


## AZ-204 Dumps

### Developing Solutions for Microsoft Azure (beta)

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



**NEW QUESTION 1**

- (Exam Topic 1)

You need to secure the Shipping Logic App. What should you use?

- A. Azure App Service Environment (ASE)
- B. Azure AD B2B integration
- C. Integration Service Environment (ISE)
- D. VNet service endpoint

**Answer: C**

**Explanation:**

Scenario: The Shipping Logic App requires secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.

You can access to Azure Virtual Network resources from Azure Logic Apps by using integration service environments (ISEs).

Sometimes, your logic apps and integration accounts need access to secured resources, such as virtual machines (VMs) and other systems or services, that are inside an Azure virtual network. To set up this access, you can create an integration service environment (ISE) where you can run your logic apps and create your integration accounts.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/connect-virtual-network-vnet-isolated-environment-overview>

**NEW QUESTION 2**

- (Exam Topic 1)

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

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

**Answer Area**

**Option**

**Value**

Tier

	▼
Standard	
Premium	

Profile

	▼
Akamai	
Microsoft	

Optimization

	▼
general web delivery	
large file download	
dynamic site acceleration	
video-on-demand media streaming	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Scenario: Shipping website

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

Tier: Standard Profile: Akamai

Optimization: Dynamic site acceleration

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

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

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

Reference:

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

**NEW QUESTION 3**

- (Exam Topic 1)

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

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

NOTE: Each correct selection is worth one point.

az webapp cors add -g shipping-apis-test-rg -n web

cors	add
config	up
deployment	remove

--

slot	http://*.wideworldimporters.com
allowed-origins	http://test-shippingapi.wideworldimporters.com
name	http://test.wideworldimporters.com
	http://www.wideworldimporters.com

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

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

Enter the full URL of the site you want to allow to access your WEB API or \* to allow all domains. Box 1: cors

Box 2: add

Box 3: allowed-origins

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

[http://donovanbrown.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-wit](http://donovanbrown.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-with-Azure-App-Service)

[h-Azure-App-Service](http://donovanbrown.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-wit)

**NEW QUESTION 4**

- (Exam Topic 1)

You need to secure the Shipping Function app.

How should you configure the app? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

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

- A. Mastered  
B. Not Mastered

Answer: A

Explanation:

Scenario: Shipping Function app: Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).

Box 1: Function

Box 2: JSON based Token (JWT)

Azure AD uses JSON based tokens (JWTs) that contain claims Box 3: HTTP

How a web app delegates sign-in to Azure AD and obtains a token

User authentication happens via the browser. The OpenID protocol uses standard HTTP protocol messages. References:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/authentication-scenarios>

NEW QUESTION 5

- (Exam Topic 3)

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

You need to review the Azure Function App code shown below. NOTE: Each correct selection is worth one point.

```
public static class 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}");
    }
}
```

	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 6

- (Exam Topic 3)

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

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

NOTE: Each correct selection is worth one point.



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

- A. Mastered  
B. Not Mastered

Answer: A

**Explanation:**

Box 1: Yes

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

Box 2: No

Box 3: Yes

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

Box 4: No

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

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

**NEW QUESTION 7**

- (Exam Topic 3)

You are creating a script that will run a large workload on an Azure Batch pool. Resources will be reused and do not need to be cleaned up after use.

You have the following parameters:

You need to write an Azure CLI script that will create the jobs, tasks, and the pool.

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

Command segments	Answer Area
<pre>az batch pool create --id mypool --vm-size Standard_A1_v2 --target-dedicated-nodes 2 --image \$image --node-agent-sku-id \$sku</pre>	
<pre>az batch job create --id myjob --pool-id mypool</pre>	<div>⬅️⬆️</div>
<pre>for i in {1..\$numberOfJobs} do</pre>	<div>➡️⬇️</div>
<pre>az batch task create --task-id mytask\$i --job-id myjob --command-line \$script</pre>	

- 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 8

- (Exam Topic 3)

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

The system performs the following operations in order:

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

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

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

Application State	Restorability
Data blob	<div> <input type="text"/> </div> <div> <input type="text"/> </div> <div> <input type="text"/> </div>
Snapshot 1	<div> <input type="text"/> </div> <div> <input type="text"/> </div> <div> <input type="text"/> </div>
Snapshot 2	<div> <input type="text"/> </div> <div> <input type="text"/> </div> <div> <input type="text"/> </div>

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Box 1: Can be restored

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

Box 2: Cannot be restored It has been deleted.

