Exam Questions AZ-204

Developing Solutions for Microsoft Azure

https://www.2passeasy.com/dumps/AZ-204/

NEW QUESTION 1
- (Exam Topic 8)
You are developing a data storage solution for a social networking app.
The solution requires a mobile app that stores user information using Azure Table Storage. You need to develop code that can insert multiple sets of user information.
How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

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

[ TableOperation | TableBatchOperaton | TableEntity | TableQuery ] op = new [ TableOperation | TableBatchOperaton | TableEntity | TableQuery ]();
. . .
table.[ ExecuteBatch | Execute | Insert | InsertOrMerge ](op);
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1, Box 2: TableBatchOperation Create the batch operation.
TableBatchOperation op = new TableBatchOperation(); Box 3: ExecuteBatch
/ Execute the batch operation. table.ExecuteBatch(op);
Note: You can insert a batch of entities into a table in one write operation. Some other notes on batch operations:
You can perform updates, deletes, and inserts in the same single batch operation. A single batch operation can include up to 100 entities.
All entities in a single batch operation must have the same partition key.
While it is possible to perform a query as a batch operation, it must be the only operation in the batch. References:
https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet

NEW QUESTION 2
- (Exam Topic 8)
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.

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 3
- (Exam Topic 8)
You are developing an Azure Web App. You configure TLS mutual authentication for the web app.
You need to validate the client certificate in the web app. 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:
Accessing the client certificate from App Service.
If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be available through the HttpRequest.ClientCertificate property. For other application stacks, the client cert will be available in your app through a base64 encoded value in the "X-ARR-ClientCert" request header. Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.
References:
https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth

NEW QUESTION 4
- (Exam Topic 8)
You are a developer for a software as a service (SaaS) company that uses an Azure Function to process orders. The Azure Function currently runs on an Azure Function app that is triggered by an Azure Storage queue.
You are preparing to migrate the Azure Function to Kubernetes using Kubernetes-based Event Driven Autoscaling (KEDA).
You need to configure Kubernetes Custom Resource Definitions (CRD) for the Azure Function.
Which CRDs should you configure? To answer, drag the appropriate CRD types to the correct locations. Each CRD type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

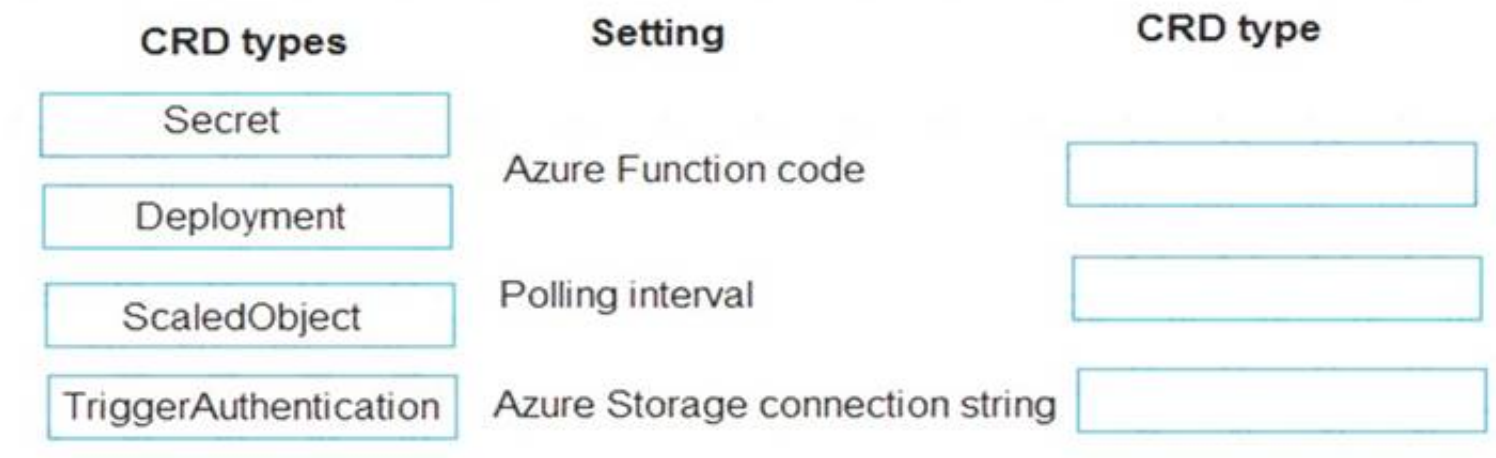

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Deployment
To deploy Azure Functions to Kubernetes use the func kubernetes deploy command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.
Box 2: ScaledObject
With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.
Example of ScaledObject with polling interval apiVersion: keda.k8s.io/v1alpha1
kind: ScaledObject metadata:
name: transformer-fn namespace: tt
labels:
deploymentName: transformer-fn spec:
scaleTargetRef: deploymentName: transformer-fn pollingInterval: 5
minReplicaCount: 0
maxReplicaCount: 100
Box 3: Secret
Store connection strings in Kubernetes Secrets. Example: to create the Secret in our demo Namespace:
create the k8s demo namespace kubectl create namespace tt
grab connection string from Azure Service Bus KEDA_SCALER_CONNECTION_STRING=$(az servicebus queue authorization-rule keys list \
-g $RG_NAME \
--namespace-name $SBN_NAME \
--queue-name inbound \
-n keda-scaler \
--query "primaryConnectionString" \
-o tsv)
create the kubernetes secret
kubectl create secret generic tt-keda-auth \
--from-literal KedaScaler=$KEDA_SCALER_CONNECTION_STRING \
--namespace tt Reference:
https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/

NEW QUESTION 5
- (Exam Topic 8)
You develop a gateway solution for a public facing news API.
The news API back end is implemented as a RESTful sen/ice and hosted in an Azure App Service instance.
You need to configure back-end authentication for the API Management service instance.
Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value 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.

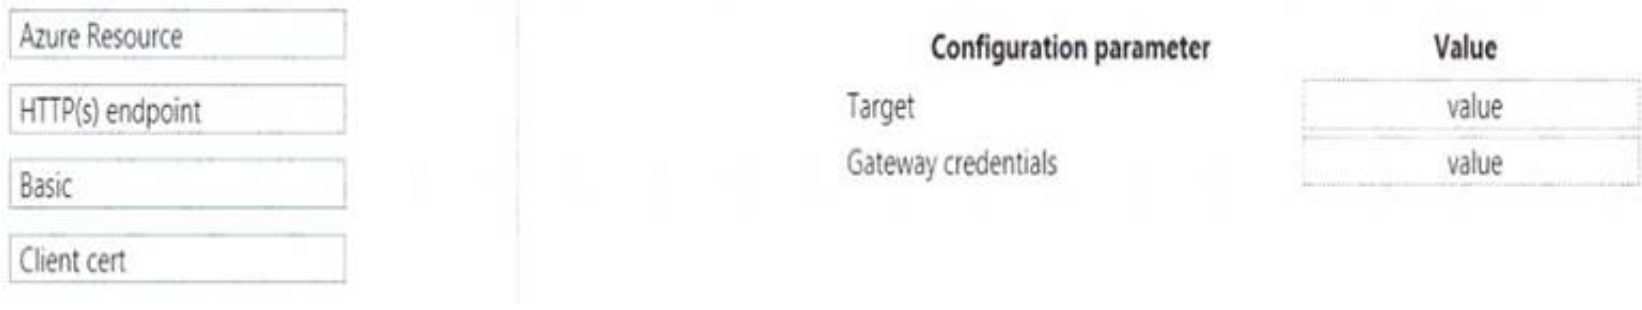

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Azure Resource Box 2: Client cert
API Management allows to secure access to the back-end service of an API using client certificates. References:
https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-ba

NEW QUESTION 6
- (Exam Topic 8)
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: Trigger the photo processing from Blob storage events. Does the solution meet the goal?

A. Yes
B. NO

Answer: B

Explanation:
You need to catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload
Note: 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 7
- (Exam Topic 8)
You are developing an Azure solution.
You need to develop code to access a secret stored in Azure Key Vault.
How should you complete the code segment? To answer, drag the appropriate code segments to the correct locations. Each code 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:
Graphical user interface, text, application Description automatically generated with medium confidence
Box 1: SecretClient
Box 2: DefaultAzureCredential
In below example, the name of your key vault is expanded to the key vault URI, in the format
"https://<your-key-vault-name>.vault.azure.net". This example is using 'DefaultAzureCredential()' class from Azure Identity Library, which allows to use the same code across different environments with different options to provide identity.
string keyVaultName = Environment.GetEnvironmentVariable("KEY_VAULT_NAME");
var kvUri = "https://" + keyVaultName + ".vault.azure.net";
var client = new SecretClient(new Uri(kvUri), new DefaultAzureCredential()); Reference:
https://docs.microsoft.com/en-us/azure/key-vault/secrets/quick-create-net

NEW QUESTION 8
- (Exam Topic 7)
You need to implement a solution to resolve the retail store location data issue.
Which three Azure Blob features should you enable? Each correct answer presents pan ol the solution. NOTE Each correct selection is worth one point

A. Immutability
B. Snapshots
C. Versioning
D. Soft delete
E. Object replication
F. Change feed

Answer: CDF

Explanation:
Scenario: You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
Before you enable and configure point-in-time restore, enable its prerequisites for the storage account: soft delete, change feed, and blob versioning.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/point-in-time-restore-manage

NEW QUESTION 9
- (Exam Topic 7)
You need to implement the corporate website. How should you configure the solution?

A. Mastered
B. Not Mastered

Answer: A

Explanation:

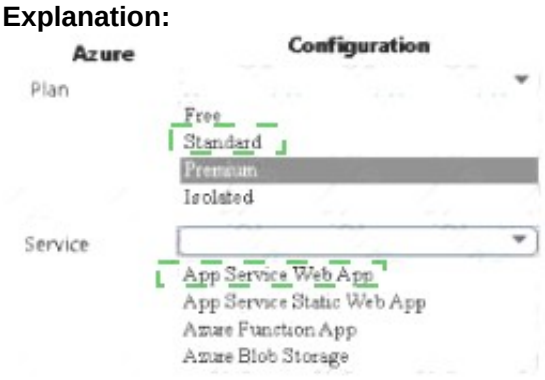

NEW QUESTION 10
- (Exam Topic 6)
You need to access data from the user claim object in the e-commerce web app. What should you do first?

A. Write custom code to make a Microsoft Graph API call from the e-commerce web app.
B. Assign the Contributor RBAC role to the e-commerce web app by using the Resource Manager create role assignment API.
C. Update the e-commerce web app to read the HTTP request header values.
D. Using the Azure CLI, enable Cross-origin resource sharing (CORS) from the e-commerce checkout API to the e-commerce web app.

Answer: C

Explanation:
Methods to Get User Identity and Claims in a .NET Azure Functions App include: ClaimsPrincipal from the Request Context
The ClaimsPrincipal object is also available as part of the request context and can be extracted from the HttpRequest.HttpContext.
User Claims from the Request Headers.
App Service passes user claims to the app by using special request headers. Reference:
https://levelup.gitconnected.com/four-alternative-methods-to-get-user-identity-and-claims-in-a-net-azurefunctio

NEW QUESTION 10
- (Exam Topic 5)
You need to ensure the security policies are met.
What code do you add at line CS07 of ConfigureSSE.ps1?

A. –PermissionsToKeys create, encrypt, decrypt
B. –PermissionsToCertificates create, encrypt, decrypt
C. –PermissionsToCertificates wrapkey, unwrapkey, get
D. –PermissionsToKeys wrapkey, unwrapkey, get

Answer: B

Explanation:
Scenario: All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
The Set-AzureRmKeyValutAccessPolicy parameter -PermissionsToKeys specifies an array of key operation permissions to grant to a user or service principal. The acceptable values for this parameter: decrypt, encrypt, unwrapKey, wrapKey, verify, sign, get, list, update, create, import, delete, backup, restore, recover, purge
Reference:
https://docs.microsoft.com/en-us/powershell/module/azurerm.keyvault/set-azurermkeyvaultaccesspolicy

NEW QUESTION 14
- (Exam Topic 4)
You need to add code at line EG15 in EventGridController.cs to ensure that the Log policy applies to all services.
How should you complete the code? To answer, drag the appropriate code segments to the correct locations. Each code 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.

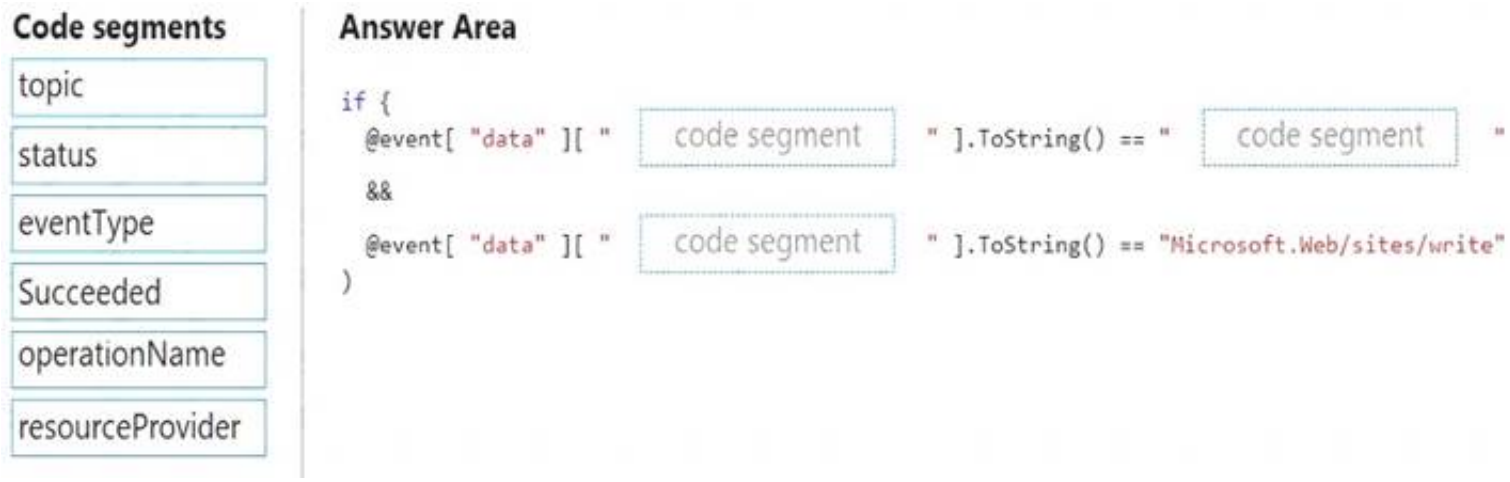

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Scenario, Log policy: All Azure App Service Web Apps must write logs to Azure Blob storage. Box 1: Status
Box 2: Succeeded
Box 3: operationName
Microsoft.Web/sites/write is resource provider operation. It creates a new Web App or updates an existing one.
Reference:
https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations

NEW QUESTION 16
- (Exam Topic 3)
You need to configure the Account Kind, Replication, and Storage tier options for the corporate website's Azure Storage account.
How should you complete the configuration? To answer, select the appropriate options in the dialog box in the answer area.
NOTE: Each correct selection is worth one point.

