

AZ-203.prepaway.premium.exam.157q

Number: AZ-203
Passing Score: 800
Time Limit: 120 min
File Version: 8.0



AZ-203

Developing Solutions for Microsoft Azure

Version 8.0

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Develop Azure Infrastructure as a Service Compute Solutions

Testlet 1

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

Background

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements

Policy service

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

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

Policies

Log Policy

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named **logdrop**. Logs must remain in the container for 15 days.

Authentication events

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

PolicyLib

You have a shared library named **PolicyLib** that contains functionality common to all ASP.NET Core web services and applications. The **PolicyLib** library must:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

- Exclude non-user actions from Application Insights telemetry.
- Provide methods that allow a web service to scale itself
- Ensure that scaling actions do not disrupt application usage

Other

Anomaly detection service

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service.

If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring

All web applications and services have health monitoring at the /health service endpoint.

Issues

Policy loss

When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue

When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency

Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

```

EventGridController.cs
EG001 public class EventGridController : Controller
EG002 {
EG003     public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG004     public IActionResult Process([FromBody] string eventsJson
EG005     {
EG006         var events = JArray.Parse(eventsJson);
EG007
EG008         foreach (var @event in events)
EG009         {
EG010             EventId.Value = @event ["id"].ToString();
EG011             if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG012             {
EG013                 SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG014             }
EG015
EG016             {
EG017                 EnsureLogging(@event["subject"].ToString());
EG018             }
EG019         }
EG020         return null;
EG021     }
EG022     private void EnsureLogging(string resource)
EG023     {
EG024         . . .
EG025     }
EG026     private async Task SendToAnomalyDetectionService(string uri)
EG027     {
EG028         var content = GetLogData(uri);
EG029         var scoreRequest = new
EG030         {
EG031             Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG032             {
EG033                 {
EG034                     "input1",
EG035                     new List<Dictionary<string, string>>()
EG036                     {
EG037                         new Dictionary<string, string>()
EG038                         {
EG039                             {
EG040                                 "logcontent", content
EG041                             }
EG042                         }
EG043                     }
EG044                 },
EG045             },
EG046             GlobalParameters = new Dictionary<string, string>() { }
EG047         };
EG048         var result = await (new HttpClient()).PostAsJsonAsync(". . .", scoreRequest);
EG049         var rawModelResult = await result.Content.ReadAsStringAsync();
EG050         var modelResult = JObject.Parse(rawModelResult);
EG051         if (modelResult["notify"].HasValues)
EG052         {
EG053             . . .
EG054         }
EG055     }
EG056     private (string name, string resourceGroup) ParseResourceId(string
resourceId)
EG057     {
EG058         . . .
EG059     }
EG060     private string GetLogData(string uri)
EG061     {
EG062         . . .
EG063     }
EG064     static string BlobStoreAccountSAS(string containerName)
EG065     {
EG066         . . .
EG067     }
EG068 }

```

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04 public string subject { get; set; }
LE05 public DateTime eventTime { get; set; }
LE06 public Dictionary<string, string> data { get; set; }
LE07 public string Serialize()
LE08 {
LE09     return JsonConvert.SerializeObject(this);
LE10 }
LE11 }
```

QUESTION 1

You need to resolve a notification latency issue.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Set Always On to **false**.
- B. Set Always On to **true**.
- C. Ensure that the Azure Function is set to use a consumption plan.
- D. Ensure that the Azure Function is using an App Service plan.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Azure Functions can run on either a Consumption Plan or a dedicated App Service Plan. If you run in a dedicated mode, you need to turn on the Always On setting for your Function App to run properly. The Function runtime will go idle after a few minutes of inactivity, so only HTTP triggers will actually "wake up" your functions. This is similar to how WebJobs must have Always On enabled.

Scenario: Notification latency: Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

Anomaly detection service: You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service.

If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

References:

<https://github.com/Azure/Azure-Functions/wiki/Enable-Always-On-when-running-on-dedicated-App-Service-Plan>

Develop Azure Infrastructure as a Service Compute Solutions

Testlet 2

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

LabelMaker app

Coho Winery produces bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends **data** to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the **LabelMaker** application to include artwork and text for custom label designs.

Requirements

Data

You identify the following requirements for data management and manipulation:

- Order data is stored as nonrelational JSON and must be queried using Structured Query Language (SQL).
- Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Security

You have the following security requirements:

- Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
- External partners must use their own credentials and authenticate with their organization's identity management solution.
- External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.
- Storage of e-commerce application settings must be maintained in Azure Key Vault.
- E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure


Active Directory (AAD).

- Conditional access policies must be applied at the application level to protect company content
- The LabelMaker applications must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.
- LabelMaker app

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture



Issues

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communications timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order .json

Relevant portions of the app files are shown below. Line numbers are included for reference only. This JSON file contains a representation of the data for an order that includes a single item.

