

# Microsoft

## Exam Questions AZ-204

Developing Solutions for Microsoft Azure



### NEW QUESTION 1

- (Topic 8)

You are developing a road tollway tracking application that sends tracking events by using Azure Event Hubs using premium tier. Each road must have a throttling policy uniquely assigned. You need to configure the event hub to allow for per-road throttling. What should you do?

- A. Ensure each road has a unique connection string.
- B. Use a unique consumer group for each road
- C. Use a unique application group for each road
- D. Ensure each road stores events in a different partition.

**Answer: D**

### NEW QUESTION 2

- (Topic 8)

The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location. You need to complete the configuration.

Which Azure CLI or PowerShell command should you run?

- A)
- ```
New-AzureRmResourceGroup
  -Name fridge-rg
  -Location fridge-loc
```
- B)
- ```
connectionStrings=$(az servicebus namespace authorization-rule keys list
  --resource-group fridge-rg
  --fridge-ns fridge-ns
  --name RootManageSharedAccessKey
  --query primaryConnectionString --output tsv)
```
- C)
- ```
New-AzureRmServiceBusQueue
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Name fridge-q
  -EnablePartitioning $False
```
- D)
- ```
New-AzureRmServiceBusNamespace
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Location fridge-loc
```

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

**Answer: C**

### NEW QUESTION 3

- (Topic 8)

You develop and deploy a web application to Azure App Service. The application accesses data stored in an Azure Storage account. The account contains several containers with several blobs with large amounts of data. You deploy all Azure resources to a single region. You need to move the Azure Storage account to the new region. You must copy all data to the new region. What should you do first?

- A. Export the Azure Storage account Azure Resource Manager template
- B. Initiate a storage account failover
- C. Configure object replication for all blobs
- D. Use the AzCopy command line tool
- E. Create a new Azure Storage account in the current region
- F. Create a new subscription in the current region

**Answer: A**

#### Explanation:

To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice and finally, delete the resources in the source region.

To get started, export, and then modify a Resource Manager template.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

**NEW QUESTION 4**

- (Topic 8)

An organization hosts web apps in Azure. The organization uses Azure Monitor. You discover that configuration changes were made to some of the web apps. You need to identify the configuration changes. Which Azure Monitor log should you review?

- A. AppServiceEnvironmentPlatformLogs
- B. AppServiceApplogs
- C. AppServiceAuditLogs
- D. AppServiceConsoleLogs

**Answer: C**

**NEW QUESTION 5**

HOTSPOT - (Topic 8)

You provisioned an Azure Cosmos DB for NoSQL account named account1 with the default consistency level.

You plan to configure the consistency level on a per request basis. The level needs to be set for consistent prefix for read and write operations to account1.

You need to identify the resulting consistency level for read and write operations. Which levels should you configure? To answer, select the appropriate options in the

answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Operation type	Resulting consistency level
Read operations	<input type="text"/> strong session consistent prefix
Write operations	<input type="text"/> strong session consistent prefix

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

**Answer Area**

Operation type	Resulting consistency level
Read operations	<input type="text"/> strong session consistent prefix
Write operations	<input type="text"/> strong session consistent prefix

**NEW QUESTION 6**

DRAG DROP - (Topic 8)

Contoso, Ltd. provides an API to customers by using Azure API Management (APIM). The API authorizes users with a JWT token.

You must implement response caching for the APIM gateway. The caching mechanism must detect the user ID of the client that accesses data for a given location and cache the response for that user ID.

You need to add the following policies to the policies file:

- a set-variable policy to store the detected user identity
- a cache-lookup-value policy
- a cache-store-value policy

• a find-and-replace policy to update the response body with the user profile information To which policy section should you add the policies? To answer, drag the appropriate sections to the correct policies. Each section may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content

NOTE: Each correct selection is worth one point

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Inbound.

A set-variable policy to store the detected user identity. Example:

```
<policies>
<inbound>
<!-- How you determine user identity is application dependent -->
<set-variable name="enduserid"
value="@ (context.Request.Headers.GetValueOrDefault("Authorization","").Split(' ')[1].AsJwt()?.Subject)" />
Box 2: Inbound
A cache-lookup-value policy Example:
```

Box 2: Inbound

A cache-lookup-value policy Example:

```
<inbound>
<base />
<cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false" downstream-caching-type="none | private | public" must-revalidate="true | false">
<vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several times -->
</cache-lookup>
</inbound>
```

Box 3: Outbound

A cache-store-value policy. Example:

```
<outbound>
<base />
<cache-store duration="3600" />
</outbound>
```

Box 4: Outbound

A find-and-replace policy to update the response body with the user profile information. Example:

```
<outbound>
<!-- Update response body with user profile-->
<find-and-replace from="$userprofile$"
to="@ ((string)context.Variables["userprofile"])" />
<base />
</outbound>
```

**NEW QUESTION 7**

- (Topic 8)

You are building a web application that performs image analysis on user photos and returns metadata containing objects identified. The image is very costly in terms of time and compute resources. You are planning to use Azure Redis Cache so duplicate uploads do not need to be reprocessed.

In case of an Azure data center outage, metadata loss must be kept to a minimum. You need to configure the Azure Redis cache instance.

Which two actions should you perform?

- A. Configure Azure Redis with rob persistence
- B. Configure second storage account far persistence.
- C. Set backup frequency to the minimum value.
- D. Configure Azure Redis with AOF persistence

**Answer: BC**

**NEW QUESTION 8**

DRAG DROP - (Topic 8)

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

- 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 9**

DRAG DROP - (Topic 8)

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		
	API	Permission	Type
<input type="text" value="client_id"/>	Azure Storage	<input type="text" value="Setting"/>	<input type="text" value="Setting"/>
<input type="text" value="profile"/>	Microsoft Graph	User.Read	<input type="text" value="Setting"/>
<input type="text" value="delegated"/>			
<input type="text" value="application"/>			
<input type="text" value="user_impersonation"/>			

- 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

**NEW QUESTION 10**

HOTSPOT - (Topic 8)

You have an Azure Web app that uses Cosmos DB as a data store. You create a CosmosDB container by running the following PowerShell script:

```
$resourceGroupName = "testResourceGroup"
$accountName = "testCosmosAccount"
$databaseName = "testDatabase"
$containerName = "testContainer"
$partitionKeyPath = "/EmployeeId"
$autoscaleMaxThroughput = 5000 New-AzCosmosDBSqlContainer
-ResourceGroupName $resourceGroupName
-AccountName $accountName
-DatabaseName $databaseName
-Name $containerName
-PartitionKeyKind Hash
-PartitionKeyPath $partitionKeyPath
-AutoscaleMaxThroughput $autoscaleMaxThroughput
```

You create the following queries that target the container:

```
SELECT * FROM c WHERE c.EmployeeId > '12345' SELECT * FROM c WHERE c.UserID = '12345'
```

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

NOTE: Each correct selection is worth one point.

	<b>Yes</b>	<b>No</b>
The minimum throughput for the container is 400 R/Us.	<input type="radio"/>	<input type="radio"/>
The first query statement is an in-partition query.	<input type="radio"/>	<input type="radio"/>
The second query statement is a cross-partition query.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: No

You set the highest, or maximum RU/s T<sub>max</sub> you don't want the system to exceed. The system automatically scales the throughput T such that  $0.1 * T_{max} \leq T \leq T_{max}$ .

In this example we have autoscaleMaxThroughput = 5000, so the minimum throughput for the container is 500 R/Us.

Box 2: No

First query: SELECT \* FROM c WHERE c.EmployeeId > '12345'

Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition. In order to be an in-partition query, the query must have an equality filter that includes the partition key:

```
SELECT * FROM c WHERE c.DeviceId = 'XMS-0001'
```

Box 3: Yes

Example of In-partition query:

Consider the below query with an equality filter on DeviceId. If we run this query on a container partitioned on DeviceId, this query will filter to a single physical partition.

```
SELECT * FROM c WHERE c.DeviceId = 'XMS-0001'
```

**NEW QUESTION 10**

- (Topic 8)

You are developing an Azure messaging solution.

You need to ensure that the solution that meets the following requirements:

- Provide transactional support
- Provide duplicate detection.
- Store the messages for an unlimited period of time

Which two technologies will meet the requirements? Each correct answer presents a complete solution NOTE Each correct selection is worth one point.

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

**Answer:** AC

**Explanation:**

The Azure Service Bus Queue and Topic has duplicate detection.

Enabling duplicate detection helps keep track of the application-controlled MessageId of all messages sent into a queue or topic during a specified time window.

Reference:

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

**NEW QUESTION 12**

- (Topic 8)

You develop a REST API. You implement a user delegation SAS token to communicate with Azure Blob storage.

The token is compromised. You need to revoke the token.

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. Revoke the delegation keys
- B. Delete the stored access policy.
- C. Regenerate the account key.
- D. Remove the role assignment for the security principle.

**Answer:** AB

**Explanation:**

A: Revoke a user delegation SAS

To revoke a user delegation SAS from the Azure CLI, call the `az storage account revoke- delegation-keys` command. This command revokes all of the user delegation keys associated with the specified storage account. Any shared access signatures associated with those keys are invalidated.

B: To revoke a stored access policy, you can either delete it, or rename it by changing the signed identifier.

Changing the signed identifier breaks the associations between any existing signatures and the stored access policy. Deleting or renaming the stored access policy immediately effects all of the shared access signatures associated with it. D18912E1457D5D1DDCBD40AB3BF70D5D

Reference:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/storage/blobs/storage-blob-user-delegationsas-create-cli.md>

<https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy#modifying-or-revoking-a-stored-access-policy>

**NEW QUESTION 14**

- (Topic 8)

You are developing a web application that runs as an Azure Web App. The web application stores data in Azure SQL Database and stores files in an Azure Storage account. The web application makes HTTP requests to external services as part of normal operations.

The web application is instrumented with Application Insights. The external services are OpenTelemetry compliant.

You need to ensure that the customer ID of the signed in user is associated with all operations throughout the overall system.

What should you do?

- A. Create a new SpanContext with the TraceRags value set to the customer ID for the signed in user.
- B. On the current SpanContext, set the Traceld to the customer ID for the signed in user.
- C. Add the customer ID for the signed in user to the CorrelationContext in the web application.
- D. Set the header Ocp-Apim-Trace to the customer ID for the signed in user.

**Answer:** C

**Explanation:**

Reference:

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

**NEW QUESTION 17**

- (Topic 8)

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 18**

- (Topic 8)

You are developing a solution that will use a multi-partitioned Azure Cosmos DB database. You plan to use the latest Azure Cosmos DB SDK for development.

The solution must meet the following requirements:

? Send insert and update operations to an Azure Blob storage account.

? Process changes to all partitions immediately.

? Allow parallelization of change processing.

You need to process the Azure Cosmos DB operations.

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 an Azure App Service API and implement the change feed estimator of the SD
- B. Scale the API by using multiple Azure App Service instances.
- C. Create a background job in an Azure Kubernetes Service and implement the change feed feature of the SDK.
- D. Create an Azure Function to use a trigger for Azure Cosmos D
- E. Configure the trigger to connect to the container.
- F. Create an Azure Function that uses a FeedIterator object that processes the change feed by using the pull model on the container

G. Use a FeedRange object to parallelize the processing of the change feed across multiple functions.

**Answer:** CD

**Explanation:**

Azure Functions is the simplest option if you are just getting started using the change feed. Due to its simplicity, it is also the recommended option for most change feed use cases. When you create an Azure Functions trigger for Azure Cosmos DB, you select the container to connect, and the Azure Function gets triggered whenever there is a change in the container. Because Azure Functions uses the change feed processor behind the scenes, it automatically parallelizes change processing across your container's partitions.

Note: You can work with change feed using the following options:

? Using change feed with Azure Functions

? Using change feed with change feed processor

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed>

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model> <https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed#azure-functions>

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model#using-feedrange-for-parallelization>

**NEW QUESTION 23**

HOTSPOT - (Topic 8)

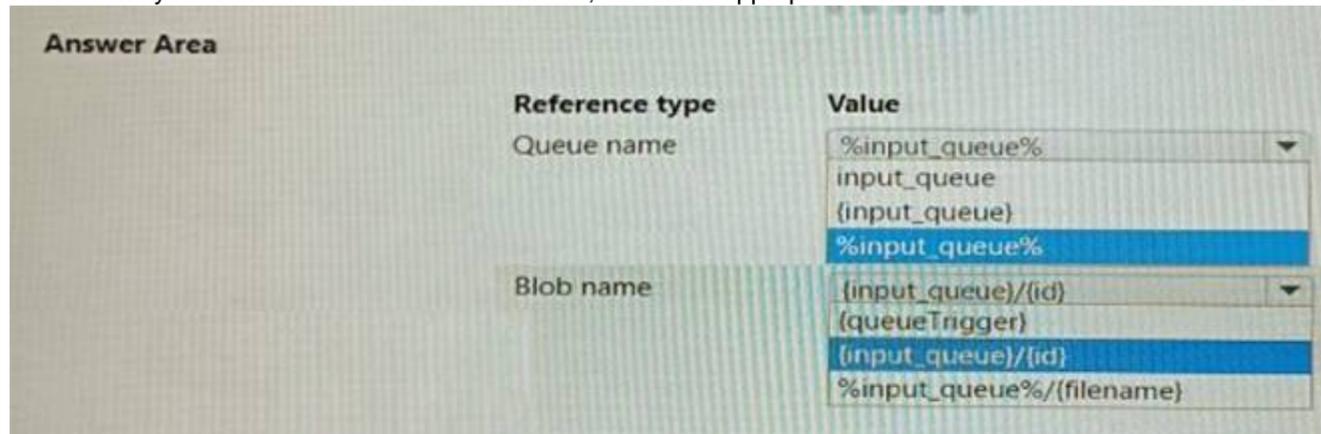
You plan to implement an Azure Functions app.