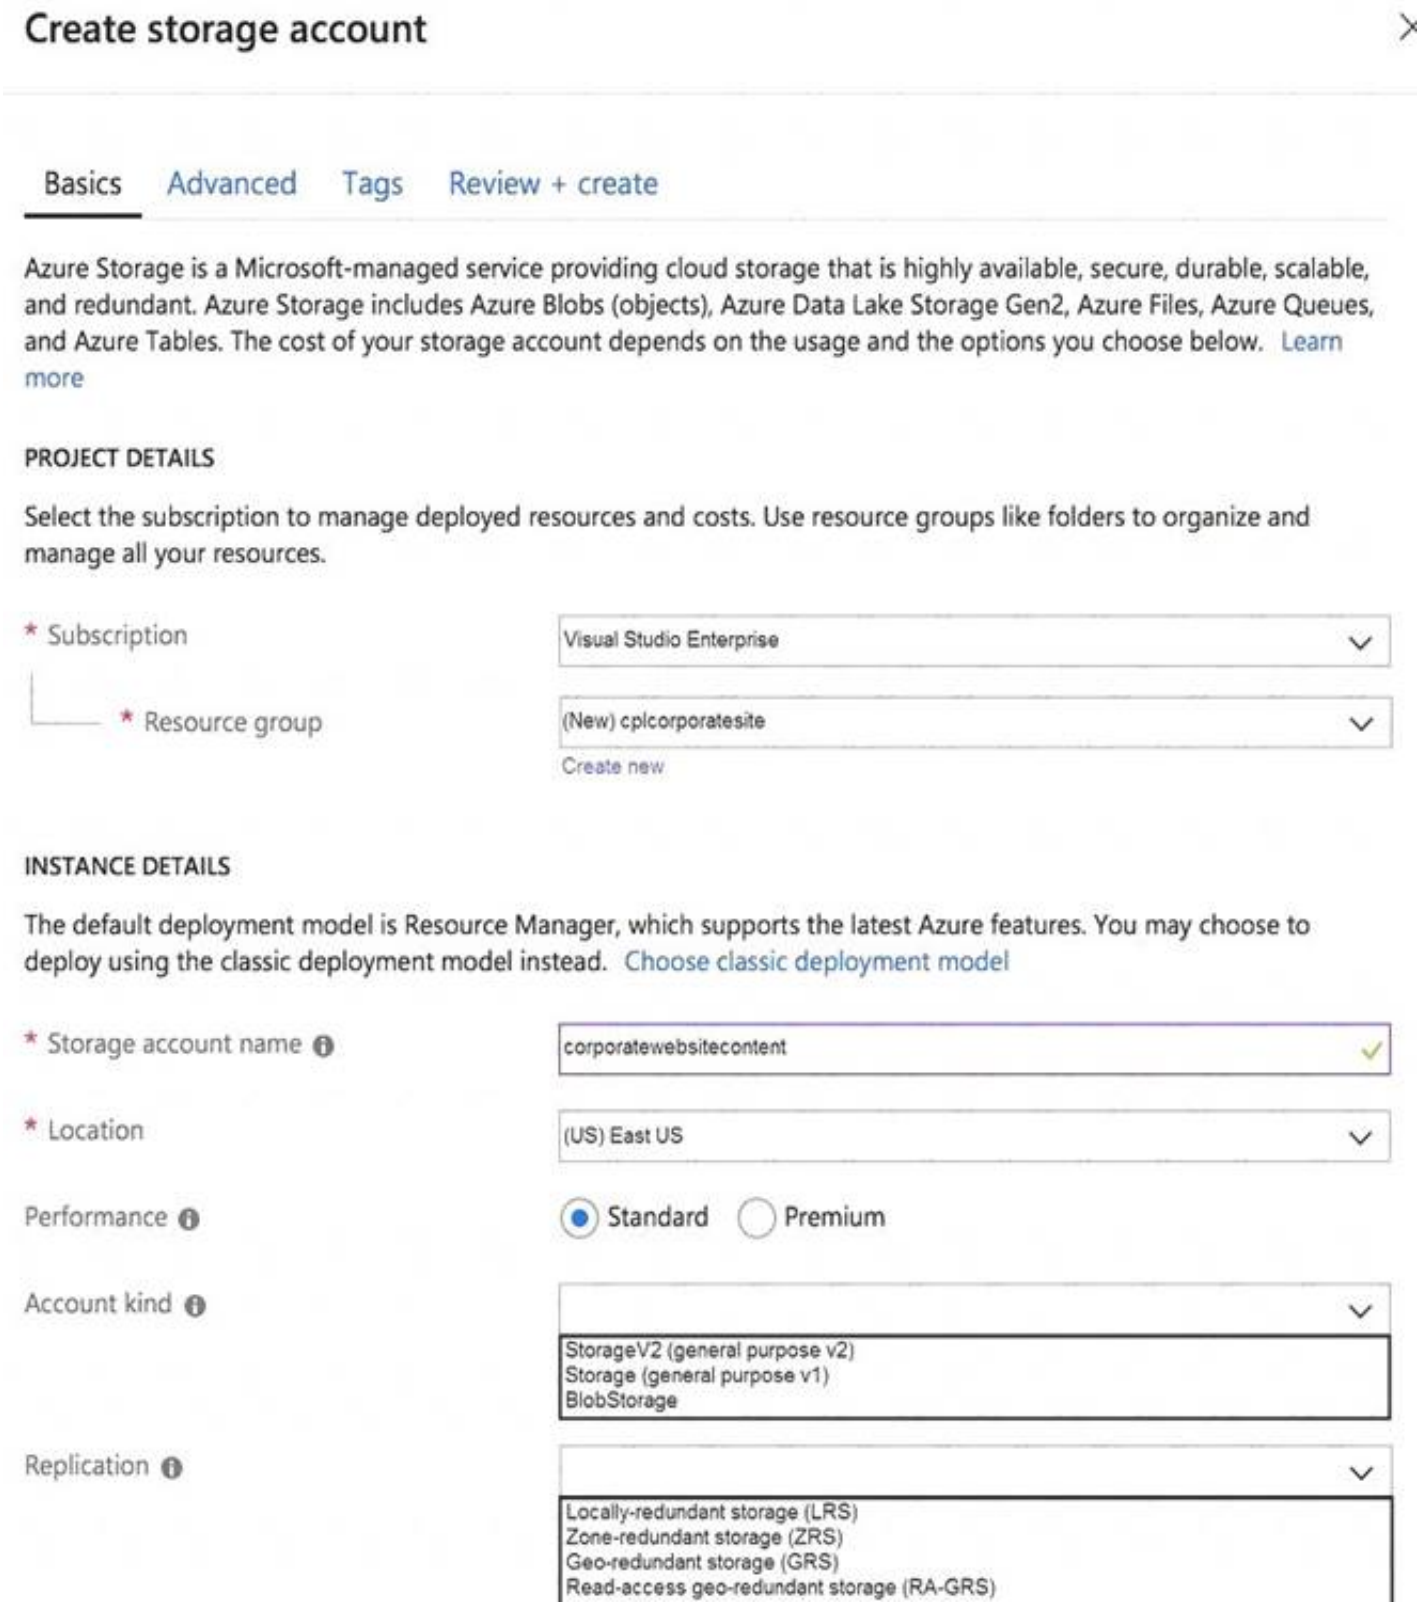

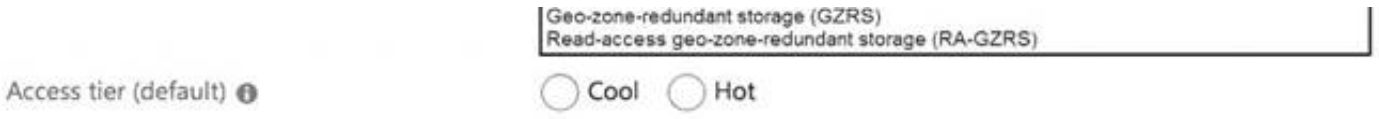

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Account Kind: StorageV2 (general-purpose v2)
Scenario: Azure Storage blob will be used (refer to the exhibit). Data storage costs must be minimized. General-purpose v2 accounts: Basic storage account type for blobs, files, queues, and tables. Recommended
for most scenarios using Azure Storage.
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-account-overview https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal

NEW QUESTION 21
- (Exam Topic 3)
You need to ensure that all messages from Azure Event Grid are processed. What should you use?

A. Azure Event Grid topic
B. Azure Service Bus topic
C. Azure Service Bus queue
D. Azure Storage queue
E. Azure Logic App custom connector

Answer: B

Explanation:
As a solution architect/developer, you should consider using Service Bus queues when:
- Your solution needs to receive messages without having to poll the queue. With Service Bus, you can achieve it by using a long-polling receive operation using the TCP-based protocols that Service Bus
supports. Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compa

NEW QUESTION 25
- (Exam Topic 3)
You need to authenticate the user to the corporate website as indicated by the architectural diagram. Which two values should you use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. ID token signature
B. ID token claims
C. HTTP response code
D. Azure AD endpoint URI
E. Azure AD tenant ID

Answer: BE

Explanation:
Claims in access tokens
JWTs (JSON Web Tokens) are split into three pieces:
- Header - Provides information about how to validate the token including information about the type of token and how it was signed.
- Payload - Contains all of the important data about the user or app that is attempting to call your service.
- Signature - Is the raw material used to validate the token.

Your client can get an access token from either the v1.0 endpoint or the v2.0 endpoint using a variety of protocols.
Scenario: User authentication (see step 5 below)
The following steps detail the user authentication process:
- The user selects Sign in in the website.
- The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
- The user signs in.
- Azure AD redirects the user's session back to the web application. The URL includes an access token.
- The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
- The back-end API validates the access token. Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies

NEW QUESTION 30
- (Exam Topic 2)
You need to deploy the CheckUserContent Azure function. The solution must meet the security and cost requirements.
Which hosting model should you use?

A. Consumption plan

B. Premium plan
C. App Service plan

Answer: A

NEW QUESTION 32
- (Exam Topic 2)
You need to ensure that validation testing is triggered per the requirements.
How should you complete the code segment? To answer, select the appropriate values in the answer area.
NOTE: Each correct selection is worth one point.

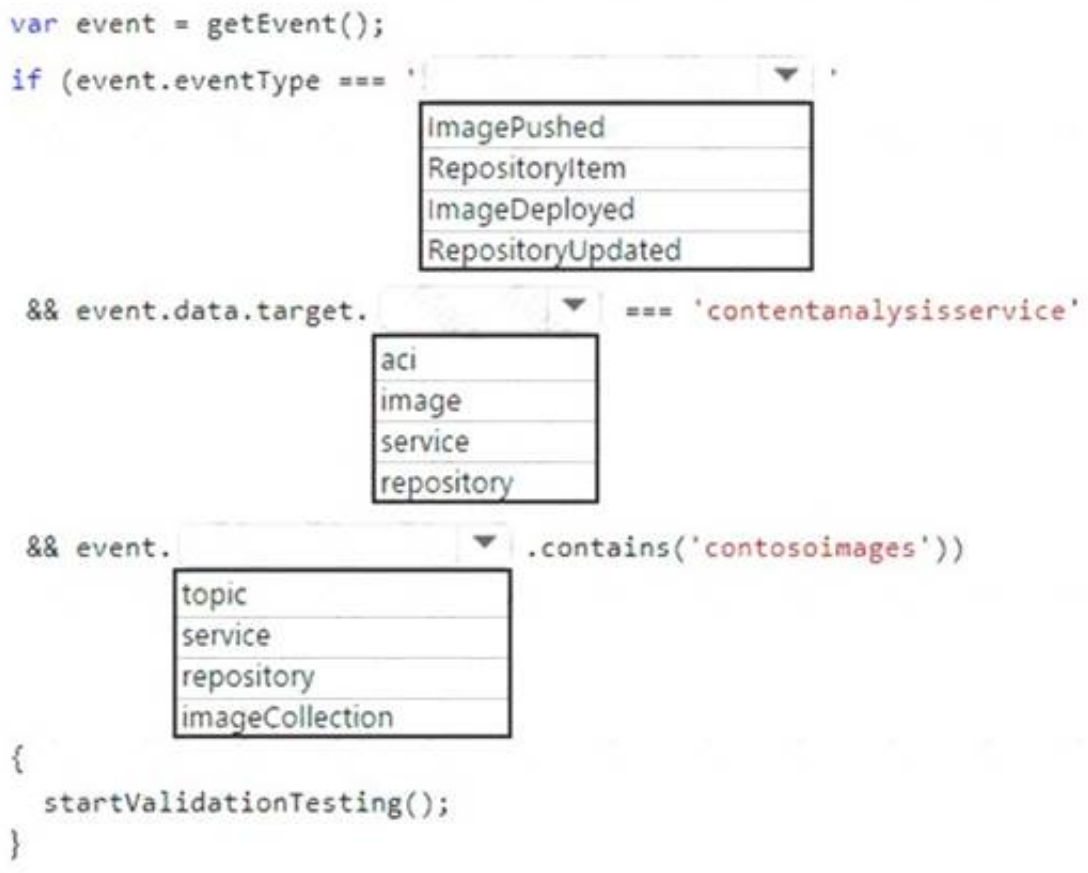

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: RepositoryUpdated
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.
Box 2: service
Box 3: imageCollection Reference:
https://docs.microsoft.com/en-us/azure/devops/notifications/oob-supported-event-types

NEW QUESTION 34
- (Exam Topic 2)
You need to add markup at line AM04 to implement the ContentReview role.
How should you complete the markup? To answer, drag the appropriate json segments to the correct locations. Each json segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

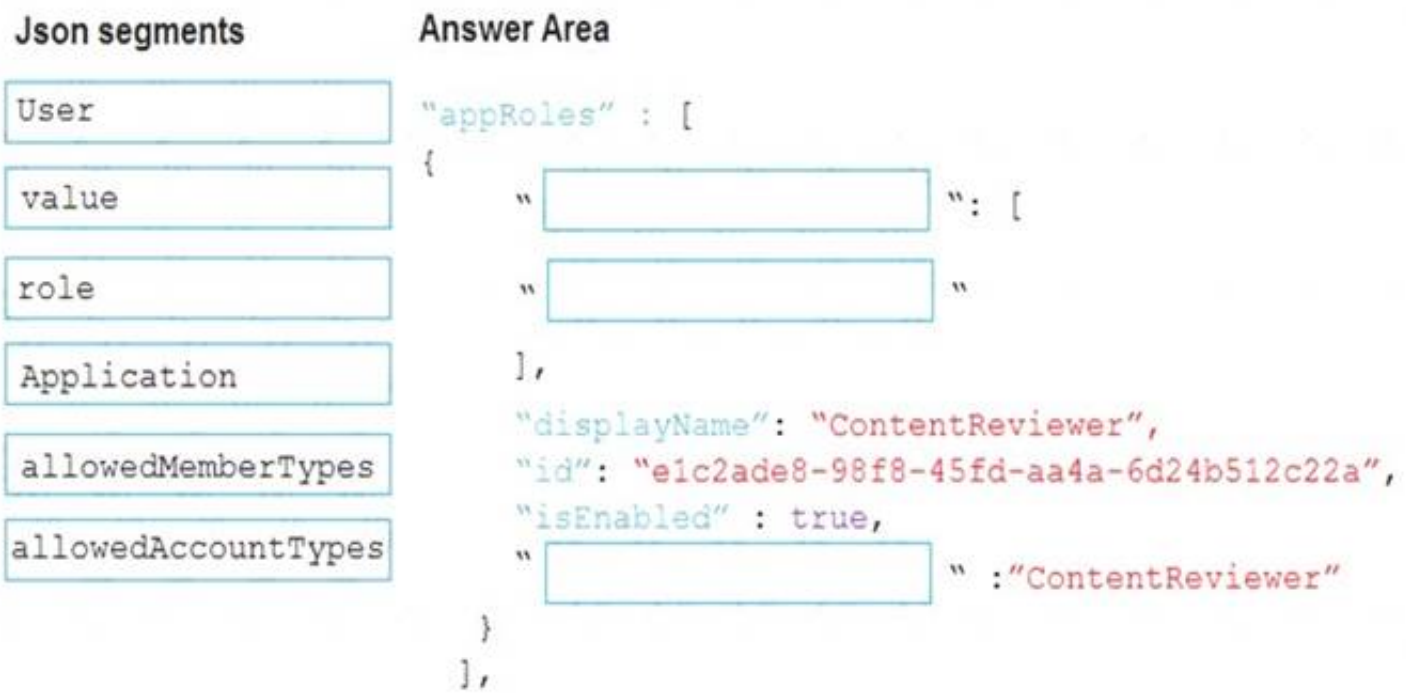

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Box 1: allowedMemberTypes
allowedMemberTypes specifies whether this app role definition can be assigned to users and groups by setting to "User", or to other applications (that are accessing this application in daemon service scenarios) by setting to "Application", or to both.
Note: The following example shows the appRoles that you can assign to users. "appId": "8763f1c4-f988-489c-a51e-158e9ef97d6a",
"appRoles": [
{
"allowedMemberTypes": ["User"
],
"displayName": "Writer",
"id": "d1c2ade8-98f8-45fd-aa4a-6d06b947c66f", "isEnabled": true,
"description": "Writers Have the ability to create tasks.", "value": "Writer"
}
],
"availableToOtherTenants": false, Box 2: User
Scenario: In order to review content a user must be part of a ContentReviewer role.
Box 3: value
value specifies the value which will be included in the roles claim in authentication and access tokens. Reference:
https://docs.microsoft.com/en-us/graph/api/resources/approle

NEW QUESTION 35
- (Exam Topic 2)
You need to monitor ContentUploadService accourding to the requirements. Which command should you use?

A. az monitor metrics alert create –n alert –g … - -scopes … - -condition "avg Percentage CPU > 8"
B. az monitor metrics alert create –n alert –g … - -scopes … - -condition "avg Percentage CPU > 800"
C. az monitor metrics alert create –n alert –g … - -scopes … - -condition "CPU Usage > 800"
D. az monitor metrics alert create –n alert –g … - -scopes … - -condition "CPU Usage > 8"

Answer: B

Explanation:
Scenario: An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU-cores
Reference:
https://docs.microsoft.com/sv-se/cli/azure/monitor/metrics/alert

NEW QUESTION 40
- (Exam Topic 1)
You need to correct the VM issues.
Which tools should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

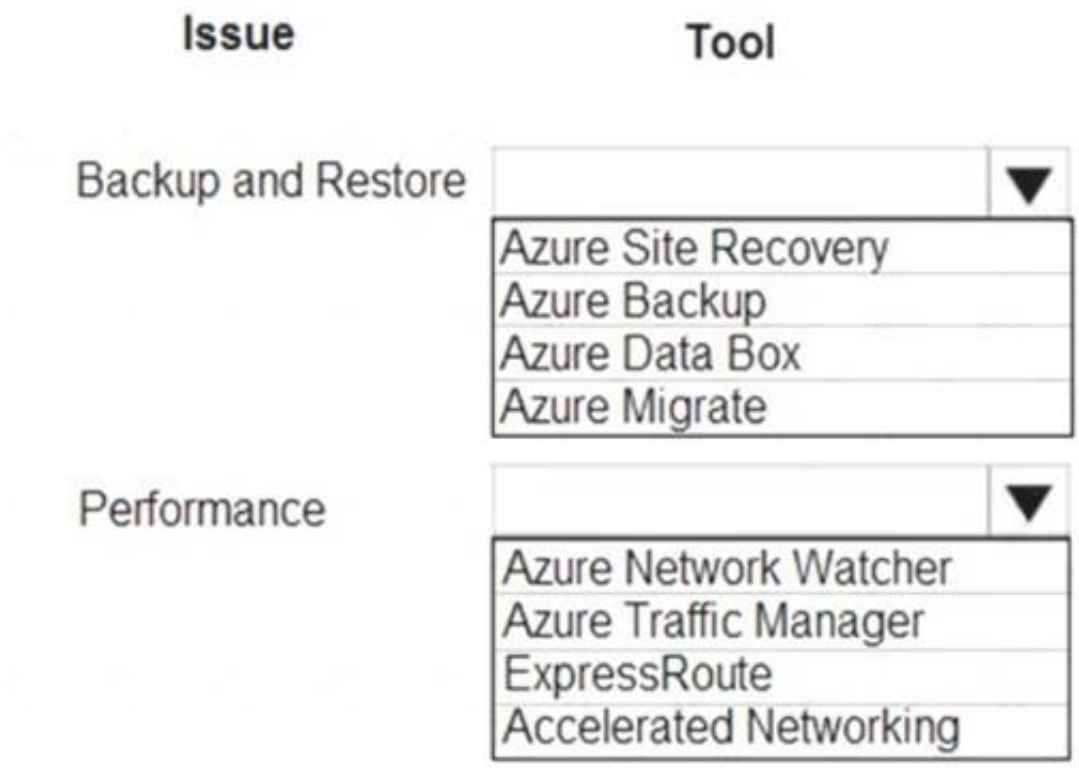

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Backup and Restore: Azure Backup
Scenario: The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.
In-Place restore of disks in IaaS VMs is a feature of Azure Backup. Performance: Accelerated Networking
Scenario: The VM shows high network latency, jitter, and high CPU utilization.

Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, greatly improving its networking performance. This high-performance path bypasses the host from the datapath, reducing latency, jitter, and CPU utilization, for use with the most demanding network workloads on supported VM types.
References:
https://azure.microsoft.com/en-us/blog/an-easy-way-to-bring-back-your-azure-vm-with-in-place-restore/

NEW QUESTION 43
- (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.

