

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 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
  >
  </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 configure Azure CDN for the Shipping web site.

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

Answer Area

Option	Value
Tier	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Standard</p> <p>Premium</p> </div> </div>
Profile	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Akamai</p> <p>Microsoft</p> </div> </div>
Optimization	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>general web delivery</p> <p>large file download</p> <p>dynamic site acceleration</p> <p>video-on-demand media streaming</p> </div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Scenario: Shipping website

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

Tier: Standard Profile: Akamai

Optimization: Dynamic site acceleration

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

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

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

Reference:

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

NEW QUESTION 4

- (Exam Topic 1)

You need to 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	Value
Plan	<div style="border: 1px solid black; padding: 5px;"> Basic Standard Premium Isolated </div>
Instance Count	<div style="border: 1px solid black; padding: 5px;"> 1 10 20 100 </div>

- 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 add -g shipping-apis-test-rg -n web
```

config

deployment

up

remove

```
-- slot http://*.wideworldimporters.com
```

allowed-origins

http://test-shippingapi.wideworldimporters.com

http://test.wideworldimporters.com

http://www.wideworldimporters.com

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

Box 2: add
Box 3: allowed-origins
Box 4: <http://testwideworldimporters.com/> References:
<http://donovanbrown.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-with-Azure-App-Service>

NEW QUESTION 6

- (Exam Topic 3)

You develop a web application.

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

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

Actions

Answer Area

Select **Manifest** from the middle-tier service registration.

In Enterprise Applications, select **New application**.


Add a Cryptographic key.

Create a new application and provide the name, account type, and redirect URL

Select the Azure AD instance.

Use an access token to access the secure resource.

In App Registrations, select **New registration**.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Register a new application using the Azure portal

- > Sign in to the Azure portal using either a work or school account or a personal Microsoft account.
- > If your account gives you access to more than one tenant, select your account in the upper right corner. Set your portal session to the Azure AD tenant that you want.
- > Search for and select Azure Active Directory. Under Manage, select App registrations.
- > Select New registration. (Step 1)
- > In Register an application, enter a meaningful application name to display to users.
- > Specify who can use the application. Select the Azure AD instance. (Step 2)
- > Under Redirect URI (optional), select the type of app you're building: Web or Public client (mobile & desktop). Then enter the redirect URI, or reply URL, for your application. (Step 3)
- > When finished, select Register.

NEW QUESTION 7

- (Exam Topic 3)

You use Azure Table storage to store customer information for an application. The data contains customer details and is partitioned by last name. You need to create a query that returns all customers with the last name Smith. Which code segment should you use?

- A. `TableQuery.GenerateFilterCondition("PartitionKey", Equals, "Smith")`
- B. `TableQuery.GenerateFilterCondition("LastName", Equals, "Smith")`
- C. `TableQuery.GenerateFilterCondition("PartitionKey", QueryComparisons.Equal, "Smith")`
- D. `TableQuery.GenerateFilterCondition("LastName", QueryComparisons.Equal, "Smith")`

Answer: C

Explanation:

Retrieve all entities in a partition. The following code example specifies a filter for entities where 'Smith' is the partition key. This example prints the fields of each

entity in the query results to the console.