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

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

### NEW QUESTION 9

- (Exam Topic 3)

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

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

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

You create the index in Azure Search.

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

- \* 1. Create a SearchIndexClient object to connect to the search index.
- \* 2. Create a DataContainer that contains the documents which must be added.

- \* 3. Create a DataSource instance and set its Container property to the DataContainer.
- \* 4. Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource. Does the solution meet the goal?

- A. Yes  
B. No

**Answer: B**

**Explanation:**

Use the following method:

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

**NEW QUESTION 10**

- (Exam Topic 3)

You are developing a project management service by using ASP.NET. The service hosts conversations, files, to-do lists, and a calendar that users can interact with at any time.

The application uses Azure Search for allowing users to search for keywords in the project data.

You need to implement code that creates the object which is used to create indexes in the Azure Search service.

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

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

**Answer: BC**

**Explanation:**

The various client libraries define classes like Index, Field, and Document, as well as operations like Indexes.Create and Documents.Search on the SearchServiceClient and SearchIndexClient classes.

Example:

The sample application we'll be exploring creates a new index named "hotels", populates it with a few documents, then executes some search queries. Here is the main program, showing the overall flow:

/ This sample shows how to delete, create, upload documents and query an index static void Main(string[] args)

```
{  
IConfigurationBuilder builder = new ConfigurationBuilder().AddJsonFile("appsettings.json"); IConfigurationRoot configuration = builder.Build();  
SearchServiceClient serviceClient = CreateSearchServiceClient(configuration); Console.WriteLine("{0}", "Deleting index...\n");  
DeleteHotelsIndexIfExists(serviceClient);  
Console.WriteLine("{0}", "Creating index...\n"); CreateHotelsIndex(serviceClient);  
ISearchIndexClient indexClient = serviceClient.Indexes.GetClient("hotels"); References:  
https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk
```

**NEW QUESTION 10**

- (Exam Topic 3)

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

Spikes in traffic have caused increases in page load times.

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

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

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



## Actions

## Answer Area

Configure the web app to the Premium App Service tier.

Configure the web app to the Standard App Service tier.

Enable autoscaling on the web-app.

Add a Scale rule.

Switch to an Azure App Services consumption plan.

Configure a Scale condition.



- A. Mastered
- B. Not Mastered

**Answer:** A

### Explanation:

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

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

First enable autoscale Step 3: Add a scale rule

Step 4: Add a Scale condition Reference:

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

### NEW QUESTION 12

- (Exam Topic 3)

You are developing a .NET Core model-view controller (MVC) application hosted on Azure for a health care system that allows providers access to their information.

You develop the following code:

```
services.AddAuthorization (options =>
{
    options.AddPolicy("ProviderPartner", policy =>
    {
        .policy.AddAuthenticationSchemes("Cookie, Bearer");
        policy.RequireAuthenticatedUser();
        policy.RequireRole("ProviderAdmin", "SysAdmin");
        policy.RequireClaim("editor", "partner");
    });
});
```

You define a role named SysAdmin.

You need to ensure that the application meets the following authorization requirements:

- Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.
- Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

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.



```
services.AddAuthorization (options =>
{
    options.AddPolicy("ProviderPartner", policy =>
    {
        policy.AddAuthenticationSchemes("Cookie, Bearer");
        policy.RequireAuthenticatedUser();
        policy.RequireRole("ProviderAdmin");
        policy.RequireClaim("editor", "partner");
    });
});
```

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

Box 1:

Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.

Box 2:

Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

**NEW QUESTION 15**

- (Exam Topic 3)

You are developing a data storage solution for a social networking app.

The solution requires a mobile app that stores user information using Azure Table Storage.

You need to develop code that can insert multiple sets of user information.

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

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

▼

op = new

▼

() ;

TableOperation

TableBatchOperaton

TableEntity

TableQuery

TableOperation

TableBatchOperaton

TableEntity

TableQuery

...

table.

▼

(op) ;

ExecuteBatch

Execute

Insert

InsertOrMerge

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

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

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

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

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

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

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

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

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

**NEW QUESTION 17**

- (Exam Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.  
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.  
Margie's Travel is an international travel and bookings management service. The company is expanding into restaurant bookings. You are tasked with implementing Azure Search for the restaurants listed in their solution.  
You create the index in Azure Search.