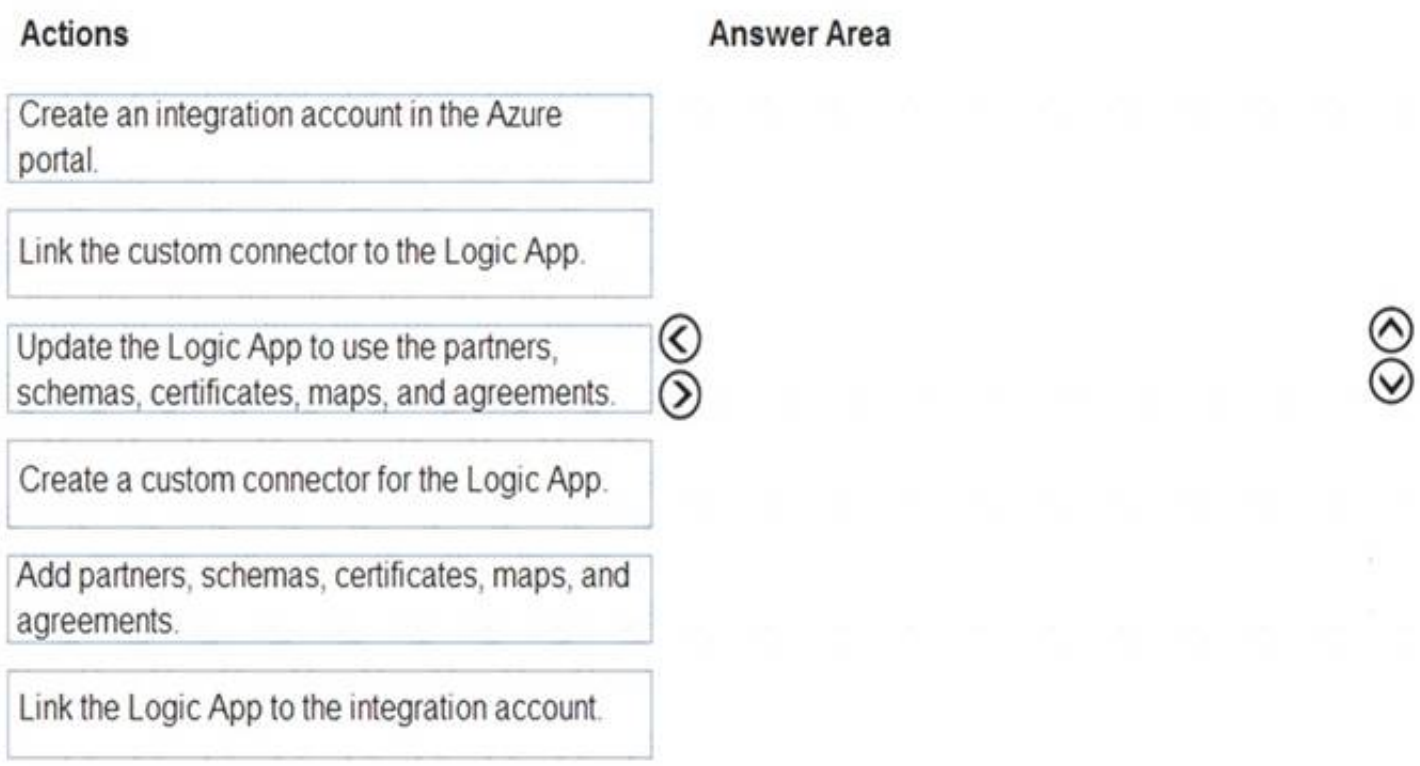

A. Mastered
B. Not Mastered

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 48
- (Exam Topic 8)
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 53
- (Exam Topic 8)
You are configuring a new development environment for a Java application.
The environment requires a Virtual Machine Scale Set (VMSS), several storage accounts, and networking components.
The VMSS must not be created until the storage accounts have been successfully created and an associated load balancer and virtual network is configured.
How should you complete the Azure Resource Manager template? 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:
Box 1: copyIndex
Notice that the name of each resource includes the copyIndex() function, which returns the current iteration in the loop. copyIndex() is zero-based.
Box 2: copy
By adding the copy element to the resources section of your template, you can dynamically set the number of resources to deploy.
Box 3: dependsOn Example:
"type": "Microsoft.Compute/virtualMachineScaleSets", "apiVersion": "2020-06-01",
"name": "[variables('namingInfix')]",
"location": "[parameters('location')]", "sku": {
"name": "[parameters('vmSku')]", "tier": "Standard",
"capacity": "[parameters('instanceCount')]"
},
"dependsOn": [
"[resourceId('Microsoft.Network/loadBalancers', variables('loadBalancerName'))]", "[resourceId('Microsoft.Network/virtualNetworks', variables('virtualNetworkName'))]"
],
Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/copy-resources https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/quick-create-template-windows

NEW QUESTION 56
- (Exam Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:
•Share session state across all ASP.NET web applications
•Support controlled, concurrent access to the same session state data for multiple readers and a single writer

•Save full HTTP responses for concurrent requests You need to store the information.
Proposed Solution: Deploy and configure Azure Cache for Redis. Update the web applications. Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
The session state provider for Azure Cache for Redis enables you to share session information between different instances of an ASP.NET web application.
The same connection can be used by multiple concurrent threads. Redis supports both read and write operations.
The output cache provider for Azure Cache for Redis enables you to save the HTTP responses generated by an ASP.NET web application.
Note: Using the Azure portal, you can also configure the eviction policy of the cache, and control access to the cache by adding users to the roles provided. These roles, which define the operations that members can perform, include Owner, Contributor, and Reader. For example, members of the Owner role have complete control over the cache (including security) and its contents, members of the Contributor role can read and write information in the cache, and members of the Reader role can only retrieve data from the cache.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching

NEW QUESTION 57
- (Exam Topic 8)
You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway.
Data from traffic sensors are stored in Azure Event Hub.
Traffic data is consumed by four departments. Each department has an Azure Web App that displays the time-series-based reports and contains a WebJob that processes the incoming data from Event Hub. All Web Apps run on App Service Plans with three instances.
Data throughout must be maximized. Latency must be minimized. You need to implement the Azure Event Hub.
Which settings should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

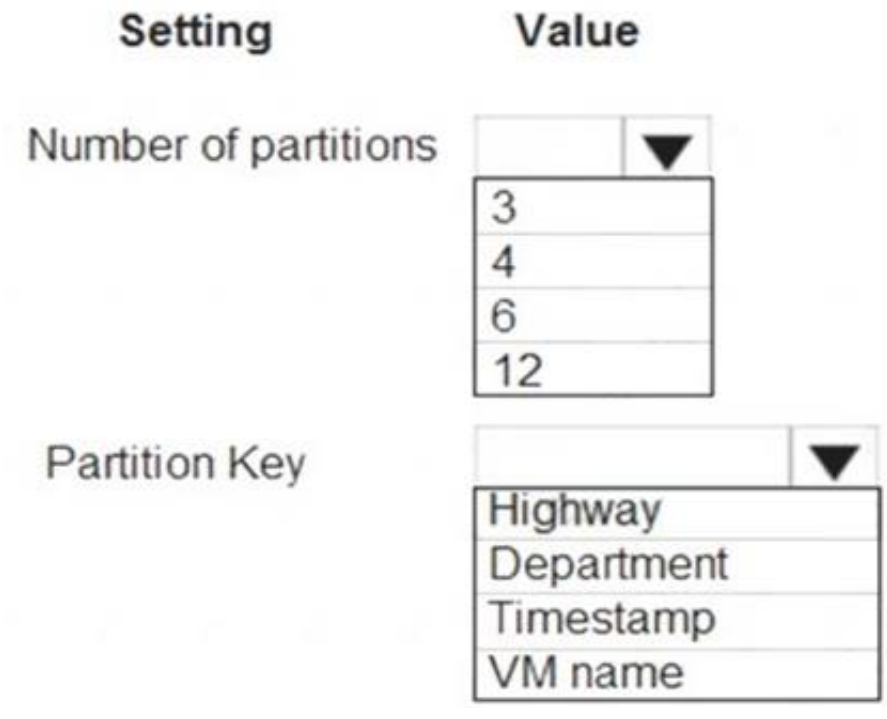

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 62
- (Exam Topic 8)
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 tor 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 SearchServiceClient 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. Set the DataSource property of the SearchServiceClient 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 63
- (Exam Topic 8)
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: Create an Azure Function app that uses the Consumption hosting model and that is triggered from the blob upload.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
In the Consumption hosting plan, resources are added dynamically as required by your functions. Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function

NEW QUESTION 66
- (Exam Topic 8)
You are developing a web service that will run on Azure virtual machines that use Azure Storage. You configure all virtual machines to use managed identities.
You have the following requirements:
> Must use only Azure Instance Metadata Service endpoints.
You need to write code to retrieve an access token to access Azure Storage. To answer, drag the appropriate code segments to the correct locations. Each code segment may be used 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.

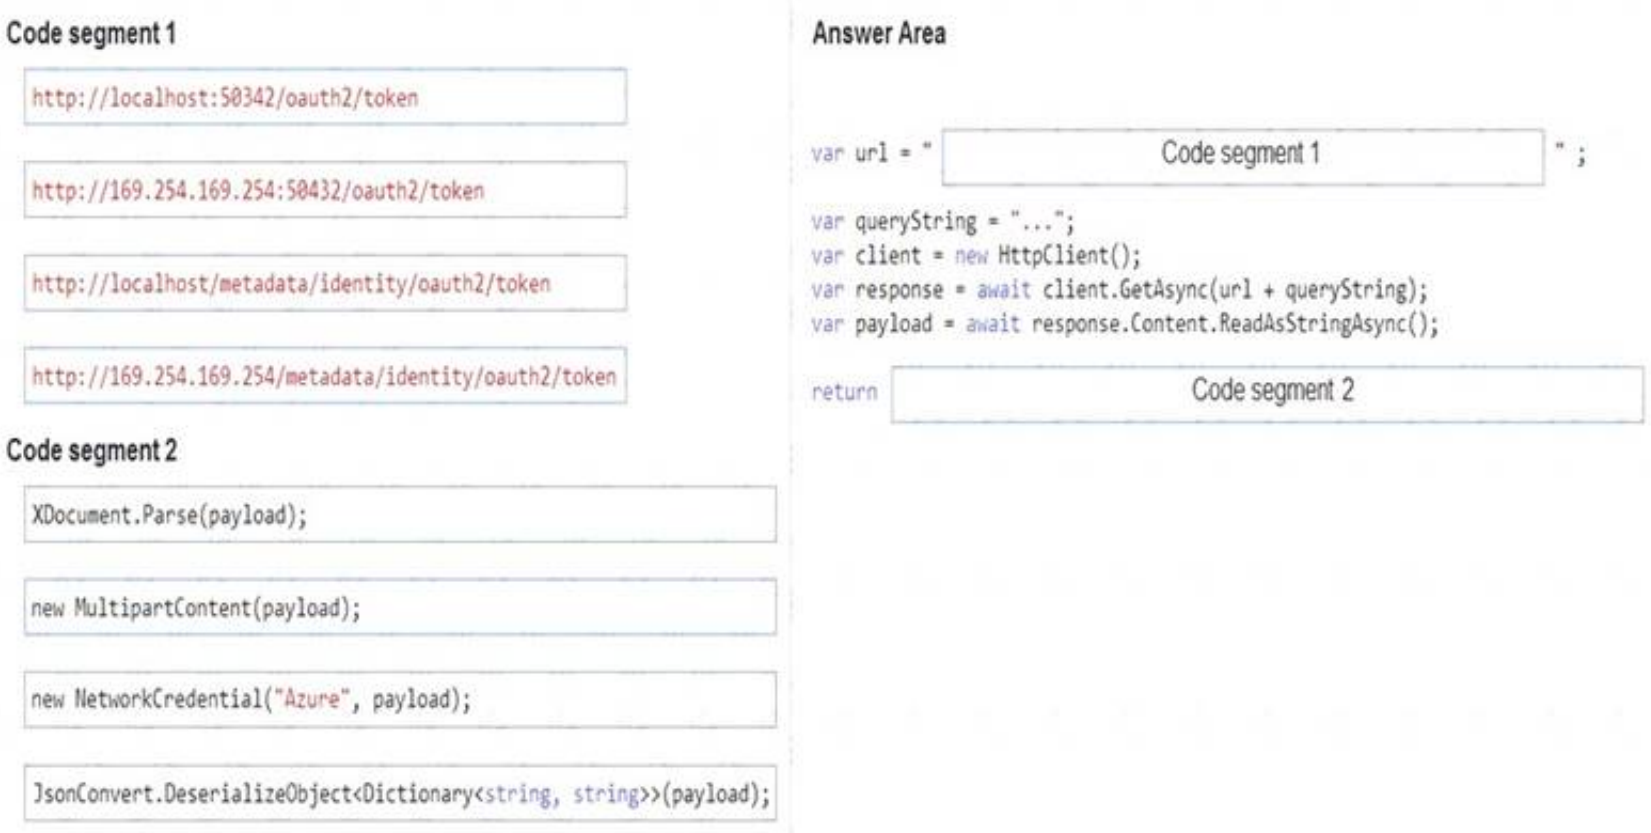

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Azure Instance Metadata Service endpoints "/oauth2/token"
Box 1: http://169.254.169.254/metadata/identity/oauth2/token
Sample request using the Azure Instance Metadata Service (IMDS) endpoint (recommended): GET 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01
&resource=https://management.azure.com/' HTTP/1.1 Metadata: true
Box 2: JsonConvert.DeserializeObject<Dictionary<string,string>>(payload); Deserialized token response; returning access code.
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-toke https://docs.microsoft.com/en-us/azure/service-fabric/how-to-managed-identity-service-fabric-app-code

NEW QUESTION 68
- (Exam Topic 8)
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 Example:
public async Task Enqueue(string payload)
{
// StartOperation is a helper method that initializes the telemetry item
// and allows correlation of this operation with its parent and children.
var operation = telemetryClient.StartOperation<DependencyTelemetry>("enqueue " + queueName); operation.Telemetry.Type = "Azure Service Bus";
operation.Telemetry.Data = "Enqueue " + queueName; var message = new BrokeredMessage(payload);
// Service Bus queue allows the property bag to pass along with the message.
// We will use them to pass our correlation identifiers (and other context)
// to the consumer.
message.Properties.Add("ParentId", operation.Telemetry.Id); message.Properties.Add("RootId", operation.Telemetry.Context.Operation.Id); Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking

NEW QUESTION 69
- (Exam Topic 8)
You are developing an application that uses Azure Storage Queues. You have the following code:

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse
(CloudConfigurationManager.GetSetting("StorageConnectionString"));
CloudQueueClient queueClient = storageAccount.CreateCloudQueueClient()

CloudQueue queue = queueClient.GetQueueReference("appqueue") ;
await queu.CreateIfNotExistsAsync() ;