The Azure Functions app has the following requirements:

- Must be triggered by a message placed in an Azure Storage queue.
- Must use the queue name set by an app setting named input-queue.
- Must create an Azure Blob Storage named the same as the content of the message.

You need to identify how to reference the queue and blob name in the function. Just file of the Azure Functions app.

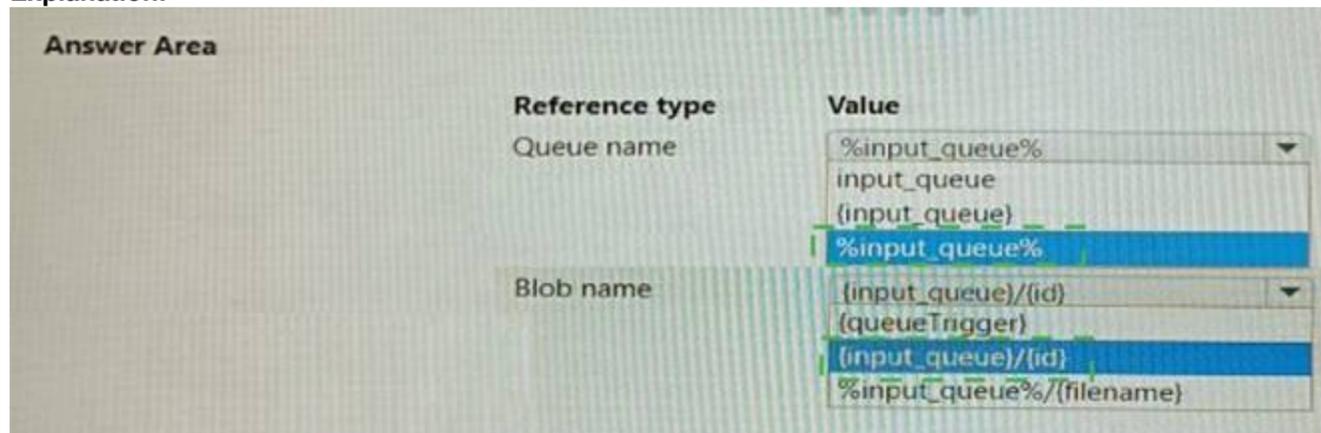
How should you reference the names? To answer, select the appropriate values in the answer area. NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 24**

- (Topic 8)

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

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

**Answer:** C

**NEW QUESTION 26**

DRAG DROP - (Topic 8)

You develop and deploy an Azure App Service ---- app. The web app accesses data in an Azure SQL database

You must update the web app to store frequently used data m a new Azure Cache for Redis Premium instance.

You need to implement the Azure Cache for Redis features.

Which feature should you implement? To answer, drag the appropriate feature to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to scroll between panes or scroll to view content.  
 NOTE: Each correct selection is worth one point.

Features	Requirement	Feature
horizontal partitionin	Create a data structure for storing collections of related items.	
channel	Create a data structure for storing the most recently accessed cache items.	
list	Send messages through a high-performance publisher/subscriber mechanism.	
set		

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Features	Requirement	Feature
horizontal partitionin	Create a data structure for storing collections of related items.	set
channel	Create a data structure for storing the most recently accessed cache items.	list
list	Send messages through a high-performance publisher/subscriber mechanism.	channel
set		

**NEW QUESTION 28**

- (Topic 8)

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

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

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

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

Solution:

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

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: A**

**Explanation:**

To configure Manifest to include Group Claims in Auth Token

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

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

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

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

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

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

References:

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

**NEW QUESTION 30**

- (Topic 8)

You are developing a mobile instant messaging app for a company. The mobile app must meet the following requirements:

- Support offline data sync.
- Update the latest messages during normal sync cycles. You need to implement Offline Data Sync.

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

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

**Answer: BE**

**Explanation:**

B: Incremental Sync: the first parameter to the pull operation is a query name that is used only on the client. If you use a non-null query name, the Azure Mobile SDK performs an incremental sync. Each time a pull operation returns a set of results, the latest updatedAt timestamp from that result set is stored in the SDK local system tables. Subsequent pull operations retrieve only records after that timestamp.

E (not D): To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. However, since the SDK adds its own sort on the updatedAt field, you cannot use a pull query that has its own orderBy clause.

References:

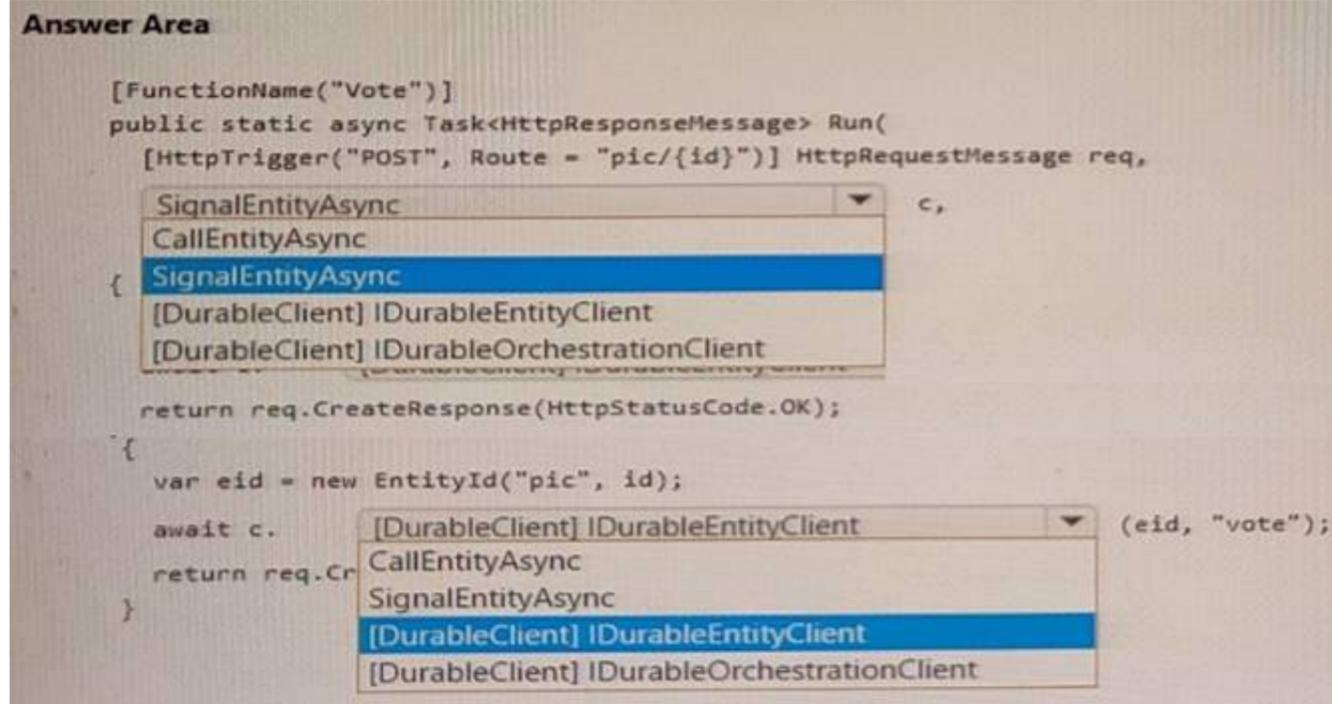
<https://docs.microsoft.com/en-us/azure/app-service-mobile/app-service-mobile-offline-data-sync>

**NEW QUESTION 32**

HOTSPOT - (Topic 8)

You are developing an online game that allows players to vote for their favorite photo that illustrates a word. The game is built by using Azure Functions and uses durable entities to track the vote count

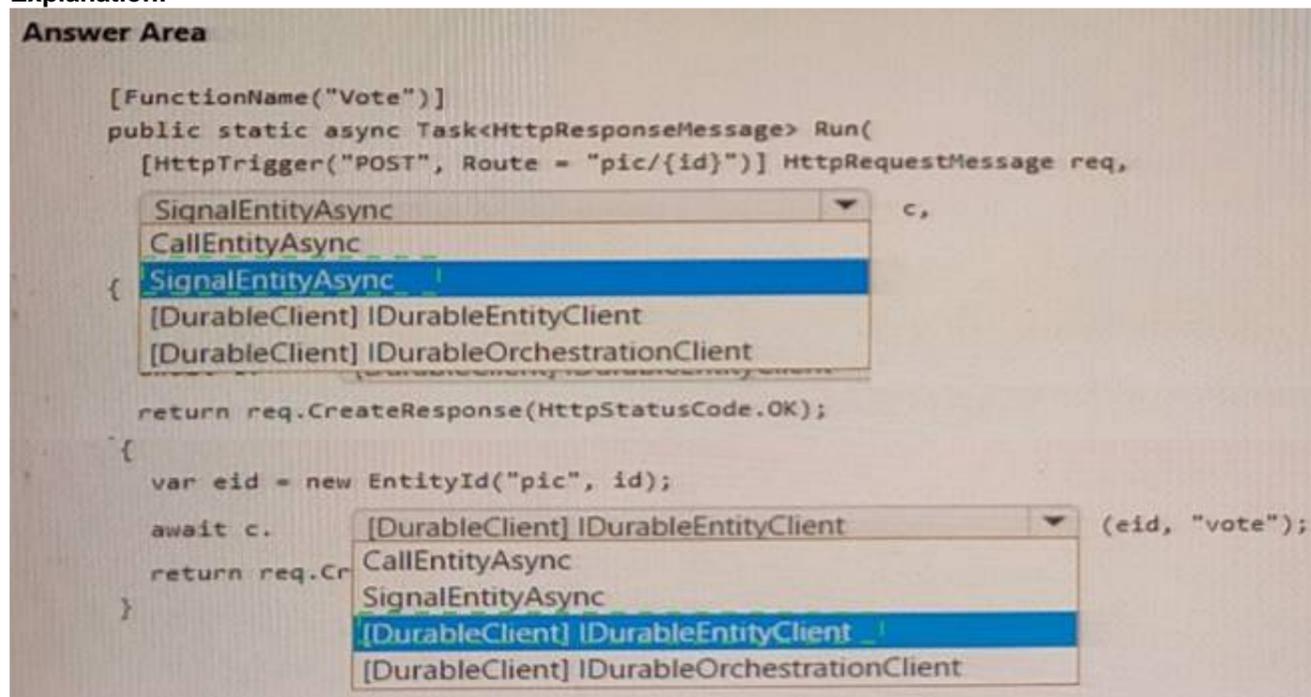
The voting window is 30 seconds. You must minimize latency. You need to implement the Azure Function for voting. How should you complete the code? To answer, select the appropriate options in the answer area.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 36

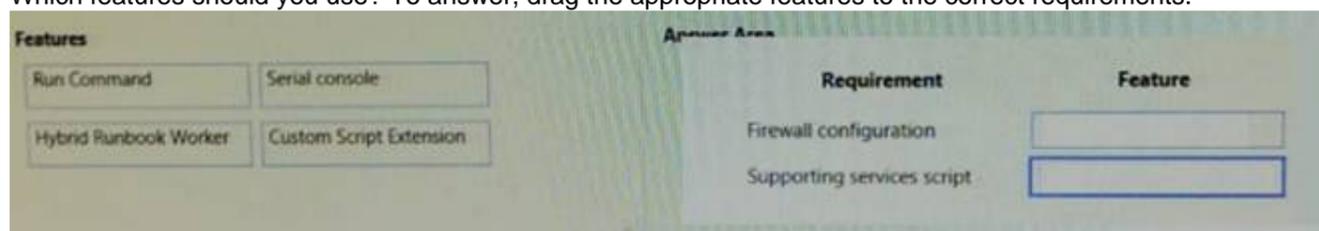
DRAG DROP - (Topic 8)

You are preparing to deploy an Azure virtual machine (VM) based application. The VMs that run the application have the following requirements:

- When a VM is provisioned the firewall must be automatically configured before it can access Azure resources.
- Supporting services must be installed by using an Azure PowerShell script that is stored in Azure Storage

You need to ensure that the requirements are met.

