

Microsoft

Exam Questions AZ-204

Developing Solutions for Microsoft Azure (beta)



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 resolve the Shipping web site error.
 How should you configure the Azure Table Storage service? To answer, select the appropriate options in the answer area.
 NOTE: Each correct selection is worth one point.

```
<?xml version="1.0" encoding="utf-8"?>
<StorageServiceProperties>
  ""
  <Cors>
    <CorsRule>
      <
        AllowedHeaders
        ExposedHeaders
        AllowedMethods
        AllowedOrigins
      >
        http://*.wideworldimporters.com
        http://test.wideworldimporters.com
        http://test-shippingapi.wideworldimporters.com
        http://www.wideworldimporters.com
      </
    >
    <AllowedMethods>
      GET,PUT
      GET
      POST
      GET,HEAD
    </AllowedMethods>
  </CorsRule>
</Cors>
</StorageServiceProperties>
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: AllowedOrigins
 A CORS request will fail if Access-Control-Allow-Origin is missing. Scenario:
 The following error message displays while you are testing the website:
 Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource. Origin 'http://testwideworldimporters.com/' is therefore not allowed access.
 Box 2: http://test-shippingapi.wideworldimporters.com Syntax: Access-Control-Allow-Origin: *
 Access-Control-Allow-Origin: <origin> Access-Control-Allow-Origin: null
 <origin> Specifies an origin. Only a single origin can be specified. Box 3: AllowedOrigins
 Box 4: POST
 The only allowed methods are GET, HEAD, and POST. In this case POST is used. "<Corsrule>" "allowedmethods" Failed to load no "Access-control-Origin" header is present References:
<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Access-Control-Allow-Origin>

NEW QUESTION 3

- (Exam Topic 1)

You need to support the message processing for the ocean transport workflow.

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

Actions

Create an integration account in the Azure portal.

Link the custom connector to the Logic App.

Update the Logic App to use the partners, schemas, certificates, maps, and agreements.

Create a custom connector for the Logic App.

Add partners, schemas, certificates, maps, and agreements.

Link the Logic App to the integration account.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 4

- (Exam Topic 1)

You need to configure Azure App Service to support the REST API requirements.

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

NOTE: Each correct selection is worth one point.

Setting

Plan

Basic

Standard

Premium

Isolated

Instance Count

1

10

20

100

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Plan: Standard
Standard support auto-scaling Instance Count: 10
Max instances for standard is 10. Scenario:
The REST API's that support the solution must meet the following requirements:
> Allow deployment to a testing location within Azure while not incurring additional costs.
> Automatically scale to double capacity during peak shipping times while not causing application downtime.
> Minimize costs when selecting an Azure payment model. References:
<https://azure.microsoft.com/en-us/pricing/details/app-service/plans/>

NEW QUESTION 5

- (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

config

deployment

add

up

remove

--

slot

allowed-origins

name

http://*.wideworldimporters.com

http://test-shippingapi.wideworldimporters.com

http://test.wideworldimporters.com

http://www.wideworldimporters.com

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Enable Cross-Origin Resource Sharing (CORS) on your Azure App Service Web App.
Enter the full URL of the site you want to allow to access your WEB API or * to allow all domains. Box 1: cors
Box 2: add
Box 3: allowed-origins
Box 4: <http://testwideworldimporters.com/> References:
[http://donovanbrown.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-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 6

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

Issue	Tool
Backup and Restore	<div>▼</div> <div> Azure Site Recovery Azure Backup Azure Data Box Azure Migrate </div>
Performance	<div>▼</div> <div> Azure Network Watcher Azure Traffic Manager ExpressRoute Accelerated Networking </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Backup and Restore: Azure Backup

Scenario: The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

In-Place restore of disks in IaaS VMs is a feature of Azure Backup. Performance: Accelerated Networking

Scenario: The VM shows high network latency, jitter, and high CPU utilization.

Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, greatly improving its networking performance. This high-performance path bypasses the host from the datapath, reducing latency, jitter, and CPU utilization, for use with the most demanding network workloads on supported VM types.

References:

<https://azure.microsoft.com/en-us/blog/an-easy-way-to-bring-back-your-azure-vm-with-in-place-restore/>

NEW QUESTION 7

- (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>▼</div> <div> Can be restored Cannot be restored </div>
Snapshot 1	<div>▼</div> <div> Can be restored Cannot be restored </div>
Snapshot 2	<div>▼</div> <div> Can be restored Cannot be restored </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 8

- (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 9

- (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 an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage

containers named Container1 and Container2. Uploading of videos occurs on an irregular basis.

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

What should you do?

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

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

Answer: B

Explanation:

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

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

-SrcContainer "ContosoUploads"

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

References:

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

NEW QUESTION 10

- (Exam Topic 3)

You are 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", "SysAdmin");
        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 14

- (Exam Topic 3)

You are developing an app that manages users for a video game. You plan to store the region, email address, and phone number for the player. Some players may not have a phone number. The player's region will be used to load-balance data.

Data for the app must be stored in Azure Table Storage.

You need to develop code to retrieve data for an individual player.

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