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

- \* 1. Create a 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 20**

- (Exam Topic 3)

You are preparing to deploy a medical records application to an Azure virtual machine (VM). The application will be deployed by using a VHD produced by an on-premises build server.

You need to ensure that both the application and related data are encrypted during and after deployment to Azure.

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
Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage.	Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage.
Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension.	
Run the Azure PowerShell command Set-AzureRmVMOsdisk.	Run the Azure PowerShell command Set-AzureRmVMOsdisk.
Encrypt the on-premises VHD by using BitLocker with a TPM. Upload the VM to Azure Storage.	
Run the Azure PowerShell command New-AzureRmVm.	Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension.

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Step 1: Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage Step 2: Run the Azure PowerShell command Set-AzureRMVMOSDisk

To use an existing disk instead of creating a new disk you can use the Set-AzureRMVMOSDisk command. Example:

```
$osDiskName = $vmname+'_osDisk'
```

```
$osDiskCaching = 'ReadWrite'
```

```
$osDiskVhdUri = "https://$storageName.blob.core.windows.net/vhds/" + $vmname + "_os.vhd"
```

```
$vm = Set-AzureRmVMOsdisk -VM $vm -VhdUri $osDiskVhdUri -name $osDiskName -Create Step 3: Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension
```

Use the Set-AzVMDiskEncryptionExtension cmdlet to enable encryption on a running IaaS virtual machine in Azure.

Incorrect:

Not TPM: BitLocker can work with or without a TPM. A TPM is a tamper resistant security chip on the system board that will hold the keys for encryption and check the integrity of the boot sequence and allows the most secure BitLocker implementation. A VM does not have a TPM.

References:

<https://www.itprotoday.com/iaaspaas/use-existing-vhd-azurerem-vm>

**NEW QUESTION 25**

- (Exam Topic 3)

You are developing an Azure Cosmos DB solution by using the Azure Cosmos DB SQL API. The data includes millions of documents. Each document may contain hundreds of properties.

The properties of the documents do not contain distinct values for partitioning. Azure Cosmos DB must scale individual containers in the database to meet the performance needs of the application by spreading the workload evenly across all partitions over time.

You need to select a partition key.

Which two partition keys can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.



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

**Answer:** AC

**Explanation:**

You can form a partition key by concatenating multiple property values into a single artificial partitionKey property. These keys are referred to as synthetic keys. Another possible strategy to distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in this way, you can perform parallel write operations across partitions.

Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys>

**NEW QUESTION 26**

- (Exam Topic 3)

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

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

How should you configure the App Service plan? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

## Answer Area

### App service plan setting

### Value

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Number of VM instances: 4

You are not charged extra for deployment slots. Pricing tier: Isolated

The App Service Environment (ASE) is a powerful feature offering of the Azure App Service that gives network isolation and improved scale capabilities. It is essentially a deployment of the Azure App Service into a subnet of a customer's Azure Virtual Network (VNet).

References:


<https://azure.microsoft.com/sv-se/blog/announcing-app-service-isolated-more-power-scale-and-ease-of-use/>

**NEW QUESTION 31**

- (Exam Topic 3)

You are validating the configuration of an Azure Search indexer.

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



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

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



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

You need to resolve the issue.

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

	Yes	No
You can resolve the issue by recreating the search index with the same settings for all fields except ItemDescription. Select the SEARCHABLE option for this field	<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 34

- (Exam Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. 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**

#### NEW QUESTION 36

- (Exam Topic 3)

Your company is developing an Azure API.

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

- All API calls must be secure.
- Callers to the API must not send credentials to the API. Which authentication mechanism should you use?

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

**Answer: C**

#### Explanation:

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

Reference:

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

#### NEW QUESTION 41

- (Exam Topic 3)

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

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

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

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

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

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

- A. Yes  
B. No

**Answer: A**

#### Explanation:

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

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

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

Reference:

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

#### NEW QUESTION 45

- (Exam Topic 3)

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

**NEW QUESTION 46**  
- (Exam Topic 3)  
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 50**

- (Exam Topic 3)

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

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

You need to run the static generation script before the website starts serving traffic.

What are two possible ways to achieve this goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

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

**Answer:** AD

**Explanation:**

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

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

command = YOUR COMMAND TO RUN FOR DEPLOYMENT

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

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

To enable your function app to run from a package, you just add a `WEBSITE_RUN_FROM_PACKAGE` setting to your function app settings.