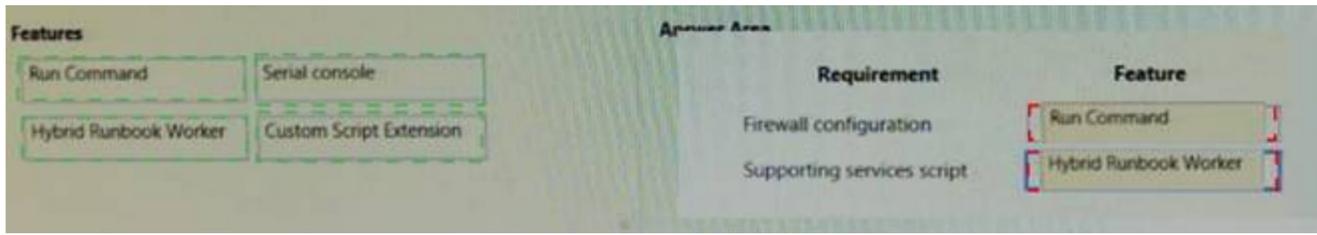
Which features should you use? To answer, drag the appropriate features to the correct requirements.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



#### NEW QUESTION 41

- (Topic 8)

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

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

You are developing an Azure 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 Service Bus Queue from the mobile application. Create an Azure Function App that uses an Azure Service Bus Queue trigger.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: A**

#### Explanation:

You can create a function that is triggered when messages are submitted to an Azure Storage queue.

Reference:

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

#### NEW QUESTION 43

- (Topic 8)

You are developing an Azure-based web application. The application goes offline periodically to perform offline data processing. While the application is offline, numerous Azure Monitor alerts fire which result in the on-call developer being paged.

The application must always log when the application is offline for any reason.

You need to ensure that the on-call developer is not paged during offline processing. What should you do?

- A. Add Azure Monitor alert processing rules to suppress notifications.
- B. Create an Azure Monitor Metric Alert.
- C. Build an Azure Monitor action group that suppresses the alerts.
- D. Disable Azure Monitor Service Health Alerts during offline processing.

**Answer: C**

#### NEW QUESTION 44

- (Topic 8)

You are developing an e-commerce solution that uses a microservice architecture.

You need to design a communication backplane for communicating transactional messages between various parts of the solution. Messages must be communicated in first-in-first-out (FIFO) order.

What should you use?

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

**Answer: C**

#### Explanation:

As a solution architect/developer, you should consider using Service Bus queues when:

- ? Your solution requires the queue to provide a guaranteed first-in-first-out (FIFO) ordered delivery.

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted>

#### NEW QUESTION 46

HOTSPOT - (Topic 8)

You are developing an application that uses Azure Storage Queues. You have the following code:

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

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

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

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

Statement	Yes	No
The code configures the lock duration for the queue.	<input type="radio"/>	<input type="radio"/>
The last message read remains in the queue after the code runs.	<input type="radio"/>	<input type="radio"/>
The storage queue remains in the storage account after the code runs.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: No

The QueueDescription.LockDuration property gets or sets the duration of a peek lock; that is, the amount of time that the message is locked for other receivers. The maximum value for LockDuration is 5 minutes; the default value is 1 minute.

Box 2: Yes

You can peek at the message in the front of a queue without removing it from the queue by calling the PeekMessage method.

Box 3: Yes

**NEW QUESTION 49**

- (Topic 8)

You need to design network connectivity for a subnet in an Azure virtual network. The subnet will contain 30 virtual machines. The virtual machines will establish outbound connections to internet hosts by using the same a pool of four public IP addresses, inbound connections to the virtual machines will be prevented. What should include in the design?

- A. Azure Private Link
- B. NAT Gateway
- C. User Defined Routes
- D. Azure Virtual WAN

Answer: D

**NEW QUESTION 50**

- (Topic 8)

You develop an ASP.NET Core app that uses Azure App Configuration. You also create an App Configuration containing 100 settings. The app must meet the following requirements:

- Ensure the consistency of all configuration data when changes to individual settings occur.
- Handle configuration data changes dynamically without causing the application to restart.
- Reduce the overall number of requests made to App Configuration APIs.

You must implement dynamic configuration updates in the app.

What are two ways to achieve this goal? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Increase the App Configuration cache expiration from the default value.
- B. Create and implement environment variables for each App Configuration store setting.
- C. Decrease the App Configuration cache expiration from the default value.
- D. Register all keys in the App Configuration stor
- E. Set the refreshAll parameter of the Register method to false.
- F. Create and register a sentinel key in the App Configuration stor
- G. Set the refreshAll parameter of the Register method to true.
- H. Create and configure Azure Key Vault
- I. Implement the Azure Key Vault configuration provider.

Answer: AE

**NEW QUESTION 51**

- (Topic 8)

You are designing a multi-tiered application that will be hosted on Azure virtual machines. The virtual machines will run Windows Server. Front-end servers will be accessible from the Internet over port 443. The other servers will NOT be directly accessible over the internet

You need to recommend a solution to manage the virtual machines that meets the following requirement

- Allows the virtual machine to be administered by using Remote Desktop.
- Minimizes the exposure of the virtual machines on the Internet Which Azure service should you recommend?

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

**Answer: C**

**NEW QUESTION 53**

HOTSPOT - (Topic 8)

A company is developing a gaming platform. Users can join teams to play online and see leaderboards that include player statistics. The solution includes an entity named Team.

You plan to implement an Azure Redis Cache instance to improve the efficiency of data operations for entities that rarely change.

You need to invalidate the cache when team data is changed.

How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
void ClearCachedTeams()
{
    
    

    ViewBag.nsg += Team data removed from cache. ";
}
```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: IDatabase cache = connection.GetDatabase();

Connection refers to a previously configured ConnectionMultiplexer.

Box 2: cache.StringSet("teams", "")

To specify the expiration of an item in the cache, use the TimeSpan parameter of StringSet.

cache.StringSet("key1", "value1", TimeSpan.FromMinutes(90));

References:

<https://azure.microsoft.com/sv-se/blog/lap-around-azure-redis-cache-preview/>

**NEW QUESTION 57**

DRAG DROP - (Topic 8)

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		
Flow	Manage all code segments from the same DevOps environment.	<input type="text"/>

- 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 62**

- (Topic 8)

You deploy an API to API Management

You must secure all operations on the API by using a client certificate.

You need to secure access to the backend service of the API by using client certificates. Which two security features can you use?

- A. Azure AD token
- B. Self-signed certificate
- C. Certificate Authority (CA) certificate
- D. Triple DES (3DES) cipher
- E. Subscription key

**Answer:** BC

**NEW QUESTION 67**

- (Topic 8)

You deploy an Azure App Service web app. You create an app registration for the app in Azure Active Directory (Azure AD) and Twitter. the app must authenticate users and must use SSL for all communications. The app must use Twitter as the identity provider. You need to validate the Azure AD request in the app code.

What should you validate?

- A. HTTP response code
- B. ID token header
- C. ID token signature
- D. Tenant ID

**Answer:** B

**NEW QUESTION 71**

HOTSPOT - (Topic 8)

You develop several Azure Grid to include hundreds of event types, such as billing, inventory, and shipping updates.

Events must be sent to a single endpoint for the Azure Functions app to process. The events must be filtered by event type before processing. You must have authorization and authentication control to partition your tenants to receive the event data.

You need to configure Azure Event Grid.

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

NOTE: Each correct selection is worth one point.

**Answer Area**

Requirement	Configuration Value
Third-party system endpoint to send events	<input type="text" value="system topic"/> <ul style="list-style-type: none"> <li>system topic</li> <li>custom topic</li> <li>event domain</li> <li>event subscription</li> </ul>
Azure Functions app endpoint to handle filtered events	<input type="text" value="event domain"/> <ul style="list-style-type: none"> <li>system topic</li> <li>custom topic</li> <li>event domain</li> <li>event subscription</li> </ul>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

Requirement	Configuration Value
Third-party system endpoint to send events	<ul style="list-style-type: none"> <li>system topic</li> <li>system topic</li> <li>custom topic</li> <li>event domain</li> <li>event subscription</li> </ul>
Azure Functions app endpoint to handle filtered events	<ul style="list-style-type: none"> <li>event domain</li> <li>system topic</li> <li>custom topic</li> <li>event domain</li> <li>event subscription</li> </ul>

**NEW QUESTION 76**

HOTSPOT - (Topic 8)

All functions in the app meet the following requirements:

- Run until either a successful run or until 10 run attempts occur.
- Ensure that there are at least 20 seconds between attempts for up to 15 minutes. You need to configure the hostjson file.

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**

```

{
  "retry": {
    "strategy": "exponentialBackoff",
    "maxRetryCount": 10,
    "healthMonitor": "singleton"
  }
}

```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

```

{
  "retry": {
    "strategy": "exponentialBackoff",
    "maxRetryCount": 10,
    "healthMonitor": "singleton"
  }
}

```

**NEW QUESTION 80**

HOTSPOT - (Topic 8)

You are developing an application that uses a premium block blob storage account. You are optimizing costs by automating Azure Blob Storage access tiers. You apply the following policy rules to the storage account. You must determine the implications of applying the rules to the data. (Line numbers are included for reference only.)

