

# Microsoft

## Exam Questions AZ-204

Developing Solutions for Microsoft Azure



**NEW QUESTION 1**

- (Exam Topic 1)

You need to update the APIs to resolve the testing error.

How should you complete the Azure CLI command? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
az webapp --resource-group shipping-apis-test-rg --name web
```

cors  
config  
deployment

add  
up  
remove

slot  
allowed-origins  
name

http://\*.wideworldimporters.com  
 http://test-shippingapi.wideworldimporters.com  
 http://test.wideworldimporters.com  
 http://www.wideworldimporters.com

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Enable Cross-Origin Resource Sharing (CORS) on your Azure App Service Web App.

Enter the full URL of the site you want to allow to access your WEB API or \* to allow all domains.

Box 1: cors

Box 2: add

Box 3: allowed-origins

Box 4: http://testwideworldimporters.com/ References:

<http://donovanbrown.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-with-Azure-App-Service>

**NEW QUESTION 2**

- (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.

| Setting             | Value  |
|---------------------|--|
| Authorization level | <div style="border: 1px solid black; padding: 5px;">           Function<br/>Anonymous<br/>Admin         </div>   |
| User claims         | <div style="border: 1px solid black; padding: 5px;">           JSON Web Token (JWT)<br/>Shared Access Signature (SAS) token<br/>API Key         </div> |
| Trigger type        | <div style="border: 1px solid black; padding: 5px;">           blob<br/>HTTP<br/>queue<br/>timer         </div>  |

- 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 3**

- (Exam Topic 1)

You need to configure Azure CDN for the Shipping web site.

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

**Answer Area**

| Option       | Value   |
|--------------|---|
| Tier         | <div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px;">▼</div> <div style="padding: 2px;">Standard</div> <div style="padding: 2px;">Premium</div> </div>   |
| Profile      | <div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px;">▼</div> <div style="padding: 2px;">Akamai</div> <div style="padding: 2px;">Microsoft</div> </div>   |
| Optimization | <div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px;">▼</div> <div style="padding: 2px;">general web delivery</div> <div style="padding: 2px;">large file download</div> <div style="padding: 2px;">dynamic site acceleration</div> <div style="padding: 2px;">video-on-demand media streaming</div> </div> |

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Scenario: Shipping website

Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.

Tier: Standard Profile: Akamai

Optimization: Dynamic site acceleration

Dynamic site acceleration (DSA) is available for Azure CDN Standard from Akamai, Azure CDN Standard from Verizon, and Azure CDN Premium from Verizon profiles.

DSA includes various techniques that benefit the latency and performance of dynamic content. Techniques include route and network optimization, TCP optimization, and more.

You can use this optimization to accelerate a web app that includes numerous responses that aren't cacheable. Examples are search results, checkout transactions, or real-time data. You can continue to use core Azure CDN caching capabilities for static data.

Reference:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-optimization-overview>

**NEW QUESTION 4**

- (Exam Topic 1)

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

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

**Answer: D**

**Explanation:**

Migrate from on-premises or cloud implementations of MongoDB to Azure Cosmos DB with minimal downtime by using Azure Database Migration Service.

Perform resilient migrations of MongoDB data at scale and with high reliability.

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

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

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

**NEW QUESTION 5**

- (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.

| Issue              | Tool   |
|--------------------|--|
| Backup and Restore | <div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; padding: 2px;">▼</div> <div style="padding: 2px;">Azure Site Recovery</div> <div style="padding: 2px;">Azure Backup</div> <div style="padding: 2px;">Azure Data Box</div> <div style="padding: 2px;">Azure Migrate</div> </div>                   |
| Performance        | <div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; padding: 2px;">▼</div> <div style="padding: 2px;">Azure Network Watcher</div> <div style="padding: 2px;">Azure Traffic Manager</div> <div style="padding: 2px;">ExpressRoute</div> <div style="padding: 2px;">Accelerated Networking</div> </div> |

- 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 6**

- (Exam Topic 2)

You need to configure the ContentUploadService deployment.

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

- A. Add the following markup to line CS23: types: Private
- B. Add the following markup to line CS24: osType: Windows
- C. Add the following markup to line CS24: osType: Linux
- D. Add the following markup to line CS23: types: Public

**Answer:** A

**Explanation:**

Scenario: All Internal services must only be accessible from Internal Virtual Networks (VNets) There are three Network Location types – Private, Public and Domain

Reference:

<https://devblogs.microsoft.com/powershell/setting-network-location-to-private/>

**NEW QUESTION 7**

- (Exam Topic 2)

You need to store the user agreements.

Where should you store the agreement after it is completed?

- A. Azure Storage queue
- B. Azure Event Hub
- C. Azure Service Bus topic
- D. Azure Event Grid topic

**Answer:** B

**Explanation:**

Azure Event Hub is used for telemetry and distributed data streaming.

This service provides a single solution that enables rapid data retrieval for real-time processing as well as repeated replay of stored raw data. It can capture the streaming data into a file for processing and analysis.

It has the following characteristics:

- > low latency
- > capable of receiving and processing millions of events per second
- > at least once delivery

Reference:

<https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services>

**NEW QUESTION 8**

- (Exam Topic 2)

You need to investigate the http server log output to resolve the issue with the ContentUploadService. Which command should you use first?

- A. az webapp log
- B. az ams live-output
- C. az monitor activity-log
- D. az container attach

**Answer: C**

**Explanation:**

Scenario: Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages. "502 bad gateway" and "503 service unavailable" are common errors in your app hosted in Azure App Service. Microsoft Azure publicizes each time there is a service interruption or performance degradation.

The az monitor activity-log command manages activity logs.

Note: Troubleshooting can be divided into three distinct tasks, in sequential order:

- > Observe and monitor application behavior
- > Collect data
- > Mitigate the issue Reference:

<https://docs.microsoft.com/en-us/cli/azure/monitor/activity-log>

**NEW QUESTION 9**

- (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.

| Json segments       | Answer Area  |
|---------------------|--|
| User                | <pre> "appRoles" : [ {   "value": [   "   ],   "displayName": "ContentReviewer",   "id": "e1c2ade8-98f8-45fd-aa4a-6d24b512c22a",   "isEnabled" : true,   "   } ], </pre> |
| value               |  |
| role                |  |
| Application         |  |
| allowedMemberTypes  |  |
| allowedAccountTypes |  |

- 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 10**

- (Exam Topic 2)

You need to add code at line AM09 to ensure that users can review content using ContentAnalysisService. How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

▼

```
"allowPublicClient":true
"oauth2Permissions":["login"]
"oauth2AllowUriPathMatching":true
"oauth2AllowIdTokenImplicitFlow":true
```

▼

```
"oauth2AllowImplicitFlow": true
"oauth2RequiredPostResponse":true
"preAuthorizedApplications":["SPA"]
"knownClientApplications":["ContentAnalysisService"]
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: "oauth2Permissions": ["login"]

oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

Box 2: "oauth2AllowImplicitFlow":true

For applications (Angular, Ember.js, React.js, and so on), Microsoft identity platform supports the OAuth 2.0 Implicit Grant flow.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest>

**NEW QUESTION 10**

- (Exam Topic 2)

You need to add YAML markup at line CS17 to ensure that the ContentUploadService can access Azure Storage access keys.