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

References:

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

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

**NEW QUESTION 53**

- (Exam Topic 3)

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	public class Startup
ServerFeatures	{
LoggerFilterOptions	... public void ConfigureServices (IServiceCollection services)
ApplicationServices	{ services.AddOptions< >().
ApplicationInsightsLoggerOptions	Configure(o => o. = true );
TrackExceptionsAsExceptionTelemetry	services.AddMvc(); }
	public void Configure (IApplicationBuilder app, IHostingEnvironment env, ILoggerFactory loggerFactory)
	{ loggerFactory.AddApplicationInsights(app. ,LogLevel.Trace); app.UseMvc(); }

- 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.computecloud.com/enabling-application-insights-trace-logging-in-asp-net-core/>

**NEW QUESTION 58**

- (Exam Topic 3)

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

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

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

You create the index in Azure Search.

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

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

Does the solution meet the goal?

- A. Yes  
B. No

**Answer: B**

**NEW QUESTION 62**

- (Exam Topic 3)

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

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

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

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



**Answer:** AD

**Explanation:**

D: The ability to specify a host override is defined in the HTTP settings and can be applied to any back-end pool during rule creation.

The ability to derive the host name from the IP or FQDN of the back-end pool members. HTTP settings also provide an option to dynamically pick the host name from a back-end pool member's FQDN if configured with the option to derive host name from an individual back-end pool member.

A (not C): SSL termination and end to end SSL with multi-tenant services.

In case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway.

Therefore, there is no need to add any authentication certificates.

**Add HTTP setting**

saiaappgw-appgw

\* Protocol

HTTP HTTPS

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

\* Port

443 ✓

\* Request timeout (seconds)

20

Override backend path

☒ Use for App service

☒ Use custom probe

OK

Reference:

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

**NEW QUESTION 67**

- (Exam Topic 3)

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

This function is maintained by your company.

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

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

## Authentication

## Type

WebHook event delivery

▼

SAS tokens

Key authentication

JWT token

Topic publishing

▼

ValidationCode handshake

ValidationURL handshake

Management Access Control

- 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 70**

- (Exam Topic 3)

You are creating a CLI script that creates an Azure web app related services in Azure App Service. The web app uses the following variables:

Variable name	Value
\$gitrepo	<code>https://github.com/Contos/webapp</code>
\$webappname	Webapp1103

You need to automatically deploy code from GitHub to the newly created web app.

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

## Answer Area

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

	▼	--name \$webappname - -resource-group myResourceGroup - -sku FREE
az webapp create		
az appservice plan create		
az webapp deployment		
az group delete		

	▼	--name \$webappname - -resource-group myResourceGroup
az webapp create		
az appservice plan create		
az webapp deployment		
az group delete		

	▼	
--repo-url \$gitrepo - -branch master - -manual-integration		
git clone \$gitrepo		
--plan \$webappname		

	▼	source config --name \$webappname
az webapp create		
az appservice plan create		
az webapp deployment		
az group delete		

--resource-group myResourceGroup	▼	
--repo-url \$gitrepo - -branch master - -manual-integration		
git clone \$gitrepo		
--plan \$webappname		

- A. Mastered  
B. Not Mastered

**Answer:** A

### Explanation:

Box 1: az appservice plan create

The azure group creates command successfully returns JSON result. Now we can use resource group to create a azure app service plan

Box 2: az webapp create Create a new web app..

Box 3: --plan \$webappname

with the serviceplan we created in step 1. Box 4: az webapp deployment

Continuous Delivery with GitHub. Example:

az webapp deployment source config --name firstsamplewebsite1 --resource-group websites--repo-url \$gitrepo

--branch master --git-token \$token

Box 5: --repo-url \$gitrepo --branch master --manual-integration Reference:

<https://medium.com/@satish1v/devops-your-way-to-azure-web-apps-with-azure-cli-206ed4b3e9b1>

## NEW QUESTION 71

- (Exam Topic 3)

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

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

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

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

Solution:

•Create a new Azure AD application's manifest, set value of the groupMembershipClaims option to All.

•In the website, use the value of the groups claim from the JWT for the user to determine permissions. Does the solution meet the goal?