CloudQueueMessage peekedMessage = await queue.PeekMessageAsync() ;
if (peekedMessage != null)
{
    Console.WriteLine("The peeked message is: {0}", peekedMessage.AsString);
}
CloudQueueMessage message = await queue.GetMessageAsync() ;
```

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

Statement	Yes	No
The code configures the lock duration for the queue.	○	○
The last message read remains in the queue after the code runs.	○	○
The storage queue remains in the storage account after the code runs.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
The QueueDescription.LockDuration property gets or sets the duration of a peek lock; that is, the amount of time that the message is locked for other receivers. The maximum value for LockDuration is 5 minutes; the default value is 1 minute.
Box 2: Yes
You can peek at the message in the front of a queue without removing it from the queue by calling the PeekMessage method.

Box 3: Yes Reference:
https://docs.microsoft.com/en-us/azure/storage/queues/storage-dotnet-how-to-use-queues https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicebus.messaging.queuedescription.lockduration

NEW QUESTION 74
- (Exam Topic 8)
You have an application that includes an Azure Web app and several Azure Function apps. Application secrets including connection strings and certificates are stored in Azure Key Vault.
Secrets must not be stored in the application or application runtime environment. Changes to Azure Active Directory (Azure AD) must be minimized.
You need to design the approach to loading application secrets. What should you do?

A. Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity.
B. Create a single Azure AD Service Principal with permission to access Key Vault and use a client secret from within the App Services to access Key Vault.
C. Create a system assigned Managed Identity in each App Service with permission to access Key Vault.
D. Create an Azure AD Service Principal with Permissions to access Key Vault for each App Service and use a certificate from within the App Services to access Key Vault.

Answer: A

Explanation:
Use Key Vault references for App Service and Azure Functions.
Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.
Reference:
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

NEW QUESTION 79
- (Exam Topic 8)
You are developing an application to use Azure Blob storage. You have configured Azure Blob storage to include change feeds.
A copy of your storage account must be created in another region. Data must be copied from the current storage account to the new storage account directly between the storage servers.
You need to create a copy of the storage account in another region and copy the data.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

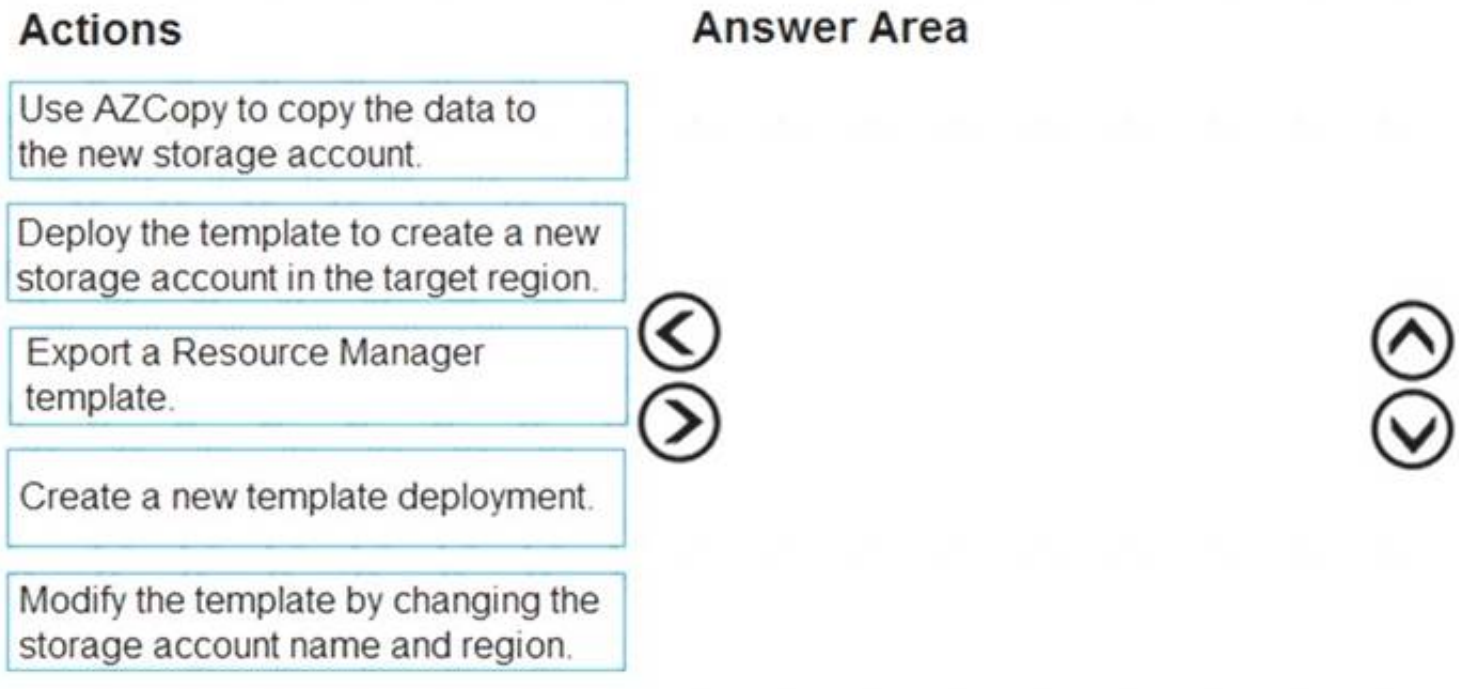

A. Mastered
B. Not Mastered

Answer: A

Explanation:
https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal#modify-the-te

NEW QUESTION 81
- (Exam Topic 8)
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 Event Grid. Configure event filtering to evaluate the device identifier. 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.
Note: An event is a lightweight notification of a condition or a state change. Event hubs is usually used reacting to status changes.
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

NEW QUESTION 83
- (Exam Topic 8)
You manage a data processing application that receives requests from an Azure Storage queue. You need to manage access to the queue. You have the following requirements:

- Provide other applications access to the Azure queue.
- Ensure that you can revoke access to the queue without having to regenerate the storage account keys.
- Specify access at the queue level and not at the storage account level.

Which type of shared access signature (SAS) should you use?

A. Service SAS with a stored access policy
B. Account SAS
C. User Delegation SAS
D. Service SAS with ad hoc SAS

Answer: A

Explanation:
A service SAS is secured with the storage account key. A service SAS delegates access to a resource in only one of the Azure Storage services: Blob storage, Queue storage, Table storage, or Azure Files.
Stored access policies give you the option to revoke permissions for a service SAS without having to regenerate the storage account keys. Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

NEW QUESTION 88
- (Exam Topic 8)
You are preparing to deploy an Azure virtual machine (VM) based application. The VMs that run the application have the following requirements:
• When a VM is provisioned the firewall must be automatically configured before it can access Azure resources.
• Supporting services must be installed by using an Azure PowerShell script that is stored in Azure Storage You need to ensure that the requirements are met.
Which features should you use? To answer, drag the appropriate features to the correct requirements.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-command

NEW QUESTION 91
- (Exam Topic 8)
You plan to deploy a new application to a Linux virtual machine (VM) that is hosted in Azure.
The entire VM must be secured at rest by using industry-standard encryption technology to address organizational security and compliance requirements.
You need to configure Azure Disk Encryption for the VM.
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.

Answer Area

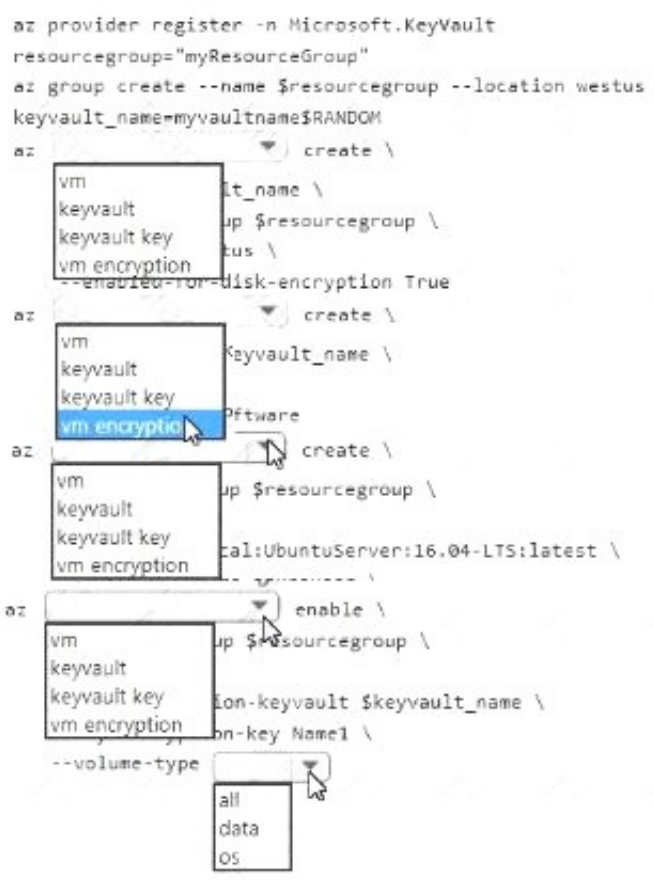

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: keyvault
Create an Azure Key Vault with az keyvault create and enable the Key Vault for use with disk encryption. Specify a unique Key Vault name for keyvault_name as follows:
keyvault_name=myvaultname$RANDOM az keyvault create \
--name $keyvault_name \
--resource-group $resourcegroup \
--location eastus \
--enabled-for-disk-encryption True Box 2: keyvault key
The Azure platform needs to be granted access to request the cryptographic keys when the VM boots to decrypt the virtual disks. Create a cryptographic key in your Key Vault with az keyvault key create. The following example creates a key named myKey:
az keyvault key create \
--vault-name $keyvault_name \
--name myKey \
--protection software Box 3: vm
Create a VM with az vm create. Only certain marketplace images support disk encryption. The following example creates a VM named myVM using an Ubuntu 16.04 LTS image:
az vm create \
--resource-group $resourcegroup \
--name myVM \
--image Canonical:UbuntuServer:16.04-LTS:latest \
--admin-username azureuser \
--generate-ssh-keys \ Box 4: vm encryption
Encrypt your VM with az vm encryption enable: az vm encryption enable \
--resource-group $resourcegroup \
--name myVM \
--disk-encryption-keyvault $keyvault_name \
--key-encryption-key myKey \
--volume-type all
Note: seems to an error in the question. Should have enable instead of create. Box 5: all
Encrypt both data and operating system.
References:
https://docs.microsoft.com/bs-latn-ba/azure/virtual-machines/linux/encrypt-disks

NEW QUESTION 93
- (Exam Topic 8)
You are implementing an order processing system. A point of sale application publishes orders to topics in an Azure Service Bus queue. The label property for the topic includes the following data:

Property	Description
ShipLocation	the country/region where the order will be shipped
CorrelationId	a priority value for the order
Quantity	a user-defined field that stores the quantity of items in an order
AuditedAt	a user-defined field that records the date an order is audited

The system has the following requirements for subscriptions

Subscription type	Comments
FutureOrders	This subscription is reserved for future use and must not receive any orders.
HighPriorityOrders	Handle all high priority orders and International orders.
InternationalOrders	Handle orders where the country/region is not United States.
HighQuantityOrders	Handle only orders with quantities greater than 100 units.
AllOrders	This subscription is used for auditing purposes. This subscription must receive every single order. AllOrders has an Action defined that updates the AuditedAt property to include the date and time it was received by the subscription.

You need to implement filtering and maximize throughput while evaluating filters.
Which filter types should you implement? To answer, drag the appropriate filter types to the correct subscriptions. Each filter 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.

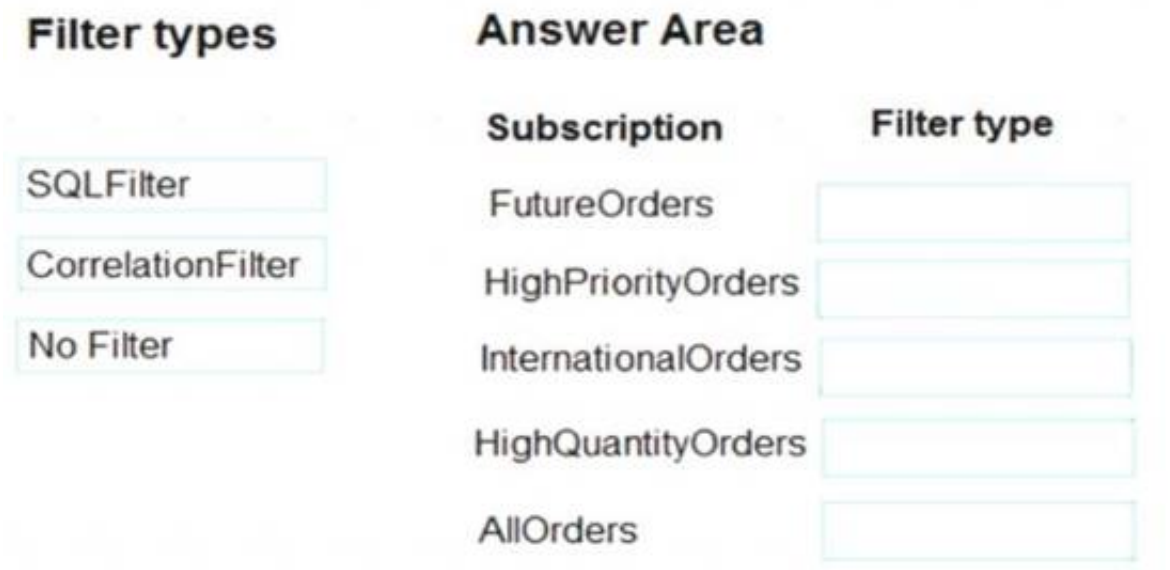

A. Mastered
B. Not Mastered

Answer: A

Explanation:
FutureOrders: SQLFilter HighPriortyOrders: CorrelationFilter CorrelationID only InternationalOrders: SQLFilter
Country NOT USA requires an SQL Filter HighQuantityOrders: SQLFilter
Need to use relational operators so an SQL Filter is needed. AllOrders: No Filter
SQL Filter: SQL Filters - A SqlFilter holds a SQL-like conditional expression that is evaluated in the broker against the arriving messages' user-defined properties and system properties. All system properties must be prefixed with sys. in the conditional expression. The SQL-language subset for filter conditions tests for the existence of properties (EXISTS), as well as for null-values (IS NULL), logical NOT/AND/OR, relational operators, simple numeric arithmetic, and simple text pattern matching with LIKE.
Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties. A common use is to match against the CorrelationId property, but the application can also choose to match against ContentType, Label, MessageId, ReplyTo, ReplyToSessionId, SessionId, To, and any user-defined properties. A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter. For string expressions, the comparison is case-sensitive. When specifying multiple match properties, the filter combines them as a logical AND condition, meaning for the filter to match, all conditions must match.
Boolean filters - The TrueFilter and FalseFilter either cause all arriving messages (true) or none of the arriving
messages (false) to be selected for the subscription. References:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters

NEW QUESTION 98
- (Exam Topic 8)
You are validating the configuration of an Azure Search indexer.
The service has been configured with an indexer that uses the Import Data option. The index is configured using options as shown in the Index Configuration exhibit. (Click the Index Configuration tab.)

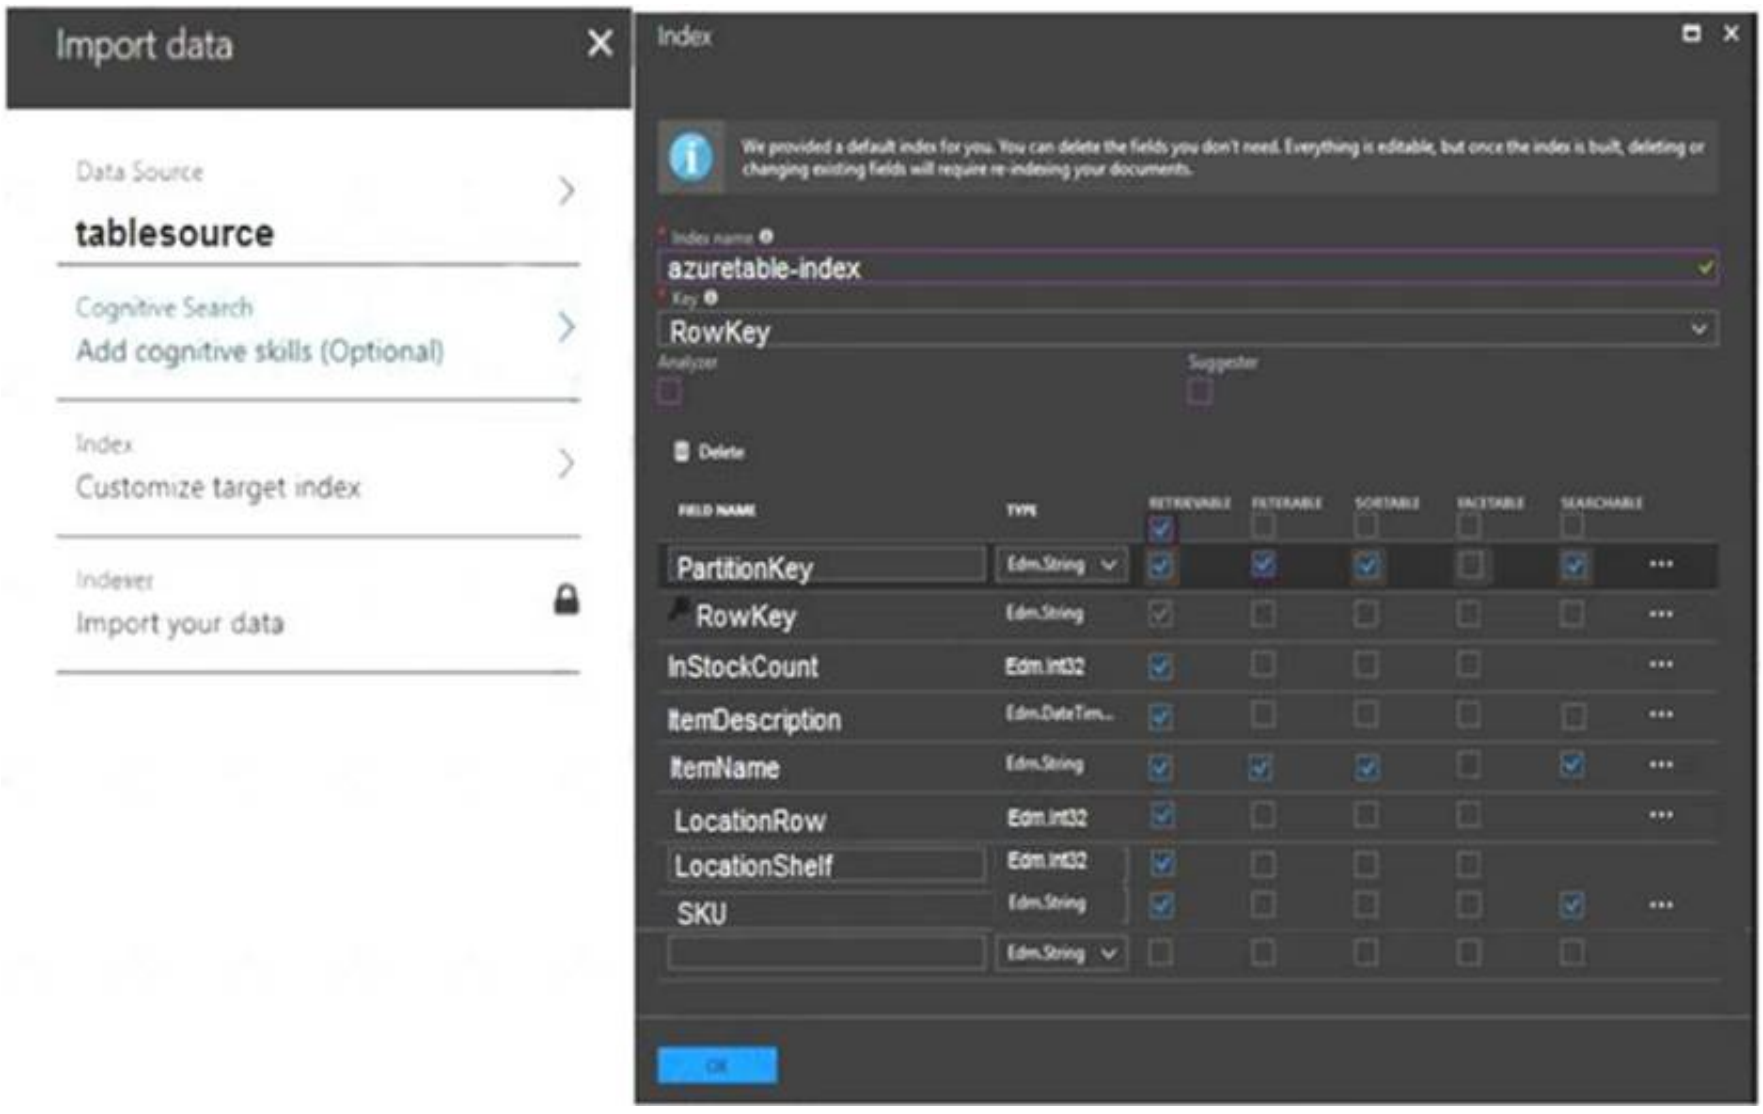

You use an Azure table as the data source for the import operation. The table contains three records with item inventory data that matches the fields in the Storage data exhibit. These records were imported when the index was created. (Click the Storage Data tab.) When users search with no filter, all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

PartitionKey	RowKey	Timestamp	InStockCount	ItemDescription	ItemName	LocationRow	LocationShelf	SKU
Food	3	2018-08-25T15:47:29.135Z	32	A box of chocolate candy bars	Choco-bar	5	3	123421
Hardware	2	2018-08-25T15:46:08.405Z	2	A bag of bolts	Bolts	1	4	678564
Hardware	1	2018-08-25T15:46:41.402Z	23	A box of nails	Nails	2	1	654365

When users search for items by description, Search explorer returns no records. The Search Explorer exhibit shows the query and results for a test. In the test, a user is trying to search for all items in the table that have a
description that contains the word bag. (Click the Search Explorer tab.) You need to resolve the issue.
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

	Yes	No
You can resolve the issue by recreating the search index with the same settings for all fields except ItemDescription. Select the SEARCHABLE option for this field	○	○
You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field.	○	○
You can resolve the issue by running the indexer.	○	○
You can resolve the issue by changing the query string in Search explorer to `bag of` to return the correct results	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes

The ItemDescription field in not searchable.
Box 2: No
The ItemDescription field in not searchable, but we would need to recreate the index.
Box 3: Yes
An indexer in Azure Search is a crawler that extracts searchable data and metadata from an external Azure data source and populates an index based on field-to-field mappings between the index and your data source. This approach is sometimes referred to as a 'pull model' because the service pulls data in without you having to write any code that adds data to an index.
Box 4: No
References:
https://docs.microsoft.com/en-us/azure/search/search-what-is-an-index https://docs.microsoft.com/en-us/azure/search/search-indexer-overview

NEW QUESTION 99
- (Exam Topic 8)
You are developing a user portal for a company.
You need to create a report for the portal that lists information about employees who are subject matter experts for a specific topic. You must ensure that administrators have full control and cosent over the data.
Which technology should you use?

A. Microsoft Graph connectors
B. Microosft graph API
C. Microsoft Graph data connect

Answer: C

NEW QUESTION 101
- (Exam Topic 8)
You are developing a serverless Java application on Azure. You create a new Azure Key Vault to work with secrets from a new Azure Functions application.
The application must meet the following requirements:

- Reference the Azure Key Vault without requiring any changes to the Java code.
- Dynamically add and remove instances of the Azure Functions host based on the number of incoming application events.
- Ensure that instances are perpetually warm to avoid any cold starts.
- Connect to a VNet.
- Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.

You need to grant the Azure Functions application access to the Azure Key Vault.
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.

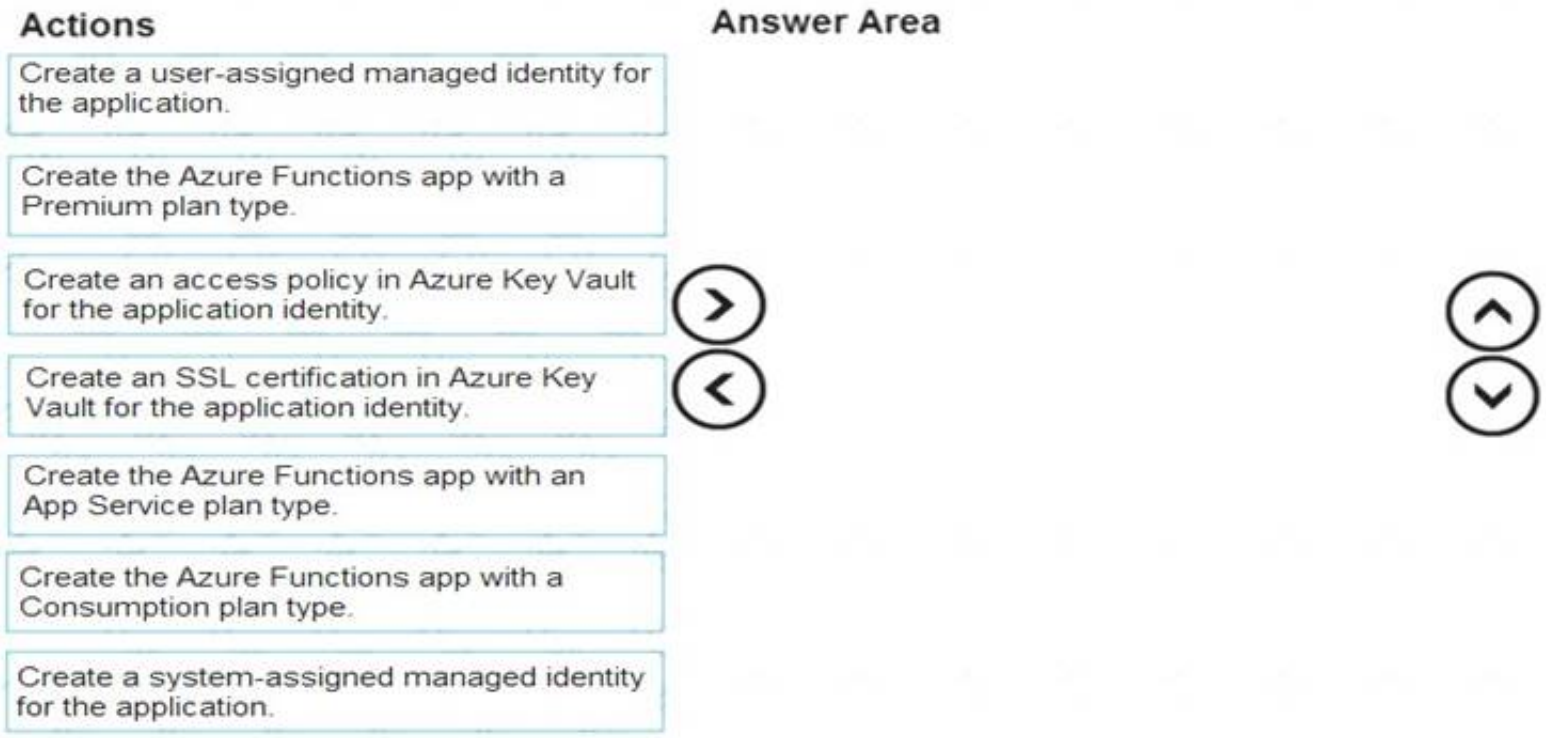

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create the Azure Functions app with a Consumption plan type. Use the Consumption plan for serverless.
Step 2: Create a system-assigned managed identity for the application. Create a system-assigned managed identity for your application.
Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.
Step 3: Create an access policy in Key Vault for the application identity.
Create an access policy in Key Vault for the application identity you created earlier. Enable the "Get" secret permission on this policy. Do not configure the "authorized application" or applicationId settings, as this is not compatible with a managed identity.
Reference:
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

NEW QUESTION 102
- (Exam Topic 8)
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

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 106
- (Exam Topic 8)
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 an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data. You need to ensure the app does not time out and processes the blob data.
Solution: Use the Durable Function async pattern to process the blob data. Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Note: Large, long-running functions can cause unexpected timeout issues. General best practices include: Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.
Reference:

https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices

NEW QUESTION 111
- (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.

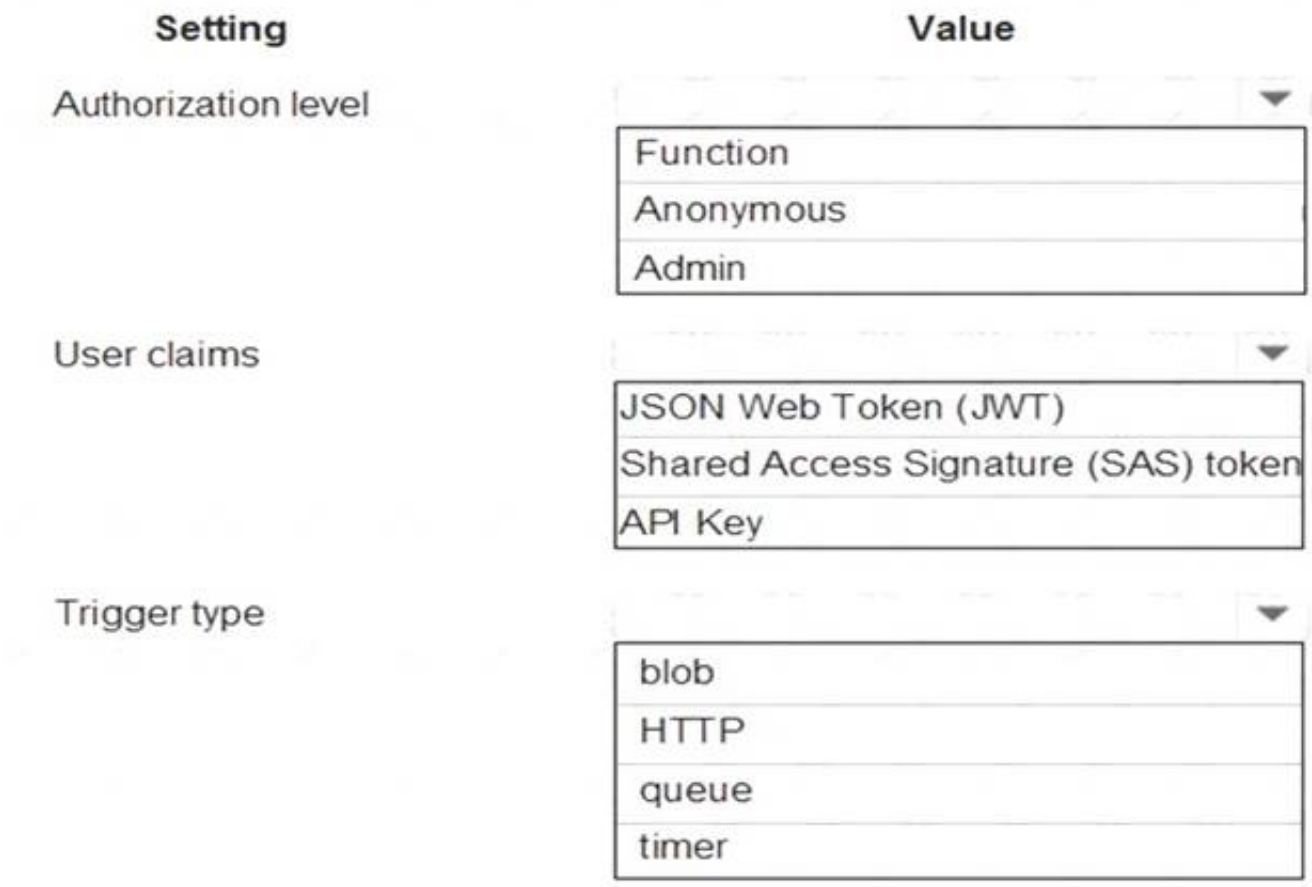

A. Mastered
B. Not Mastered

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 113
- (Exam Topic 8)
You are maintaining an existing application that uses an Azure Blob GPv1 Premium storage account. Data older than three months is rarely used.
Data newer than three months must be available immediately. Data older than a year must be saved but does not need to be available immediately.
You need to configure the account to support a lifecycle management rule that moves blob data to archive storage for data not modified in the last year.
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.

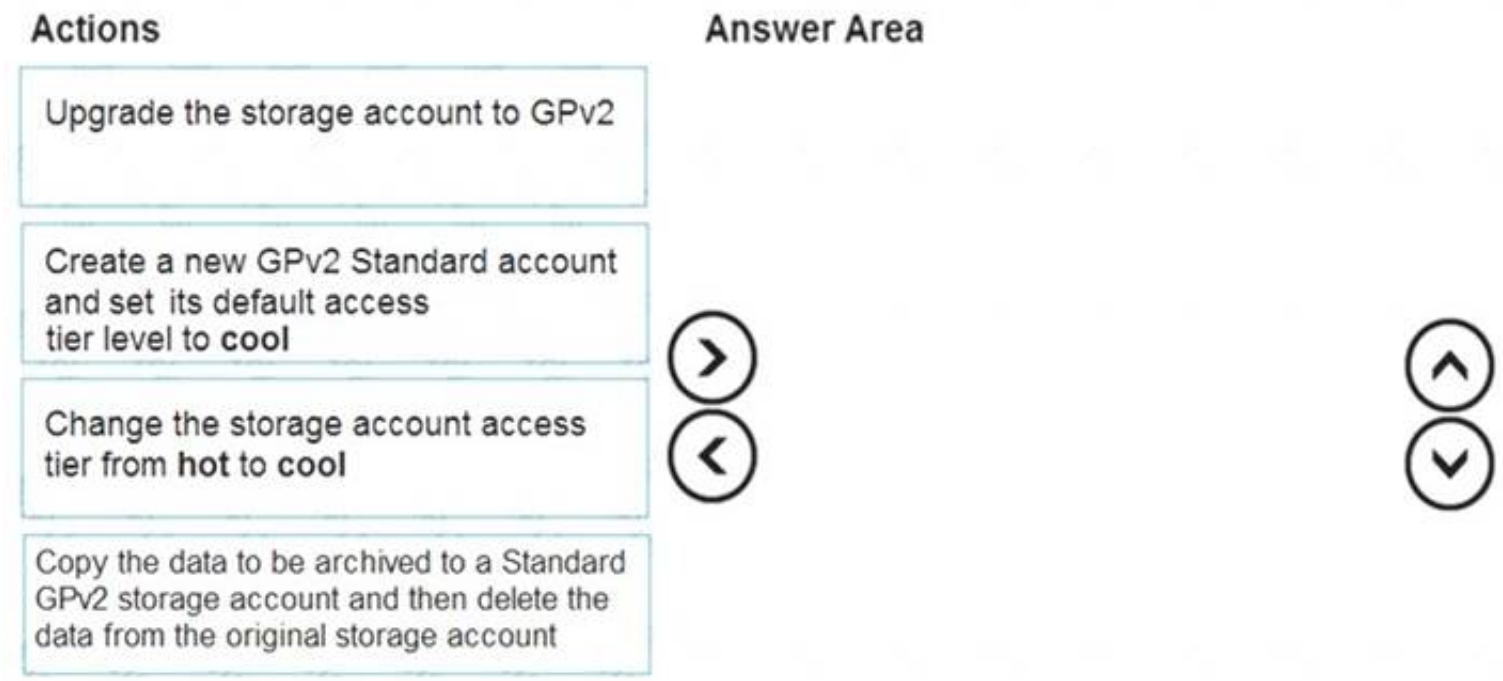

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Upgrade the storage account to GPv2
Object storage data tiering between hot, cool, and archive is supported in Blob Storage and General Purpose v2 (GPv2) accounts. General Purpose v1 (GPv1) accounts don't support tiering.
You can easily convert your existing GPv1 or Blob Storage accounts to GPv2 accounts through the Azure portal.
Step 2: Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account