How should you complete the YAML markup? To answer, drag the appropriate YAML segments to the correct locations. Each YAML segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

**YAML segments**

secret

envVar

secretValues

volumes

volumeMounts

environmentVariables

**Answer Area**

```
YAML segment :
- mountPath: /mnt/secrets
  name: accesskey
YAML segment :
- name: accesskey
YAML segment :
  key: TXkgZmlyc3Qgc2VjcmV0IEZPTwo=
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: volumeMounts Example:

volumeMounts:

- mountPath: /mnt/secrets name: secretvolume1 volumes:

- name: secretvolume1 secret:

mysecret1: TXkgZmlyc3Qgc2VjcmV0IEZPTwo= Box 2: volumes

Box 3: secret Reference:

<https://docs.microsoft.com/en-us/azure/container-instances/container-instances-volume-secret>

**NEW QUESTION 14**

- (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.
- > 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 18**

- (Exam Topic 3)

You need to investigate the Azure Function app error message in the development environment. What should you do?

- A. Connect Live Metrics Stream from Application Insights to the Azure Function app and filter the metrics.
- B. Create a new Azure Log Analytics workspace and instrument the Azure Function app with Application Insights.
- C. Update the Azure Function app with extension methods from Microsoft.Extensions.Logging to log events by using the log instance.
- D. Add a new diagnostic setting to the Azure Function app to send logs to Log Analytics.

**Answer:** A

**Explanation:**

Azure Functions offers built-in integration with Azure Application Insights to monitor functions.

The following areas of Application Insights can be helpful when evaluating the behavior, performance, and errors in your functions:

Live Metrics: View metrics data as it's created in near real-time. Failures

Performance Metrics Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-monitoring>

**NEW QUESTION 21**

- (Exam Topic 3)

You need to correct the Azure Logic app error message.

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

NOTE: Each correct selection is worth one point.

| Setting              | Value   |
|----------------------|---|
| authentication level | <div style="border: 1px solid #ccc; padding: 5px;">                     anonymous<br/>                     function<br/>                     admin                 </div> |
| managed identity     | <div style="border: 1px solid #ccc; padding: 5px;">                     system-assigned<br/>                     user-assigned                 </div>                     |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Scenario: You test the Logic app in a development environment. The following error message displays:

'400 Bad Request'

Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function. Note: If the inbound call's request body doesn't match your schema, the trigger returns an HTTP 400 Bad Request error.

Box 1: function

If you have an Azure function where you want to use the system-assigned identity, first enable authentication for Azure functions.

Box 2: system-assigned

Your logic app or individual connections can use either the system-assigned identity or a single user-assigned identity, which you can share across a group of logic apps, but not both.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/create-managed-service-identity>

**NEW QUESTION 22**

- (Exam Topic 4)

You need to ensure that the solution can meet the scaling requirements for Policy Service. Which Azure Application Insights data model should you use?

- A. an Application Insights dependency
- B. an Application Insights event
- C. an Application Insights trace
- D. an Application Insights metric

**Answer: D**

**Explanation:**

Application Insights provides three additional data types for custom telemetry:

Trace - used either directly, or through an adapter to implement diagnostics logging using an instrumentation framework that is familiar to you, such as Log4Net or System.Diagnostics.

Event - typically used to capture user interaction with your service, to analyze usage patterns. Metric - used to report periodic scalar measurements.

Scenario:

Policy service must use Application Insights to automatically scale with the number of policy actions that it is performing.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/data-model>

**NEW QUESTION 26**

- (Exam Topic 4)

You need to insert code at line LE03 of LoginEvent.cs to ensure that all authentication events are processed correctly.

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: id

id is a unique identifier for the event.

Box 2: eventType

eventType is one of the registered event types for this event source.

Box 3: dataVersion

dataVersion is the schema version of the data object. The publisher defines the schema version.

Scenario: Authentication events are used to monitor users signing in and signing out. All authentication events must be processed by Policy service. Sign outs must be processed as quickly as possible.

The following example shows the properties that are used by all event publishers: [

```
{
  "topic": string, "subject": string, "id": string,
  "eventType": string, "eventTime": string, "data": {
    object-unique-to-each-publisher
  },
  "dataVersion": string, "metadataVersion": string
}
```

Reference:

<https://docs.microsoft.com/en-us/azure/event-grid/event-schema>

**NEW QUESTION 28**

- (Exam Topic 5)

You need to resolve the log capacity issue. What should you do?

- A. Create an Application Insights Telemetry Filter
- B. Change the minimum log level in the host.json file for the function
- C. Implement Application Insights Sampling
- D. Set a LogCategoryFilter during startup

**Answer: C**

**Explanation:**

Scenario, the log capacity issue: Developers report that the number of log message in the trace output for the processor is too high, resulting in lost log messages. Sampling is a feature in Azure Application Insights. It is the recommended way to reduce telemetry traffic and storage, while preserving a statistically correct analysis of application data. The filter selects items that are related, so that you can navigate between items when you are doing diagnostic investigations. When metric counts are presented to you in the portal, they are renormalized to take account of the sampling, to minimize any effect on the statistics. Sampling reduces traffic and data costs, and helps you avoid throttling. Reference: <https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling>

**NEW QUESTION 31**

- (Exam Topic 5)

You need to add code at line PC26 of Processing.cs to ensure that security policies are met.

How should you complete the code that you will add at line PC26? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
var resolver = new KeyVaultKeyResolver(_keyVaultClient);
var keyBundle = await _keyVaultClient.GetKeyAsync("...", "...");
```

▼

```
var key = keyBundle.Key;
var key = keyBundle.KeyIdentifier.Identifier;
var key = await resolver.ResolveKeyAsync("encrypt", null);
var key = await resolver.ResolveKeyAsync(keyBundle.KeyIdentifier.Identifier, CancellationToken.None);
```

▼

```
var x = keyBundle.Managed;
var x = AuthenticationScheme.SharedKey;
var x = new BlobEncryptionPolicy(key, resolver);
var x = new DeleteRetentionPolicy (Enabled = key.Kid != null);
```

▼

```
cloudBlobClient.AuthenticationScheme = x;
cloudBlobClient.DefaultRequestOptions.RequireEncryption = x;
cloudBlobClient.DefaultRequestOptions.EncryptionPolicy = x;
cloudBlobClient.SetServiceProperties(new ServiceProperties(deleteRetentionPolicy:x));
```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: var key = await Resolver.ResolveKeyAsyn(keyBundle,KeyIdentifier.CancellationToken.None); Box 2: var x = new BlobEncryptionPolicy(key,resolver);

Example:

// We begin with cloudKey1, and a resolver capable of resolving and caching Key Vault secrets.

BlobEncryptionPolicy encryptionPolicy = new BlobEncryptionPolicy(cloudKey1, cachingResolver); client.DefaultRequestOptions.EncryptionPolicy = encryptionPolicy;

Box 3: cloudblobClient. DefaultRequestOptions.EncryptionPolicy = x; Reference:

<https://github.com/Azure/azure-storage-net/blob/master/Samples/GettingStarted/EncryptionSamples/KeyRotatio>

**NEW QUESTION 32**

- (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/azurermskeyvault/set-azurermskeyvaultaccesspolicy>

**NEW QUESTION 37**

- (Exam Topic 5)

You need to add code at line PC32 in Processing.cs to implement the GetCredentials method in the Processing class.