```
Construct the query operation for all customer entities where PartitionKey="Smith".
TableQuery<CustomerEntity> query = new
TableQuery<CustomerEntity>().Where(TableQuery.GenerateFilterCondition("PartitionKey",
QueryComparisons.Equal, "Smith"));
```

References:

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

NEW QUESTION 8

- (Exam Topic 3)

You are developing a microservices solution. You plan to deploy the solution to a multinode Azure Kubernetes Service (AKS) cluster.

You need to deploy a solution that includes the following features:

- > reverse proxy capabilities
- > configurable traffic routing
- > TLS termination with a custom certificate

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

NOTE: Each correct selection is worth one point.

Answer Area

Components	Action	Component
Helm		
Draft	Deploy solution.	<input type="text"/>
Brigade	View cluster and external IP addressing.	<input type="text"/>
KubeCtl	Implement a single, public IP endpoint that is routed to multiple microservices.	<input type="text"/>
Ingress Controller		
CoreDNS		
Virtual Kubelet		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Helm

To create the ingress controller, use Helm to install nginx-ingress. Box 2: kubectl

To find the cluster IP address of a Kubernetes pod, use the kubectl get pod command on your local machine, with the option -o wide .

Box 3: Ingress Controller

An ingress controller is a piece of software that provides reverse proxy, configurable traffic routing, and TLS termination for Kubernetes services. Kubernetes ingress resources are used to configure the ingress rules and routes for individual Kubernetes services.

Reference:

<https://docs.microsoft.com/bs-cyrl-ba/azure/aks/ingress-basic> <https://www.digitalocean.com/community/tutorials/how-to-inspect-kubernetes-networking>

NEW QUESTION 9

- (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

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

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

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

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



- 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 10

- (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 style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; padding: 2px;">Can be restored</div> <div style="padding: 2px;">Cannot be restored</div> </div>
Snapshot 1	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; padding: 2px;">Can be restored</div> <div style="padding: 2px;">Cannot be restored</div> </div>
Snapshot 2	<div style="border: 1px solid black; padding: 2px;"> <div style="border-bottom: 1px solid black; padding: 2px;">Can be restored</div> <div style="padding: 2px;">Cannot be restored</div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Can be restored

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

Box 2: Cannot be restored It has been deleted.

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

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

NEW QUESTION 10

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead use Azure Cache for Redis.

Note: Azure Cache for Redis provides a session state provider that you can use to store your session state in-memory with Azure Cache for Redis instead of a SQL Server database. To use the caching session state provider, first configure your cache, and then configure your ASP.NET application for cache using the Azure Cache for Redis Session State NuGet package.

References:

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

NEW QUESTION 11

- (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, , string pk, string rk)
    {
        
        
        
        
        
        
        
        
        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 12

- (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(
    CloudConfigurationManager.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 13

- (Exam Topic 3)

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

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

You need to create the custom role.

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

Item	Value
Powershell command	<div style="border: 1px solid black; padding: 5px;"> <p>Get-AzureRmRoleDefinition-Name "Reader" ConvertTo-Json Out-File C:\SupportRole.json</p> <p>Get-AzureRmRoleDefinition-Name "Operator" ConvertTo-Json Out-File C:\SupportRole.json</p> <p>Set-AzureRmRoleDefinition-Name "Reader" Input-File C:\SupportRole.json</p> <p>Set-AzureRmRoleDefinition Input-File C:\SupportRole.json</p> </div>
Actions section	<div style="border: 1px solid black; padding: 5px;"> <p>"/read*, *Microsoft.Support/*"</p> <p>"/read*</p> <p>"/read*, *Microsoft.Support/*"</p> <p>"/read*</p> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 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.

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

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

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

Solution: Provision an Azure Event Grid. Configure event filtering to evaluate the device identifier. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Note: An event is a lightweight notification of a condition or a state change. Event hubs is usually used reacting to status changes.

Reference:

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

NEW QUESTION 20

- (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 placed 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. ();

        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.)));

            outputFileList.Add(new OutputFile(fileTask.Output,
                new OutputFileDestination(failedContainer),
                new OutputFileUploadOptions(OutputFileUploadCondition.)));

            task. = outputFileList;

            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 24

- (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 28

- (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 30

- (Exam Topic 3)

You have a web service that is used to pay for food deliveries. The web service uses Azure Cosmos DB as the data store.

You plan to add a new feature that allows users to set a tip amount. The new feature requires that a property named tip on the document in Cosmos DB must be present and contain a numeric value.

There are many existing websites and mobile apps that use the web service that will not be updated to set the tip property for some time.

How should you complete the trigger?

NOTE: Each correct selection is worth one point.

```
function ensureTip() {
var r = 

|                             |
|-----------------------------|
| _.value();                  |
| _.readDocument('item');     |
| getContext().getRequest();  |
| getContext().getResponse(); |


var i = r.getBody();


|                                            |
|--------------------------------------------|
| if (!("tip" in i)) {                       |
| if (request.getValue("tip") === null){     |
| if (isNaN(i)["tip"]    i["tip"]=== null) { |
| if (typeof_.pluck("tip") == 'number') {    |
| i["tip"] = 0;                              |
| }                                          |



|                      |
|----------------------|
| r.setBody(i);        |
| r.setValue(i);       |
| _.upsertDocument(i); |
| _.replaceDocument(i) |


}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: getContext().getRequest(); Box 2: if(isNaN(i)["tip"] ..

In JavaScript, there are two ways to check if a variable is a number :

isNaN() – Stands for “is Not a Number”, if variable is not a number, it return true, else return false. typeof – If variable is a number, it will returns a string named “number”.