```

01 {
02   "id" : 1,
03   "customers" : [
04     {
05       "familyName" : "Doe",
06       "givenName" : "John",
07       "customerid" : 5
08     }
09   ],
10   "line_items" : [
11     {
12       "fulfillable_quantity" : 1,
13       "id" : 6,
14       "price" : "199.99",
15       "product_id" : 7513594,
16       "quantity" : 1,
17       "requires_shipping" : true,
18       "sku" : "SFC-342-N" ,
19       "title" : "Surface Go",
20       "vendor" : "Microsoft" ,
21       "name" : "Surface Go - 8GB",
22       "taxable" : true,
23       "tax_lines" : [
24         {
25           "title" : "State Tax",
26           "price" : "3.93",
27           "rate" : 0.06
28         }
29       ],
30       "total_discount" : "5.00"

```

art.com/

```

31   "discount_allocations" : [
32     {

```

QUESTION 1

HOTSPOT

You need to ensure that you can deploy the LabelMaker application.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

az create --name --location eastus

az create --resource-group CohoWineryLabelMaker --name

LabelMakerCluster --node-count 5 --enable-addons

Correct Answer:

Answer Area

az create --name --location eastus

az create --resource-group CohoWineryLabelMaker --name

LabelMakerCluster --node-count 5 --enable-addons

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: group

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.

The following example creates a resource group named myResourceGroup in the westeurope location.

az group create --name myResourceGroup --location westeurope

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Box 2: CohoWinterLabelMaker

Use the resource group named, which is used in the second command.

Box 3: aks

The command `az aks create`, is used to create a new managed Kubernetes cluster.

Box 4: monitoring

Scenario: LabelMaker app

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Develop Azure Infrastructure as a Service Compute Solutions

Testlet 3

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment

Windows Server 2016 virtual machine

This virtual machine (VM) runs Biz Talk Server 2016. The VM runs the following workflows:

- Ocean Transport – This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.
- Inland Transport – This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

- Container API – This API provides container information including weight, contents, and other attributes.
- Location API – This API provides location information regarding shipping ports of call and truck stops.
- Shipping REST API – This API provides shipping information for use and display on the shipping website.

Shipping Data

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

Shipping Web Site

The site displays shipping container tracking information and container contents. The site is located at <http://shipping.wideworldimporters.com>


Proposed solution

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:



Shipping Logic App

The Shipping Logic app must meet the following requirements:

- Support the ocean transport and inland transport workflows by using a Logic App.
- Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
- Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
- Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app

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

REST APIs

The REST API's that support the solution must meet the following requirements:

- Secure resources to the corporate VNet.
- 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.

Shipping data

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

Shipping website

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

Issues

Windows Server 2016 VM

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

Shipping website and REST APIs

The following error message displays while you are testing the website:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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.

QUESTION 1

You need to support the requirements for the Shipping Logic App.

What should you use?

- A. Azure Active Directory Application Proxy
- B. Point-to-Site (P2S) VPN connection
- C. Site-to-Site (S2S) VPN connection
- D. On-premises Data Gateway

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Before you can connect to on-premises data sources from Azure Logic Apps, download and install the on-premises data gateway on a local computer. The gateway works as a bridge that provides quick data transfer and encryption between data sources on premises (not in the cloud) and your logic apps.

The gateway supports BizTalk Server 2016.

Note: Microsoft have now fully incorporated the Azure BizTalk Services capabilities into Logic Apps and Azure App Service Hybrid Connections.

Logic Apps Enterprise Integration pack bring some of the enterprise B2B capabilities like AS2 and X12, EDI standards support

Scenario: The Shipping Logic app must meet the following requirements:

- Support the ocean transport and inland transport workflows by using a Logic App.
- Support industry standard protocol X12 message format for various messages including vessel content details and arrival notices.
- Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
- Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-gateway-install>

QUESTION 2

HOTSPOT

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.

Hot Area:

Answer Area

Setting	Value
Plan	▼ Basic Standard Premium Isolated
Instance Count	▼ 1 10 20 100