```
public class PlayerEntity : TableEntity
{
    public PlayerEntity()
    {
    }
    public PlayerEntity(string region, string email)
    {
        PartitionKey =  ;
        RowKey =  ;
    }
    public string Phone { get; set; }
}

public class Player
{
    protected PlayerEntity player;
    async void GetPlayer(string cs,  table, string pk, string rk)
    {
        
        TableEntity query = TableEntity.Retrieve<PlayerEntity>(pk, rk);
        TableOperation query = TableOperation.Retrieve<PlayerEntity>(pk, rk);
        TableResult query = TableQuery.Retrieve<PlayerEntity>(pk, rk);
        TableResultSegment query = TableResult.Retrieve<PlayerEntity>(pk, rk);

        
        TableEntity data = await table.ExecuteAsync(query);
        TableOperation data = await table.ExeucteAsync(query);
        TableQuery data = await table.ExecuteAsync(query);
        TableResult data = await table.ExecuteAsync(query);
        player = data.Result as PlayerEntity;
    }
}
```

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Box 1: region

The player's region will be used to load-balance data. Choosing the PartitionKey.

The core of any table's design is based on its scalability, the queries used to access it, and storage operation requirements. The PartitionKey values you choose will dictate how a table will be partitioned and the type of queries that can be used. Storage operations, in particular inserts, can also affect your choice of PartitionKey values.

Box 2: email

Not phone number some players may not have a phone number. Box 3: CloudTable

Box 4 : TableOperation query =.. Box 5: TableResult

References:

<https://docs.microsoft.com/en-us/rest/api/storageservices/designing-a-scalable-partitioning-strategy-for-azure-ta>

NEW QUESTION 18

- (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 23

- (Exam Topic 3)

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

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

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

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

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

Solution: Provision an Azure Service Bus. Configure a topic to receive the device data by using a correlation filter.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

A message is raw data produced by a service to be consumed or stored elsewhere. The Service Bus is for high-value enterprise messaging, and is used for order processing and financial transactions.

Reference:

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

NEW QUESTION 24

- (Exam Topic 3)

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

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

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

You need to resolve the error.

What should you do?

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

Answer: C

Explanation:

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

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

NEW QUESTION 28

- (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'
```

```
$osDiskVhdUri = "https://$stname.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-azure-vm>

NEW QUESTION 30

- (Exam Topic 3)

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

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

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

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

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

Solution: Provision an Azure Notification Hub. Register all devices with the hub. Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Instead use an Azure Service Bus, which is used order processing and financial transactions.

Reference:

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

NEW QUESTION 31

- (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 33

- (Exam Topic 3)

You develop a news and blog content delivery app for Windows devices.

A notification must arrive on a user's device when there is a new article available for them to view. You need to implement push notifications.

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

```
string notificationHubName = "contoso_hub";
string notificationHubConnection = "connection_string";

[Box 1] hub=
[Box 2]
[Box 3]
(notificationHubConnection, notificationHubName);
string windowsToastPayload =
@"<toast><visual><binding template=""ToastText01""><text id=""1"">" +
@"New item to view" + @"</text></binding></visual></toast>";
try
{
    var result=
        await hub.[Box 4] (windowsToastPayload);
    . . .
}
catch (System.Exception ex)
{
    . . .
}
. . .
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: NotificationHubClient
Box 2: NotificationHubClient
Box 3: CreateClientFromConnectionString
// Initialize the Notification Hub
NotificationHubClient hub = NotificationHubClient.CreateClientFromConnectionString(listenConnString, hubName);
Box 4: SendWindowsNativeNotificationAsync Send the push notification.
var result = await hub.SendWindowsNativeNotificationAsync(windowsToastPayload);
References:
<https://docs.microsoft.com/en-us/azure/notification-hubs/notification-hubs-push-notification-registration-manag>
<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/app-service-mobile/app-service-mobile-windo>

NEW QUESTION 35

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

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

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

Answer Area

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

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

            outputFileList.Add(new OutputFile(fileTask.Output,
                new OutputFileDestination(failedContainer),
                new OutputFileUploadOptions(OutputFileUploadCondition.▼)));
            TaskSuccess
            TaskFailure
            TaskCompletion

            task.▼ = outputFileList;
            OutputFiles
            FilesToStage
            ResourceFiles
            StageFiles

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

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Box 1: CreateJob

Box 2: TaskSuccess

TaskSuccess: Upload the file(s) only after the task process exits with an exit code of 0.

Incorrect: TaskCompletion: Upload the file(s) after the task process exits, no matter what the exit code was. Box 3: TaskFailure

TaskFailure: Upload the file(s) only after the task process exits with a nonzero exit code. Box 4: OutputFiles

To specify output files for a task, create a collection of OutputFile objects and assign it to the CloudTask.OutputFiles property when you create the task.