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.

| Code segments   | Answer Area                                       |
|---|---|
| MSITokenProvider("...", null)                           | var tp = new [code segment]                       |
| tp.GetAccessTokenAsync("...")                           | var t = new TokenCredential(await [code segment]) |
| AzureServiceTokenProvider()                             | return new StorageCredentials(t);                 |
| StringTokenProvider("storage", "msi")                   |   |
| tp.GetAuthenticationHeaderAsync(CancellationToken.None) |   |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: AzureServiceTokenProvider()  
 Box 2: tp.GetAccessTokenAsync("...")

Acquiring an access token is then quite easy. Example code: private async Task<string> GetAccessTokenAsync()  
 {  
 var tokenProvider = new AzureServiceTokenProvider();  
 return await tokenProvider.GetAccessTokenAsync("https://storage.azure.com/");  
 }

Reference:  
<https://joonasw.net/view/azure-ad-authentication-with-azure-storage-and-managed-service-identity>

**NEW QUESTION 39**

- (Exam Topic 5)

You need to resolve the capacity issue. What should you do?

- A. Convert the trigger on the Azure Function to an Azure Blob storage trigger
- B. Ensure that the consumption plan is configured correctly to allow scaling
- C. Move the Azure Function to a dedicated App Service Plan
- D. Update the loop starting on line PC09 to process items in parallel

**Answer:** D

**Explanation:**

If you want to read the files in parallel, you cannot use foreach. Each of the async callback function calls does return a promise. You can await the array of promises that you'll get with Promise.all.

Scenario: Capacity issue: During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

```

PC08 var container = await GetCloudBlobContainer();
PC09 foreach (var fileItem in await ListFiles())
PC10 {
PC11     var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12     var ms = new MemoryStream();
PC13     await file.DownloadToStreamAsync(ms);
PC14     var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15     await blob.UploadFromStreamAsync(ms);
PC16
PC17 }
```

Reference:  
<https://stackoverflow.com/questions/37576685/using-async-await-with-a-foreach-loop>

**NEW QUESTION 42**

- (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 45

- (Exam Topic 7)

You are developing applications for a company. You plan to host the applications on Azure App Services. The company has the following requirements:

- > Every five minutes verify that the websites are responsive.
- > Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly.
- > Generate alerts if a website is experiencing issues.
- > If a website fails to load, the system must attempt to reload the site three more times. You need to implement this process with the least amount of effort.

What should you do? D18912E1457D5D1DDCBD40AB3BF70D5D

- A. Create a Selenium web test and configure it to run from your workstation as a scheduled task.
- B. Set up a URL ping test to query the home page.
- C. Create an Azure function to query the home page.
- D. Create a multi-step web test to query the home page.
- E. Create a Custom Track Availability Test to query the home page.

**Answer: D**

#### Explanation:

You can monitor a recorded sequence of URLs and interactions with a website via multi-step web tests. Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-multistep>

#### NEW QUESTION 47

- (Exam Topic 7)

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 Azure solutions.

You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager. You need to obtain an Azure Resource Manager access token.

Solution: Use an X.509 certificate to authenticate the VM with Azure Resource Manager. Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

#### Explanation:

Instead run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm>

#### NEW QUESTION 50

- (Exam Topic 7)

D18912E1457D5D1DDCBD40AB3BF70D5D

You are building a website that uses Azure Blob storage for data storage. You configure Azure Blob storage lifecycle to move all blobs to the archive tier after 30 days.

Customers have requested a service-level agreement (SLA) for viewing data older than 30 days. You need to document the minimum SLA for data recovery.

Which SLA should you use?

- A. at least two days
- B. between one and 15 hours
- C. at least one day
- D. between zero and 60 minutes

**Answer: B**

#### Explanation:

The archive access tier has the lowest storage cost. But it has higher data retrieval costs compared to the hot and cool tiers. Data in the archive tier can take several hours to retrieve depending on the priority of the rehydration. For small objects, a high priority rehydrate may retrieve the object from archive in under 1 hour.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal>

#### NEW QUESTION 52

- (Exam Topic 7)

You develop an ASP.NET Core MVC application. You configure the application to track webpages and custom events.

You need to identify trends in application usage.

Which Azure Application Insights Usage Analysis features should you use? To answer, drag the appropriate features to the correct requirements. Each feature 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.

| Requirement  | Feature              |
|--|----------------------|
| Which pages visited by users most often correlate to a product purchase?                       | <input type="text"/> |
| How does load time of the product display page affect a user's decision to purchase a product? | <input type="text"/> |
| Which events most influence a user's decision to continue to use the application?              | <input type="text"/> |
| Are there places in the application that users often perform repetitive actions?               | <input type="text"/> |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box1: Users Box 2: Impact

One way to think of Impact is as the ultimate tool for settling arguments with someone on your team about how slowness in some aspect of your site is affecting whether users stick around. While users may tolerate a certain amount of slowness, Impact gives you insight into how best to balance optimization and performance to maximize user conversion.

Box 3: Retention

The retention feature in Azure Application Insights helps you analyze how many users return to your app, and how often they perform particular tasks or achieve goals. For example, if you run a game site, you could compare the numbers of users who return to the site after losing a game with the number who return after winning. This knowledge can help you improve both your user experience and your business strategy.

Box 4: User flows

The User Flows tool visualizes how users navigate between the pages and features of your site. It's great for answering questions like:

How do users navigate away from a page on your site? What do users click on a page on your site?

Where are the places that users churn most from your site?

Are there places where users repeat the same action over and over?

**NEW QUESTION 53**

- (Exam Topic 7)

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

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

You need to create the custom role.

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

| Item               | Value   |
|--------------------|---|
| Powershell command | <input type="text"/> <ul style="list-style-type: none"> <li>Get-AzureRmRoleDefinition-Name "Reader"   ConvertTo-Json   Out-File C:\SupportRole.json</li> <li>Get-AzureRmRoleDefinition-Name "Operator"   ConvertTo-Json   Out-File C:\SupportRole.json</li> <li>Set-AzureRmRoleDefinition-Name "Reader"   Input-File C:\SupportRole.json</li> <li>Set-AzureRmRoleDefinition Input-File C:\SupportRole.json</li> </ul> |
| Actions section    | <input type="text"/> <ul style="list-style-type: none"> <li>"/read", "Microsoft.Support/"</li> <li>"/read"</li> <li>"* *Microsoft.Support/*"</li> <li>""</li> </ul>   |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

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

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

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

**NEW QUESTION 54**

- (Exam Topic 7)

You are developing Azure WebJobs.

You need to recommend a WebJob type for each scenario.

Which WebJob type should you recommend? To answer, drag the appropriate WebJob types to the correct scenarios. Each WebJob type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

| WebJob types                        | Scenario  | WebJob type              |
|-------------------------------------|---|--------------------------|
| <input type="checkbox"/> Triggered  | Run on all instances that the web app runs on. Optionally restrict the WebJob to a single instance. | <input type="checkbox"/> |
| <input type="checkbox"/> Continuous | Run on a single instance that Azure select for load balancing.                                      | <input type="checkbox"/> |
|                                     | Supports remote debugging   | <input type="checkbox"/> |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Continuous

Continuous runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.

Box 2: Triggered

Triggered runs on a single instance that Azure selects for load balancing. Box 3: Continuous

Continuous supports remote debugging. Note:

The following table describes the differences between continuous and triggered WebJobs.

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

**References:**

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

**NEW QUESTION 55**

- (Exam Topic 7)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 60**

- (Exam Topic 7)

You develop a web application.