Step 3: Change the storage account access tier from hot to cool Note: Hot - Optimized for storing data that is accessed frequently.
Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.
Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.
Only the hot and cool access tiers can be set at the account level. The archive access tier can only be set at the blob level.
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers

NEW QUESTION 118
- (Exam Topic 8)
You develop and deploy an ASP.NET web app to Azure App Service. You use Application Insights telemetry to monitor the app.
You must test the app to ensure that the app is available and responsive from various points around the world and at regular intervals. If the app is not responding, you must send an alert to support staff.
You need to configure a test for the web app.
Which two test types can you use? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. integration
B. multi-step web
C. URL ping
D. unit
E. load

Answer: BC

Explanation:
There are three types of availability tests:
- URL ping test: a simple test that you can create in the Azure portal.
- Multi-step web test: A recording of a sequence of web requests, which can be played back to test more complex scenarios. Multi-step web tests are created in Visual Studio Enterprise and uploaded to the portal for execution.
- Custom Track Availability Tests: If you decide to create a custom application to run availability tests, the TrackAvailability() method can be used to send the results to Application Insights.

Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability

NEW QUESTION 123
- (Exam Topic 8)
You are developing an app that manages users for a video game. You plan to store the region, email address, and phone number for the player. Some players may not have a phone number. The player's region will be used to load-balance data.
Data for the app must be stored in Azure Table Storage.
You need to develop code to retrieve data for an individual player.
How should you complete the code? 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: region
The player's region will be used to load-balance data. Choosing the PartitionKey.
The core of any table's design is based on its scalability, the queries used to access it, and storage operation requirements. The PartitionKey values you choose will dictate how a table will be partitioned and the type of queries that can be used. Storage operations, in particular inserts, can also affect your choice of PartitionKey values.
Box 2: email
Not phone number some players may not have a phone number. Box 3: CloudTable
Box 4 : TableOperation query =.. Box 5: TableResult
References:
https://docs.microsoft.com/en-us/rest/api/storageservices/designing-a-scalable-partitioning-strategy-for-azure-ta

NEW QUESTION 126
- (Exam Topic 8)
You are developing an ASP.NET Core website that uses Azure FrontDoor. The website is used to build custom weather data sets for researchers. Data sets are downloaded by users as Comma Separated Value
(CSV) files. The data is refreshed every 10 hours.
Specific files must be purged from the FrontDoor cache based upon Response Header values. You need to purge individual assets from the Front Door cache.
Which type of cache purge should you use?

A. single path
B. wildcard
C. root domain

Answer: A

Explanation:
These formats are supported in the lists of paths to purge:
- Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example, /pictures/strasbourg.png;
- Wildcard purge: Asterisk (*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /*, for example, /pictures/*.
- Root domain purge: Purge the root of the endpoint with "/" in the path. Reference:
https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching

NEW QUESTION 130
- (Exam Topic 8)

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 a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.
You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.
You need to configure authorization. Solution:

- Configure and use Integrated Windows Authentication in the website.
- In the website, query Microsoft Graph API to load the group to which the user is a member. Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Microsoft Graph is a RESTful web API that enables you to access Microsoft Cloud service resources.
Instead in the Azure AD application's manifest, set value of the groupMembershipClaims option to All. In the website, use the value of the groups claim from the JWT for the user to determine permissions.
Reference:
https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/

NEW QUESTION 132
- (Exam Topic 8)
You are designing a multi-tiered application that will be hosted on Azure virtual machines. The virtual machines will run Windows Server. Front-end servers will be accessible from the Internet over port 443. The other servers will NOT be directly accessible over the internet
You need to recommend a solution to manage the virtual machines that meets the following requirement
• Allows the virtual machine to be administered by using Remote Desktop.
• Minimizes the exposure of the virtual machines on the Internet Which Azure service should you recommend?

A. Azure Bastion
B. Service Endpoint
C. Azure Private Link
D. Azure Front Door

Answer: C

NEW QUESTION 137
- (Exam Topic 8)
You have an existing Azure storage account that stores large volumes of data across multiple containers.
You need to copy all data from the existing storage account to a new storage account. The copy process must meet the following requirements:

- Automate data movement.
- Minimize user input required to perform the operation.
- Ensure that the data movement process is recoverable.

What should you use?

A. AzCopy
B. Azure Storage Explorer
C. Azure portal
D. .NET Storage Client Library

Answer: A

Explanation:
You can copy blobs, directories, and containers between storage accounts by using the AzCopy v10 command-line utility.
The copy operation is synchronous so when the command returns, that indicates that all files have been copied. Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy

NEW QUESTION 138
- (Exam Topic 8)
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 an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data. You need to ensure the app does not time out and processes the blob data.
Solution: Pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
Large, long-running functions can cause unexpected timeout issues. General best practices include: Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue

trigger function. This approach lets you defer the actual work and return an immediate response.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices

NEW QUESTION 142
- (Exam Topic 8)
You are developing a mobile app that uses an API which stores geospabal data in Azure Cosmos D& The app will be used to find restaurants in a particular area and related information including food types, menu information and the optimal route to a selected restaurant from the user's current location.
Which Azure Cosmos DB API should you use for the API?

A. MongoDB
B. Gremlin
C. Cassandra
D. Core

Answer: A

NEW QUESTION 143
- (Exam Topic 8)
You need to implement the Azure Function for delivery driver profile information.
Which configurations should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Configuration	Value
Code library	Microsoft Authentication Library (MSAL) Microsoft Azure Key Vault SDK Azure Identity library
API	Microsoft Graph Azure Active Directory Graph Azure Key Vault

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Code Library: MSAL
API: Microsoft Graph
https://docs.microsoft.com/en-us/azure/active-directory/develop/msal-overview

NEW QUESTION 145
- (Exam Topic 8)
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.

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

NOTE: Each correct selection is worth one point.

	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 147
- (Exam Topic 8)
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 = [ _.value(); | _.readDocument('item'); | getContext().getRequest(); | getContext().getResponse(); ]
var i = r.getBody();
[ if (!("tip" in i)) { | if (request.getValue("tip") === null){ | if (isNaN(i)["tip"] || i["tip"]=== null) { | if (typeof_.pluck("tip") == 'number') { ]
    i["tip"] = 0;
}
[ 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 149
- (Exam Topic 8)
Your company is designing an application named App1 that will use data from Azure SQL Database. App1 will be accessed over the internet by many users.
You need to recommend a solution for improving the performance ofApp1. What should you include in the recommendation?

A. Azure HPC cache
B. ExpressRoute
C. a CON profile
D. Azure Cache for Redis

Answer: D

NEW QUESTION 154
- (Exam Topic 8)
You are developing a solution to store documents in Azure Blob storage. Customers upload documents to multiple containers. Documents consist of PDF, CSV, Microsoft Office format, and plain text files.
The solution must process millions of documents across hundreds of containers. The solution must meet the following requirements:
* Document must the categorized by a customer identifier as they are uploaded to the storage account.
* Allow filtering by the customer identifier.
* Allow searching of information contained within a document.
* Minimize costs.
You created and configure a standard general-purpose v2 storage account to support the solution. You need to implement the solution.
NOTE: Each correct selection is worth one point.
Answer Area

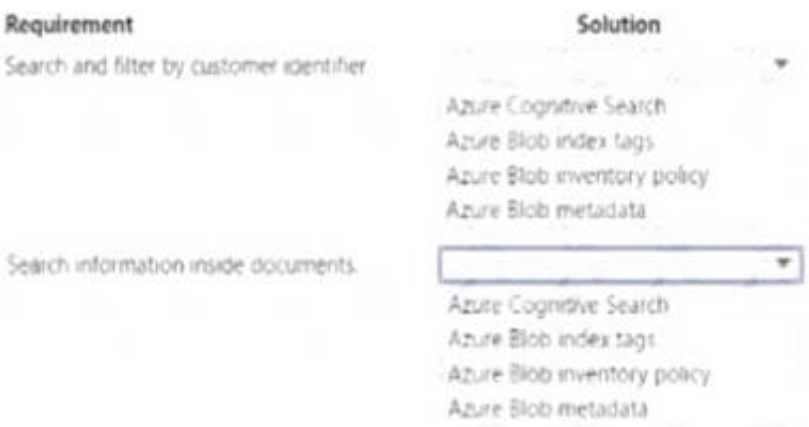

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Azure Blob Index tags:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-index-how-to?tabs=azure-portal
Azure Cognitive Search: Search inside documents

NEW QUESTION 155
- (Exam Topic 8)
You plan to create a Docker image that runs as ASP.NET Core application named ContosoApp. You have a setup script named setupScript.ps1 and a series of application files including ContosoApp.dll.
You need to create a Dockerfile document that meets the following requirements:
•Call setupScript.ps1 when the container is built.
•Run ContosoApp.dll when the container starts.
The Docker document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.
Which four commands should you use to develop the solution? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

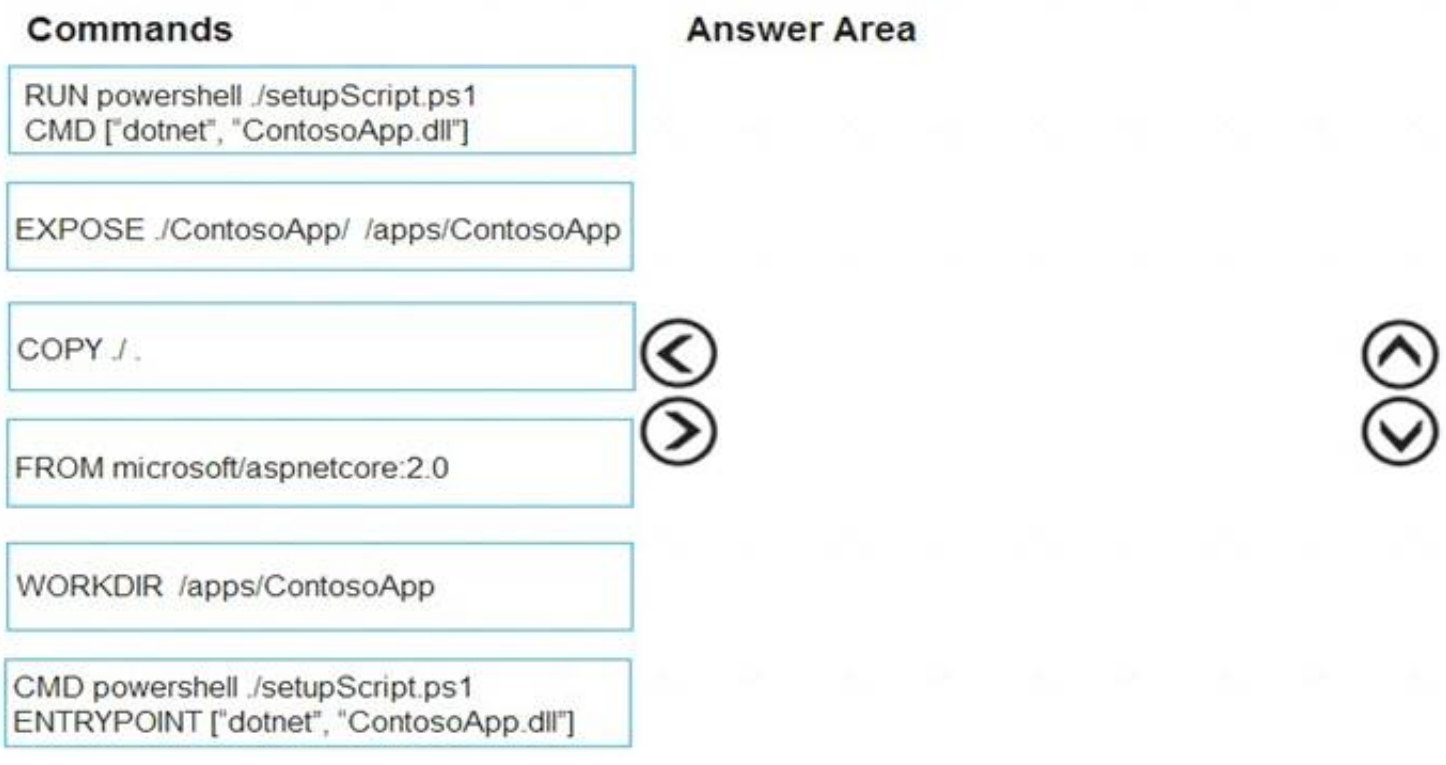

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: WORKDIR /apps/ContosoApp Step 2: COPY ./The Docker document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.
Step 3: EXPOSE ./ContosoApp/ /app/ContosoApp Step 4: CMD powershell ./setupScript.ps1 ENTRYPOINT ["dotnet", "ContosoApp.dll"]
You need to create a Dockerfile document that meets the following requirements:

- Call setupScript.ps1 when the container is built.
- Run ContosoApp.dll when the container starts. References:

https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image

NEW QUESTION 160
- (Exam Topic 8)
You are developing an application that needs access to an Azure virtual machine (VM). The access lifecycle for the application must be associated with the VM service instance. You need to enable managed identity for the VM.
How should you complete the PowerShell segment? To answer, select the appropriate options in the answer area.
NOTE Each correct selection is worth one point.

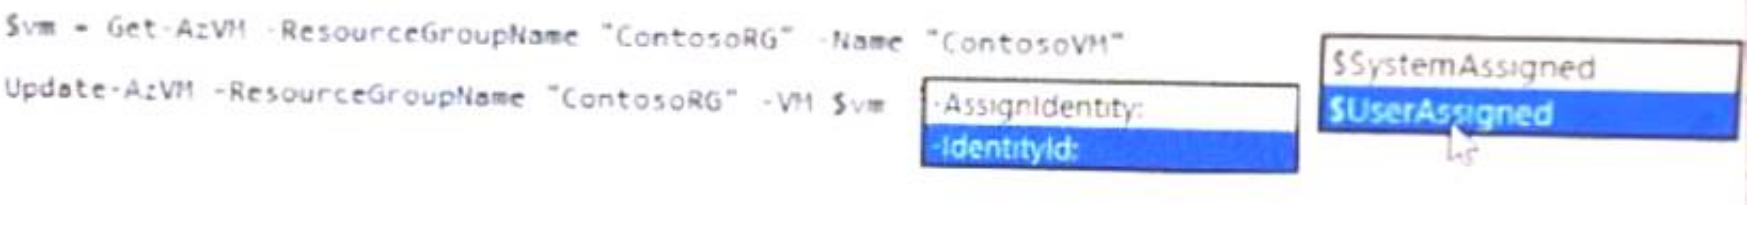

A. Mastered
B. Not Mastered

Answer: A

Explanation:
$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM
Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -AssignIdentity:$SystemAssigned https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-power

NEW QUESTION 161
- (Exam Topic 8)
You are developing an Azure messaging solution.
You need to ensure that the solution that meets the following requirements:
• Provide transactional support
• Provide duplicate detection.
• Store the messages for an unlimited period of time
Which two technologies will meet the requirements? Each correct answer presents a complete solution NOTE Each correct selection is worth one point.

A. Azure Service Bus Queue
B. Azure Storage Queue
C. Azure Service Bus Topic D Azure Event Hub

Answer: AC

Explanation:
The Azure Service Bus Queue and Topic has duplicate detection.
Enabling duplicate detection helps keep track of the application-controlled MessageId of all messages sent into a queue or topic during a specified time window.

Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/duplicate-detection

NEW QUESTION 166
- (Exam Topic 8)
You are a developing a SaaS application that stores data as key value pairs.
You must make multiple editions of the application available. In the lowest cost edition, the performance must be best-effort, and there is no regional failover.
In higher cos! editions customers must be able to select guaranteed performance and support for multiple regions. Azure costs must be minimized.
Which Azure Cosmos OB API should you use for the application?

A. Core
B. MongoDB
C. Cassandra
D. Table API

Answer: C

NEW QUESTION 168
- (Exam Topic 8)
You are developing an application that includes two Docker containers. The application must meet the following requirements
- The containers must not run as root.
- The containers must be deployed to Azure Container Instances by using a YAML file.
- The containers must share a lifecycle, resources, local network and storage volume.
- The storage volume must persist through container crashes.
- The storage volume must be destroyed on stop or restart of the containers. You need to configure Azure Container Instances for the application.

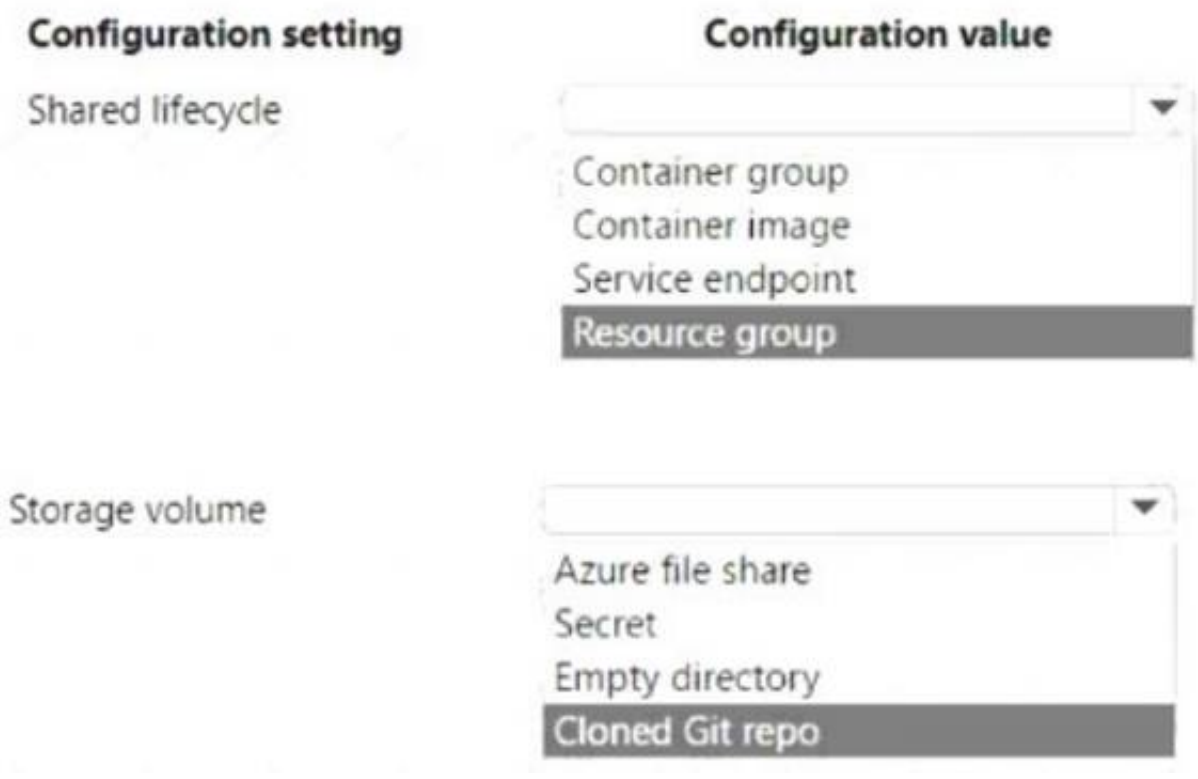

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 170
- (Exam Topic 8)
Your company purchases an Azure subscription and plans to migrate several on-premises virtual machines to Azure. You need to design the infrastructure required (or the Azure virtual machines solution. What should you include in the design?

A. the number of Azure Storage accounts
B. the settings of the Azure virtual networks
C. the size of the virtual machines
D. the number of Azure regions

Answer: C

NEW QUESTION 173
- (Exam Topic 8)
You are developing an ASP.NET Core Web API web service that uses Azure Application Insights to monitor performance and track events.
You need to enable logging and ensure that log messages can be correlated to events tracked by Application Insights.
How should you complete the code? To answer, drag the appropriate code segments to the correct locations. Each code 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:
Box 1: ApplicationInsightsLoggerOptions
If you want to include the EventId and EventName properties, then add the following to the ConfigureServices method:
services AddOptions<ApplicationInsightsLoggerOptions>() Configure(o => o.IncludeEventId = true);
Box 2: IncludeEventID
Box 3: ApplicationServices
In Asp.Net core apps it turns out that trace logs do not show up in Application Insights out of the box. We need to add the following code snippet to our Configure method in Startup.cs:
loggerFactory.AddApplicationInsights(app.ApplicationServices, logLevel); References:
https://blog.computedcloud.com/enabling-application-insights-trace-logging-in-asp-net-core/

NEW QUESTION 177
- (Exam Topic 8)
A company is developing a gaming platform. Users can join teams to play online and see leaderboards that include player statistics. The solution includes an entity named Team.
You plan to implement an Azure Redis Cache instance to improve the efficiency of data operations for entities that rarely change.
You need to invalidate the cache when team data is changed.
How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

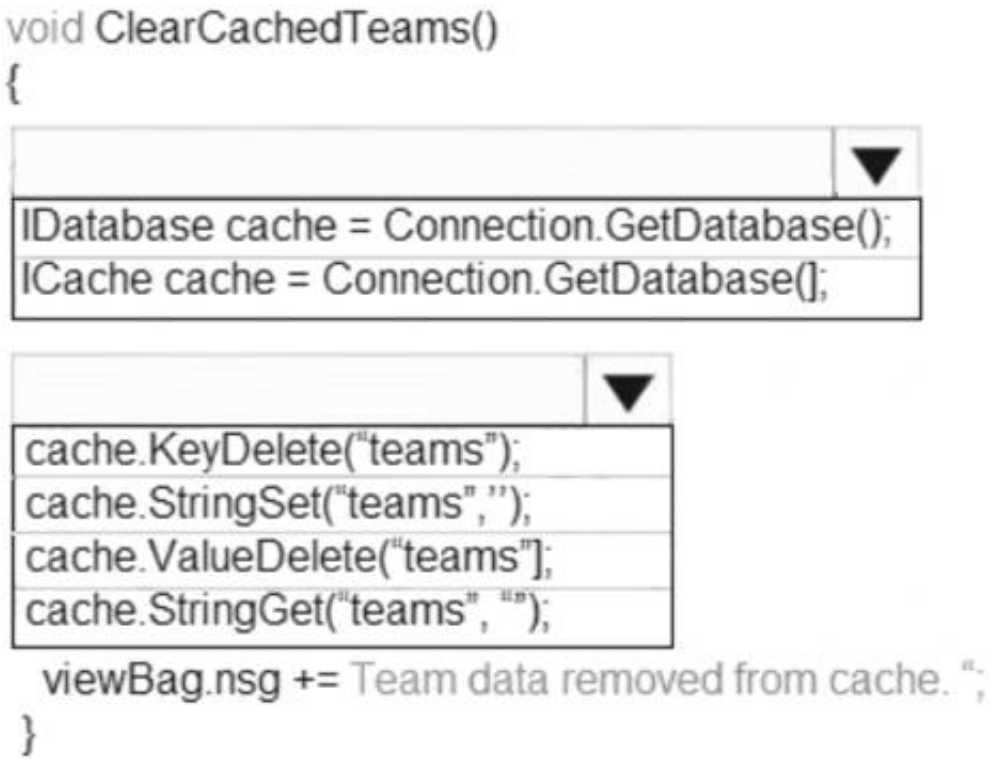

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: IDatabase cache = connection.GetDatabase();
Connection refers to a previously configured ConnectionMultiplexer. Box 2: cache.StringSet("teams","")
To specify the expiration of an item in the cache, use the TimeSpan parameter of StringSet. cache.StringSet("key1", "value1", TimeSpan.FromMinutes(90));
References:
https://azure.microsoft.com/sv-se/blog/lap-around-azure-redis-cache-preview/

NEW QUESTION 182
- (Exam Topic 8)
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.

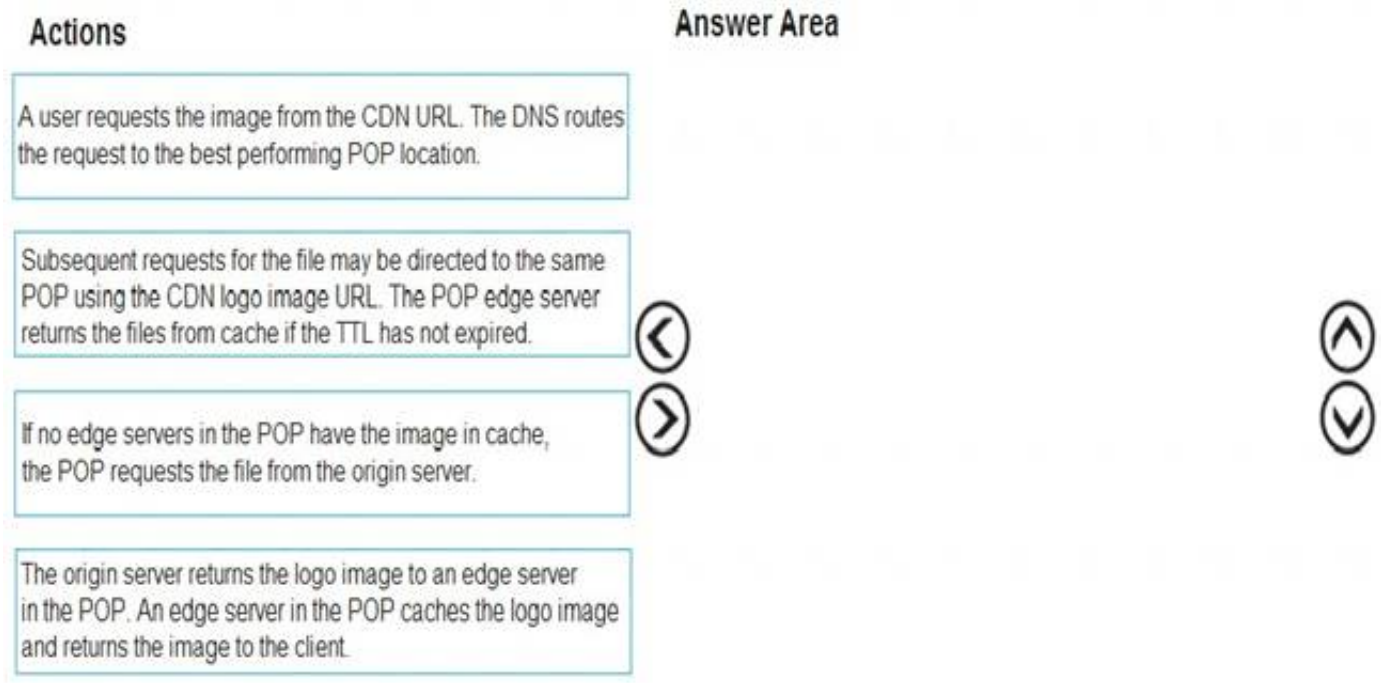

A. Mastered
B. Not Mastered

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 184
- (Exam Topic 8)
You create the following PowerShell script:

```
$source = New-AzScheduledQueryRuleSource -Query 'Heartbeat | where TimeGenerated > ago(1h)' -DataSourceId "contoso"
$schedule = New-AzScheduledQueryRuleSchedule -FrequencyInMinutes 60 -TimeWindowInMinutes 60
$triggerCondition = New-AzScheduledQueryRuleTriggerCondition -ThresholdOperator "LessThan" -Threshold 5
$aznsActionGroup = New-AzScheduledQueryRuleAznsActionGroup -ActionGroup "contoso" -EmailSubject "Custom email subject"
 -CustomWebhookPayload "{ '"alert'":'"#alertrulename'", '"IncludeSearchResults'":true }"
$alertingAction = New-AzScheduledQueryRuleAlertingAction -AznsAction $aznsActionGroup -Severity "3" -Trigger $triggerCondition
New-AzScheduledQueryRule -ResourceGroupName "contoso" -Location "eastus" -Action $alertingAction -Enabled $true
 -Description "Alert description" -Schedule $schedule -Source $source -Name "Alert Name"
```

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

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	○	○
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	○	○
The log alert is scheduled to run every two hours.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated
Box 1: No
The AzScheduledQueryRuleSource is Heartbeat, not CPU.
Box 2: Yes
The AzScheduledQueryRuleSource is Heartbeat!
Note: New-AzScheduledQueryRuleTriggerCondition creates an object of type Trigger Condition. This object is to be passed to the command that creates Alerting Action object.
Box 3: No
The schedule is 60 minutes, not two hours.
-FrequencyInMinutes: The alert frequency.
-TimeWindowInMinutes: The alert time window
The New-AzScheduledQueryRuleSchedule command creates an object of type Schedule. This object is to be passed to the command that creates Log Alert Rule.
Reference:
https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryrule https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryruletriggercondition

NEW QUESTION 189
- (Exam Topic 8)
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 194
- (Exam Topic 8)
You are building a web application that performs image analysis on user photos and returns metadata containing objects identified. The image analysis is very costly in terms of time and compute resources. You are planning to use Azure Redo Cache so Cache uploads do not need to be reprocessed.
In case of an Azure data center outage metadata loss must be kept to a minimum. You need to configure the Azure Redis cache instance.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection in worth one point.

A. Configure Azure Redis with persistence
B. Configure second storage account for persistence
C. Set backup frequency to the minimum value
D. Configure Azure Redis with RDS persistence

Answer: BD

NEW QUESTION 196
- (Exam Topic 8)
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 198
- (Exam Topic 8)
You are developing a project management service by using ASP.NET. The service hosts conversations, files, to-do lists, and a calendar that users can interact with at any time.
The application uses Azure Search for allowing users to search for keywords in the project data.
You need to implement code that creates the object which is used to create indexes in the Azure Search service.
Which two objects should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. SearchService
B. SearchIndexClient
C. SearchServiceClient
D. SearchCredentials

Answer: BC

Explanation:
The various client libraries define classes like Index, Field, and Document, as well as operations like Indexes.Create and Documents.Search on the SearchServiceClient and SearchIndexClient classes.
Example:
The sample application we'll be exploring creates a new index named "hotels", populates it with a few documents, then executes some search queries. Here is the main program, showing the overall flow:
/ This sample shows how to delete, create, upload documents and query an index static void Main(string[] args)
{
IConfigurationBuilder builder = new ConfigurationBuilder().AddJsonFile("appsettings.json"); IConfigurationRoot configuration = builder.Build();
SearchServiceClient serviceClient = CreateSearchServiceClient(configuration); Console.WriteLine("{0}", "Deleting index...\n");
DeleteHotelsIndexIfExists(serviceClient);
Console.WriteLine("{0}", "Creating index...\n"); CreateHotelsIndex(serviceClient);
ISearchIndexClient indexClient = serviceClient.Indexes.GetClient("hotels"); References:
https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk

NEW QUESTION 203
- (Exam Topic 8)
A company is developing a mobile app for field service employees using Azure App Service Mobile Apps as the backend.
The company’s network connectivity varies throughout the day. The solution must support offline use and synchronize changes in the background when the app is online app.

You need to implement the solution.
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.