```

01 {
02   "rules":
03     {
04       "name": "agingDataRule",
05       "enabled": true,
06       "type": "Lifecycle",

```

**Answer Area**

	Yes	No
Block blobs prefixed with container1/salesorders or container2/inventory which have not been modified in over 60 days are moved to cool storage. Blobs that have not been modified in 120 days are moved to the archive tier.	<input type="radio"/>	<input type="radio"/>
Blobs are moved to cool storage if they have not been accessed for 30 days.	<input type="radio"/>	<input checked="" type="radio"/>
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.	<input type="radio"/>	<input type="radio"/>
All block blobs older than 730 days will be deleted.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

- \* 1. Yes
- \* 2. Yes
- \* 3. Yes
- \* 4. No

<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#move-aging-data-to-a-cooler-tier>

**NEW QUESTION 84**

- (Topic 8)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:

- Share session state across all ASP.NET web applications
- Support controlled, concurrent access to the same session state data for multiple readers and a single writer
- Save full HTTP responses for concurrent requests You need to store the information.

Proposed Solution: 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 85**

- (Topic 8)

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

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

You are developing an Azure 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 Service Bus Queue from the mobile application. Create an Azure Windows VM that is triggered from Azure Service Bus Queue.

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 88**

- (Topic 8)

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

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

**Answer:** ABC

#### NEW QUESTION 92

- (Topic 8)

You are developing a software solution for an autonomous transportation system. The solution uses large data sets and Azure Batch processing to simulate navigation sets for entire fleets of vehicles.

You need to create compute nodes for the solution on Azure Batch. What should you do?

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

**Answer:** D

#### Explanation:

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

Note:

Step 1: Create a pool of compute nodes. When you create a pool, you specify the number of compute nodes for the pool, their size, and the operating system. When each task in your job runs, it's assigned to execute on one of the nodes in your pool.

Step 2 : Create a job. A job manages a collection of tasks. You associate each job to a specific pool where that job's tasks will run.

Step 3: Add tasks to the job. Each task runs the application or script that you uploaded to process the data files it downloads from your Storage account. As each task completes, it can upload its output to Azure Storage.

#### NEW QUESTION 97

- (Topic 8)

You are creating an app that will use CosmosDB for data storage. The app will process batches of relational data.

You need to select an API for the app. Which API should you use?

- A. MongoDBAPI
- B. Table API
- C. SQL API
- D. Cassandra API

**Answer:** C

#### Explanation:

For relational data you will need the SQL API

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>

#### NEW QUESTION 102

- (Topic 8)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:

- Share session state across all ASP.NET web applications
- Support controlled, concurrent access to the same session state data for multiple readers and a single writer
- Save full HTTP responses for concurrent requests You need to store the information.

Proposed Solution: Deploy and configure an Azure Database for PostgreSQL. Update the web applications.

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:

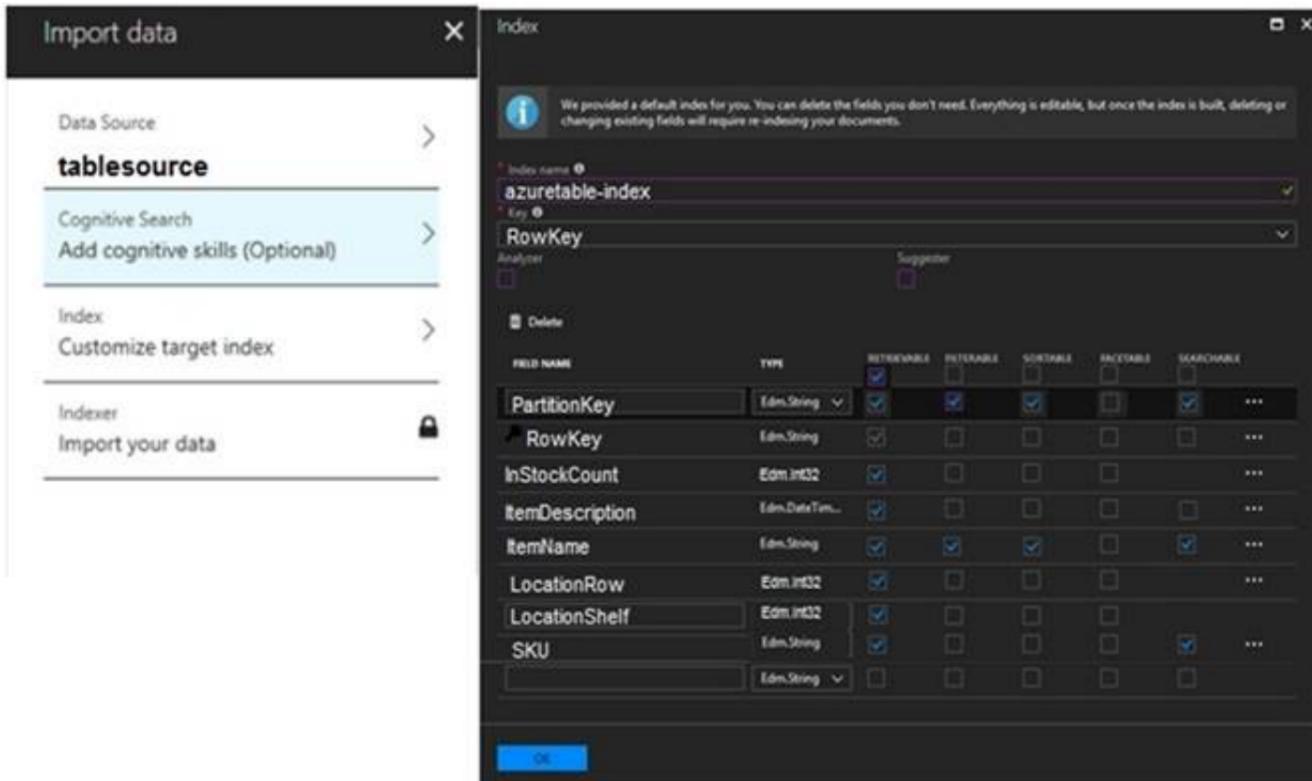
<https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching#managing-concurrency-in-a-cache>

#### NEW QUESTION 105

HOTSPOT - (Topic 8)

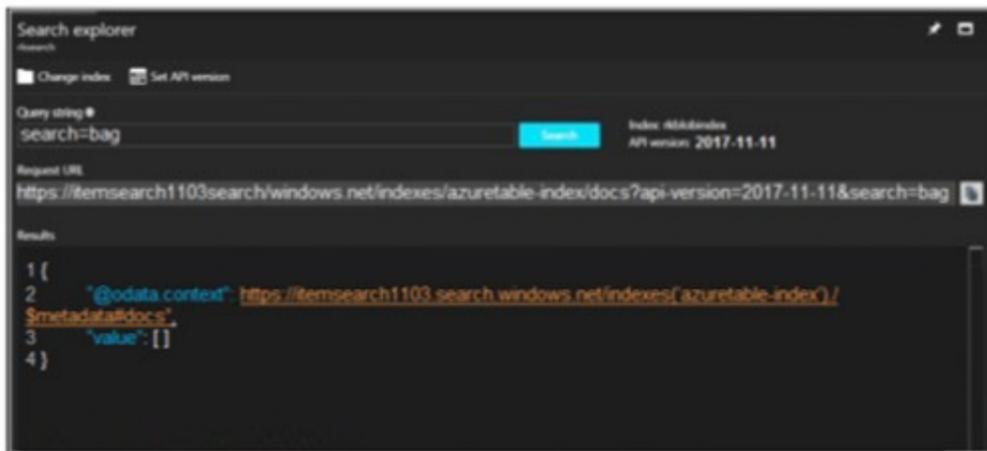
You are validating the configuration of an Azure Search indexer.

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



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

Category	RowKey	Timestamp	InStockCount	ItemDescription	ItemName	LocationRow	LocationShelf	SKU
Food	3	2018-08-25T15:47:26.135Z	12	A box of chocolate candy bars	Choco-bar	3	123421	
Hardware	2	2018-08-25T15:46:08.409Z	2	A bag of bolts	Bolts	1	4	678944
Hardware	1	2018-08-25T15:46:41.403Z	23	A box of nails	Nails	2	1	654365



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

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Yes

The ItemDescription field is not searchable.

Box 2: No

The ItemDescription field is not searchable, but we would need to recreate the index.

Box 3: Yes

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

Box 4: No References:

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

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

**NEW QUESTION 106**

DRAG DROP - (Topic 8)

You develop and deploy several APIs to Azure API Management. You create the following policy fragment named APICounts:

```
<fragment>
  <emit-metric value="1" namespace="custom-metrics">
    <dimension name="User ID" />
    <dimension name="Operation ID" />
    <dimension name="API ID" />
    <dimension name="Client IP" value="@(<context.Request.IpAddress>)" />
  </emit-metric>
</fragment>
```

The policy fragment must be reused across various scopes and APIs. The policy fragment must be applied to all APIs and run when a calling system invokes any API.

You need to implement the policy fragment.

How should you complete the policy segment? To answer, drag the appropriate XML elements to the correct targets. Each XML element 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.

XML elements	Answer Area
name	<pre>&lt;policies&gt;   &lt; [ ] &gt;     &lt; [ ] = "APICounts" /&gt;     &lt;base /&gt;   &lt;/ [ ] &gt;   . . . &lt;/policies&gt;</pre>
inbound	
outbound	
set-variable	
fragment-id	
include-fragment	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

<https://learn.microsoft.com/en-us/azure/api-management/include-fragment-policy>

**NEW QUESTION 108**

- (Topic 8)

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

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

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

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

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

Solution: Use the Azure Blob Storage change feed to trigger photo processing. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

The change feed is a log of changes that are organized into hourly segments but appended to and updated every few minutes. These segments are created only when there are blob change events that occur in that hour.

Instead catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.

Reference:

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

**NEW QUESTION 112**

HOTSPOT - (Topic 8)

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

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

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

How should you complete the commands? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup
az [ ] create --location centralus --name $resourcegroupname
group create --name $webappname --resource-group $resourcegroupname
webapp - -sku S3
appservice plan create --name $webappname --resource-group $resourcegroupname
webapp deployment slot \ - -plan $webappname
webapp deployment source create --name $webappname --resource-group $resourcegroupname
\ - -slot staging
az [ ] config --name $webappname --resource-group $resourcegroupname
group \ - -slot staging - -repo-url
webapp $gitrepo - -branch master - -manual-integration
appservice plan
webapp deployment slot
webapp deployment source
az [ ]
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
az [ ]
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
az [ ]
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
az [ ]
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
```

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: group  
 # Create a resource group.  
 az group create --location westeurope --name myResourceGroup  
 Box 2: appservice plan  
 # Create an App Service plan in STANDARD tier (minimum required by deployment slots). az appservice plan create --name \$webappname --resource-group myResourceGroup -- sku S1  
 Box 3: webapp  
 # Create a web app.  
 az webapp create --name \$webappname --resource-group myResourceGroup \ --plan \$webappname  
 Box 4: webapp deployment slot  
 #Create a deployment slot with the name "staging".  
 az webapp deployment slot create --name \$webappname --resource-group myResourceGroup \ --slot staging  
 Box 5: webapp deployment source  
 # Deploy sample code to "staging" slot from GitHub.  
 az webapp deployment source config --name \$webappname --resource-group myResourceGroup \ --slot staging --repo-url \$gitrepo --branch master --manual-integration  
 References:  
<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment>

**NEW QUESTION 115**

- (Topic 8)  
 You are building a B2B web application that uses Azure B2B collaboration for authentication. Paying customers authenticate to Azure B2B using federation. The application allows users to sign up for trial accounts using any email address. When a user converts to a paying customer, the data associated with the trial should be kept, but the user must authenticate using federation. You need to update the user in Azure Active Directory (Azure AD) when they convert to a paying customer. Which Graph API parameter is used to change authentication from one-time password to federation?

- A. usrFlowType
- B. Status
- C. invittdUstr
- D. resetRedemption

Answer: B

**NEW QUESTION 117**

- (Topic 8)  
 Your company has several containers based on the following operating systems:

- Windows Server 2019 Nano Server
- Windows Server 2019 Server Core
- Windows Server 2022 Nano Server
- Windows Server 2022 Server Core
- Linux

You plan to migrate the containers to an Azure Kubernetes cluster. What is the minimum number of node pools that the cluster must have?

- A. 1
- B. 2
- C. 3
- D. 6

**Answer: C**

**NEW QUESTION 122**

HOTSPOT - (Topic 8)

You plan to deploy a new application to a Linux virtual machine (VM) that is hosted in Azure.

The entire VM must be secured at rest by using industry-standard encryption technology to address organizational security and compliance requirements.

You need to configure Azure Disk Encryption for the VM.

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

NOTE: Each correct selection is worth one point.

**Answer Area**

```

az provider register -n Microsoft.KeyVault
resourcegroup="myResourceGroup"
az group create --name $resourcegroup --location westus
keyvault_name=myvaultname$RANDOM
az 
vm  \
keyvault  \
keyvault key  \
vm encryption 
az 
vm  \
keyvault  \
keyvault key  \
vm encryption 
az 
vm  \
keyvault  \
keyvault key  \
vm encryption 
az 
vm  \
keyvault  \
keyvault key  \
vm encryption 
--volume-type 

```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: keyvault

Create an Azure Key Vault with az keyvault create and enable the Key Vault for use with disk encryption. Specify a unique Key Vault name for keyvault\_name as follows:

```

keyvault_name=myvaultname$RANDOM az keyvault create \
--name $keyvault_name \
--resource-group $resourcegroup \
--location eastus \
--enabled-for-disk-encryption True

```

Box 2: keyvault key

The Azure platform needs to be granted access to request the cryptographic keys when the VM boots to decrypt the virtual disks. Create a cryptographic key in your Key Vault with az keyvault key create. The following example creates a key named myKey:

```

az keyvault key create \
--vault-name $keyvault_name \
--name myKey \
--protection software

```

Box 3: vm

Create a VM with az vm create. Only certain marketplace images support disk encryption. The following example creates a VM named myVM using an Ubuntu

```
16.04 LTS image:
az vm create \
--resource-group $resourcegroup \
--name myVM \
--image Canonical:UbuntuServer:16.04-LTS:latest \
--admin-username azureuser \
--generate-ssh-keys \
Box 4: vm encryption
```

Encrypt your VM with az vm encryption enable:

```
az vm encryption enable \
--resource-group $resourcegroup \
--name myVM \
--disk-encryption-keyvault $keyvault_name \
--key-encryption-key myKey \
--volume-type all
```

Note: seems to an error in the question. Should have enable instead of create. Box 5: all Encrypt both data and operating system.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/virtual-machines/linux/encrypt-disks>

### NEW QUESTION 124

- (Topic 8)

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

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

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

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

You need to configure authorization. Solution:

? 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 129

HOTSPOT - (Topic 8)

You are developing an Azure Function App by using Visual Studio. The app will process orders input by an Azure Web App. The web app places the order information into Azure Queue Storage.

You need to review the Azure Function App code shown below.

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

NOTE:Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: No

ExpirationTime - The time that the message expires.

InsertionTime - The time that the message was added to the queue.

Box 2: Yes

maxDequeueCount - The number of times to try processing a message before moving it to the poison queue. Default value is 5.

Box 3: Yes

When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them. By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.

Box 4: Yes References:

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

**NEW QUESTION 134**

DRAG DROP - (Topic 8)

A company has multiple warehouse. Each warehouse contains IoT temperature devices which deliver temperature data to an Azure Service Bus queue.

You need to send email alerts to facility supervisors immediately if the temperature at a warehouse goes above or below specified threshold temperatures.

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

**Actions**

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

**Answer Area**

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Create a blank Logic app. Create and configure a Logic App.

Step 2: Add a logical app trigger that fires when one or more messages arrive in the queue. Configure the logic app trigger.

Under Triggers, select When one or more messages arrive in a queue (auto-complete). Step 3: Add an action that reads IoT temperature data from the Service Bus queue

Step 4: Add a condition that compares the temperature against the upper and lower thresholds.

Step 5: Add an action that sends an email to specified personnel if the temperature is outside of those thresholds

**NEW QUESTION 139**

- (Topic 8)

You have an application that includes an Azure Web app and several Azure Function apps. Application secrets including connection strings and certificates are stored in Azure Key Vault.