You need to register the application with an active Azure Active Directory (Azure AD) tenant.

Which three actions should you perform in sequence? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

**Actions**

**Answer Area**

- Select **Manifest** from the middle-tier service registration.
- In Enterprise Applications, select **New application**.
- Add a Cryptographic key.
- Create a new application and provide the name, account type, and redirect URL
- Select the Azure AD instance.
- Use an access token to access the secure resource.
- In App Registrations, select **New registration**.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Register a new application using the Azure portal

- > Sign in to the Azure portal using either a work or school account or a personal Microsoft account.
- > If your account gives you access to more than one tenant, select your account in the upper right corner.

Set your portal session to the Azure AD tenant that you want.

- > Search for and select Azure Active Directory. Under Manage, select App registrations.
- > Select New registration. (Step 1)
- > In Register an application, enter a meaningful application name to display to users.
- > Specify who can use the application. Select the Azure AD instance. (Step 2)
- > Under Redirect URI (optional), select the type of app you're building: Web or Public client (mobile & desktop). Then enter the redirect URI, or reply URL, for your application. (Step 3)
- > When finished, select Register.

**NEW QUESTION 61**

- (Exam Topic 7)

You are developing an application to securely transfer data between on-premises file systems and Azure Blob storage. The application stores keys, secrets, and certificates in Azure Key Vault. The application uses the Azure Key Vault APIs.

The application must allow recovery of an accidental deletion of the key vault or key vault objects. Key vault objects must be retained for 90 days after deletion.

You need to protect the key vault and key vault objects.

Which Azure Key Vault feature should you use? To answer, drag the appropriate features to the correct actions. Each feature 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.

**Features**

**Answer Area**

- Access policy
- Purge protection
- Soft delete
- Shared access signature

| Action  | Feature |
|---|---------|
| Enable retention period and accidental deletion.  | Feature |
| Enforce retention period and accidental deletion. | Feature |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Soft delete

When soft-delete is enabled, resources marked as deleted resources are retained for a specified period (90 days by default). The service further provides a mechanism for recovering the deleted object, essentially undoing the deletion.

Box 2: Purge protection

Purge protection is an optional Key Vault behavior and is not enabled by default. Purge protection can only be enabled once soft-delete is enabled.

When purge protection is on, a vault or an object in the deleted state cannot be purged until the retention period has passed. Soft-deleted vaults and objects can still be recovered, ensuring that the retention policy will be followed.

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/general/soft-delete-overview>

**NEW QUESTION 64**

- (Exam Topic 7)

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**

```

{
  ...
  "resources": [
    {
      "apiVersion": "2016-01-01",
      "type": "Microsoft.Storage/storageAccounts",
      "name": "[concat(
        'storage-', copyIndex(), 'storage', uniqueString(resourceGroup().id))]",
      "location": "[resourceGroup().location]",
      "sku": {
        "name": "Standard_LRS"
      },
      "kind": "Storage",
      "properties": {},
      "copy": {
        "name": "storagesetup",
        "count": 3
      },
      "dependsOn": [
        "[variables('loadBalancerName')]",
        "[variables('virtualNetworkName')]",
        "storagesetup"
      ]
    },
    {
      "apiVersion": "2015-06-15",
      "type": "Microsoft.Compute/virtualMachines",
      "name": "[concat('VM', uniqueString(resourceGroup().id))]",
      "copy": {
        "name": "VM",
        "count": 1
      },
      "dependsOn": [
        "[variables('loadBalancerName')]",
        "[variables('virtualNetworkName')]",
        "storagesetup"
      ]
    }
  ],
  "outputs": {}
}
    
```

- 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 68**

- (Exam Topic 7)

You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway. Data from traffic sensors are stored in Azure Event Hub.

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

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

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

NOTE: Each correct selection is worth one point.

| Setting              | Value  |
|----------------------|--|
| Number of partitions | <input type="text" value="6"/> <ul style="list-style-type: none"> <li>3</li> <li>4</li> <li>6</li> <li>12</li> </ul>                                   |
| Partition Key        | <input type="text" value="Highway"/> <ul style="list-style-type: none"> <li>Highway</li> <li>Department</li> <li>Timestamp</li> <li>VM name</li> </ul> |

- 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 70**

- (Exam Topic 7)

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

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

You need to configure the App Service plan for the Web App.

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

**Answer Area**

| App service plan setting | Value  |
|--------------------------|--|
| Number of VM instances   | <input type="text" value="4"/> <ul style="list-style-type: none"> <li>2</li> <li>4</li> <li>8</li> <li>16</li> </ul>                                     |
| Pricing tier             | <input type="text" value="Isolated"/> <ul style="list-style-type: none"> <li>Isolated</li> <li>Standard</li> <li>Premium</li> <li>Consumption</li> </ul> |

- 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/>

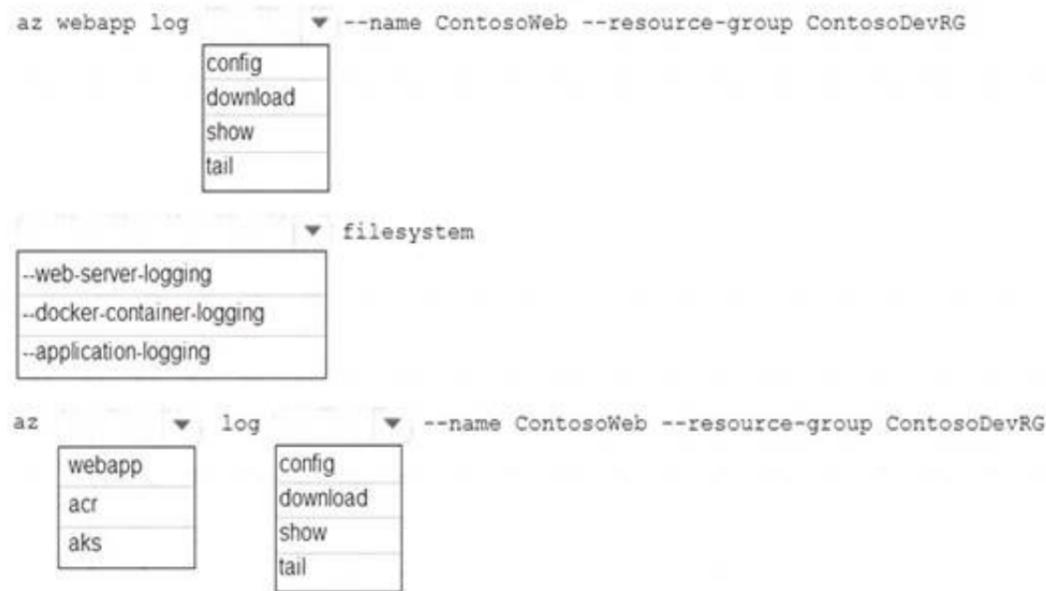
**NEW QUESTION 71**

- (Exam Topic 7)

You plan to deploy a web app to App Service on Linux. You create an App Service plan. You create and push a custom Docker image that image that contains the web app to Azure Container Registry.

You need to access the console logs generated from inside the container in real-time.

How should you complete the Azure CLI command? 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: config  
 To Configure logging for a web app use the command: az webapp log config  
 Box 2: --docker-container-logging Syntax include:  
 az webapp log config [--docker-container-logging {filesystem, off}]  
 Box 3: webapp  
 To download a web app's log history as a zip file use the command: az webapp log download  
 Box 4: download References:  
<https://docs.microsoft.com/en-us/cli/azure/webapp/log>

**NEW QUESTION 76**

- (Exam Topic 7)

You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (AAD) for authentication. You need to implement multifactor authentication for the website.