Correct Answer:

Answer Area

Setting	Value
Plan	▼ Basic Standard Premium Isolated
Instance Count	▼ 1 10 20 100

Section: (none)

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation

Explanation/Reference:

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/>

Develop Azure Infrastructure as a Service Compute Solutions

Question Set 4

QUESTION 1

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.EnableJobAsync(String, IEnumerable<BatchClientBehavior>, Cancellation.Token)
- B. JobOperations.CreateJob()
- C. CloudJob.Enable(IEnumerable<BatchClientBehavior>)
- D. JobOperations.EnableJob(String, IEnumerable<BatchClientBehavior>)
- E. CloudJob.CommitAsync(IEnumerable<BatchClientBehavior>, Cancellation.Token)

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 2

DRAG DROP

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.

Select and Place:

Answer Area

WebJob types	Scenario	WebJob type
Triggered	Run on all instances that the web app runs on. Optionally restrict the WebJob to a single instance.	
Continuous	Run on a single instance that Azure select for load balancing.	
	Supports remote debugging	

Correct Answer:

Answer Area

WebJob types	Scenario	WebJob type
Triggered	Run on all instances that the web app runs on. Optionally restrict the WebJob to a single instance.	Continuous
Continuous	Run on a single instance that Azure select for load balancing.	Triggered
	Supports remote debugging	Continuous

Section: (none)

Explanation

Explanation/Reference:

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://vceoreteconvert.com/>

<https://t.me/learningnets>

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

QUESTION 3

DRAG DROP

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

You need to create compute nodes for the solution on Azure Batch.

What should you do?

Put the actions in the correct order.

Select and Place:

Select these

- In the Azure CLI, run the command: `az batch account create`
- In Azure CLI, run the command: `az batch task create`
- In Azure CLI, run the command: `az batch pool create`
- In Azure CLI, run the command: `az batch job create`

Place here

Correct Answer:

Select these

Place here

- In the Azure CLI, run the command: `az batch account create`
- In Azure CLI, run the command: `az batch pool create`
- In Azure CLI, run the command: `az batch job create`
- In Azure CLI, run the command: `az batch task create`

Section: (none)

Explanation

Explanation/Reference:

Explanation:

With the Azure CLI:

Step 1: In the Azure CLI, run the command: `az batch account create`
First we create a batch account.

Step 2: In Azure CLI, run the command: `az batch pool create`
Now that you have a Batch account, create a sample pool of Linux compute nodes using the `az batch pool create` command.

Step 3: In Azure CLI, run the command: `az batch job create`
Now that you have a pool, create a job to run on it. A Batch job is a logical group for one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. Create a Batch job by using the `az batch job create` command.

Step 4: In Azure CLI, run the command: `az batch task create`
Now use the `az batch task create` command to create some tasks to run in the job.

References:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

<https://docs.microsoft.com/en-us/azure/batch/quick-create-cli>

QUESTION 4
DRAG DROP

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.

Select and Place:

Command segments

Answer Area

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



Correct Answer:

Command segments

Answer Area

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



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

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Step 1: az group create

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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/>

QUESTION 5

DRAG DROP

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.

Select and Place:

Actions

Encrypted the on-premises VHD by using BitLocker with a TPM.
Upload the VM to Azure Storage.

Run the Azure PowerShell command `Set-AzureRmVMDiskEncryptionExtension`.

Run the Azure PowerShell command `New-AzureRmVM`.

Encrypt the on-premises VHD by using BitLocker without a TPM.
Upload the VM to Azure Storage.

Run the Azure PowerShell command `Set-AzureRmVMDiskEncryptionExtension`.

Answer Area



Correct Answer:

Actions

Encrypted the on-premises VHD by using BitLocker with a TPM.
Upload the VM to Azure Storage.

Run the Azure PowerShell command `Set-AzureRmVMOSDisk`.

Run the Azure PowerShell command `New-AzureRmVM`.

Encrypt the on-premises VHD by using BitLocker without a TPM.
Upload the VM to Azure Storage.

Run the Azure PowerShell command `Set-AzureRmVMDiskEncryptionExtension`.

Answer Area

Encrypt the on-premises VHD by using BitLocker without a TPM.
Upload the VM to Azure Storage.

Run the Azure PowerShell command `Set-AzureRmVMOSDisk`.

Run the Azure PowerShell command `Set-AzureRmVMDiskEncryptionExtension`.

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Step 1: Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage

Step 2: Run the Azure PowerShell command `Set-AzureRmVMOSDisk`

To use an existing disk instead of creating a new disk you can use the `Set-AzureRmVMOSDisk` command.

Example:

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

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

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