```
var client = new MobileServiceClient("MOBILE_APP_URL");
var store = new MobileServiceSQLiteStore
(Constants.OfflineDbPath);
store.DefineTable<TodoItem>();
await client.SyncContext.IntitalizeAsync(store);
```

[▼]
- var todoTable = client.GetSyncTable<TodoItem>();
- var todoTable = client.GetTable<TodoItem>();
- var todoTable = client.SyncTable;
- var todoTable = client.Table;

```
await client.SyncContext.PushAsync();
```

[▼]
- await todoTable.PullAsync("allTodoItems",todoTable.CreateQuery());
- await todoTable.UpdateAsync();
- todoTable.PullAsync("allTodoItems", todoTable.CreateQuery());
- todoTable.UpdateAsync();

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: var todoTable = client GetSyncTable<TodoItem>()
To setup offline access, when connecting to your mobile service, use the method GetSyncTable instead of GetTable (example):
IMobileServiceSyncTable todoTable = App.MobileService.GetSyncTable(); / Box 2: await todoTable.PullAsync("allTodoItems",todo.Table.CreateQuery());
Your app should now use IMobileServiceSyncTable (instead of IMobileServiceTable) for CRUD operations. This will save changes to the local database and also keep a log of the changes. When the app is ready to synchronize its changes with the Mobile Service, use the methods PushAsync and PullAsync (example):
await App.MobileService.SyncContext.PushAsync(); await todoTable.PullAsync();
References:
https://azure.microsoft.com/es-es/blog/offline-sync-for-mobile-services/

NEW QUESTION 208
- (Exam Topic 8)
You are developing an application to store and retrieve data in Azure Blob storage. The application will be hosted in an on-premises virtual machine (VM). The VM is connected to Azure by using a Site-to-Site VPN gateway connection. The application is secured by using Azure Active Directory (Azure AD) credentials.
The application must be granted access to the Azure Blob storage account with a start time, expiry time, and read permissions. The Azure Blob storage account access must use the Azure AD credentials of the application to secure data access. Data access must be able to be revoked if the client application security is breached.
You need to secure the application access to Azure Blob storage.
Which security features should you use? To answer select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Component	Security Feature
Application (Client)	[▼] Storage Account Access Key / System-assigned Managed Identity / Shared access signature (SAS) token
Azure Storage (Server)	[▼] Stored Access Policy / User-assigned Managed Identity / Cross-Origin Resource Sharing (CORS)

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text, letter Description automatically generated
Box 1: Shared access signature (SAS) token
When your application design requires shared access signatures for access to Blob storage, use Azure AD credentials to create a user delegation SAS when possible for superior security.

Box 2: Stored access policy
Stored access policies give you the option to revoke permissions for a service SAS without having to regenerate the storage account keys.
A shared access signature can take one of the following two forms:

- Service SAS with stored access policy. A stored access policy is defined on a resource container, which can be a blob container, table, queue, or file share. The stored access policy can be used to manage constraints for one or more service shared access signatures. When you associate a service SAS with a stored access policy, the SAS inherits the constraints – the start time, expiry time, and permissions – defined for the stored access policy.
- Ad hoc SAS. Reference:

https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

NEW QUESTION 210
- (Exam Topic 8)
You develop Azure Web Apps for a commercial diving company. Regulations require that all divers fill out a health questionnaire every 15 days after each diving job starts.
You need to configure the Azure Web Apps so that the instance count scales up when divers are filling out the questionnaire and scales down after they are complete.
You need to configure autoscaling.
What are two possible autoscaling configurations to achieve this goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Predictive autoscaling
B. CPU usage-based autoscaling
C. Recurrence profile
D. Fixed date profile

Answer: AC

NEW QUESTION 213
- (Exam Topic 8)
You are developing an application to retrieve user profile information. The application will use the Microsoft Graph SDK.
The app must retrieve user profile information by using a Microsoft Graph API call. You need to call the Microsoft Graph API from the application.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

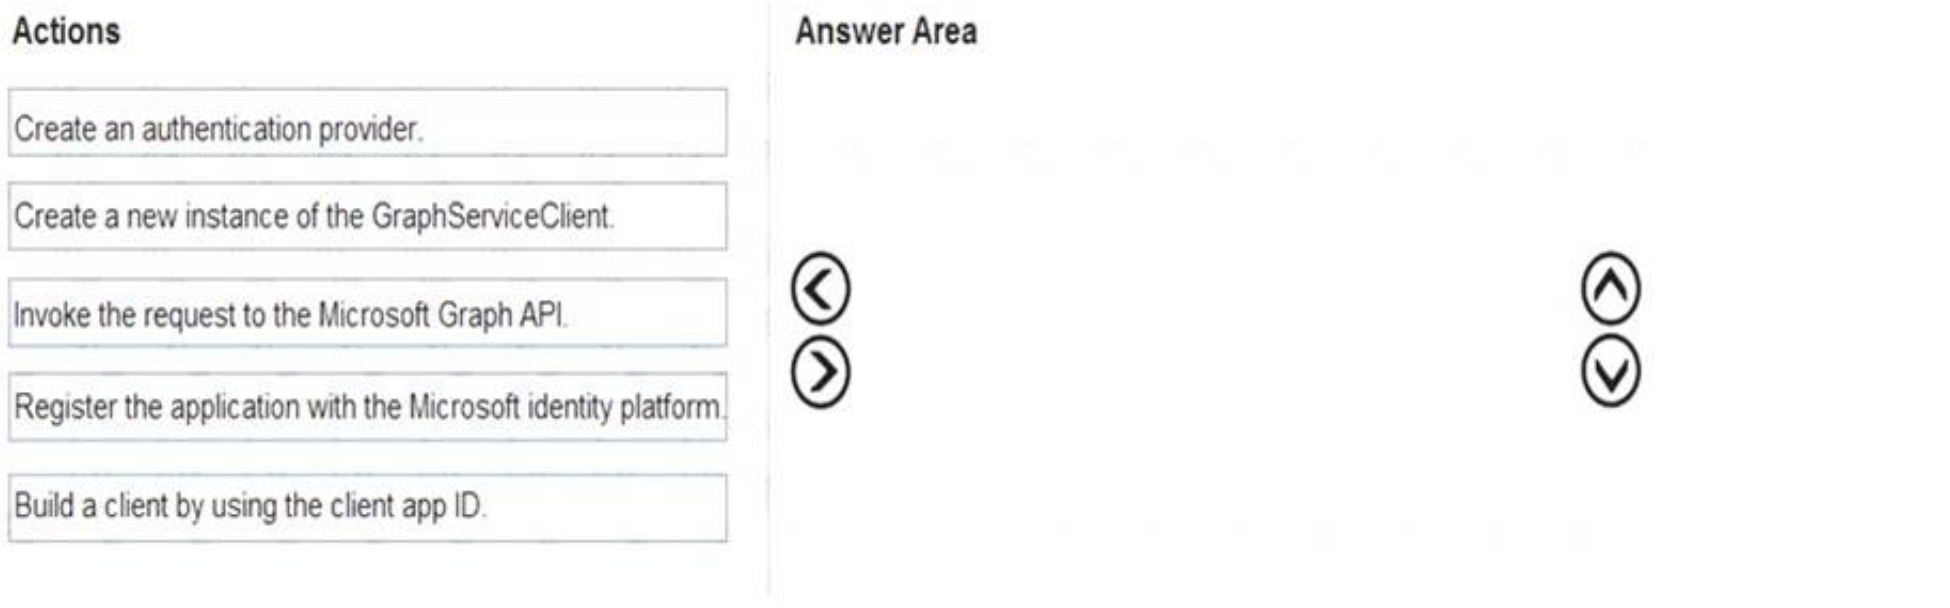

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Table Description automatically generated
Step 1: Register the application with the Microsoft identity platform.
To authenticate with the Microsoft identity platform endpoint, you must first register your app at the Azure app registration portal
Step 2: Build a client by using the client app ID Step 3: Create an authentication provider
Create an authentication provider by passing in a client application and graph scopes. Code example:
DeviceCodeProvider authProvider = new DeviceCodeProvider(publicClientApplication, graphScopes);
// Create a new instance of GraphServiceClient with the authentication provider. GraphServiceClient graphClient = new GraphServiceClient(authProvider);
Step 4: Create a new instance of the GraphServiceClient Step 5: Invoke the request to the Microsoft Graph API Reference:
https://docs.microsoft.com/en-us/graph/auth-v2-service https://docs.microsoft.com/en-us/graph/sdks/create-client

NEW QUESTION 218
- (Exam Topic 8)
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:

- Orders are sent to all available drivers in an area.
- Only orders for the selected restaurants will appear for the driver.
- 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.

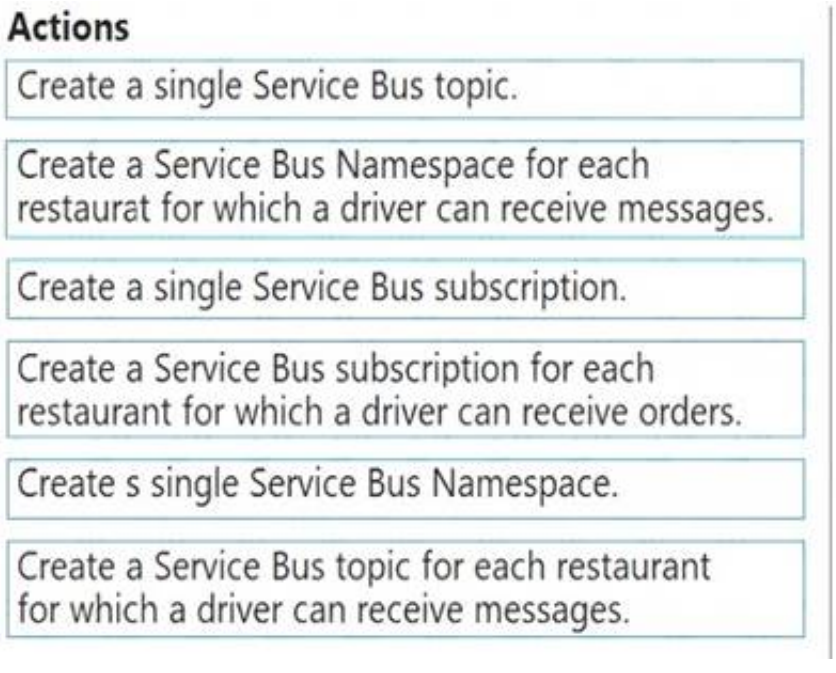

A. Mastered
B. Not Mastered

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

NEW QUESTION 220
- (Exam Topic 8)
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 and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.
You need to ensure that scripts run and resources are available before a swap operation occurs. Solution: Enable auto swap for the Testing slot. Deploy the app to the Testing slot.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.
Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.