What should you do? Each correct answer presents part of the solution. NOTE; Each correct selection is worth one point.

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

**Answer: AE**

**Explanation:**

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

**NEW QUESTION 79**

- (Exam Topic 7)

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

- Code should be scalable and allow for growth.
- Each record must consist of a playedId, gameId, score, and time played.
- When users reach a new high score, the system will save the new score using the SaveScore function below.
- Each game is assigned an Id based on the series title.

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

```

01 public void SaveScore(string gameId, string playerId, int score, long timePlayed)
02 {
03     CloudStorageAccount storageAccount = CloudStorageAccount.Parse(connectionString);
04     CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
05     CloudTable table = tableClient.GetTableReference("scoreTable");
06     table.CreateIfNotExists();
07     var scoreRecord = new PlayerScore(gameId, playerId, score, timePlayed);
08     TableOperation insertOperation = TableOperation.Insert(scoreRecord);
09     table.Execute(insertOperation);
10 }
11 public class PlayerScore : TableEntity
12 {
13     public PlayerScore(string gameId, string playerId, int score, long timePlayed)
14     {
15         this.PartitionKey = gameId;
16         this.RowKey = playerId;
17         Score = score;
18         TimePlayed = timePlayed;
19     }
20     public int Score { get; set; }
21     public long TimePlayed { get; set; }
22 }
    
```

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

| PartitionKey | RowKey | Email              |
|--------------|--------|--------------------|
| Harp         | Walter | wharp@contoso.com  |
| Smith        | Steve  | ssmith@contoso.com |
| Smith        | Jeff   | jsmith@contoso.com |

```

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

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

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Yes  
 Code for CosmosDB, example:  
 // Parse the connection string and return a reference to the storage account. CloudStorageAccount storageAccount = CloudStorageAccount.Parse(CloudConfigurationManager.GetSetting("StorageConnectionString"));  
 // Create the table client.  
 CloudTableClient tableClient = storageAccount.CreateCloudTableClient();  
 // Retrieve a reference to the table.  
 CloudTable table = tableClient.GetTableReference("people");  
 // Create the TableOperation object that inserts the customer entity. TableOperation insertOperation = TableOperation.Insert(customer1);  
 Box 2: No  
 A new record will always be added as TableOperation.Insert is used, instead of TableOperation.InsertOrReplace.  
 Box 3: No  
 No partition key is used. Box 4: Yes  
 References:  
<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet>

**NEW QUESTION 81**

- (Exam Topic 7)

You are developing an application that use an Azure blob named data to store application data. The application creates blob snapshots to allow application state to be reverted to an earlier state. The Azure storage account has soft deleted enabled.

The system performs the following operations in order:

- The blob is updated
- Snapshot 1 is created.
- Snapshot 2 is created.
- Snapshot 1 is deleted.

A system error then deletes the data blob and all snapshots. You need to determine which application states can be restored.

What is the restorability of the application data? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

| Application State | Restorability   |
|-------------------|---|
| Data blob         | <div style="border: 1px solid black; padding: 2px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> <span style="border-bottom: 1px solid black; width: 80%;"></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>Can be restored</p> <p>Cannot be restored</p> </div> </div> |
| Snapshot 1        | <div style="border: 1px solid black; padding: 2px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> <span style="border-bottom: 1px solid black; width: 80%;"></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>Can be restored</p> <p>Cannot be restored</p> </div> </div> |
| Snapshot 2        | <div style="border: 1px solid black; padding: 2px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> <span style="border-bottom: 1px solid black; width: 80%;"></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>Can be restored</p> <p>Cannot be restored</p> </div> </div> |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Can be restored

When enabled, soft delete enables you to save and recover your data when blobs or blob snapshots are deleted. This protection extends to blob data that is erased as the result of an overwrite.

Box 2: Cannot be restored It has been deleted.

Box 3: Can be restored It has not been deleted. References:

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

**NEW QUESTION 82**

- (Exam Topic 7)

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.

Solution: Enable Application Request Routing (ARR). Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead deploy and configure Azure Cache for Redis. Update the web applications. Reference:

<https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching#managing-concurrency-in-a-cache>

**NEW QUESTION 84**

- (Exam Topic 7)

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: Add the web applications to Docker containers. Deploy the containers. Deploy the containers to Azure Kubernetes Service (AKS).

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead use Azure Cache for Redis.

Note: Azure Cache for Redis provides a session state provider that you can use to store your session state in-memory with Azure Cache for Redis instead of a SQL Server database. To use the caching session state

provider, first configure your cache, and then configure your ASP.NET application for cache using the Azure Cache for Redis Session State NuGet package.

References:

<https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-aspnet-session-state-provider>

**NEW QUESTION 89**

- (Exam Topic 7)

You are developing a .NET application that communicates with Azure Storage. A message must be stored when the application initializes.

You need to implement the message.

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

**Answer Area**

```

CloudStorageAccount storageAccount = CloudStorageAccount.Parse(CloudConfigurationManager.GetSetting
("StorageConnectionString"));
CloudQueueClient pVar1 = storageAccount.CreateCloudQueueClient();
CloudTableClient pVar2 = pVar1.CreateCloudTableClient();
CloudQueue GetQueueReference();
CloudTable GetTableReference();
CloudQueueClient tExistsAsync();
CloudTableClient CreateCloudQueueClient("contoso-storage");
CloudQueue CreateCloudTableClient();
CloudTable GetQueueReference();
CloudTable GetTableReference();
    
```

A. Mastered

B. Not Mastered

**Answer:** A

**Explanation:**  
 Answer Area

```

CloudStorageAccount storageAccount = CloudStorageAccount.Parse(CloudConfigurationManager.GetSetting
("StorageConnectionString"));
CloudQueueClient pVar1 = storageAccount.CreateCloudQueueClient ();
CloudTableClient pVar2 = pVar1.CreateCloudTableClient ();
CloudQueue GetQueueReference ();
CloudTable GetTableReference ();
CloudQueueClient tExistsAsync ();
CloudQueueClient CreateCloudQueueClient ("contoso-storage");
CloudTableClient CreateCloudTableClient ();
CloudQueue GetQueueReference ();
CloudTable GetTableReference ();
    