- A. Yes  
B. No

**Answer:** A

### Explanation:

To configure Manifest to include Group Claims in Auth Token

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

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

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

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

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



Now your application will include group claims in your manifest and you can use this fact in your code. References:  
<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

### NEW QUESTION 73

- (Exam Topic 3)

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

You need to identify trends in application usage.

Which Azure Application Insights Usage Analysis features should you use? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Box1: Users Box 2: Impact

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

Box 3: Retention

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

Box 4: User flows

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

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

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

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

### NEW QUESTION 77

- (Exam Topic 3)

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

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

You are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.

You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.

You need to implement a solution to receive the device data.

Solution: Provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.

- A. Yes
- B. No

**Answer:** A

#### Explanation:

References:

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

### NEW QUESTION 81

- (Exam Topic 3)

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

Metric source

Storage queue

Service Bus queue

Current resource

Storage queue (classic)

Resource type

Service Bus Namespaces

Resource

MessageQueue1103

Queues

itemqueue

Criteria

Metric name

Message Count

Active Message Count

Time grain statistic

1 minute time grain

Total

Maximum

Average

Count

Greater than

Greater than or equal to

Less than

Less than or equal to

Threshold

1000

Action

Operation

Increase count by

Increase count to

Decrease count by

Decrease count to

Instance count

1

Cool down (minutes)

5

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Service bus queue

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

Box 2: ActiveMessage Count

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

Box 4: Less than or equal to

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

Box 5: Decrease count by

#### NEW QUESTION 86

- (Exam Topic 3)

You are configuring a development environment for your team. You deploy the latest Visual Studio image from the Azure Marketplace to your Azure subscription. The development environment requires several software development kits (SDKs) and third-party components to support application development across the organization. You install and customize the deployed virtual machine (VM) for your development team. The customized VM must be saved to allow provisioning of a new team member development environment.

You need to save the customized VM for future provisioning.

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

#### Answer Area

Action	Tool or service
Generalize the VM.	<div><div>Azure PowerShell</div><div>Visual Studio command prompt</div><div>Azure Migrate</div><div>Azure Backup</div></div>
Store images.	<div><div>Azure Blob Storage</div><div>Azure Data Lake Storage</div><div>Azure File Storage</div><div>Azure Table Storage</div></div>

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Box 1: Azure Powershell

Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks.

Before you begin, make sure that you have the latest version of the Azure PowerShell module. You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image. Box 2: Azure Blob Storage

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource#create-an-image-of-a>

#### NEW QUESTION 90

- (Exam Topic 3)

You have an Azure App Services Web App. Azure SQL Database instance. Azure Storage Account and an Azure Redis Cache instance in a resource group.

A developer must be able to publish code to the web app. You must grant the developer the Contributor role to the web app

You need to grant the role.

What two commands can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

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

**Answer:** AB

#### Explanation:

References:

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

#### NEW QUESTION 94

- (Exam Topic 3)

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

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

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

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

## Azure CLI Commands

## Answer Area

az group create

az group update

az webapp update

az webapp create

az appservice plan create

⬅

➡

⬆

⬆

- A. Mastered
- B. Not Mastered

**Answer:** A

### Explanation:

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

Step 1: az group create

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

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

Step 3: az webapp create

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

References:

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

### NEW QUESTION 97

- (Exam Topic 3)

A company backs up all manufacturing data to Azure Blob Storage. Admins move blobs from hot storage to archive tier storage every month.

You must automatically move blocks to Archive tier after they have not been accessed for 180 days. The path for any item that is not archived must be placed in an existing queue. This operation must be performed automatically once a month. You set the value of TierAgeInDays to 180.

How should you configure the Logic App? To answer, drag the appropriate triggers or action blocks to the correct trigger or action slots. Each trigger or action block may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

### Triggers and Action Blocks

**Insert Entity**

\*Table: processing

\*Entity: Path

Show advanced options

**Tier blob**

If blob is older than the defined value, tier it to Cool or Archive tier

\*Blob path: Path

\*Blob Tier: Archive

**When there are messages in a queue**

\*Queue Name: processing

Show advanced options

Connected to tableStorageAccountConnection. Change connection.


**Recurrence**

\*Interval: 1

\*Frequency: Month

Show advanced options

### Answer Area



- A. Mastered  
B. Not Mastered

**Answer: A**

#### Explanation:

Box 1: Recurrence Box 2: Insert Entity

Box 3 (if true): Tier Blob Box 4: (if false):

Leave blank. References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations>

#### NEW QUESTION 102

- (Exam Topic 3)

You are working for Contoso, Ltd.