Secrets must not be stored in the application or application runtime environment. Changes to Azure Active Directory (Azure AD) must be minimized.

You need to design the approach to loading application secrets. What should you do?

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

**Answer: C**

**Explanation:**

Use Key Vault references for App Service and Azure Functions.

Key Vault references currently only support system-assigned managed identities. User- assigned identities cannot be used.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references>

**NEW QUESTION 144**

- (Topic 8)

Note: This question is part of a series of questions that present the same scenario.

Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You are developing an Azure 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 148**

DRAG DROP - (Topic 8)

You are implementing an order processing system. A point of sale application publishes orders to topics in an Azure Service Bus queue. The label property for the topic includes the following data:

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

The system has the following requirements for subscriptions

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

You need to implement filtering and maximize throughput while evaluating filters.

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

NOTE: Each correct selection is worth one point.

## Filter types

- SQLFilter
- CorrelationFilter
- No Filter

## Answer Area

Subscription	Filter type
FutureOrders	
HighPriorityOrders	
InternationalOrders	
HighQuantityOrders	
AllOrders	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

FutureOrders: SQLFilter  
 HighPriorityOrders: CorrelationFilter  
 CorrelationID only  
 InternationalOrders: SQLFilter  
 Country NOT USA requires an SQL Filter  
 HighQuantityOrders: SQLFilter

Need to use relational operators so an SQL Filter is needed. AllOrders: No Filter

SQL Filter: SQL Filters - A SqlFilter holds a SQL-like conditional expression that is evaluated in the broker against the arriving messages' user-defined properties and system properties. All system properties must be prefixed with sys. in the conditional expression. The SQL-language subset for filter conditions tests for the existence of properties (EXISTS), as well as for null-values (IS NULL), logical NOT/AND/OR, relational operators, simple numeric arithmetic, and simple text pattern matching with LIKE.

Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties. A common use is to match against the CorrelationId property, but the application can also choose to match against ContentType, Label, MessageId, ReplyTo, ReplyToSessionId, SessionId, To, and any user-defined properties. A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter. For string expressions, the comparison is case-sensitive. When specifying multiple match properties, the filter combines them as a logical AND condition, meaning for the filter to match, all conditions must match.

Boolean filters - The TrueFilter and FalseFilter either cause all arriving messages (true) or none of the arriving messages (false) to be selected for the subscription.

References:

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

**NEW QUESTION 153**

- (Topic 8)

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

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

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

You create the index in Azure Search.

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

Solution:

- \* 1 Create a SearchIndexClient object to connect to the search index
- \* 2. Create an IndexBatch that contains the documents which must be added.
- \* 3. Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

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

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

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

```
var batch = IndexBatch.Upload(hotels);
* 3. The next step is to populate the newly-created index Example:
var batch = IndexBatch.Upload(hotels);
try
{
indexClient.Documents.Index(batch);
}
}
References:
https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk
```

**NEW QUESTION 154**

DRAG DROP - (Topic 8)

You are developing an ASP.NET Core Web API web service that uses Azure Application Insights to monitor performance and track events. You need to enable logging and ensure that log messages can be correlated to events tracked by Application Insights. How should you complete the code? To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

Code segments	Answer Area
IncludeEventId	<pre>public class Startup {     . . .     public void ConfigureServices (IServiceCollection services)     {         services.AddOptions&lt; &gt;().         Configure(o =&gt; o. = true );         services.AddMvc();     }     public void Configure (IApplicationBuilder app,         IHostingEnvironment env, ILoggerFactory loggerFactory)     {         loggerFactory.AddApplicationInsights(app. ,LogLevel.Trace);         app.UseMvc();     } }</pre>
ServerFeatures	
LoggerFilterOptions	
ApplicationServices	
ApplicationInsightsLoggerOptions	
TrackExceptionsAsExceptionTelemetry	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: ApplicationInsightsLoggerOptions  
 If you want to include the EventId and EventName properties, then add the following to the ConfigureServices method:  
 services.AddOptions<ApplicationInsightsLoggerOptions>().Configure(o => o.IncludeEventId = true);  
 Box 2: IncludeEventID  
 Box 3: ApplicationServices  
 In Asp.Net core apps it turns out that trace logs do not show up in Application Insights out of the box. We need to add the following code snippet to our Configure method in Startup.cs:  
 loggerFactory.AddApplicationInsights(app.ApplicationServices, logLevel);  
 References:  
<https://blog.computedcloud.com/enabling-application-insights-trace-logging-in-asp-net-core/>

**NEW QUESTION 158**

- (Topic 8)

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

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

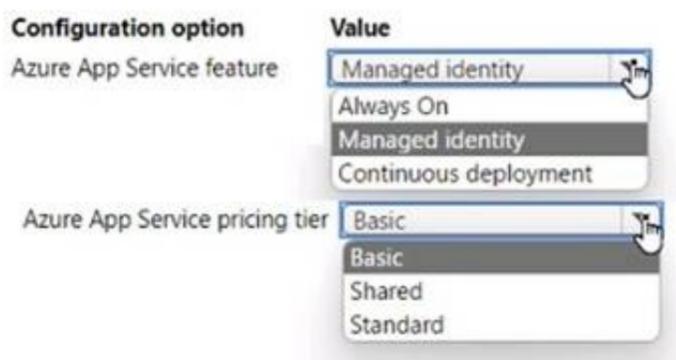
**Answer: AC**

**NEW QUESTION 159**

HOTSPOT - (Topic 8)

You have an App Service plan named aspl based on the Free pricing tier. You plan to use aspl to implement an Azure Function app with a queue trigger. Your solution must minimize cost. You need to identify the configuration options that will meet the requirements. Which value should you configure? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

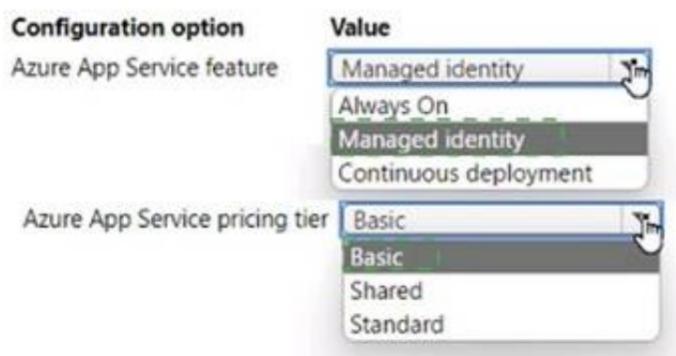


- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**



**NEW QUESTION 161**

- (Topic 8)

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

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

You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Enable auto swap for the Testing slot. Deploy the app to the Testing slot. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

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

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

Reference:

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

**NEW QUESTION 162**

DRAG DROP - (Topic 8)

You plan to create a Docker image that runs as ASP.NET Core application named ContosoApp. You have a setup script named setupScript.ps1 and a series of application files including ContosoApp.dll.

You need to create a Dockerfile document that meets the following requirements:

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

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

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

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

- 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 167**

HOTSPOT - (Topic 8)

You are using Azure Front Door Service.

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

You need to determine the root cause for the issue.

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

NOTE:Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: No

Front Door can dynamically compress content on the edge, resulting in a smaller and faster response to your clients. All files are eligible for compression.

However, a file must be of a MIME type that is eligible for compression list.

Box 2: No

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

Box 3: Yes

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

**NEW QUESTION 171**

DRAG DROP - (Topic 8)

You are maintaining an existing application that uses an Azure Blob GPv1 Premium storage account. Data older than three months is rarely used. Data newer than three months must be available immediately. Data older than a year must be saved but does not need to be available immediately. You need to configure the account to support a lifecycle management rule that moves blob data to archive storage for data not modified in the last year. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

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	⬅

- 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.

**NEW QUESTION 176**

HOTSPOT - (Topic 8)

You implement an Azure solution to include Azure Cosmos DB, the latest Azure Cosmos DB SDK, and the Azure Cosmos DB for NoSQL API. You also implement a change feed processor on a new container instance by using the Azure Functions trigger for Azure Cosmos DB.

A large batch of documents continues to fail when reading one of the documents in the batch. The same batch of documents is continuously retried by the triggered function and a new batch of documents must be read.

You need to implement the change feed processor to read the documents.

Which feature should you implement? To answer, select the appropriate features in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Feature
Read a new batch of documents while keeping track of the failing batch of documents.	<ul style="list-style-type: none"> <li>Change feed estimator</li> <li>Lease container</li> <li>Dead-letter queue</li> <li>Life-cycle notifications</li> <li>Change feed estimator</li> </ul>
Handle errors in the change feed processor.	<ul style="list-style-type: none"> <li>Dead-letter queue</li> <li>Lease container</li> <li>Dead-letter queue</li> <li>Life-cycle notifications</li> <li>Change feed estimator</li> </ul>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

<p><b>Requirement</b>                  Read a new batch of documents while keeping track of the failing batch of documents.</p>	<p><b>Feature</b></p> <div style="border: 1px solid gray; padding: 2px;"> <p>Change feed estimator</p> <p>Lease container</p> <p>Dead-letter queue</p> <p>Life-cycle notifications</p> <p>Change feed estimator</p> </div> <div style="border: 1px solid gray; padding: 2px; margin-top: 5px;"> <p>Dead-letter queue</p> <p>Lease container</p> <p>Dead-letter queue</p> <p>Life-cycle notifications</p> <p>Change feed estimator</p> </div>
<p>Handle errors in the change feed processor.</p>	

**NEW QUESTION 177**

- (Topic 8)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.

You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:

- Share session state across all ASP.NET web applications
- Support controlled, concurrent access to the same session state data for multiple readers and a single writer
- Save full HTTP responses for concurrent requests You need to store the information.

Proposed Solution: Deploy and configure Azure Cache for Redis. Update the web applications.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: A**

**Explanation:**

The session state provider for Azure Cache for Redis enables you to share session information between different instances of an ASP.NET web application. The same connection can be used by multiple concurrent threads. Redis supports both read and write operations.

The output cache provider for Azure Cache for Redis enables you to save the HTTP responses generated by an ASP.NET web application.

Note: Using the Azure portal, you can also configure the eviction policy of the cache, and control access to the cache by adding users to the roles provided. These roles, which define the operations that members can perform, include Owner, Contributor, and Reader. For example, members of the Owner role have complete control over the cache (including security) and its contents, members of the Contributor role can read and write information in the cache, and members of the Reader role can only retrieve data from the cache.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching>

**NEW QUESTION 178**

HOTSPOT - (Topic 8)

You are creating an app that uses Event Grid to connect with other services. Your app's event data will be sent to a serverless function that checks compliance. This function is maintained by your company.

You write a new event subscription at the scope of your resource. The event must be invalidated after 3 specific period of time. You need to configure Event Grid to ensure security.

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

<b>Authentication</b>	<b>Type</b>
WebHook event delivery	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>SAS tokens</p> <p>Key authentication</p> <p>JWT token</p> </div> </div>
Topic publishing	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span></span> <span>▼</span> </div> <div style="padding: 2px;"> <p>ValidationCode handshake</p> <p>ValidationURL handshake</p> <p>Management Access Control</p> </div> </div>

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: SAS tokens

Custom topics use either Shared Access Signature (SAS) or key authentication. Microsoft recommends SAS, but key authentication provides simple programming,

and is compatible with many existing webhook publishers.

In this case we need the expiration time provided by SAS tokens.

Box 2: ValidationCode handshake

Event Grid supports two ways of validating the subscription: ValidationCode handshake (programmatic) and ValidationURL handshake (manual).

If you control the source code for your endpoint, this method is recommended.

#### NEW QUESTION 181

- (Topic 8)

A company is developing a solution that allows smart refrigerators to send temperature information to a central location. You have an existing Service Bus. The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location.

You need to complete the configuration.

Which Azure CLI or PowerShell command should you run?

- A. `az servicebus namespace create --resource-group fridge-rg --name fridge-ns --location fridge-loc`
- B. `az servicebus queue create --resource-group fridge-rg --namespace-name fridge-ns --name fridge-q`
- C. `connectionString=$(az servicebus namespace authorization-rule keys list --resource-group fridge-rg --fridge-ns fridge-ns --name RootManageSharedAccessKey --query primaryConnectionString --output tsv)`
- D. `az group create --name fridge-rg --location fridge-log`

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

**Answer: B**

#### Explanation:

A service bus instance has already been created (Step 2 below). Next is step 3, Create a Service Bus queue.