```

**NEW QUESTION 94**

- (Exam Topic 7)

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: Disable auto swap. Update the app with a method named statuscheck to run the scripts. Re-enable auto swap and deploy the app to the Production slot. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**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 97**

- (Exam Topic 7)

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

Will each protection method meet the requirement? To answer, drag the appropriate responses to the correct protection methods. Each response may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Yes

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

Box 2: No

Box 3: Yes

In SQL Database, the VIEW permissions are not granted by default to the public fixed database role. This enables certain existing, legacy tools (using older

versions of DacFx) to work properly. Consequently, to work with encrypted columns (even if not decrypting them) a database administrator must explicitly grant the two VIEW permissions.

Box 4: No

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

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

**NEW QUESTION 100**

- (Exam Topic 7)

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 Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.

You have the following requirements:

- > Queue size must not grow larger than 80 gigabytes (GB).
- > Use first-in-first-out (FIFO) ordering of messages.
- > Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application. Create an Azure Function App that uses an Azure Storage Queue trigger.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Create an Azure Function App that uses an Azure Service Bus Queue trigger. Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function>

**NEW QUESTION 105**

- (Exam Topic 7)

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

Spikes in traffic have caused increases in page load times.

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

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

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

| Actions   | Answer Area |
|---|-------------|
| Configure the web app to the Premium App Service tier.  |             |
| Configure the web app to the Standard App Service tier. |             |
| Enable autoscaling on the web-app.                      | ⏪           |
| Add a Scale rule.                                       | ⏩           |
| Switch to an Azure App Services consumption plan.       | ⏴           |
| Configure a Scale condition.                            | ⏵           |

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Step 1: Configure the web app to the Standard App Service Tier

The Standard tier supports auto-scaling, and we should minimize the cost. Step 2: Enable autoscaling on the web app

First enable autoscale Step 3: Add a scale rule

Step 4: Add a Scale condition Reference:

<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-autoscale-get-started>

**NEW QUESTION 110**

- (Exam Topic 7)

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

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

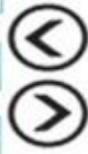
You need to create the application and perform an initial deployment.

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

**Azure CLI Commands**

**Answer Area**

- az group create
- az group update
- az webapp update
- az webapp create
- az appservice plan create



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

Step 1: az group create

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

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

Step 3: az webapp create

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

References:

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

**NEW QUESTION 114**

- (Exam Topic 7)

You develop a gateway solution for a public facing news API.

The news API back end is implemented as a RESTful service 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.

| Values           | Answer Area             |                      |
|------------------|-------------------------|----------------------|
|                  | Configuration parameter | Value                |
| Azure Resource   | Target                  | <input type="text"/> |
| HTTP(s) endpoint | Gateway credentials     | <input type="text"/> |
| Basic            |                         |                      |
| Client cert      |                         |                      |

- 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.

Reference:

<https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-ba>

**NEW QUESTION 116**

- (Exam Topic 7)

You are developing a solution that uses the Azure Storage Client library for .NET. You have the following code: (Line numbers are included for reference only.)

```

01 CloudBlockBlob src = null;
02 try
03 {
04     src = container.ListBlobs().OfType<CloudBlockBlob>().FirstOrDefault();
05     var id = await src.AcquireLeaseAsync(null);
06     var dst = container.GetBlockBlobReference(src.Name);
07     string cpid = await dst.StartCopyAsync(src);
08     await dst.FetchAttributeAsync();
09     return id;
10 }
11 catch (Exception e)
12 {
13     throw;
14 }
15 finally
16 {
17     if (src != null)
18         await src.FetchAttributesAsync();
19     if (src.Properties.LeaseState != LeaseState.Available)
20         await src.BreakLeaseAsync(new TimeSpan(0));
21 }

```

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 creates an infinite lease            | <input type="radio"/> | <input type="radio"/> |
| The code at line 06 always creates a new blob | <input type="radio"/> | <input type="radio"/> |
| The finally block releases the lease          | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Yes

AcquireLeaseAsync does not specify leaseTime.

leaseTime is a TimeSpan representing the span of time for which to acquire the lease, which will be rounded down to seconds. If null, an infinite lease will be acquired. If not null, this must be 15 to 60 seconds.

Box 2: No

The GetBlockBlobReference method just gets a reference to a block blob in this container. Box 3: Yes

The BreakLeaseAsync method initiates an asynchronous operation that breaks the current lease on this container.

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync> <https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.getblockblobreference> <https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync>

**NEW QUESTION 117**

- (Exam Topic 7)

You are writing code to create and run an Azure Batch job.

You have created a pool of compute nodes.

You need to choose the right class and its method to submit a batch job to the Batch service. Which method should you use?

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

**Answer: C**

**Explanation:**

A Batch job is a logical grouping of one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. The app uses the BatchClient.JobOperations.CreateJob method to create a job on your pool.

The Commit method submits the job to the Batch service. Initially the job has no tasks.

```

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

```

References:

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

**NEW QUESTION 120**

- (Exam Topic 7)

You are developing an application that uses Azure Blob storage.

The application must read the transaction logs of all the changes that occur to the blobs and the blob metadata in the storage account for auditing purposes. The changes must be in the order in which they occurred, include only create, update, delete, and copy operations and be retained for compliance reasons.

You need to process the transaction logs asynchronously. What should you do?

- A. Process all Azure Blob storage events by using Azure Event Grid with a subscriber Azure Function app.
- B. Enable the change feed on the storage account and process all changes for available events.
- C. Process all Azure Storage Analytics logs for successful blob events.
- D. Use the Azure Monitor HTTP Data Collector API and scan the request body for successful blob events.

**Answer: B**

**Explanation:**

:

Change feed support in Azure Blob Storage

The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes. Client applications can read these logs at any time, either in streaming or in batch mode. The change feed enables you to build efficient and scalable solutions that process change events that occur in your Blob Storage account at a low cost.

Reference:

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

**NEW QUESTION 123**

- (Exam Topic 7)

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 128**

- (Exam Topic 7)

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 129**

- (Exam Topic 7)

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

| Policy section | Answer Area | Policy             | Policy section |
|----------------|-------------|--------------------|----------------|
| Inbound        |             | Set-variable       | policy section |
| Outbound       |             | Cache-lookup-value | policy section |
|                |             | Cache-store-value  | policy section |
|                |             | Find-and-replace   | policy section |

- 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 />
```

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 131**

- (Exam Topic 7)

You are developing a Java application that uses Cassandra to store key and value data. You plan to use a new Azure Cosmos DB resource and the Cassandra API in the application. You create an Azure Active Directory (Azure AD) group named Cosmos DB Creators to enable provisioning of Azure Cosmos accounts, databases, and containers.

The Azure AD group must not be able to access the keys that are required to access the data. You need to restrict access to the Azure AD group.

Which role-based access control should you use?

- A. DocumentDB Accounts Contributor
- B. Cosmos Backup Operator
- C. Cosmos DB Operator
- D. Cosmos DB Account Reader

**Answer:** C

**Explanation:**

Azure Cosmos DB now provides a new RBAC role, Cosmos DB Operator. This new role lets you provision Azure Cosmos accounts, databases, and containers, but can't access the keys that are required to access the data. This role is intended for use in scenarios where the ability to grant access to Azure Active Directory service principals to manage deployment operations for Cosmos DB is needed, including the account, database, and containers.

Reference:

<https://azure.microsoft.com/en-us/updates/azure-cosmos-db-operator-role-for-role-based-access-control-rbac-is>

**NEW QUESTION 136**

- (Exam Topic 7)

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: Configure the app to use an App Service hosting plan and enable the Always On setting. 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 141**

- (Exam Topic 7)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You are developing a solution that will be deployed to an Azure Kubernetes Service (AKS) cluster. The solution will include a custom VNet, Azure Container Registry images, and an Azure Storage account.

The solution must allow dynamic creation and management of all Azure resources within the AKS cluster. You need to configure an AKS cluster for use with the Azure APIs.

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead create an AKS cluster that supports network policy. Create and apply a network to allow traffic only from within a defined namespace

References:

<https://docs.microsoft.com/en-us/azure/aks/use-network-policies>

**NEW QUESTION 142**

- (Exam Topic 7)

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:** B

**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 144**

- (Exam Topic 7)

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 145**

- (Exam Topic 7)

You are creating a hazard notification system that has a single signaling server which triggers audio and visual alarms to start and stop.