```
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostName="[app hostname]" />
<add initializationPage="/Home/About" hostName="[app hostname]" />
</applicationInitialization>
</system.webServer>
```

Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps

NEW QUESTION 225
- (Exam Topic 8)
A company has multiple warehouse. Each warehouse contains IoT temperature devices which deliver temperature data to an Azure Service Bus queue.
You need to send email alerts to facility supervisors immediately if the temperature at a warehouse goes above or below specified threshold temperatures.
Which five 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.

Actions	Answer Area
Add a logic app trigger that fires when one or more messages arrive in the queue.	
Add a Recurrence trigger that schedules the app to run every 15 minutes.	
Add an action that sends an email to specified personnel if the temperature is outside of those thresholds.	
Add a trigger that reads IoT temperature data from a Service Bus queue.	
Add a logic app action that fires when one or more messages arrive in the queue.	
Add a condition that compares the temperature against the upper and lower thresholds.	
Create a blank Logic app.	
Add an action that reads IoT temperature data from the Service Bus queue.	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create a blank Logic app. Create and configure a Logic App.
Step 2: Add a logical app trigger that fires when one or more messages arrive in the queue. Configure the logic app trigger.
Under Triggers, select When one or more messages arrive in a queue (auto-complete). Step 3: Add an action that reads IoT temperature data from the Service Bus queue
Step 4: Add a condition that compares the temperature against the upper and lower thresholds.
Step 5: Add an action that sends an email to specified personnel if the temperature is outside of those thresholds
Reference:
https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-monitoring-notifications-with-azure-logic-apps

NEW QUESTION 226
- (Exam Topic 8)
You are developing an Azure function that connects to an Azure SQL Database instance. The function is triggered by an Azure Storage queue.
You receive reports of numerous System.InvalidOperationExceptions with the following message: “Timeout expired. The timeout period elapsed prior to obtaining a connection from the pool. This may have occurred because all pooled connections were in use and max pool size was reached.”
You need to prevent the exception. What should you do?

A. In the host.json file, decrease the value of the batchSize option
B. Convert the trigger to Azure Event Hub
C. Convert the Azure Function to the Premium plan
D. In the function.json file, change the value of the type option to queueScaling

Answer: A

Explanation:
With the Premium plan the max outbound connections per instance is unbounded compared to the 600 active (1200 total) in a Consumption plan.
Note: The number of available connections is limited partly because a function app runs in a sandbox environment. One of the restrictions that the sandbox imposes on your code is a limit on the number of outbound connections, which is currently 600 active (1,200 total) connections per instance. When you reach this limit, the functions runtime writes the following message to the logs: Host thresholds exceeded: Connections.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/manage-connections https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#service-limits

NEW QUESTION 227
- (Exam Topic 8)
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 and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.
You need to ensure that scripts run and resources are available before a swap operation occurs.
Solution: Update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:

Specify custom warm-up.
Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostName="[app hostname]" />
<add initializationPage="/Home/About" hostName="[app hostname]" />
</applicationInitialization>
</system.webServer> Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps

NEW QUESTION 232
- (Exam Topic 8)
You are developing an e-commerce solution that uses a microservice architecture.
You need to design a communication backplane for communicating transactional messages between various parts of the solution. Messages must be communicated in first-in-first-out (FIFO) order.
What should you use?

A. Azure Storage Queue
B. Azure Event Hub
C. Azure Service Bus
D. Azure Event Grid

Answer: C

Explanation:
As a solution architect/developer, you should consider using Service Bus queues when:
- Your solution requires the queue to provide a guaranteed first-in-first-out (FIFO) ordered delivery. Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compa

NEW QUESTION 236
- (Exam Topic 8)
You provision virtual machines (VMs) as development environments. One VM does not have host.
The VM is stuck in a Windows update process. You attach the OS disk for the affected VM to a recovery VM. You need to correct the issue.
In which order should you perform the actions' 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:
Remove the update that causes the problem
- Take a snapshot of the OS disk of the affected VM as a backup.
- Attach the OS disk to a recovery VM.
- Once the OS disk is attached on the recovery VM, run diskmgmt.msc to open Disk Management, and ensure the attached disk is ONLINE.
- (Step 1) Open an elevated command prompt instance (Run as administrator). Run the following command to get the list of the update packages that are on the attached OS disk:
dism /image:<Attached OS disk>:\ /get-packages > c:\temp\Patch_level
- (Step 2) Open the C:\temp\Patch_level.txt file, and then read it from the bottom up. Locate the update that's in Install Pending or Uninstall Pending state.
- Remove the update that caused the problem:
dism /Image:<Attached OS disk>:\ /Remove-Package /PackageName:<PACK
- (Step 4) Detach the OS disk and recreate the VM. Then check whether the issue is resolved. Reference:
https://docs.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshoot-stuck-updating-boot-error

NEW QUESTION 238
- (Exam Topic 8)
You are implementing an Azure API app that uses built-in authentication and authorization functionality. All app actions must be associated with information about the current user.
You need to retrieve the information about the current user.
What are two possible ways to achieve the goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. HTTP headers
B. environment variables
C. /.auth/me HTTP endpoint
D. /.auth/login endpoint

Answer: AC

Explanation:
A: After App Service Authentication has been configured, users trying to access your API are prompted to sign in with their organizational account that belongs to the same Azure AD as the Azure AD application used to secure the API. After signing in, you are able to access the information about the current user through the HttpContext.Current.User property.
C: While the server code has access to request headers, client code can access GET /.auth/me to get the same access tokens (
References:
https://docs.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-auth-aad https://docs.microsoft.com/en-us/sharepoint/dev/spfx/web-parts/guidance/connect-to-api-secured-with-aad

NEW QUESTION 239
- (Exam Topic 8)
You are developing an Azure Static Web app that contains training materials for a tool company. Each tool's training material is contained In a static web page that Is linked from the tool's publicly available description page.
A user must be authenticated using Azure AD prior to viewing training.
You need to ensure That the user can view training material pages after authentication.
How should you complete the configuration file? 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:
Graphical user interface, application Description automatically generated

NEW QUESTION 243
- (Exam Topic 8)
You are developing an ASP.NET Core web application. You plan to deploy the application to Azure Web App for Containers.
The application needs to store runtime diagnostic data that must be persisted across application restarts. You have the following code:

```
public void SaveDiagData(string data)
{
    var path = Environment.GetEnvironmentVariable("DIAGDATA")
    File.WriteAllText(Path.Combine(path, "data"), data);
}
```

You need to configure the application settings so that diagnostic data is stored as required.
How should you configure the web app's settings? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

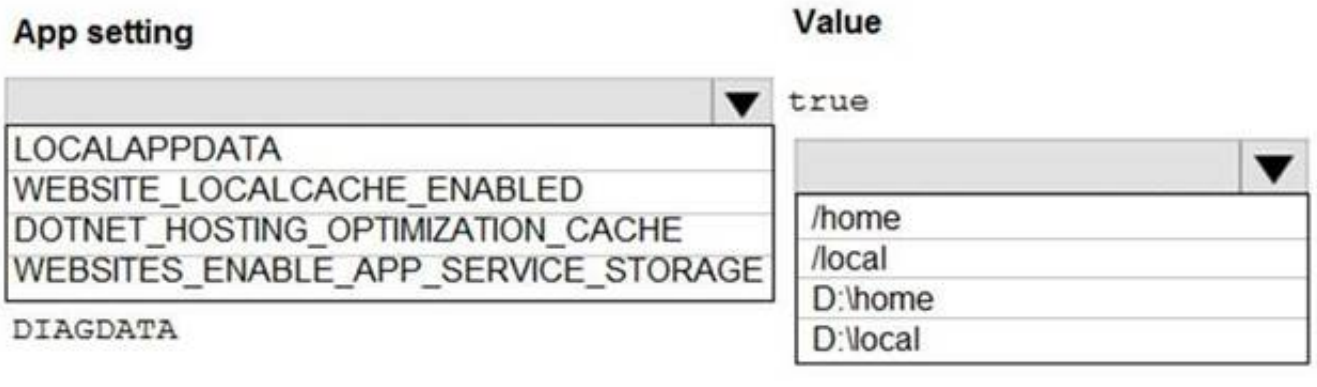

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Box 1: If WEBSITES_ENABLE_APP_SERVICE_STORAGE
If WEBSITES_ENABLE_APP_SERVICE_STORAGE setting is unspecified or set to true, the /home/ directory will be shared across scale instances, and files written will persist across restarts
Box 2: /home nce:
https://docs.microsoft.com/en-us/azure/app-service/containers/app-service-linux-faq

NEW QUESTION 245
- (Exam Topic 8)
You are developing a web application that uses Azure Cache for Redis. You anticipate that the cache will frequently fill and that you will need to evict keys.
You must configure Azure Cache for Redis based on the following predicted usage pattern: A small subset of elements will be accessed much more often than the rest.
You need to configure the Azure Cache for Redis to optimize performance for the predicted usage pattern.
Which two eviction policies will achieve the goal?
NOTE: Each correct selection is worth one point.

A. noeviction
B. allkeys-lru
C. volatile-lru
D. allkeys-random
E. volatile-ttl
F. volatile-random

Answer: BD

Explanation:
B: The allkeys-lru policy evict keys by trying to remove the less recently used (LRU) keys first, in order to make space for the new data added. Use the allkeys-lru policy when you expect a power-law distribution in the popularity of your requests, that is, you expect that a subset of elements will be accessed far more often than the rest.
C: volatile-lru: evict keys by trying to remove the less recently used (LRU) keys first, but only among keys that have an expire set, in order to make space for the new data added.
Note: The allkeys-lru policy is more memory efficient since there is no need to set an expire for the key to be evicted under memory pressure.
Reference: https://redis.io/topics/lru-cache

NEW QUESTION 250
- (Exam Topic 8)
You are developing a web application by using the Azure SDK. The web application accesses data m a zone-redundant BlockBlobStorage storage account
The application must determine whether the data has changed since the application last reao the data. Update operations must use the latest data changes when writing data to the storages.....
You need to implement the update operations.
Which values should you use? To answer, select the appropriate option m the answer area. NOTE Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
See the answer in below image.

Answer Area

Code evaluation	Value
HTTP Header value	VersionId
Conditional header	If-Match

NEW QUESTION 255
- (Exam Topic 8)
You are developing a web application that makes calls to the Microsoft Graph API. You register the application in the Azure portal and upload a valid X509 certificate.
You create an appsettings.json file containing the certificate name, client identifier for the application, and the tenant identifier of the Azure active Directory (Azure AD). You create a method named ReadCertificate to return the X509 certificate by name.
You need to implement code that acquires a token by using the certificate.
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.

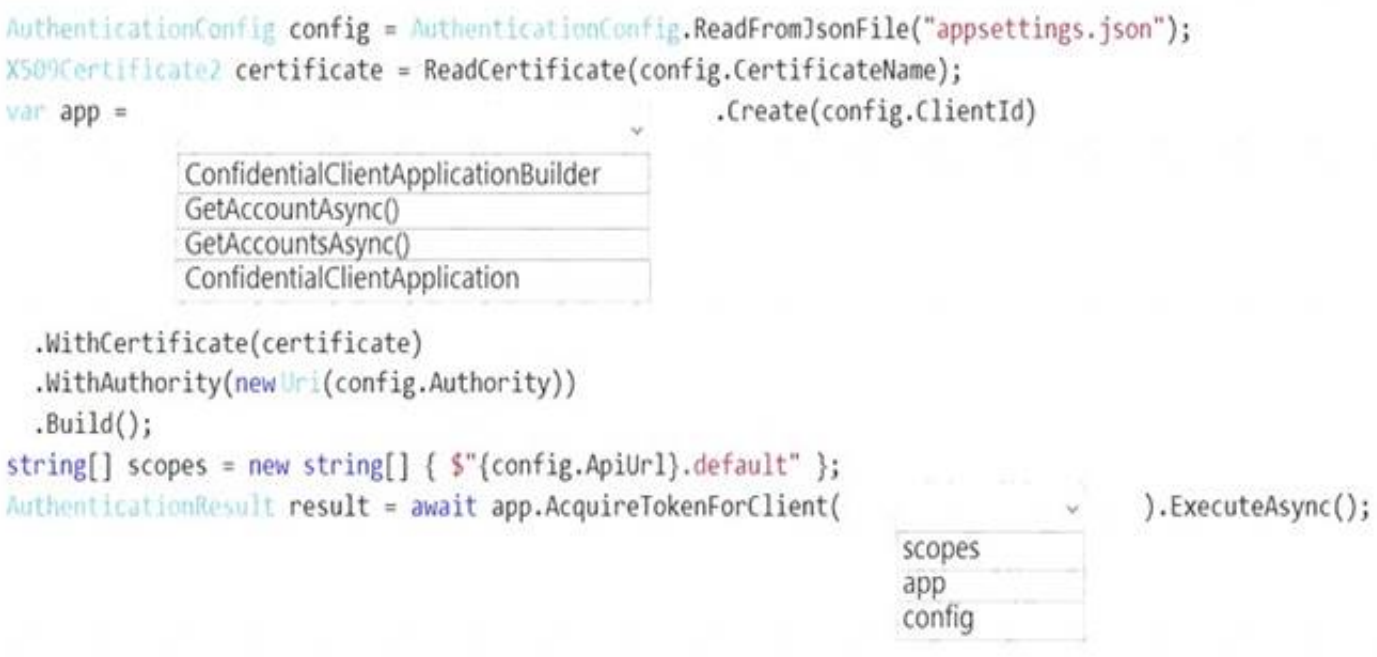

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-app-configuration?tabs=dotne https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-acquire-token?tabs=dotnet#ac

NEW QUESTION 256
- (Exam Topic 8)
You ate developing a web application that uses the Microsoft identity platform to authenticate users and resources. The web application calls several REST APIs. The APIs require an access token from the Microsoft identity platform. You need to request a token.
Which three properties should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Application secret
B. Redirect URI/URL
C. Application name
D. Supported account type
E. Application ID

Answer: ABE

NEW QUESTION 260
- (Exam Topic 8)
You are developing a web application that will use Azure Storage. Older data will be less frequently used than more recent data.
You need to configure data storage for the application. You have the following requirements:
- Retain copies of data for five years.
- Minimize costs associated with storing data that is over one year old.
- Implement Zone Redundant Storage for application data.

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

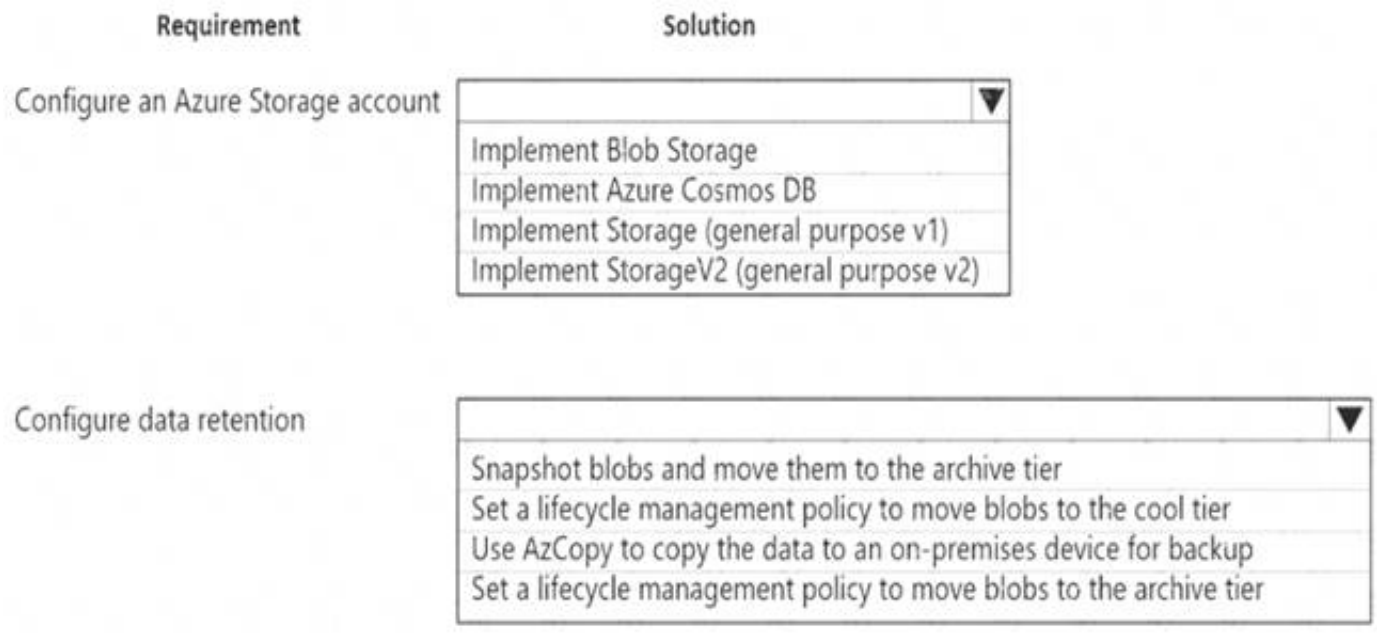

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated with medium confidence
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers
https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy?toc=/azure/storage/blobs/toc.json

NEW QUESTION 265
- (Exam Topic 8)
You develop a web app that uses tier D1 app service plan by using the Web Apps feature of Microsoft Azure App Service.
Spikes in traffic have caused increases in page load times.
You need to ensure that the web app automatically scales when CPU load is about 85 percent and minimize costs.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

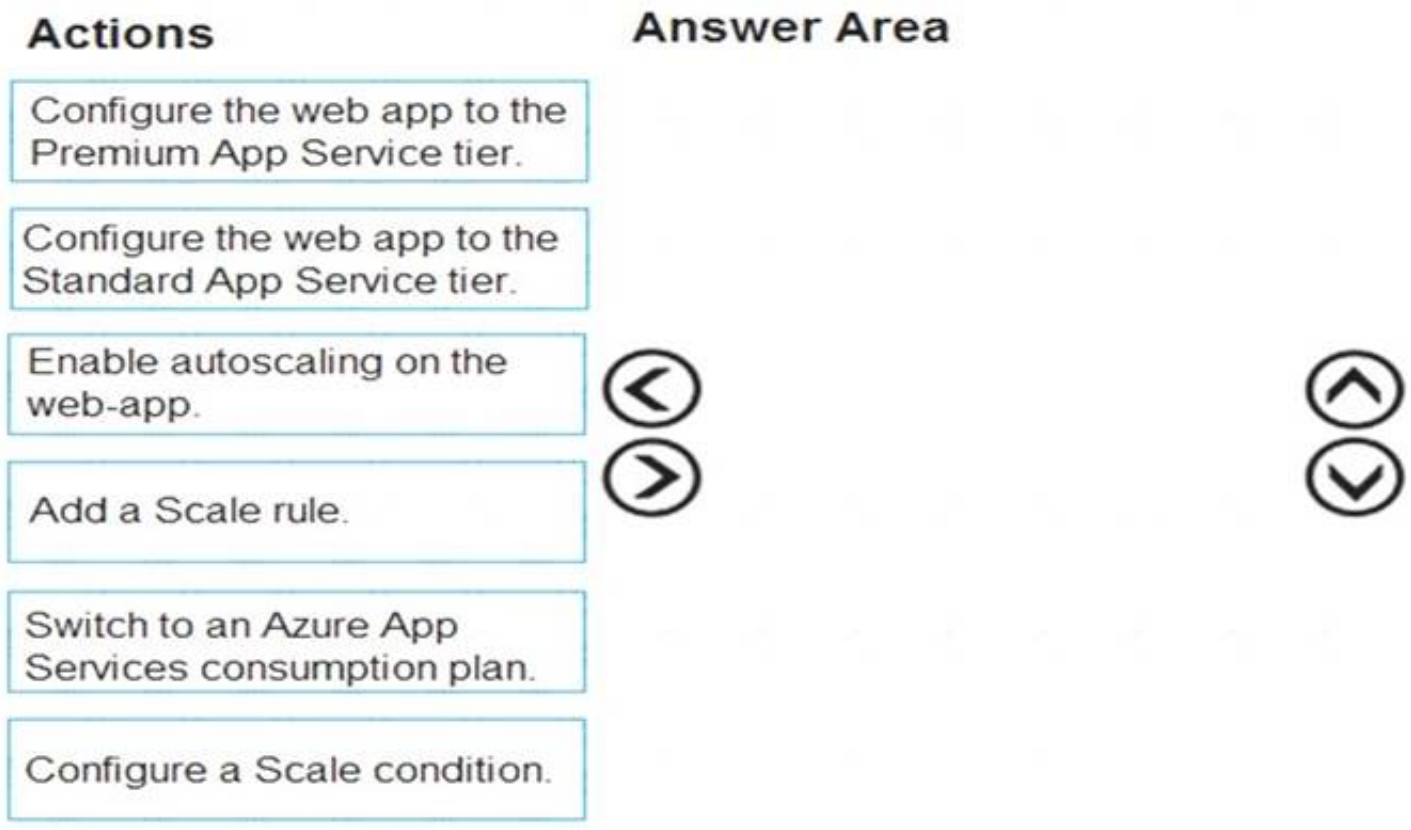

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 270
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AZ-204 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AZ-204 Product From:

https://www.2passeasy.com/dumps/AZ-204/

Money Back Guarantee

AZ-204 Practice Exam Features:

* AZ-204 Questions and Answers Updated Frequently

* AZ-204 Practice Questions Verified by Expert Senior Certified Staff

* AZ-204 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AZ-204 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year