References:

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

NEW QUESTION 37

- (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 40

- (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 41

- (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.

Answer Area

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 43

- (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	<div><div></div><div>2</div><div>4</div><div>8</div><div>16</div></div>
Pricing tier	<div><div></div><div>Isolated</div><div>Standard</div><div>Premium</div><div>Consumption</div></div>

- 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 44

- (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: Convert the Azure Storage account to a BlobBlobStorage storage account. Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Not necessary to convert the account, instead move photo processing to an Azure Function triggered from the blob upload..

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.

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 47

- (Exam Topic 3)

You develop software solutions for a mobile delivery service. You are developing a mobile app that users can use to order from a restaurant in their area. The app uses the following workflow:

* 1. A driver selects the restaurants for which they will deliver orders.

* 2. Orders are sent to all available drivers in an area.

* 3. Only orders for the selected restaurants will appear for the driver.

* 4. The first driver to accept an order removes it from the list of available orders. You need to implement an Azure Service Bus solution.

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

Actions

- Create a Service Bus topic for each restaurant for which a driver can receive messages.
- Create a single Service Bus topic.
- Create a single Service Bus subscription.
- Create a single Service Bus Namespace.
- Create a Service Bus Namespace for each restaurant for which a driver can receive messages.
- Create a Service Bus subscription for each restaurant for which a driver can receive orders.

Answer area

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Box 1: Create a single Service Bus Namespace

To begin using Service Bus messaging entities in Azure, you must first create a namespace with a name that is unique across Azure. A namespace provides a scoping container for addressing Service Bus resources within your application.

Box 2: Create a Service Bus Topic for each restaurant for which a driver can receive messages. Create topics.

Box 3: Create a Service Bus subscription for each restaurant for which a driver can receive orders. Topics can have multiple, independent subscriptions.

References:

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

NEW QUESTION 51

- (Exam Topic 3)

You are using Azure Front Door Service.

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

You need to determine the root cause for the issue.

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

NOTE: Each correct selection is worth one point.

Answer Area

Statement	Yes	No
The file MIME type is supported by the service.	<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 Reference:
<https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching>

NEW QUESTION 53

- (Exam Topic 3)
Note: This question is part of a series of questions that present the same scenario. Each question in the sries 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 DataContamer
* 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 54

- (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.