```
$vm = Set-AzureRmVMOSDisk -VM $vm -VhdUri $osDiskVhdUri -name $osDiskName -Create
```

Step 3: Run the Azure PowerShell command `Set-AzureRmVMDiskEncryptionExtension`

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

Incorrect:

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

References:

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

QUESTION 6

DRAG DROP

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 Dockerfile 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.

Select and Place:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Commands**Answer Area**

```
RUN powershell ./setupScript.ps1  
CMD ["dotnet", "ContosoApp.dll"]
```

```
EXPOSE ./ContosoApp/ /apps/ContosoApp
```

```
COPY . .
```

```
FROM microsoft/aspnetcore:2.0
```

```
WORKDIR /apps/ContosoApp
```

```
CMD powershell ./setupScript.ps1  
ENTRYPOINT ["dotnet", "ContosoApp.dll"]
```

**Correct Answer:****Commands****Answer Area**

```
RUN powershell ./setupScript.ps1  
CMD ["dotnet", "ContosoApp.dll"]
```

```
EXPOSE ./ContosoApp/ /apps/ContosoApp
```

```
COPY . .
```

```
FROM microsoft/aspnetcore:2.0
```

```
WORKDIR /apps/ContosoApp
```

```
CMD powershell ./setupScript.ps1  
ENTRYPOINT ["dotnet", "ContosoApp.dll"]
```