Note:

Steps:

Step 1: # Create a resource group resourceGroupName="myResourceGroup"

`az group create --name $resourceGroupName --location eastus`

Step 2: # Create a Service Bus messaging namespace with a unique name namespaceName=myNameSpace\$RANDOM

`az servicebus namespace create --resource-group $resourceGroupName --name`

`$namespaceName --location eastus`

Step 3: # Create a Service Bus queue

`az servicebus queue create --resource-group $resourceGroupName --namespace-name`

`$namespaceName --name BasicQueue`

Step 4: # Get the connection string for the namespace

`connectionString=$(az servicebus namespace authorization-rule keys list --resource-group`

`$resourceGroupName --namespace-name $namespaceName --name RootManageSharedAccessKey --query primaryConnectionString --output tsv)`

References:

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

#### NEW QUESTION 184

- (Topic 8)

You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published. You must ensure that the website remains available and responsive while minimizing cost. You need to deploy the website. What should you do?

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

**Answer: C**

#### Explanation:

Windows Azure Web Sites (WAWS) offers 3 modes: Standard, Free, and Shared.

Standard mode carries an enterprise-grade SLA (Service Level Agreement) of 99.9% monthly, even for sites with just one instance.

Standard mode runs on dedicated instances, making it different from the other ways to buy Windows Azure Web Sites.

**NEW QUESTION 187**

HOTSPOT - (Topic 8)

You are developing a web application that uses the Microsoft identify platform for user and resource authentication. The web application calls several REST APIs. You are implementing various authentication and authorization flows for the web application. You need to validate the claims in the authentication token.

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

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Token type
Identify users for the application by using a JWT token that contains claims.	<input type="text" value="Access"/> <input checked="" type="checkbox"/> ID <input type="checkbox"/> Refresh <input type="checkbox"/> SAML
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	<input type="text" value="Access"/>
Provide the web application with long-term access to resources on behalf of users without requiring interaction with those users.	<input type="text" value="Access"/> <input type="checkbox"/> ID <input type="checkbox"/> Refresh <input type="checkbox"/> SAML
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	<input type="text" value="Access"/> <input type="checkbox"/> ID <input checked="" type="checkbox"/> Refresh <input type="checkbox"/> SAML

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Requirement	Token type
Identify users for the application by using a JWT token that contains claims.	<input type="text" value="Access"/> <input checked="" type="checkbox"/> ID <input type="checkbox"/> Refresh <input type="checkbox"/> SAML
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	<input type="text" value="Access"/>
Provide the web application with long-term access to resources on behalf of users without requiring interaction with those users.	<input type="text" value="Access"/> <input type="checkbox"/> ID <input checked="" type="checkbox"/> Refresh <input type="checkbox"/> SAML
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	<input type="text" value="Access"/> <input type="checkbox"/> ID <input type="checkbox"/> Refresh <input checked="" type="checkbox"/> SAML

**NEW QUESTION 190**

HOTSPOT - (Topic 8)

You are building an application that stores sensitive customer data in Azure Blob storage.

The data must be encrypted with a key that is unique for each customer.

If the encryption key has been corrupted it must not be used for encryption. You need to ensure that the blob is encrypted.

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

```

from azure.storage.blob import BlobServiceClient

from azure.storage.blob.aio import BlobType x = BlobType(key, verify)
from azure.storage.blob import BlobSasPermissions x = BlobSasPermissions.from_string(key + verify)
from azure.storage.blob import CustomerProvidedEncryptionKey x = CustomerProvidedEncryptionKey(key, verify)
from azure.core.configuration import Configuration x = Configuration(key, verify)

if x.tag == verify:
    creds = creds
if x.makeitrans == verify:
    if x.EncryptionKeyHash == verify:
    if x.proxy_policy == verify:

bsc = BlobServiceClient("", credential = creds)
c = bsc.get_blob_client("con", blob)

c.upload_blob(data, pe=x)
c.upload_blob(data, bt=x)
c.upload_blob(data, bsp=x)
c.upload_blob(data, cpk=x)
    
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area

```

from azure.storage.blob import BlobServiceClient

from azure.storage.blob.aio import BlobType x = BlobType(key, verify)
from azure.storage.blob import BlobSasPermissions x = BlobSasPermissions.from_string(key + verify)
from azure.storage.blob import CustomerProvidedEncryptionKey x = CustomerProvidedEncryptionKey(key, verify)
from azure.core.configuration import Configuration x = Configuration(key, verify)

if x.tag == verify:
    creds = creds
if x.makeitrans == verify:
    if x.EncryptionKeyHash == verify:
    if x.proxy_policy == verify:

bsc = BlobServiceClient("", credential = creds)
c = bsc.get_blob_client("con", blob)

c.upload_blob(data, pe=x)
c.upload_blob(data, bt=x)
c.upload_blob(data, bsp=x)
c.upload_blob(data, cpk=x)
    
```

**NEW QUESTION 192**

- (Topic 8)

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

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

You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.

The app continues to time out after four minutes. The app must process the blob data. You need to ensure the app does not time out and processes the blob data.

Solution: Use the Durable Function async pattern to process the blob data. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.

Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

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

**NEW QUESTION 196**

- (Topic 8)

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 198**

- (Topic 8)

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

After you answer a question in this 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: Create an AKS cluster that supports network policy. Create and apply a network to allow traffic only from within a defined namespace.

Does the solution meet the goal?

A. Yes

B. No

**Answer:** A

**Explanation:**

When you run modern, microservices-based applications in Kubernetes, you often want to control which components can communicate with each other. The principle of least privilege should be applied to how traffic can flow between pods in an Azure Kubernetes Service (AKS) cluster. Let's say you likely want to block traffic directly to back-end applications. The Network Policy feature in Kubernetes lets you define rules for ingress and egress traffic between pods in a cluster.

References:

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

**NEW QUESTION 201**

- (Topic 8)

You develop a solution that uses Azure Virtual Machines (VMs).

The VMs contain code that must access resources in an Azure resource group. You grant the VM access to the resource group in Resource Manager.

You need to obtain an access token that uses the VMs system-assigned managed identity. Which two actions should you perform? Each correct answer presents part of the solution.

A. Use PowerShell on a remote machine to make a request to the local managed identity for Azure resources endpoint.

B. Use PowerShell on the VM to make a request to the local managed identity for Azure resources endpoint.

C. From the code on the V

D. call Azure Resource Manager using an access token.

E. From the code on the V

F. call Azure Resource Manager using a SAS token.

G. From the code on the V

H. generate a user delegation SAS token.

**Answer:** BC

**NEW QUESTION 206**

- (Topic 8)

You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage containers named Container1 and Container2. Uploading of videos occurs on an irregular basis.

You need to copy specific blobs from Container1 to Container2 in real time when specific requirements are met, excluding backup blob copies. What should you do?

A. Download the blob to a virtual machine and then upload the blob to Container2.

B. Run the Azure PowerShell command Start-AzureStorageBlobCopy.

C. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API.

D. Use AzCopy with the Snapshot switch blobs to Container2.

**Answer:** B

**Explanation:**

The Start-AzureStorageBlobCopy cmdlet starts to copy a blob. Example 1: Copy a named blob

C:\PS>Start-AzureStorageBlobCopy -SrcBlob "ContosoPlanning2015" -DestContainer "ContosoArchives" -SrcContainer "ContosoUploads"

This command starts the copy operation of the blob named ContosoPlanning2015 from the container named ContosoUploads to the container named ContosoArchives.

References:

<https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps-6.13.0>

**NEW QUESTION 208**

- (Topic 8)

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

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

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

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

Solution:

- \* 1. Create a SearchServiceClient object to connect to the search index.
- \* 2. Create a DataContainer that contains the documents which must be added.
- \* 3. Create a DataSource instance and set its Container property to the DataContainer.
- \* 4. Set the DataSource property of the SearchServiceClient

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Use the following method:

- \* 1. Create a SearchIndexClient object to connect to the search index
- \* 2. Create an IndexBatch that contains the documents which must be added.
- \* 3. Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.

References:

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

**NEW QUESTION 210**

- (Topic 8)

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

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

You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.

The app continues to time out after four minutes. The app must process the blob data. You need to ensure the app does not time out and processes the blob data.

Solution: 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 214**

- (Topic 8)

You are preparing to deploy an ASP.NET Core website to an Azure Web App from a GitHub repository. The website includes static content generated by a script.

You plan to use the Azure Web App continuous deployment feature.

You need to run the static generation script before the website starts serving traffic. What are two possible ways to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

**Answer: AD**

**Explanation:**

A: To customize your deployment, include a .deployment file in the repository root.

You just need to add a file to the root of your repository with the name .deployment and the content:

```
[config]
command = YOUR COMMAND TO RUN FOR DEPLOYMENT
```

this command can be just running a script (batch file) that has all that is required for your deployment, like copying files from the repository to the web root directory for example.

D: In Azure, you can run your functions directly from a deployment package file in your function app. The other option is to deploy your files in the d:\home\site\wwwroot directory of your function app (see A above).

To enable your function app to run from a package, you just add a WEBSITE\_RUN\_FROM\_PACKAGE setting to your function app settings.

Note: The host.json metadata file contains global configuration options that affect all functions for a function app.

References:

<https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script>

<https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package>

**NEW QUESTION 215**

HOTSPOT - (Topic 8)

An organization deploys a Mob storage account. Users take multiple snapshots of the blob storage account over time. You need to delete all snapshots or the blob storage account. You must not delete the blob storage account itself. 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

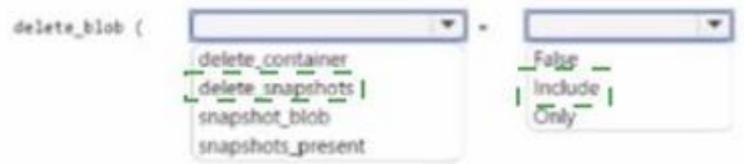


- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Answer Area



**NEW QUESTION 219**

- (Topic 8)

You are developing a software solution for an autonomous transportation system. The solution uses large data sets and Azure Batch processing to simulate navigation sets for entire fleets of vehicles.

You need to create compute nodes for the solution on Azure Batch. What should you do?

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

**Answer: B**

**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.

**NEW QUESTION 224**

HOTSPOT - (Topic 8)

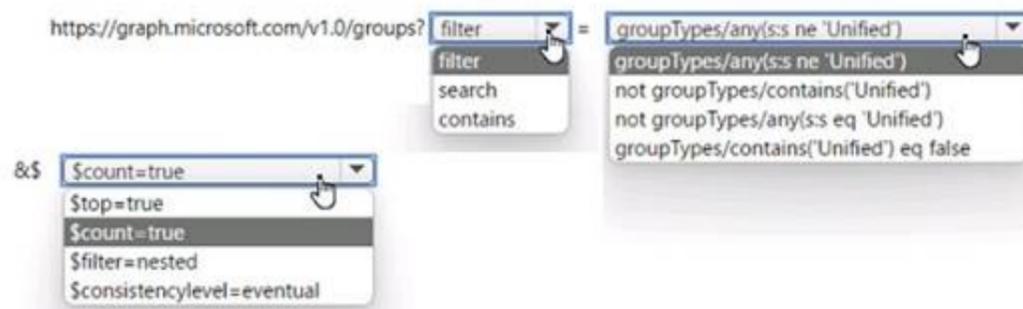
You develop a web app that interacts with Azure Active Directory (Azure AD) groups by using Microsoft Graph.

You build a web page that shows all Azure AD groups that are not of the type 'Unified'. You need to build the Microsoft Graph query for the page.

How should you complete the query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

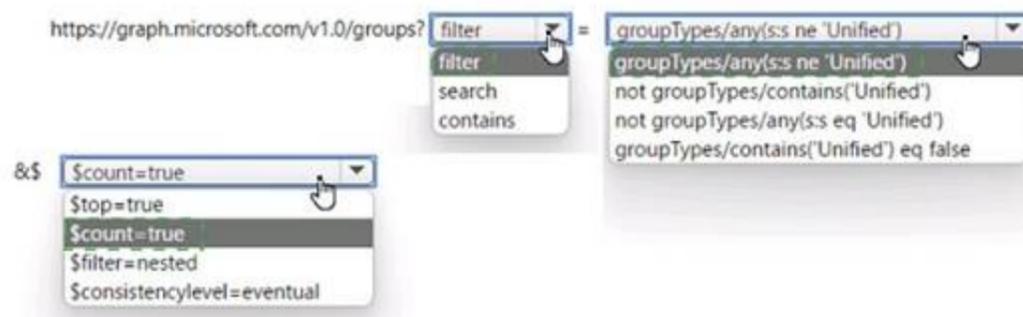


- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

**Answer Area**



**NEW QUESTION 228**

DRAG DROP - (Topic 8)

You are developing an application to use Azure Blob storage. You have configured Azure Blob storage to include change feeds.

A copy of your storage account must be created in another region. Data must be copied from the current storage account to the new storage account directly between the storage servers.

You need to create a copy of the storage account in another region and copy the data.

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

**Actions**

**Answer Area**

- Use AZCopy to copy the data to the new storage account.
- Deploy the template to create a new storage account in the target region.
- Export a Resource Manager template.
- Create a new template deployment.
- Modify the template by changing the storage account name and region.



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

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

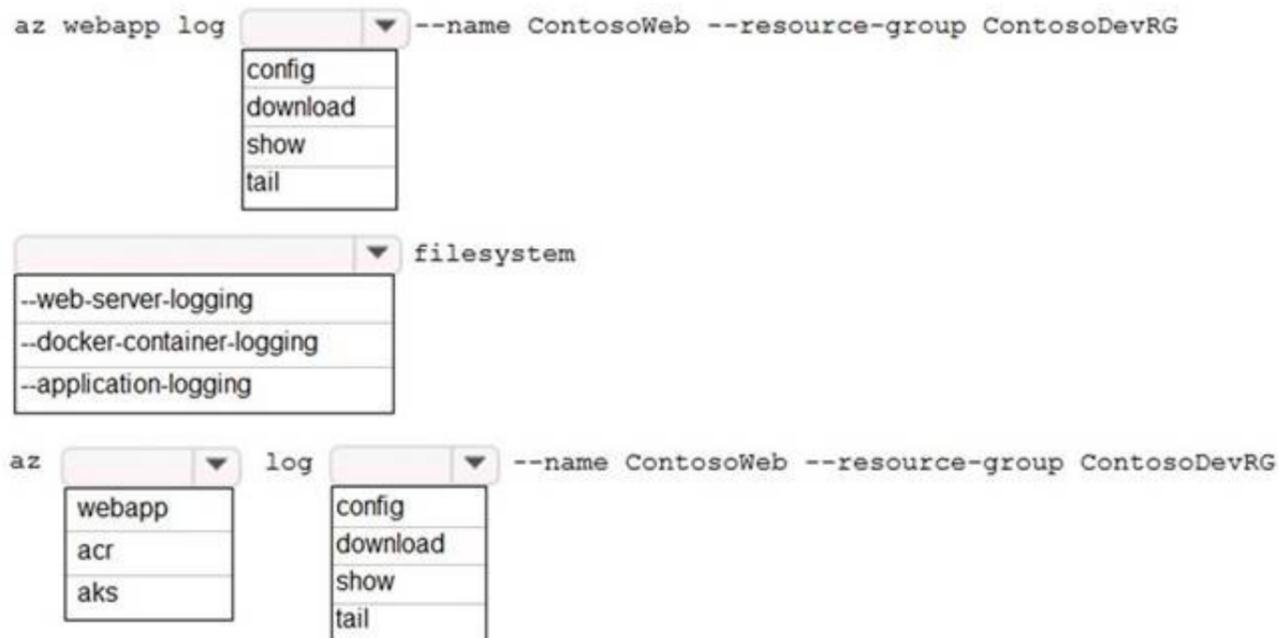
**NEW QUESTION 233**

HOTSPOT - (Topic 8)

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 237**

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

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

Answer: A

**NEW QUESTION 241**

DRAG DROP - (Topic 8)  
 You are creating a script that will run a large workload on an Azure Batch pool. Resources will be reused and do not need to be cleaned up after use.  
 You have the following parameters:

Parameter name	Description
\$script	the script that will run across the batch pool
\$image	the image that pool worker processes will use
\$sku	the node agent SKU Id
\$numberOfJobs	the number of jobs to run

You need to write an Azure CLI script that will create the jobs, tasks, and the pool.  
 In which order should you arrange the commands to develop the solution? To answer, move the appropriate commands from the list of command segments to the answer area and arrange them in the correct order.

**Command segments**