Box 3:r.setBody(i);

// update the item that will be created References:

<https://docs.microsoft.com/bs-latn-ba/azure/cosmos-db/how-to-write-stored-procedures-triggers-udfs>

<https://mkyong.com/javascript/check-if-variable-is-a-number-in-javascript/>

NEW QUESTION 35

- (Exam Topic 3)

You have an app that stores player scores for an online game. The app stores data in Azure tables using a class named PlayerScore as the table entity. The table is populated with 100,000 records.

You are reviewing the following section of code that is intended to retrieve 20 records where the player score exceeds 15,000. (Line numbers are included for reference only.)

```

1 public void GetScore(string playerId, int score, string gameName)
2 {
3     TableQuery<DynamicTableEntity> query = new TableQuery<DynamicTableEntity>().Select(new string[] { "Score" })
        .Where(TableQuery.GenerateFilterConditionForInt("Score", QueryComparisons.GreaterThanOrEqual, 15000)).Take
(20);
4     EntityResolver<KeyValuePair<string, int?>> resolver =
        (partitionKey, rowKey, ts, props, etag) => new KeyValuePair<string, int?>(rowKey, props["Score"].Int32Value);
5     foreach (var scoreItem in scoreTable.ExecuteQuery(query, resolver, null, null))
6     {
7         Console.WriteLine($"{scoreItem.Key} {scoreItem.Value}");
8     }

9 public class PlayerScore : TableEntity
10 {
11     public PlayerScore(string gameId, string playerId, int score, long timePlayed)
12     {
13         PartitionKey = gameId;
14         RowKey = playerId;
15         Score = score;
16         TimePlayed = timePlayed;
17     }
18     public int Score { get; set; }
19     public long TimePlayed { get; set; }
20 }

```

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

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

```

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

```

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

	Yes	No
The code queries the Azure table and retrieves the TimePlayed property from the table	<input type="radio"/>	<input type="radio"/>
The code will display a maximum of twenty records.	<input type="radio"/>	<input type="radio"/>
All records will be sent to the client. The client will display records for scores greater than or equal to 15,000.	<input type="radio"/>	<input type="radio"/>
The scoreItem.Key property of the KeyValuePairs that ExecuteQuery returns will contain a value for PlayerID.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

Box 2: Yes

The TableQuery.Take method defines the upper bound for the number of entities the query returns. Example:

query.Take(10);

Box 3: Yes

Box 4: Yes

References:

<https://www.vkinfotek.com/azureqa/how-do-i-query-azure-table-storage-using-tablequery-class.html>

NEW QUESTION 40

- (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 44

- (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 style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="padding: 2px;">2</td></tr> <tr><td style="padding: 2px;">4</td></tr> <tr><td style="padding: 2px;">8</td></tr> <tr><td style="padding: 2px;">16</td></tr> </table> </div>	2	4	8	16
2					
4					
8					
16					
Pricing tier	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td style="padding: 2px;">Isolated</td></tr> <tr><td style="padding: 2px;">Standard</td></tr> <tr><td style="padding: 2px;">Premium</td></tr> <tr><td style="padding: 2px;">Consumption</td></tr> </table> </div>	Isolated	Standard	Premium	Consumption
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 48

- (Exam Topic 3)

You provide an Azure API Management managed web service to clients. The back end web service implements HTTP Strict Transport Security (HSTS).

Every request to the backend service must include a valid HTTP authorization header. You need to configure the Azure API Management instance with an authentication policy. Which two policies can you use? Each correct answer presents a complete solution NOTE: Each correct selection is worth one point.

- A. Certificate Authentication
- B. Basic Authentication
- C. OAuth Client Credential Grant
- D. Digest Authentication

Answer: AC

NEW QUESTION 49

- (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 51

- (Exam Topic 3)

A company is implementing a publish-subscribe (Pub/Sub) messaging component by using Azure Service Bus. You are developing the first subscription application.

In the Azure portal you see that messages are being sent to the subscription for each topic. You create and initialize a subscription client object by supplying the correct details, but the subscription application is still not consuming the messages.

You need to complete the source code of the subscription client What should you do?

- A. `await subscriptionClient.CloseAsync();`
- B. `await subscriptionClient.AddRuleAsync(new RuleDescription(RuleDescription.DefaultRuleName, new TrueFilter()));`
- C. `subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);`
- D. `subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName);`

Answer: C

Explanation:

Using topic client, call RegisterMessageHandler which is used to receive messages continuously from the entity. It registers a message handler and begins a new thread to receive messages. This handler is waited on every time a new message is received by the receiver.