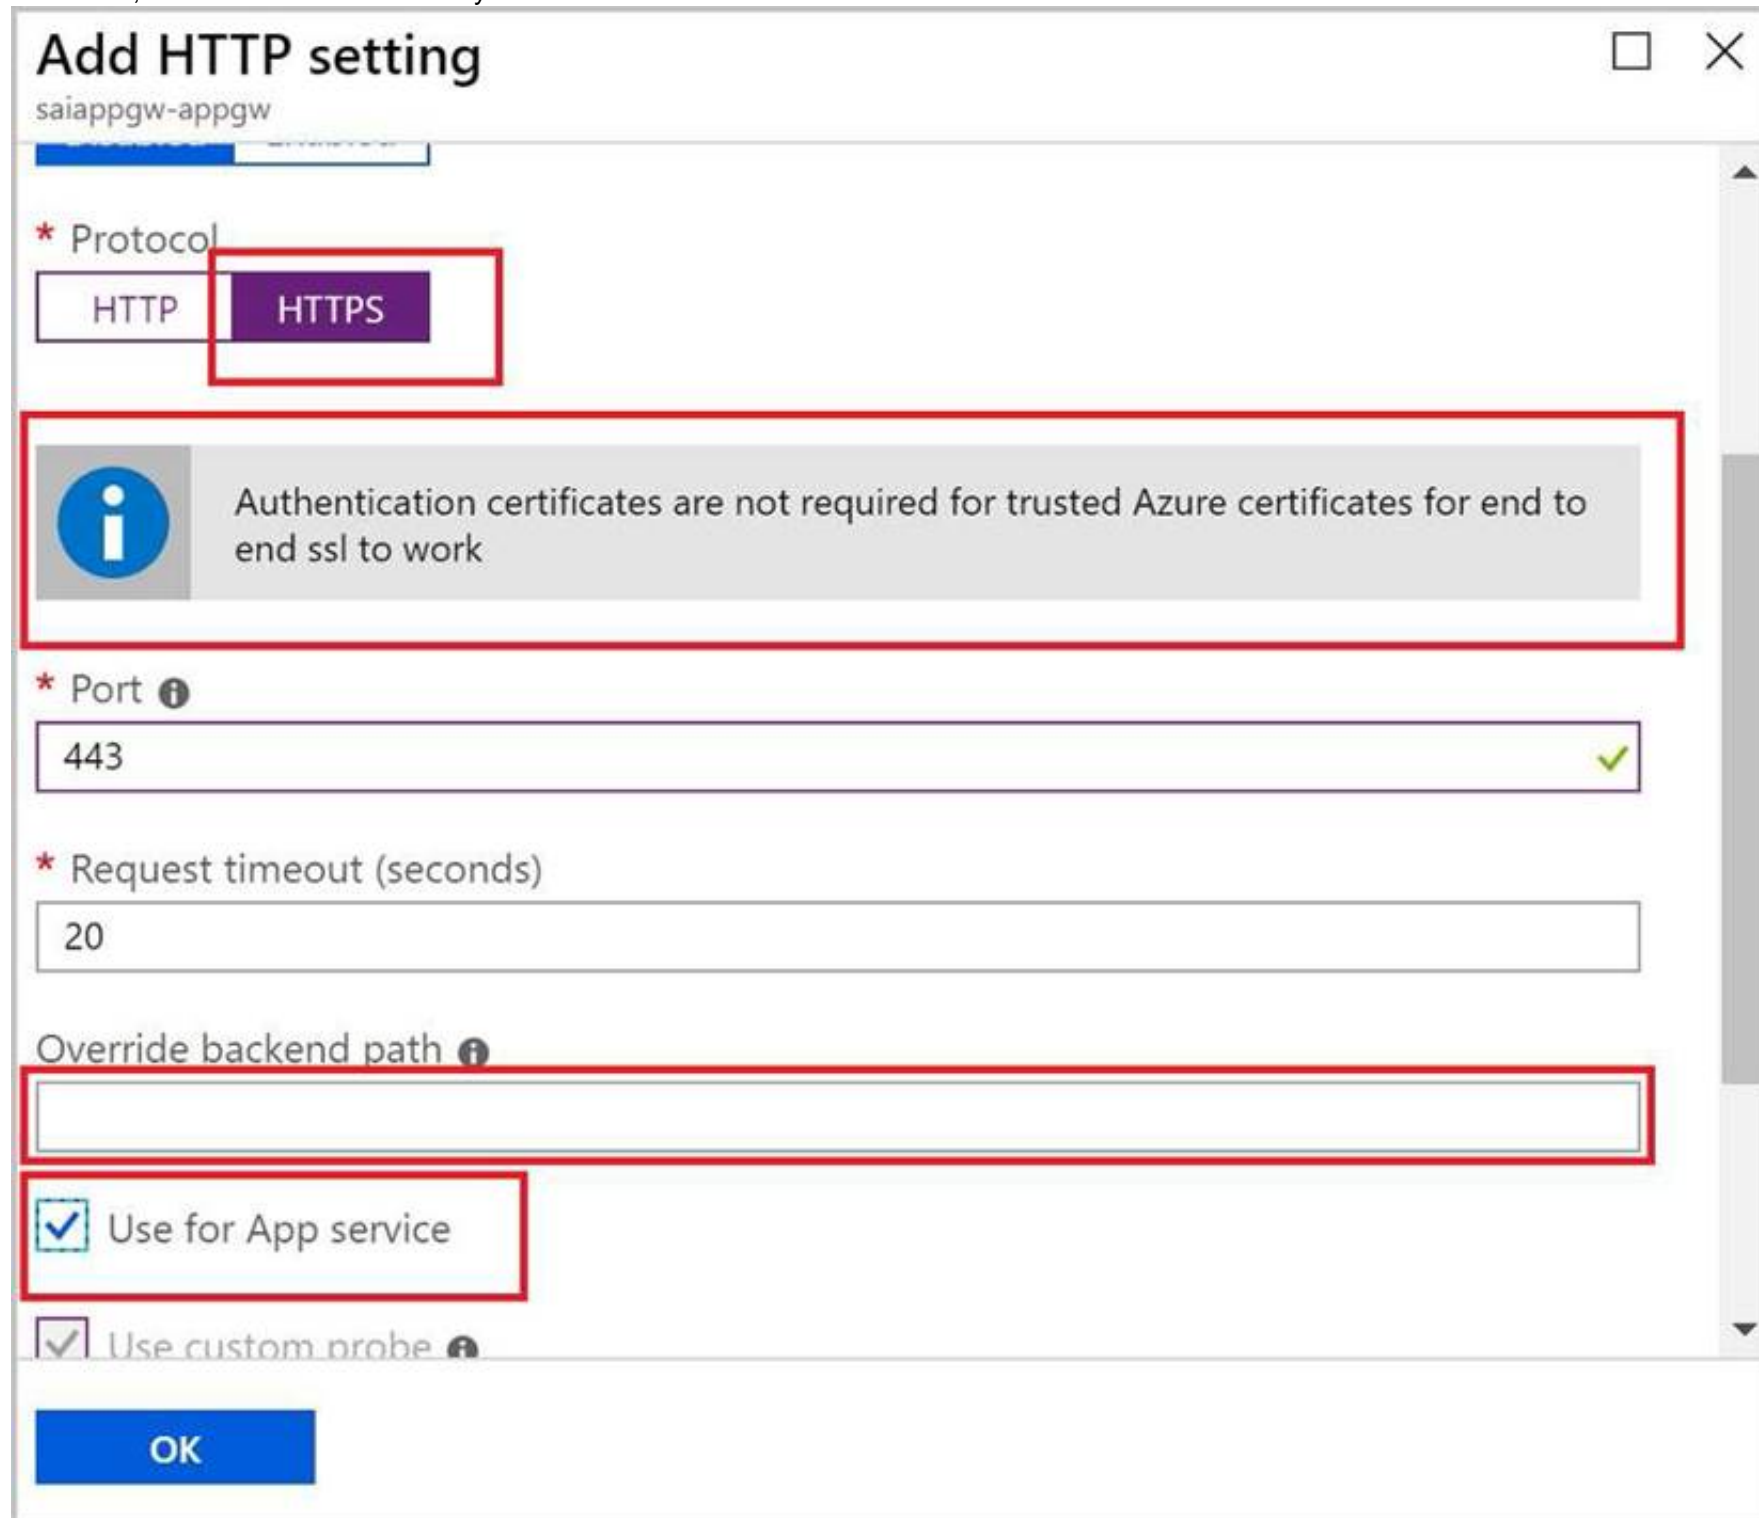
- A. In the Azure Application Gateway’s HTTP setting, enable the Use for App service setting.
- B. Convert the web app to run in an Azure App service environment (ASE).
- C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application gateway.
- D. In the Azure Application 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