You define an API Policy object by using the following XML markup:



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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 107

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

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

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

```
resourceGroupName- +airlineResourceGroup'
name- +docdb-airline-reservations'
databaseName- 'docdb-tickets-database'
collectionName- 'docdb-tickets-collection'
consistencyLevel-
```

▼
Strong
Eventual
ConsistentPrefix
BoundedStaleness

```
az cosmosdb create \
--name $name \
```

▼
--enable-virtual-network true\
--enable-automatic-failover true\
--kind 'GlobalDocumentDB' \
--kind 'MongoDB'\

```
--resource group $resourceGroupName \
--max interval 5 \
```

▼
--locations 'southcentralus'
--locations 'eastus'
--locations'southcentralus=0 eastus=1 westus=2'
--locations 'southcentralus=0'

```
--default-consistency-level - $consistencylevel
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: BoundedStaleness

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

The number of versions (K) of the item

The time interval (T) by which the reads might lag behind the writes Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels> <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/manage-with-cli.md>

**NEW QUESTION 110**

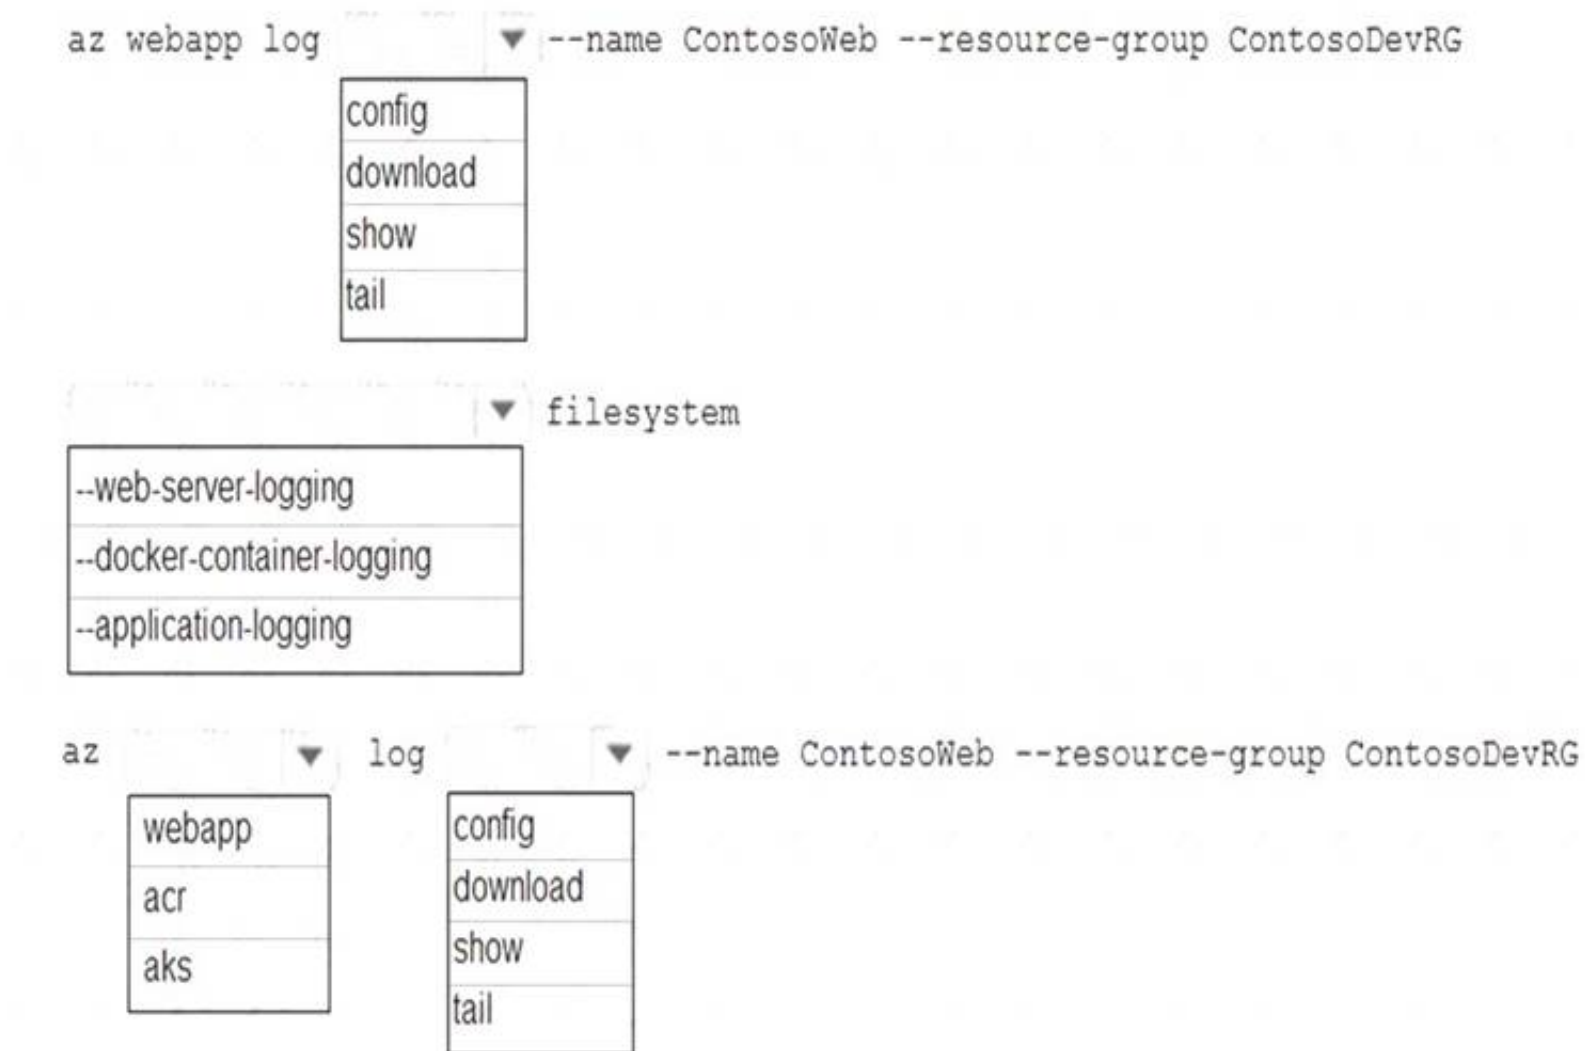
- (Exam Topic 3)

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 112**

- (Exam Topic 3)

You develop Azure solutions.

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

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

**Answer:** C

**Explanation:**

Example:

// Create a new instance of the Cosmos Client

this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)

//ADD THIS PART TO YOUR CODE

await this.CreateDatabaseAsync(); Reference:

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

**NEW QUESTION 113**

- (Exam Topic 3)

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

- The most recent patient status details must be retrieved even if multiple users in different locations have updated the patient record.
- Patient health monitoring data retrieved must be the current version or the prior version.
- After a patient is discharged and all charges have been assessed, the patient billing record contains the final charges.