```
WORKDIR /apps/ContosoApp
```

```
COPY . .
```



```
EXPOSE ./ContosoApp/ /apps/ContosoApp
```

```
CMD powershell ./setupScript.ps1  
ENTRYPOINT ["dotnet", "ContosoApp.dll"]
```

**Section: (none)****Explanation****Explanation/Reference:**

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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>

QUESTION 7

DRAG DROP

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

You have the following parameters:

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

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

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

Select and Place:

Command segments

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

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

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

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

Answer Area



Correct Answer:

Command segments

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

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

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

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

Answer Area

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

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

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

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

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 8

HOTSPOT

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Box 1: No
ExpirationTime - The time that the message expires.
InsertionTime - The time that the message was added to the queue.

Box 2: Yes
maxDequeueCount - The number of times to try processing a message before moving it to the poison queue. Default value is 5.

Box 3: Yes
When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them. By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.

Box 4: Yes

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

QUESTION 9

DRAG DROP

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.

Select and Place:

Azure CLI Commands

Answer Area



Correct Answer:

Azure CLI Commands

Answer Area



Section: (none)

Explanation

Explanation/Reference:

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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>

QUESTION 10

DRAG DROP

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

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

You need to ensure that the requirements are met.

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

Select and Place:

Features	Answer Area	
	Requirement	Feature
Run Command	Firewall configuration	
Serial console	Supporting services script	
Hybrid Runbook Worker		
Custom Script Extension		

Correct Answer:

Features	Answer Area	
	Requirement	Feature
Serial console	Firewall configuration	Run Command
	Supporting services script	Hybrid Runbook Worker
Custom Script Extension		

Section: (none)

Explanation

Explanation/Reference:

References:

<https://docs.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker>

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-command>

QUESTION 11

DRAG DROP

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.

Select and Place:

Components	Answer Area	
	Action	Component
Helm	Deploy solution.	Component
Draft	View cluster and external IP addressing.	Component
Brigade	Implement a single, public IP endpoint that is routed to multiple microservices.	Component
KubeCtl		
Ingress Controller		
CoreDNS		
Virtual Kubelet		

Correct Answer:

Components	Answer Area	
	Action	Component
	Deploy solution.	Helm
Draft	View cluster and external IP addressing.	KubeCtl
Brigade	Implement a single, public IP endpoint that is routed to multiple microservices.	Ingress Controller
CoreDNS		
Virtual Kubelet		

Section: (none)

Explanation

Explanation/Reference:

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.

Incorrect Answers:

Virtual Kubelet: Virtual Kubelet is an open-source Kubernetes kubelet implementation that masquerades as a kubelet. This allows Kubernetes nodes to be backed by Virtual Kubelet providers such as serverless cloud container platforms.

CoreDNS: CoreDNS is a flexible, extensible DNS server that can serve as the Kubernetes cluster DNS. Like Kubernetes, the CoreDNS project is hosted by the CNCF.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Reference:

<https://docs.microsoft.com/bs-cyrl-ba/azure/aks/ingress-basic>

<https://www.digitalocean.com/community/tutorials/how-to-inspect-kubernetes-networking>

QUESTION 12

HOTSPOT

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.

Hot Area:

Answer Area

Action	Tool or service
Generalize the VM.	Azure Power Shell Visual Studio command prompt Azure Migrate Azure Backup
Store images.	Azure Blob Storage Visual Data Lake Storage Azure File Storage Azure Table Storage

Correct Answer:

Answer Area

Action	Tool or service
Generalize the VM.	Azure Power Shell Visual Studio command prompt Azure Migrate Azure Backup
Store images.	Azure Blob Storage Visual Data Lake Storage Azure File Storage Azure Table Storage

Section: (none)

Explanation

Explanation/Reference:

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-vm-using-powershell>

QUESTION 13

DRAG DROP

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.

Select and Place:

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

Correct Answer:

Code segments

Ingress

Deployment

ingress.class

Answer Area

apiVersion: v1

kind: Service

metadata:

name: web-app

annotations:

service.beta.kubernetes.azure-load-balancer-internal : "true"

spec:

type: LoadBalancer

ports:

- port: 80

selector:

app: web-app

Section: (none)

Explanation

Explanation/Reference:

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>

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

QUESTION 14

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

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

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

The solution must allow dynamic creation and management of all Azure resources within the AKS cluster.

You need to configure an AKS cluster for use with the Azure APIs.

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

References:

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

QUESTION 15

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

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

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

The solution must allow dynamic creation and management of all Azure resources within the AKS cluster.

You need to configure an AKS cluster for use with the Azure APIs.

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation:

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

References:

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

QUESTION 16

HOTSPOT

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
<code>fileTasks</code>	a list of tasks to be run
<code>jobId</code>	the identifier that must be assigned to the job
<code>outputContainerSasUrl</code>	a storage SAS URL to store successfully converted files
<code>failedContainerSasUrl</code>	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.

Hot Area:

Answer Area

```
public List<CloudTask> StartTasks(List<FileTask> fileTasks, string jobId,
    string outputContainerSasUrl, string failedContainerSasUrl, string poolId)
{
    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;

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

GetJob
GetTask
EnableJob
CreateJob

TaskSuccess
TaskFailure
TaskCompletion


TaskSuccess
TaskFailure
TaskCompletion


OutputFiles
FilesToStage
ResourceFiles
StageFiles


Correct Answer:

Answer Area

```
public List<CloudTask> StartTasks(List<FileTask> fileTasks, string jobId,
    string outputContainersSasUrl, string failedContainersSasUrl, string poolId)
{
    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(outputContainersSasUrl);
            OutputFileBlobContainerDestination failedContainer =
                new OutputFileBlobContainerDestination (failedContainersSasUrl);
            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;

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

Section: (none)

Explanation

Explanation/Reference:

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://vceoreteconvert.com/>

<https://t.me/learningnets>

<https://docs.microsoft.com/en-us/azure/batch/batch-task-output-files>

QUESTION 17

DRAG DROP

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

You need to create compute nodes for the solution on Azure Batch.

What should you do?

Put the actions in the correct order.

Select and Place:

Select these

In the Azure portal, add a Job

In the Azure portal, create a Batch account.

In the Azure portal, create tasks

In the Azure portal, create a pool of compute nodes

Place here

Correct Answer:

Select these

Place here

In the Azure portal, create a Batch account.

In the Azure portal, create a pool of compute nodes

In the Azure portal, add a Job

In the Azure portal, create tasks

Section: (none)

Explanation

Explanation/Reference:

Explanation:

With the Azure Portal:

Step 1: In the Azure portal, create a Batch account.

First we create a batch account.

Step 2: In the Azure portal, create a pool of compute nodes

Now that you have a Batch account, create a sample pool of Windows compute nodes for test purposes.

Step 3: In the Azure portal, add a Job.

Now that you have a pool, create a job to run on it. A Batch job is a logical group for one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. Initially the job has no tasks.

Step 4: In the Azure portal, create tasks

Now create sample tasks to run in the job. Typically you create multiple tasks that Batch queues and distributes to run on the compute nodes.

References:

<https://docs.microsoft.com/en-us/azure/batch/quick-create-portal>

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

QUESTION 18
DRAG DROP

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

You need to create compute nodes for the solution on Azure Batch.

What should you do?

Put the actions in the correct order.

Select and Place:

Select these	Place here
In a .NET method, call the method: BatchClient.PoolOperations.Create Pool	
In a .NET method, call the method: BatchClient.PoolOperations.CreateJob	
In the Azure portal, create a Batch account	
In a .NET method, call the method: BatchClient.JobOperations.AddTask	

Correct Answer:

Select these	Place here
	In the Azure portal, create a Batch account
	In a .NET method, call the method: BatchClient.PoolOperations.Create Pool
	In a .NET method, call the method: BatchClient.PoolOperations.CreateJob
	In a .NET method, call the method: BatchClient.JobOperations.AddTask

Section: (none)
Explanation

Explanation/Reference:
Explanation:

With .NET:

Step 1: In the Azure portal, create a Batch account.
First we create a batch account.

Step 2: In a .NET method, call the method: BatchClient.PoolOperations.CreatePool
Now that you have a Batch account, create a sample pool of Windows compute nodes for test purposes. To create a Batch pool, the app uses the BatchClient.PoolOperations.CreatePool method to set the number of nodes, VM size, and a pool configuration.

Step 3: In a .NET method, call the method: BatchClient.PoolOperations.CreateJob
Now that you have a pool, create a job to run on it. A Batch job is a logical group for one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. Initially the job has no tasks. The app uses the BatchClient.JobOperations.CreateJob method to create a job on your pool.

Step 4: In a .NET method, call the method: `batchClient.JobOperations.AddTask`
Now create sample tasks to run in the job. Typically you create multiple tasks that Batch queues and distributes to run on the compute nodes. The app adds tasks to the job with the `AddTask` method, which queues them to run on the compute nodes.
For example: `batchClient.JobOperations.AddTask(JobId, tasks);`

References:

<https://docs.microsoft.com/en-us/azure/batch/quick-create-portal>

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

QUESTION 19

DRAG DROP

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

You need to create compute nodes for the solution on Azure Batch.

What should you do?

Put the actions in the correct order.

Select and Place:

Select these	Place here
In Python, implement the class: TaskAddParameter	
In Python, implement the class: JobAddParameter	
In the Azure portal, create a Batch account.	
In Python, implement the class: PoolAddParameter	

Correct Answer:

Select these	Place here
	In the Azure portal, create a Batch account.
	In Python, implement the class: PoolAddParameter
	In Python, implement the class: JobAddParameter
	In Python, implement the class: TaskAddParameter

Section: (none)

Explanation

Explanation/Reference:

Explanation:

With Python:

Step 1: In the Azure portal, create a Batch account.

First we create a batch account.

Step 2: In Python, implement the class: `PoolAddParameter`

Now that you have a Batch account, create a sample pool of Windows compute nodes for test purposes. To create a Batch pool, the app uses the `PoolAddParameter` class to set the number of nodes, VM size, and a pool configuration.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Step 3: In Python, implement the class: JobAddParameter

Now that you have a pool, create a job to run on it. A Batch job is a logical group for one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. Initially the job has no tasks. The app uses the JobAddParameter class to create a job on your pool.

Step 4: In Python, implement the class: TaskAddParameter

Now create sample tasks to run in the job. Typically you create multiple tasks that Batch queues and distributes to run on the compute nodes. The app creates a list of task objects using the TaskAddParameter class.

References:

<https://docs.microsoft.com/en-us/azure/batch/quick-create-portal>

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

Develop Azure Platform as a Service Compute Solutions

Testlet 1

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

Background

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web API project that is deployed as an Azure Web App.

Overall architecture

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing

Employees may upload receipts in two ways:

- Uploading using an Azure Files mounted folder
- Uploading using the web application

Data Storage

Receipt and employee information is stored in an Azure SQL database.

Documentation

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for Azure File upload, and instructions on how to configure the mounted folder.

Solution details

Users table

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging

Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled. Application Insights must always contain all log messages.

Requirements

Receipt processing

Concurrent processing of a receipt must be prevented.

Disaster recovery

Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security

- Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
- All certificates and secrets used to secure data must be stored in Azure Key Vault.
- You must adhere to the Least Privilege Principal and provide privileges which are essential to perform the intended function.
- All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI)
- Receipt data must always be encrypted at rest.
- All data must be protected in transit.
- User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment with the remaining parts obscured.
- In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues

Upload format issue

Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal Server error page.

Capacity issue

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Log capacity issue

Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Processing.cs