saippgw-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 56

- (Exam Topic 3)

Fourth Coffee has an ASP.NET Core web app that runs in Docker. The app is mapped to the www.fourthcoffee.com domain.

Fourth Coffee is migrating this application to Azure.

You need to provision an App Service Web App to host this docker image and map the custom domain to the App Service web app.

A resource group named FourthCoffeePublicWebResourceGroup has been created in the WestUS region that contains an App Service Plan named AppServiceLinuxDockerPlan.

Which order should the CLI commands be used to develop the solution? To answer, move all of the Azure CLI command from the list of commands to the answer area and arrange them in the correct order.

Azure CLI commands

```
az webapp config hostname add
--webapp-name $appName
--resource-group fourthCoffeePublicWebResourceGroup\
--hostname $fqdn

#/bin/bash
appName="FourthCoffeePublicWeb$random".
location "WestUS"
dockerHubContainerPath="FourthCoffee/publicweb:v1"
fqdn=http://www.fourthcoffee.com>www.fourthcoffee.com

az webapp create
--name $appName
--plan AppServiceLinuxDockerPlan
--resource-group fourthCoffeePublicWebResourceGroup

az webapp config container set
--docker-custom-image-name $dockerHibContainerPath
--name $appName
--resource-group fourthCoffeePublicWebResourceGroup
```

Answer area

⏪

⏩

⏴

⏵

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 61

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

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

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

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 64

- (Exam Topic 3)

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 68

- (Exam Topic 3)

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

Actions

A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.

Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the files from cache if the TTL has not expired.

If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.

The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: A user requests the image..
A user requests a file (also called an asset) by using a URL with a special domain name, such as <endpoint name>.azureedge.net. This name can be an endpoint hostname or a custom domain. The DNS routes the request to the best performing POP location, which is usually the POP that is geographically closest to the user.

Step 2: If no edge servers in the POP have the..
If no edge servers in the POP have the file in their cache, the POP requests the file from the origin server. The origin server can be an Azure Web App, Azure Cloud Service, Azure Storage account, or any publicly accessible web server.

Step 3: The origin server returns the..
The origin server returns the file to an edge server in the POP.
An edge server in the POP caches the file and returns the file to the original requestor (Alice). The file remains cached on the edge server in the POP until the time-to-live (TTL) specified by its HTTP headers expires. If the origin server didn't specify a TTL, the default TTL is seven days.

Step 4: Subsequent requests for..
Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.
If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.

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

NEW QUESTION 69

- (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

Generalize the VM.

Store images.

Tool or service

Azure PowerShell

Visual Studio command prompt

Azure Migrate

Azure Backup

Azure Blob Storage

Azure Data Lake Storage

Azure File Storage

Azure Table Storage

- 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 74

- (Exam Topic 3)

A company is developing a mobile app for field service employees using Azure App Service Mobile Apps as the backend.

The company's network connectivity varies throughout the day. The solution must support offline use and synchronize changes in the background when the app is online app.

You need to implement the solution.

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

NOTE: Each correct selection is worth one point.