```
az batch pool create
--id mypool --vm-size Standard_A1_v2
--target-dedicated-nodes 2
--image $image
--node-agent-sku-id $sku
```

```
az batch job
create
--id myjob
--pool-id mypool
```

```
for i in {1..$numberOfJobs}
do
```

```
az batch task create
--task-id mytask$i
--job-id myjob
--command-line $script
```

**Answer Area**

⬅

➡

⬆

⬇

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Step 1: az batch pool create  
 # Create a new Linux pool with a virtual machine configuration. az batch pool create \  
 --id mypool \  
 --vm-size Standard\_A1 \  
 --target-dedicated 2 \  
 --image canonical:ubuntu:16.04-LTS \  
 --node-agent-sku-id "batch.node.ubuntu 16.04"  
 Step 2: az batch job create

```
# Create a new job to encapsulate the tasks that are added. az batch job create \
--id myjob \
--pool-id mypool
Step 3: az batch task create
# Add tasks to the job. Here the task is a basic shell command. az batch task create \
--job-id myjob \
--task-id task1 \
--command-line "/bin/bash -c 'printenv AZ_BATCH_TASK_WORKING_DIR'" Step 4: for i in {1..$numberOfJobs} do
References:
https://docs.microsoft.com/bs-latn-ba/azure/batch/scripts/batch-cli-sample-run-job
```

**NEW QUESTION 242**

DRAG DROP - (Topic 8)

You are developing several microservices named serviceA, serviceB, and serviceC. You deploy the microservices to a new Azure Container Apps environment. You have the following requirements.

- The microservices must persist data to storage.
  - serviceA must persist data only visible to the current container and the storage must be restricted to the amount of disk space available in the container
  - serviceB must persist data for the lifetime of the replica and allow multiple containers in the replica to mount the same storage location.
  - serviceC must persist data beyond the lifetime of the replica while allowing multiple containers to access the storage and enable per object permissions.
- You need to configure storage for each microservice.

Storage types	Answer Area								
Azure Blob Storage	<table border="1"> <thead> <tr> <th>Microservice</th> <th>Storage type</th> </tr> </thead> <tbody> <tr> <td>serviceA</td> <td></td> </tr> <tr> <td>serviceB</td> <td></td> </tr> <tr> <td>serviceC</td> <td></td> </tr> </tbody> </table>	Microservice	Storage type	serviceA		serviceB		serviceC	
Microservice		Storage type							
serviceA									
serviceB									
serviceC									
Azure Files storage									
Ephemeral volume									
Container file system									

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Storage types	Answer Area								
Azure Blob Storage	<table border="1"> <thead> <tr> <th>Microservice</th> <th>Storage type</th> </tr> </thead> <tbody> <tr> <td>serviceA</td> <td>Ephemeral volume</td> </tr> <tr> <td>serviceB</td> <td>Container file system</td> </tr> <tr> <td>serviceC</td> <td>Azure Files storage</td> </tr> </tbody> </table>	Microservice	Storage type	serviceA	Ephemeral volume	serviceB	Container file system	serviceC	Azure Files storage
Microservice		Storage type							
serviceA		Ephemeral volume							
serviceB		Container file system							
serviceC	Azure Files storage								
Azure Files storage									
Ephemeral volume									
Container file system									

**NEW QUESTION 245**

HOTSPOT - (Topic 8)

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

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

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

You need to configure Azure CDN caching rules.

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

NOTE:Each correct selection is worth one point.

Setting	Action
Caching behavior	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #ccc; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">Bypass cache</div> <div style="padding: 2px;">Override</div> <div style="padding: 2px;">Set if missing</div> </div>
Cache expiration duration	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #ccc; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">1 second</div> <div style="padding: 2px;">1 minute</div> <div style="padding: 2px;">1 hour</div> <div style="padding: 2px;">1 day</div> </div>
Query string caching behavior	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #ccc; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">Ignore query strings</div> <div style="padding: 2px;">Bypass caching for query strings</div> <div style="padding: 2px;">Cache every unique URL</div> </div>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Override

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

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

Incorrect:

Bypass cache: Do not cache and ignore origin-provided cache-directive headers.

Box 2: 1 hour

All media content must expire from the cache after one hour.

Box 3: Cache every unique URL

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

**NEW QUESTION 250**

- (Topic 8)

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

After you answer a question in this 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 253**

HOTSPOT - (Topic 8)

You are developing a content management application for technical manuals. The application is deployed as an Azure Static Web app.

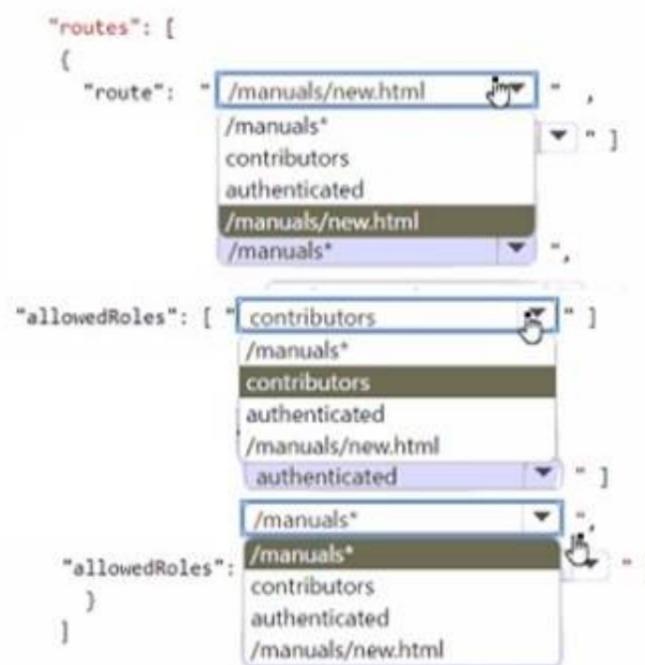
Authenticated users can view pages under /manuals but only contributors can access the page /manuals/new.html.

You need to configure the routing for the web app.

How should you complete the configuration? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**