`subscriptionClient.RegisterMessageHandler(ReceiveMessagesAsync, messageHandlerOptions);` References:

<https://www.c-sharpcorner.com/article/azure-service-bus-topic-and-subscription-pub-sub/>

NEW QUESTION 53

- (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 58

- (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 60

- (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 62

- (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

Comparison

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 67

- (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.	<ul style="list-style-type: none"> Azure PowerShell Visual Studio command prompt Azure Migrate Azure Backup
Store images.	<ul style="list-style-type: none"> 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 70

- (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 style="border: 1px solid black; padding: 2px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="border-top: 1px solid black; border-bottom: 1px solid black; padding: 2px;">3</div> <div style="border-top: 1px solid black; border-bottom: 1px solid black; padding: 2px;">4</div> <div style="border-top: 1px solid black; border-bottom: 1px solid black; padding: 2px;">6</div> <div style="border-top: 1px solid black; border-bottom: 1px solid black; padding: 2px;">12</div> </div>
Partition Key	<div style="border: 1px solid black; padding: 2px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="border-top: 1px solid black; border-bottom: 1px solid black; padding: 2px;">Highway</div> <div style="border-top: 1px solid black; border-bottom: 1px solid black; padding: 2px;">Department</div> <div style="border-top: 1px solid black; border-bottom: 1px solid black; padding: 2px;">Timestamp</div> <div style="border-top: 1px solid black; border-bottom: 1px solid black; padding: 2px;">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 71

- (Exam Topic 3)

You have an application that provides weather forecasting data to external partners. You use Azure API Management to publish APIs. You must change the behavior of the API to meet the following requirements:

- Support alternative input parameters.
- Remove formatting text from responses.
- Provide additional context to back-end services.

Which types of policies should you implement? To answer, drag the policy types to the correct scenarios. Each policy 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.

Policy types	Requirement	Policy type
Inbound	Rewrite the request URL to match to the format expected by the web service.	policy type
Outbound	Remove formatting text from responses.	policy type
Backend	Forward the user ID that is associated with the subscription key for the original request to the back-end service.	policy type

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Policy types	Requirement	Policy type
Inbound	Rewrite the request URL to match to the format expected by the web service.	Outbound
Outbound	Remove formatting text from responses.	Inbound
Backend	Forward the user ID that is associated with the subscription key for the original request to the back-end service.	Backend

NEW QUESTION 73

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

```
YAML:
apiVersion: v1 kind: Service metadata:
name: internal-app annotations:
service.beta.kubernetes.io/azure-load-balancer-internal: "true" spec:
type: LoadBalancer ports:
- port: 80 selector:
app: internal-app
References:
https://docs.microsoft.com/en-us/azure/aks/internal-lb
```

NEW QUESTION 77

- (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 79

- (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.
- The code has a bug and will insert an additional copy of the Game record with a new Id.
- The code has a bug and will insert the wrong gameId value.
- There is a valid many-to-many relationship between Players and Games.

A. Mastered

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: keyvault

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

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

Box 2: keyvault key
The Azure platform needs to be granted access to request the cryptographic keys when the VM boots to decrypt the virtual disks. Create a cryptographic key in your Key Vault with az keyvault key create. The following example creates a key named myKey:

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

Box 3: vm
Create a VM with az vm create. Only certain marketplace images support disk encryption. The following example creates a VM named myVM using an Ubuntu 16.04 LTS image:

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

Box 4: vm encryption
Encrypt your VM with az vm encryption enable: az vm encryption enable \

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

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

References:

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

NEW QUESTION 87

- (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
<div style="border: 1px solid black; padding: 2px; display: inline-block;">Triggered</div>	Run on all instances that the web app runs on. Optionally restrict the WebJob to a single instance.	<input type="text"/>
<div style="border: 1px solid black; padding: 2px; display: inline-block;">Continuous</div>	Run on a single instance that Azure select for load balancing.	<input type="text"/>
	Supports remote debugging	<input type="text"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Continuous

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

Box 2: Triggered

Triggered runs on a single instance that Azure selects for load balancing.

Box 3: Continuous
Continuous supports remote debugging. Note:

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

Continuous

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.

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

Supports remote debugging.

Triggered

Starts only when triggered manually or on a schedule.

Runs on a single instance that Azure selects for load balancing.

Doesn't support remote debugging.

References:

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

NEW QUESTION 88

- (Exam Topic 3)

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

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

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

Command segments

Answer Area

- az aks get-credentials
- az appservice plan create
- az aks create
- az group create
- kubectl apply



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: az group create

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

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

Step 2 : az aks create

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

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

Step 4: az aks get-credentials

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

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

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

NEW QUESTION 91

- (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 92

.....

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>