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

▼

var todoTable = client.GetSyncTable<TodoItem>();
var todoTable = client.GetTable<TodoItem>();
var todoTable = client.SyncTable;
var todoTable = client.Table;

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

▼

await todoTable.PullAsync("allTodoItems",todoTable.CreateQuery());
await todoTable.UpdateAsync();
todoTable.PullAsync("allTodoItems", todoTable.CreateQuery());
todoTable.UpdateAsync();

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: var todoTable = client GetSyncTable<TodoItem>()

To setup offline access, when connecting to your mobile service, use the method GetSyncTable instead of GetTable (example):

IMobileServiceSyncTable todoTable = App.MobileService.GetSyncTable(); / Box 2: await todoTable.PullAsync("allTodoItems",todo.Table.CreateQuery());

Your app should now use IMobileServiceSyncTable (instead of IMobileServiceTable) for CRUD operations. This will save changes to the local database and also keep a log of the changes. When the app is ready to synchronize its changes with the Mobile Service, use the methods PushAsync and PullAsync (example):

await App.MobileService.SyncContext.PushAsync(); await todoTable.PullAsync();

References:

<https://azure.microsoft.com/es-es/blog/offline-sync-for-mobile-services/>

NEW QUESTION 77

- (Exam Topic 3)

You are developing a new page for a website that uses Azure Cosmos DB for data storage. The feature uses documents that have the following format:

You must display data for the new page in a specific order. You create the following query for the page:

You need to configure a Cosmos DB policy to the support the query.

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

NOTE: Each correct selection is worth one point.

JSON segments	Answer Area
<div>orderBy</div>	{ "automatic": true, "indexingMode": "Consistent", "includedPaths": [{ "path": "/*" }], "excludedPaths": [], "compositeIndexes": [{ "path": "/name", "order": "descending" }, { "path": "/city", "order": " <div></div> " }] }
<div>sortOrder</div>	
<div>ascending</div>	
<div>descending</div>	
<div>compositeIndexes</div>	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: compositeIndexes

You can order by multiple properties. A query that orders by multiple properties requires a composite index. Box 2: descending

Example: Composite index defined for (name ASC, age ASC):

It is optional to specify the order. If not specified, the order is ascending.

```
{
  "automatic":true, "indexingMode":"Consistent", "includedPaths":[
    {
      "path":"/*"
    }
  ],
  "excludedPaths":[], "compositeIndexes":[ [
    {
      "path":"/name",
    },
    {
      "path":"/age",
    }
  ]
]
```

NEW QUESTION 82

- (Exam Topic 3)

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: AC

Explanation:

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

Reference:

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

NEW QUESTION 84

- (Exam Topic 3)

You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway.

Data from traffic sensors are stored in Azure Event Hub.

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

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

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

NOTE: Each correct selection is worth one point.

Setting	Value
Number of partitions	<div><div></div><div>▼</div><div>3</div><div>4</div><div>6</div><div>12</div></div>
Partition Key	<div><div></div><div>▼</div><div>Highway</div><div>Department</div><div>Timestamp</div><div>VM name</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: 6

The number of partitions is specified at creation and must be between 2 and 32. There are 6 highways.

Box 2: Highway References:

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

NEW QUESTION 88

- (Exam Topic 3)

You are writing code to create and run an Azure Batch job. You have created a pool of compute nodes.

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

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

Answer: C

Explanation:

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

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

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

References:

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

NEW QUESTION 92

- (Exam Topic 3)

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

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

NOTE: Each correct selection is worth one point.

Code segments

Ingress

Service

LoadBalancer

Deployment

ingress.class

azure-load-balancer-internal

Answer Area

apiVersion: v1

kind: Code segment

metadata:

name: web-app

annotations:

service.beta.kubernetes. Code segment : "true"

spec:

type: Code segment

ports:

- port: 80

selector:

app: web-app

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
To create an internal load balancer, create a service manifest named internal-lb.yaml with the service type LoadBalancer and the azure-load-balancer-internal annotation as shown in the following example:
YAML:
apiVersion: v1 kind: Service metadata:
name: internal-app annotations:
service.beta.kubernetes.io/azure-load-balancer-internal: "true" spec:
type: LoadBalancer ports:
- port: 80 selector:
app: internal-app
References:
<https://docs.microsoft.com/en-us/azure/aks/internal-lb>

NEW QUESTION 93
- (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

az group create

az group update

az webapp update

az webapp create

az appservice plan create

Answer Area

⬅

➡

⬆

⬆

- 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 96

- (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.