You provision a Cosmos DB NoSQL database and set the default consistency level for the database account to Strong. You set the value for Indexing Mode to Consistent.

You need to minimize latency and any impact to the availability of the solution. You must override the default consistency level at the query level to meet the required consistency guarantees for the scenarios.

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

NOTE: Each correct selection is worth one point.



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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Strong

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

Box 2: Bounded staleness

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

The number of versions (K) of the item

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

Box 3: Eventual  
Eventual: There's no ordering guarantee for reads. In the absence of any further writes, the replicas eventually converge.

**NEW QUESTION 116**

- (Exam Topic 3)

You must implement Application Insights instrumentation capabilities utilizing the Azure Mobile Apps SDK to provide meaningful analysis of user interactions with a mobile app.

You need to capture the data required to implement the Usage Analytics feature of Application Insights. Which three data values should you capture? Each correct answer presents part of the solution

NOTE: Each correct selection is worth one point.

- A. Trace
- B. Session Id
- C. Exception
- D. User Id
- E. Events

**Answer:** ADE

**Explanation:**

Application Insights is a service for monitoring the performance and usage of your apps. This module allows you to send telemetry of various kinds (events, traces, etc.) to the Application Insights service where your data can be visualized in the Azure Portal.

Application Insights manages the ID of a session for you. References: <https://github.com/microsoft/ApplicationInsights-Android>

**NEW QUESTION 118**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

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

\* One year free update

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

\* Trusted by Millions

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

\* Shop Securely

All transactions are protected by VeriSign!

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

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