```

"routes": [
  {
    "route": "/manuals/new.html",
    "allowedRoles": [
      "contributors"
    ]
  }
]

```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area



**NEW QUESTION 255**

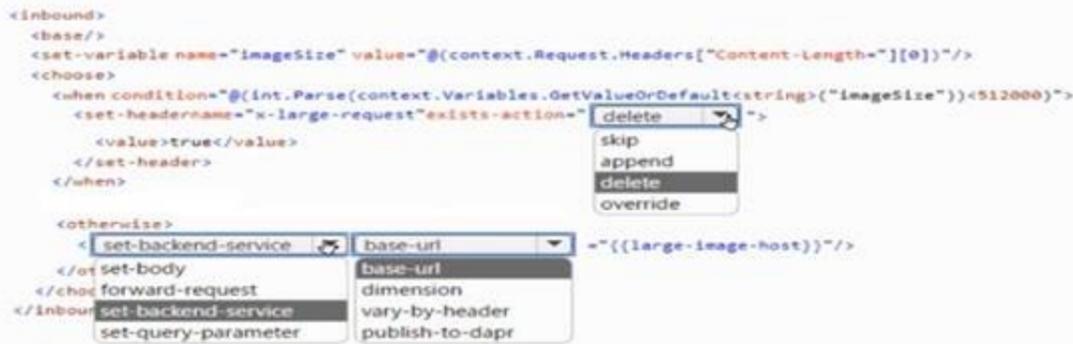
HOTSPOT - (Topic 8)

You develop an image upload service that is exposed using Azure API Management. Images are analyzed after upload for automatic tagging. Images over 500 KB are processed by a different backend that offers a lower tier of service that costs less money. The lower tier of service is denoted by a header named `x-large-request`. Images over 500 KB must never be processed by backends for smaller images and must always be charged the lower price. You need to implement API Management policies to ensure that images are processed correctly.

How should you complete the API Management inbound policy? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

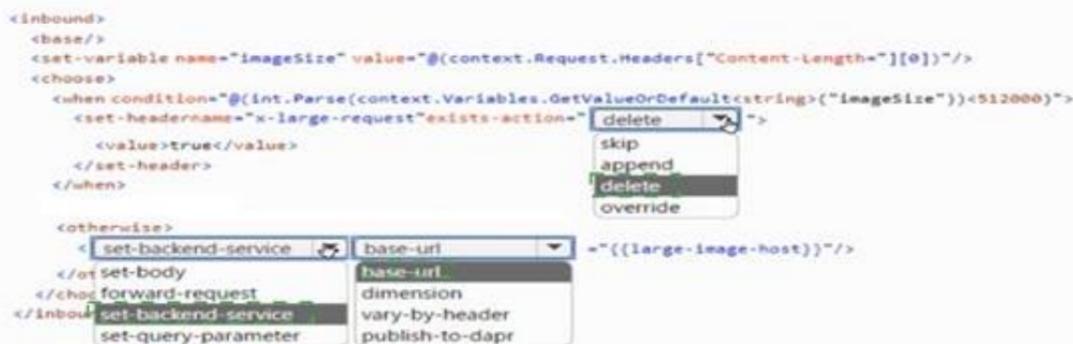


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



**NEW QUESTION 256**

HOTSPOT - (Topic 8)

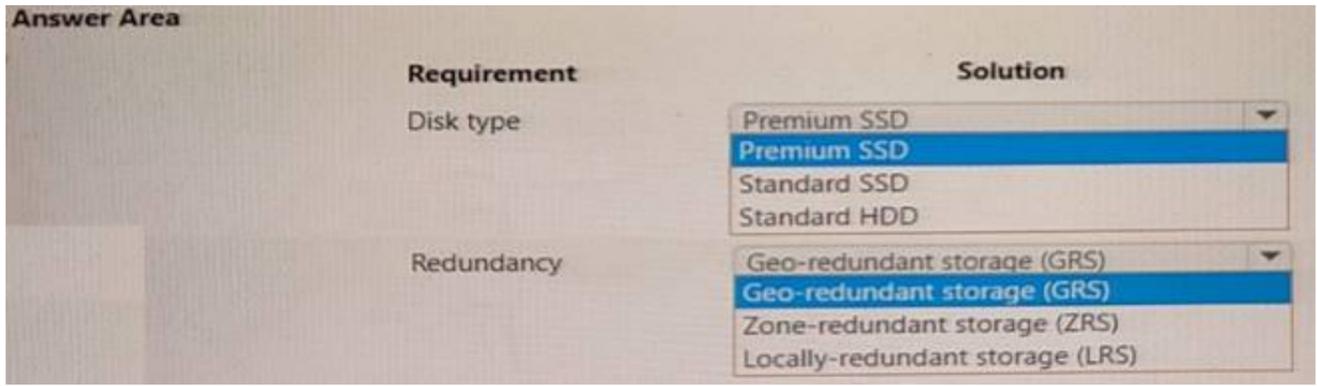
You are building a software-as-a-service (SaaS) application that analyzes DNA data that will run on Azure virtual machines (VMs) in an availability zone. The data is stored on managed disks attached to the VM. The performance of the analysis is determined by the speed of the disk attached to the VM.

You have the following requirements:

- The application must be able to quickly revert to the previous day's data if a systemic error is detected.
- The application must minimize downtime in the case of an Azure datacenter outage.

You need to provision the managed disk for the VM to maximize performance while meeting the requirements. Which type of Azure Managed Disk should you use? To answer, select the appropriate options in the answer area.

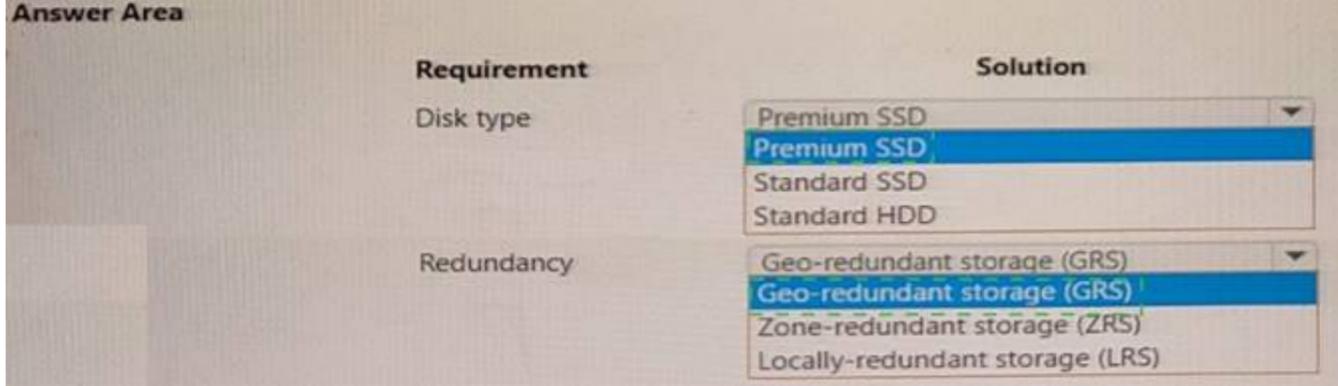
NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 261**

HOTSPOT - (Topic 8)

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  
 [Dropdown: Storage queue, Service Bus queue, Current resource, Storage queue (classic)]

Resource type  
 [Dropdown: Service Bus Namespaces]

Resource  
 [Dropdown: MessageQueue1103]

\* Queues  
 [Dropdown: itemqueue]

Criteria

\* Metric name  
 [Dropdown: Message Count, Active Message Count]

1 minute time grain

\* Time grain statistic [Info icon]  
 [Dropdown: Total, Maximum, Average, Count]

\* Operator

- 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 262**

HOTSPOT - (Topic 8)

You are developing an Azure-hosted e-commerce web application. The application will use Azure Cosmos DB to store sales orders. You are using the latest SDK to manage the sales orders in the database.

You create a new Azure Cosmos DB instance. You include a valid endpoint and valid authorization key to an appSettings.json file in the code project.

You are evaluating the following application code: (Line number are included for reference only.)

```

01 using System;
02 using System.Threading.Tasks;
03 using Microsoft.Azure.Cosmos;
04 using Microsoft.Extensions.Configuration;
05 using Newtonsoft.Json;
06 namespace SalesOrders
07 {
08     public class SalesOrder
09     {
10         ...
11     }
12     internal class ManageSalesOrders
13     {
14         private static async Task GenerateSalesOrders()
15         {
16             IConfigurationRoot configuration = new ConfigurationBuilder().AddJsonFile("appSettings.json").Build();
17             string endpoint = configuration["EndPointUrl"];
18             string authKey = configuration["AuthorizationKey"];
19             using CosmosClient client = new CosmosClient(endpoint, authKey);
20             Database database = null;
21             using (await client.GetDatabase("SalesOrders").DeleteStreamAsync() { })
22                 database = await client.CreateDatabaseIfNotExistsAsync("SalesOrders");
23             Container container1 = await database.CreateContainerAsync(id: "Container1", partitionKeyPath: "/AccountNumber");
24             Container container2 = await database.CreateContainerAsync(id: "Container2", partitionKeyPath: "/AccountNumber");
25             SalesOrder salesOrder1 = new SalesOrder() { AccountNumber = "123456" };
26             await container1.CreateItemAsync(salesOrder1, new PartitionKey(salesOrder1.AccountNumber));
27             SalesOrder salesOrder2 = new SalesOrder() { AccountNumber = "654321" };
28             await container1.CreateItemAsync(salesOrder2, new PartitionKey(salesOrder2.AccountNumber));
29             SalesOrder salesOrder3 = new SalesOrder() { AccountNumber = "109876" };
30             await container2.CreateItemAsync(salesOrder3, new PartitionKey(salesOrder3.AccountNumber));
31             _ = await database.CreateUserAsync("User1");
32             User user1 = database.GetUser("User1");
33             _ = await user1.ReadAsync();
34         }
35     }
36 }
    
```

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

NOTE:Each correct selection is worth one point.

Statements	Yes	No
A database named SalesOrders is created. The database will include two containers.	<input type="radio"/>	<input type="radio"/>
Container1 will contain two items.	<input type="radio"/>	<input type="radio"/>
Container2 will contain one item.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Yes  
 The createDatabaseIfNotExistsAsync method checks if a database exists, and if it doesn't, create it.  
 The Database.CreateContainerAsync method creates a container as an asynchronous operation in the Azure Cosmos service.  
 Box 2: Yes  
 The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service.  
 Box 3: Yes

**NEW QUESTION 264**

HOTSPOT - (Topic 8)

You develop and deploy a web app to Azure App service. The web app allows users to authenticate by using social identity providers through the Azure B2C service. All user profile information is stored in Azure B2C.

You must update the web app to display common user properties from Azure B2C to include the following information:

- ? Email address
- ? Job title

- ? First name
- ? Last name
- ? Office Location

You need to implement the user properties in the web app.

**Requirement**

API to access user properties

**Value**

Microsoft Graph  
 Azure AD Graph  
 Azure Key Vault  
 Azure AD entitlement management

Code library to interface to Azure AD B2C

Microsoft Authentication Library (MSAL)  
 Microsoft Azure Key Vault SDK  
 Azure Identity library

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Requirement**

API to access user properties

**Value**

Microsoft Graph  
 Azure AD Graph  
 Azure Key Vault  
 Azure AD entitlement management

Code library to interface to Azure AD B2C

Microsoft Authentication Library (MSAL)  
 Microsoft Azure Key Vault SDK  
 Azure Identity library

**NEW QUESTION 268**

- (Topic 8)

You develop Azure Web Apps for a commercial diving company. Regulations require that all divers fill out a health questionnaire every 15 days after each diving job starts.

You need to configure the Azure Web Apps so that the instance count scales up when divers are filling out the questionnaire and scales down after they are complete.

You need to configure autoscaling.

What are two possible autoscaling configurations to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

**Answer:** AD

**NEW QUESTION 269**

- (Topic 8)

You are developing a web application that uses Azure Cache for Redis. You anticipate that the cache will frequently fill and that you will need to evict keys.

You must configure Azure Cache for Redis based on the following predicted usage pattern: A small subset of elements will be accessed much more often than the rest.

You need to configure the Azure Cache for Redis to optimize performance for the predicted usage pattern.

Which two eviction policies will achieve the goal?

NOTE: Each correct selection is worth one point.

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

F. volatile-random

**Answer:** BD

**Explanation:**

B: The allkeys-lru policy evict keys by trying to remove the less recently used (LRU) keys first, in order to make space for the new data added. Use the allkeys-lru policy when you expect a power-law distribution in the popularity of your requests, that is, you expect that a subset of elements will be accessed far more often than the rest.

C: volatile-lru: evict keys by trying to remove the less recently used (LRU) keys first, but only among keys that have an expire set, in order to make space for the new data added.

Note: The allkeys-lru policy is more memory efficient since there is no need to set an expire for the key to be evicted under memory pressure.

Reference: <https://redis.io/topics/lru-cache>

**NEW QUESTION 274**

DRAG DROP - (Topic 8)

You develop an Azure solution that uses Cosmos DB.

The current Cosmos DB container must be replicated and must use a partition key that is optimized for queries.

You need to implement a change feed processor solution.

Which change feed processor components should you use? To answer, drag the appropriate components to the correct requirements. Each component 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 the content.

NOTE: Each correct selection is worth one point.

Components	Requirement	Component
Host	Store the data from which the change feed is generated.	Component
Delegate	Coordinate processing of the change feed across multiple workers.	Component
Lease container	Use the change feed processor to listen for changes.	Component
Monitored container	Handle each batch of changes.	Component

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: The monitored container

The monitored container has the data from which the change feed is generated. Any inserts and updates to the monitored container are reflected in the change feed of the container.

Box 2: The lease container

The lease container acts as a state storage and coordinates processing the change feed across multiple workers. The lease container can be stored in the same account as the monitored container or in a separate account.

Box 3: The host: A host is an application instance that uses the change feed processor to listen for changes. Multiple instances with the same lease configuration can run in parallel, but each instance should have a different instance name.

Box 4: The delegate

The delegate is the code that defines what you, the developer, want to do with each batch of changes that the change feed processor reads.

**NEW QUESTION 278**

- (Topic 8)

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

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

You develop 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: Run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

Get an access token using the VM's system-assigned managed identity and use it to call Azure Resource Manager

You will need to use PowerShell in this portion.

? In the portal, navigate to Virtual Machines and go to your Windows virtual machine and in the Overview, click Connect.

? Enter in your Username and Password for which you added when you created the

Windows VM.

? Now that you have created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.

? Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.

Example:

```
$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}
```

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

**NEW QUESTION 281**

- (Topic 8)

You develop and deploy an ASP.NET web app to Azure App Service. You use Application Insights telemetry to monitor the app.

You must test the app to ensure that the app is available and responsive from various points around the world and at regular intervals. If the app is not responding, you must send an alert to support staff.

You need to configure a test for the web app.

Which two test types can you use? Each correct answer presents a complete solution.

NOTE:Each correct selection is worth one point.

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

**Answer:** BC

**Explanation:**

There are three types of availability tests:

? URL ping test: a simple test that you can create in the Azure portal.

? Multi-step web test: A recording of a sequence of web requests, which can be played back to test more complex scenarios. Multi-step web tests are created in Visual Studio Enterprise and uploaded to the portal for execution.

? Custom Track Availability Tests: If you decide to create a custom application to run availability tests, the TrackAvailability() method can be used to send the results to Application Insights.

Reference:

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

**NEW QUESTION 284**

DRAG DROP - (Topic 8)

An organization plans to deploy Azure storage services.

You need to configure shared access signature (SAS) for granting access to Azure Storage.

Which SAS types should you use? To answer, drag the appropriate SAS types to the correct requirements. Each SAS 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.

SAS types	Requirement	SAS type
Account-level	Delegate access to resources in one or more of the storage services	
Service-level	Delegate access to a resource in a single storage service	
User delegation	Secure a resource by using Azure AD credentials	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

SAS types	Requirement	SAS type
Account-level	Delegate access to resources in one or more of the storage services	Account-level
Service-level	Delegate access to a resource in a single storage service	Service-level
User delegation	Secure a resource by using Azure AD credentials	User delegation

**NEW QUESTION 285**

.....

## 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](#)