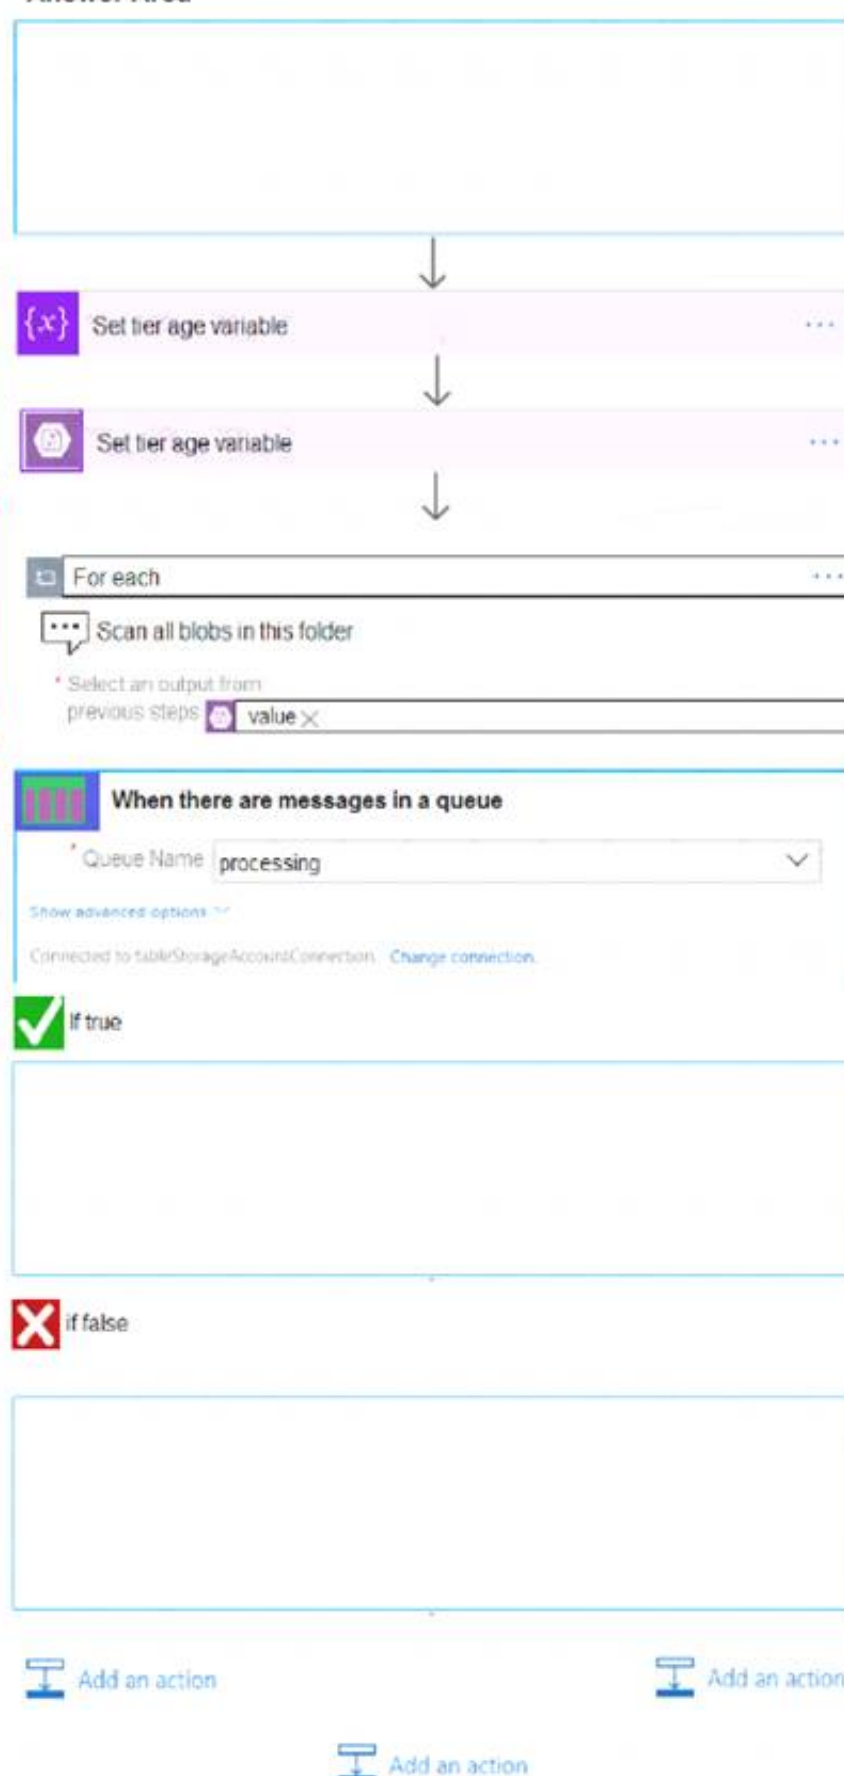
Recurrence

*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 97

- (Exam Topic 3)

You are working for a company that designs mobile applications. They maintain a server where player records are assigned to their different games. The tracking system is new and in development.

The application uses Entity Framework to connect to an Azure Database. The database holds a Player table and Game table.

When adding a player, the code should insert a new player record, and add a relationship between an existing game record and the new player record.

The application will call CreatePlayerWithGame with the correct gameId and the playerId to start the process. (Line numbers are included for reference only.)

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

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

	Yes	No
The code will successfully insert a player record.	<input type="radio"/>	<input type="radio"/>
The code has a bug and will insert an additional copy of the Game record with a new Id.	<input type="radio"/>	<input type="radio"/>
The code has a bug and will insert the wrong gameld value.	<input type="radio"/>	<input type="radio"/>
There is a valid many-to-many relationship between Players and Games.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Many-to-many relationships without an entity class to represent the join table are not yet supported. However, you can represent a many-to-many relationship by including an entity class for the join table and mapping two separate one-to-many relationships.

```
protected override void OnModelCreating(ModelBuilder modelBuilder)
{
    modelBuilder.Entity<PostTag>().HasKey(t => new { t.PostId, t.TagId }); modelBuilder.Entity<PostTag>() HasOne(pt => pt.Post)
    WithMany(p => p.PostTags) HasForeignKey(pt => pt.PostId); modelBuilder.Entity<PostTag>() HasOne(pt => pt.Tag) WithMany(t => t.PostTags) HasForeignKey(pt
    => pt.TagId);
}
```

NEW QUESTION 99

- (Exam Topic 3)

You manage several existing Logic Apps.