```
PC01 public static class Processing
PC02 {
PC03     public static class Function
PC04     {
PC05         [FunctionName ("IssueWork")]
PC06         public static async Task Run ([TimerTrigger("0 */5" ***)] TimerInfo timer, ILogger log)
PC07         {
PC08             var container = await GetCloudBlobContainer();
PC09             foreach (var fileItem in await ListFiles())
PC10             {
PC11                 var file = new CloudFile (fileItem.StorageUri.PrimaryUri);
PC12                 var ms = new MemoryStream();
PC13                 await file.DownloadToStreamAsync(ms);
PC14                 var blob = container.GetBlockBlobReference (fileItem.Uri.ToString());
PC15                 await blob.UploadFromStreamAsync(ms);
PC16             }
PC17         }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob (CloudBlockBlob sourceBlob)
PC20     {
PC21         . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25         var cloudBlobClient = new CloudBlobClient (new Uri(" . . ."), await GetCredentials());
PC26
PC27         await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28         return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32         . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36         . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(" . . .");
PC38 }
PC39 }
```

Database.cs

```

DB01 public class Database
DB02 {
DB03     private string ConnectionString =
DB04
DB05     public async Task<object> LoadUserDetails(string userId)
DB06     {
DB07
DB08         return await policy.ExecuteAsync (async () =>
DB09         {
DB10             using (var connection = new SqlConnection (ConnectionString))
DB11             {
DB12                 await connection.OpenAsync();
DB13                 using (var command = new SqlCommand("_", connection))
DB14                 using (var reader = command.ExecuteReader())
DB15                 {
DB16                     -
DB17                 }
DB18             }
DB19         }
DB20     }
DB21 }

```

ReceiptUploader.cs

```

RU01 public class ReceiptUploader
RU02 {
RU03     public async Task UploadFile(string file, byte[] binary)
RU04     {
RU05         var httpClient = new HttpClient();
RU06         var response = await httpClient.PutAsync( "...", new ByteArrayContent(binary));
RU07         while (ShouldRetry (response))
RU08         {
RU09             response = await httpClient.PutAsync ( "...", new ByteArrayContent(binary));
RU10         }
RU11     }
RU12     private bool ShouldRetry(HttpResponseMessage response)
RU13     {
RU14
RU15     }
RU16 }

```

ConfigureSSE.ps1

```

CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "... " -AccountName "... "
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "... "
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "... "
CS04 Set-AzureRmKeyVaultAccessPolicy'
CS05 -VaultName $keyVault.VaultName'
CS06 -ObjectId $storageAccount.Identity.PrincipalId'
CS07
CS08
CS09 Set-AzureRmStorageAccount"
CS10 -ResourceGroupName $storageAccount.ResourceGroupName'
CS11 -AccountName $storageAccount.StorageAccountName'
CS12 -EnableEncryptionService File `
CS13 -KeyvaultEncryption'
CS14 -KeyName $key.Name
CS15 -KeyVersion $key.Version'
CS16 -KeyVaultUri $keyVault.VaultUri

```

QUESTION 1

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

HOTSPOT

You need to configure retries in the LoadUserDetails function in the Database class without impacting user experience.

What code should you insert on line DB07?

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
var policy = 

|                      |   |
|----------------------|---|
|                      | ▼ |
| Policy               |   |
| RetryPolicy          |   |
| RetryOptions         |   |
| ReconnectRetryPolicy |   |

  
.Handle<Exception> ()
```