You implement Azure Service Bus to publish alarms. Each alarm controller uses Azure Service Bus to receive alarm signals as part of a transaction. Alarm events must be recorded for audit purposes. Each transaction record must include information about the alarm type that was activated.

You need to implement a reply trail auditing solution.

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

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

**Answer:** AD

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads>

**NEW QUESTION 147**

- (Exam Topic 7)

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 app with a method named statuscheck to run the scripts. Update the app settings for the app. Set the

WEBSITE\_SWAP\_WARMUP\_PING\_PATH and WEBSITE\_SWAP\_WARMUP\_PING\_STATUSES with a path to the new method and appropriate response codes.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

These are valid warm-up behavior options, but are not helpful in fixing swap problems.

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 151**

- (Exam Topic 7)

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.

| Actions  | Answer Area |
|--|-------------|
| Upgrade the storage account to GPv2  |             |
| Create a new GPv2 Standard account and set its default access tier level to cool   | ⤵           |
| Change the storage account access tier from hot to cool  | ⤴           |
| Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account | ⤴<br>⤵      |

- 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 154**

- (Exam Topic 7)

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

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

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

**Answer:** A

**Explanation:**

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

References:

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

**NEW QUESTION 157**

- (Exam Topic 7)

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:

- Create a new Azure AD application. In the application's manifest, define application roles that match the required permission levels for the application.
- Assign the appropriate Azure AD group to each role. In the website, use the value of the roles claim from the JWT for the user to determine permissions.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

To configure Manifest to include Group Claims in Auth Token

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

- Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.
- Locate the "groupMembershipClaims" setting. Set its value to either "SecurityGroup" or "All". To help you decide which:
  - "SecurityGroup" - groups claim will contain the identifiers of all security groups of which the user is a member.
  - "All" - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member

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

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

**NEW QUESTION 158**

- (Exam Topic 7)

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

A rule already exists to scale up the App Service when the average queue length of unprocessed and valid queue messages is greater than 1000.

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

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

**Answer Area**

**Scale rule** [X]

Metric source  
 Storage queue  
 Service Bus queue  
 Current resource  
 Storage queue (classic)

Resource type  
 Service Bus Namespaces

Resource  
 MessageQueue1103

Queues  
 itemqueue

Criteria  
 \* Metric name  
 Message Count  
 Active Message Count

\* Time grain statistic 1 minute time grain  
 Total  
 Maximum  
 Average  
 Count

Greater than  
 Greater than or equal to  
 Less than  
 Less than or equal to

\* Threshold  
 1000

Action  
 \* Operation  
 Increase count by  
 Increase count to  
 Decrease count by  
 Decrease count to

\* Instance count  
 1

\* Cool down (minutes)  
 5

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Service bus queue

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

Box 2: ActiveMessage Count

ActiveMessageCount: Messages in the queue or subscription that are in the active state and ready for delivery. Box 3: Count

Box 4: Less than or equal to

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

Box 5: Decrease count by

**NEW QUESTION 160**

- (Exam Topic 7)

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

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

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

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

How should you configure the application? To answer, drag the appropriate setting to the correct location. Each setting can be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

| Settings           | Answer Area |
|--------------------|-------------|
| client_id          |             |
| profile            |             |
| delegated          |             |
| application        |             |
| user_impersonation |             |

| API             | Permission | Type    |
|-----------------|------------|---------|
| Azure Storage   | Setting    | Setting |
| Microsoft Graph | User.Read  | Setting |

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: user\_impersonation

Box 2: delegated Example:

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

Example

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

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

**NEW QUESTION 164**

- (Exam Topic 7)

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 169**

- (Exam Topic 7)

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

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

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

NOTE: Each correct selection is worth one point.

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

**Answer: AC**

**Explanation:**

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

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

**NEW QUESTION 170**

- (Exam Topic 7)

You have a web app named MainApp. You are developing a triggered App Service background task by using the WebJobs SDK. This task automatically invokes a function code whenever any new data is received in a queue.

You need to configure the services.

Which service should you use for each scenario? To answer, drag the appropriate services to the correct scenarios. Each service may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

| Services   | Scenario   | Service              |
|------------|--|----------------------|
| Logic Apps | Process a queue data item.                                 | <input type="text"/> |
| WebJobs    | Manage all code segments from the same DevOps environment. | <input type="text"/> |
| Flow       |  |                      |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: WebJobs

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

Box 2: Flow

**NEW QUESTION 171**

- (Exam Topic 7)

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

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

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

| Command segments          | Answer Area          |
|---------------------------|----------------------|
| az aks get-credentials    |                      |
| az appservice plan create | <input type="text"/> |
| az aks create             | <input type="text"/> |
| az group create           | <input type="text"/> |
| kubectl apply             | <input type="text"/> |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: az group create

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

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

Step 2 : az aks create

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

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

Step 4: az aks get-credentials

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

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

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

**NEW QUESTION 176**

- (Exam Topic 7)

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.

| Property                    | Value   |
|-----------------------------|---|
| Client certificate location | <input type="text"/><br>HTTP request header<br>Client cookie<br>HTTP message body<br>URL query string |
| Encoding type               | <input type="text"/><br>HTML<br>URL<br>Unicode<br>Base64  |

- 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 179**

- (Exam Topic 7)

You are preparing to deploy an application to an Azure Kubernetes Service (AKS) cluster. The application must only be available from within the VNet that includes the cluster. You need to deploy the application.

How should you complete the deployment YAML? To answer, drag the appropriate YAML segments to the correct locations. Each YAML segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

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

YAML:

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

References:

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

**NEW QUESTION 183**

- (Exam Topic 7)

Your company is developing an Azure API.

You need to implement authentication for the Azure API. You have the following requirements:

- > All API calls must be secure.
- > Callers to the API must not send credentials to the API.

Which authentication mechanism should you use?

- A. Basic
- B. Anonymous
- C. Managed identity
- D. Client certificate

**Answer: C**

**Explanation:**

Use the authentication-managed-identity policy to authenticate with a backend service using the managed identity of the API Management service. This policy essentially uses the managed identity to obtain an access token from Azure Active Directory for accessing the specified resource. After successfully obtaining the token, the policy will set the value of the token in the Authorization header using the Bearer scheme.

Reference:

<https://docs.microsoft.com/bs-cyrl-ba/azure/api-management/api-management-authentication-policies>

**NEW QUESTION 185**

- (Exam Topic 7)

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.

| Commands  | Answer Area |
|---|-------------|
| RUN powershell .\setupScript.ps1<br>CMD ["dotnet", "ContosoApp.dll"]        |             |
| EXPOSE ./ContosoApp/ /apps/ContosoApp                                       |             |
| COPY . .  | ⬅           |
| FROM microsoft/aspnetcore:2.0   | ➡           |
| WORKDIR /apps/ContosoApp  |             |
| CMD powershell .\setupScript.ps1<br>ENTRYPOINT ["dotnet", "ContosoApp.dll"] | ⬆<br>⬇      |

- 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 187**

- (Exam Topic 7)

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

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

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

You assign users to RBAC role.

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

How should you configure the application? To answer, drag the appropriate setting to the correct location. Each setting may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: user\_impersonation

Box 2: delegated Example:

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

Box 3: delegated Example

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

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

**NEW QUESTION 191**

- (Exam Topic 7)

You develop and deploy an Azure Logic App that calls an Azure Function app. The Azure Function App includes an OpenAPI (Swagger) definition and uses an Azure Blob storage account. All resources are secured by using Azure Active Directory (Azure AD).