You need to change definitions, add new logic, and optimize these apps on a regular basis.

What should you use? To answer, drag the appropriate tools to the correct functionalities. Each tool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Enterprise Integration Pack

After you create an integration account that has partners and agreements, you are ready to create a business to business (B2B) workflow for your logic app with the Enterprise Integration Pack.

Box 2: Code View Editor

To work with logic app definitions in JSON, open the Code View editor when working in the Azure portal or in Visual Studio, or copy the definition into any editor that you want.

Box 3: Logical Apps Designer

You can build your logic apps visually with the Logic Apps Designer, which is available in the Azure portal through your browser and in Visual Studio.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-b2b> <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions> <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-overview>

NEW QUESTION 103

- (Exam Topic 3)

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

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

Answer: D

Explanation:

Sampling is an effective way to reduce charges and stay within your monthly quota.

You can set sampling manually, either in the portal on the Usage and estimated costs page; or in the ASP.NET SDK in the .config file; or in the Java SDK in the ApplicationInsights.xml file, to also reduce the network traffic.

Adaptive sampling is the default for the ASP.NET SDK. Adaptive sampling automatically adjusts to the volume of telemetry that your app sends. It operates automatically in the SDK in your web app so that telemetry traffic on the network is reduced.

References:

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

NEW QUESTION 106

- (Exam Topic 3)

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()
{
    

IDatabase cache = Connection.GetDatabase();
        ICache cache = Connection.GetDatabase();



cache.KeyDelete("teams");
        cache.StringSet("teams", "");
        cache.ValueDelete("teams");
        cache.StringGet("teams", "");


    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 111

- (Exam Topic 3)

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

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"]

Answer Area

- 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 114

- (Exam Topic 3)
You are developing Azure WebJobs.
You need to recommend a WebJob type for each scenario.
Which WebJob type should you recommend? To answer, drag the appropriate WebJob types to the correct scenarios. Each WebJob type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

WebJob types

Scenario

WebJob type

Triggered

Continuous

Run on all instances that the web app runs on.
Optionally restrict the WebJob to a single instance.

Run on a single instance that Azure select for load balancing.

Supports remote debugging

- 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 118

- (Exam Topic 3)
ASP.NET Core API app by using C#. The API app will allow users to authenticate by using Twitter and Azure Active Directory (Azure AD). Users must be authenticated before calling API methods. You must log the user's name for each method call. You need to configure the API method calls. Which values should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Code segment	Value
Attribute	<div><div></div><div>Authorize</div><div>AllowAnonymous</div><div>AutoValidateAntiforgeryToken</div></div>
Request Header	<div><div></div><div>X-MS-CLIENT-PRINCIPAL-NAME</div><div>Proxy-Authorization</div><div>X-Forwarded-For</div><div>X-MS-CLIENT-PRINCIPAL-ID</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Authorize
Box 2: X-MS-CLIENT-PRINCIPAL-NAME
App Service passes user claims to your application by using special headers. External requests aren't allowed to set these headers, so they are present only if set by App Service. Some example headers include:
X-MS-CLIENT-PRINCIPAL-NAME X-MS-CLIENT-PRINCIPAL-ID
Here's the set of headers you get from Easy Auth for a Twitter authenticated user:
{
"cookie": "AppServiceAuthSession=Lx43...xHDTA==", "x-ms-client-principal-name": "evilSnobu",
"x-ms-client-principal-id": "35....", "x-ms-client-principal-idp": "twitter",
"x-ms-token-twitter-access-token": "35...Dj",
"x-ms-token-twitter-access-token-secret": "OK3...Jx",
}

References:
<https://docs.microsoft.com/en-us/azure/app-service/app-service-authentication-how-to>

NEW QUESTION 123

- (Exam Topic 3)
You are creating a hazard notification system that has a single signaling server which triggers audio and visual alarms to start and stop. You implement Azure Service Bus to publish alarms. Each alarm controller uses Azure Service Bus to receive alarm signals as part of a transaction. Alarm events must be recorded for audit purposes. Each transaction record must include information about the alarm type that was activated. You need to implement a reply trail auditing solution. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

- E. Assign the value of the hazard message MessageId property to the SequenceNumber property.
- F. Assign the value of the hazard message MessageId property to the CorrelationId property.

Answer: AB

NEW QUESTION 128

- (Exam Topic 3)

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

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

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

Answer: AE

Explanation:

References:

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

NEW QUESTION 130

.....

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](#)