The Logic App must use Azure Monitor logs to record and store information about runtime data and events. The logs must be stored in the Azure Blob storage account.

You need to set up Azure Monitor logs and collect diagnostics data for the Azure Logic App.

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.

| Actions   | Answer Area |
|---|-------------|
| Create action groups and alert rules.               |             |
| Create a Log Analytics workspace.                   |             |
| Install the Logic Apps Management solution.         | ⬅           |
| Add a diagnostic setting to the Azure Function App. | ➡           |
| Create an Azure storage account.                    |             |
| Add a diagnostic setting to the Azure Logic App.    | ⬆           |
|   | ⬇           |

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Text Description automatically generated

Step 1: Create a Log Analytics workspace

Before you start, you need a Log Analytics workspace. Step 2: Install the Logic Apps Management solution

To set up logging for your logic app, you can enable Log Analytics when you create your logic app, or you can install the Logic Apps Management solution in your Log Analytics workspace for existing logic apps.

Step 3: Add a diagnostic setting to the Azure Logic App Set up Azure Monitor logs

- > In the Azure portal, find and select your logic app.
- > On your logic app menu, under Monitoring, select Diagnostic settings > Add diagnostic setting. Reference:  
<https://docs.microsoft.com/en-us/azure/logic-apps/monitor-logic-apps-log-analytics>

**NEW QUESTION 196**

- (Exam Topic 7)

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 197**

- (Exam Topic 7)

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 Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.

You have the following requirements:

- > Queue size must not grow larger than 80 gigabytes (GB).
- > Use first-in-first-out (FIFO) ordering of messages.
- > Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application. Create an Azure VM that is triggered from Azure Storage Queue events.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Don't use a VM, instead create an Azure Function App that uses an Azure Service Bus Queue trigger. Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function>

**NEW QUESTION 198**

- (Exam Topic 7)

You develop and deploy a Java RESTful API to Azure App Service.

You open a browser and navigate to the URL for the API. You receive the following error message:

```
Failed to load http://api.azurewebsites.net:6000/#/api/Products: No 'Access-Control-Allow-Origin' header is present on the requested resource. Origin 'http://localhost:6000' is therefore not allowed access
```

You need to resolve the error. What should you do?

- A. Bind an SSL certificate
- B. Enable authentication
- C. Enable CORS
- D. Map a custom domain
- E. Add a CDN

**Answer:** C

**Explanation:**

We need to enable Cross-Origin Resource Sharing (CORS). References:

<https://medium.com/@xinganwang/a-practical-guide-to-cors-51e8fd329a1f>

**NEW QUESTION 199**

- (Exam Topic 7)

You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms. If the stored intake forms are downloaded from storage by a third party, the content of the forms must not be compromised.

You need to store the intake forms according to the requirements. Solution:

- > Create a Azure Key Vault key named skey.
- > Encrypt the intake forms using the public key portion of skey.
- > Store the encrypted data in Azure Blob storage

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**NEW QUESTION 204**

- (Exam Topic 7)

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 206**

- (Exam Topic 7)

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 208**

- (Exam Topic 7)

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.

| Actions  | Answer Area |
|--|-------------|
| Create a user-assigned managed identity for the application.                 |             |
| Create the Azure Functions app with a Premium plan type.                     |             |
| Create an access policy in Azure Key Vault for the application identity.     | ➤           |
| Create an SSL certification in Azure Key Vault for the application identity. | ⬅           |
| Create the Azure Functions app with an App Service plan type.                | ⬆           |
| Create the Azure Functions app with a Consumption plan type.                 | ⬆           |
| Create a system-assigned managed identity for the application.               |             |

- 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 212**

- (Exam Topic 7)

A company uses Azure SQL Database to store data for an app. The data includes sensitive information.

You need to implement measures that allow only members of the managers group to see sensitive information. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Include the managers group.
- B. Exclude the managers group.
- C. Exclude the administrators group.
- D. Navigate to the following URL:  
`PUT https://management.azure.com/subscriptions/00000000-1111-2222-3333-444444444444  
/resourceGroups/rg01/providers/Microsoft.Sql/servers/server01/databases/customers  
/transparentDataEncryption/current?api-version=2014-04-01`
- E. Run the following Azure PowerShell command:  
`New-AzureRmSqlDatabaseDataMaskingRule -SchemaName "dbo" -TableName "customers" -  
-ColumnName "ssn" -MaskingFunction "Default"`

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

**Answer:** BE

**Explanation:**

Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.

SQL users excluded from masking - A set of SQL users or AAD identities that get unmasked data in the SQL query results.

Note: The New-AzureRmSqlDatabaseDataMaskingRule cmdlet creates a data masking rule for an Azure SQL database.

References:

<https://docs.microsoft.com/en-us/powershell/module/azurermsql/new-azurermsqldatabasedatamaskingrule?view>

**NEW QUESTION 213**

- (Exam Topic 7)

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 Hub. Configure the machine identifier as the partition key and enable capture.

- A. Yes
- B. No

**Answer:** A

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-programming-guide>

**NEW QUESTION 215**

- (Exam Topic 7)

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 218**

- (Exam Topic 7)

You plan to create a Docker image that runs an 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 setupScripts.ps1 when the container is built.
- > Run ContosoApp.dll when the container starts.

The Dockerfile document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.

Which five 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.

**Commands**

- FROM microsoft/aspnetcore:latest
- WORKDIR /apps/ContosoApp
- CMD ["dotnet", "ContosoApp.dll"]
- COPY ./ .
- RUN powershell ./setupScript.ps1

**Answer Area**

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: CMD [..]

Cmd starts a new instance of the command interpreter, Cmd.exe. Syntax: CMD <string>

Specifies the command you want to carry out.

Box 2: FROM microsoft/aspnetcore-build:latest

Box 3: WORKDIR /apps/ContosoApp

Box 4: COPY ./ .

Box 5: RUN powershell ./setupScript.ps1

**NEW QUESTION 223**

- (Exam Topic 7)

You are developing an Azure App Service REST API.

The API must be called by an Azure App Service web app. The API must retrieve and update user profile information stored in Azure Active Directory (Azure AD).

You need to configure the API to make the updates.

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

- A. Microsoft Graph API
- B. Microsoft Authentication Library (MSAL)
- C. Azure API Management
- D. Microsoft Azure Security Center
- E. Microsoft Azure Key Vault SDK

**Answer:** AC

**Explanation:**

A: You can use the Azure AD REST APIs in Microsoft Graph to create unique workflows between Azure AD resources and third-party services.

Enterprise developers use Microsoft Graph to integrate Azure AD identity management and other services to automate administrative workflows, such as employee onboarding (and termination), profile maintenance, license deployment, and more.

C: API Management (APIM) is a way to create consistent and modern API gateways for existing back-end services.

API Management helps organizations publish APIs to external, partner, and internal developers to unlock the potential of their data and services.

Reference:

<https://docs.microsoft.com/en-us/graph/azuread-identity-access-management-concept-overview>

**NEW QUESTION 226**

- (Exam Topic 7)

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: C**

**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 228**

- (Exam Topic 7)

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

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

op = new ( );

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

...

table. (op);

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

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

**NEW QUESTION 232**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **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

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The AZ-204 Practice Test Here](#)**