	▼
.Retry(3);	
.CircuitBreaker(3, TimeSpan.FromMilliseconds(100));	
.WaitAndRetryAsync(3, i=> TimeSpan.FromMilliseconds(100));	

Correct Answer:

Answer Area

```
var policy = 

|                      |   |
|----------------------|---|
|                      | ▼ |
| Policy               |   |
| RetryPolicy          |   |
| RetryOptions         |   |
| ReconnectRetryPolicy |   |

  
.Handle<Exception> ()
```

	▼
.Retry(3);	
.CircuitBreaker(3, TimeSpan.FromMilliseconds(100));	
.WaitAndRetryAsync(3, i=> TimeSpan.FromMilliseconds(100));	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Box 1: Policy

```
RetryPolicy retry = Policy  
.Handle<HttpRequestException>()  
.Retry(3);
```

The above example will create a retry policy which will retry up to three times if an action fails with an exception handled by the Policy.

Box 2: `WaitAndRetryAsync(3, i => TimeSpan.FromMilliseconds(100 * Math.Pow(2, i - 1)))`;

A common retry strategy is exponential backoff: this allows for retries to be made initially quickly, but then at progressively longer intervals, to avoid hitting a subsystem with repeated frequent calls if the subsystem may be struggling.

Example:

```
Policy  
.Handle<SomeExceptionType>()  
.WaitAndRetry(3, retryAttempt =>  
TimeSpan.FromSeconds(Math.Pow(2, retryAttempt))  
);
```

References:

<https://github.com/App-vNext/Polly/wiki/Retry>

Develop Azure Platform as a Service Compute Solutions

Testlet 2

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

LabelMaker app

Coho Winery produces bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends **data** to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the **LabelMaker** application to include artwork and text for custom label designs.

Requirements

Data

You identify the following requirements for data management and manipulation:

- Order data is stored as nonrelational JSON and must be queried using Structured Query Language (SQL).
- Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Security

You have the following security requirements:

- Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
- External partners must use their own credentials and authenticate with their organization's identity management solution.
- External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.
- Storage of e-commerce application settings must be maintained in Azure Key Vault.
- E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).


- Conditional access policies must be applied at the application level to protect company content
- The LabelMaker applications must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

LabelMaker app

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture



Issues

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communications timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order .json

Relevant portions of the app files are shown below. Line numbers are included for reference only. This JSON file contains a representation of the data for an order that includes a single item.

```

01 {
02   "id" : 1,
03   "customers" : [
04     {
05       "familyName" : "Doe",
06       "givenName" : "John",
07       "customerid" : 5
08     }
09   ],
10   "line_items" : [
11     {
12       "fulfillable_quantity" : 1,
13       "id" : 6,
14       "price" : "199.99",
15       "product_id" : 7513594,
16       "quantity" : 1,
17       "requires_shipping" : true,
18       "sku" : "SFC-342-N" ,
19       "title" : "Surface Go",
20       "vendor" : "Microsoft" ,
21       "name" : "Surface Go - 8GB",
22       "taxable" : true,
23       "tax_lines" : [
24         {
25           "title" : "State Tax",
26           "price" : "3.93",
27           "rate" : 0.06
28         }
29       ],
30       "total_discount" : "5.00"

```

art.com/

```

31   "discount_allocations" : [
32     {

```

QUESTION 1

You need to implement the e-commerce checkout API.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Set the function template's Mode property to **Webhook** and the Webhook type property to **Generic JSON**.
- B. Create an Azure Function using the HTTP POST function template.
- C. In the Azure Function App, enable Cross-Origin Resource Sharing (CORS) with all origins permitted.
- D. In the Azure Function App, enable Managed Service Identity (MSI).
- E. Set the function template's Mode property to **Webhook** and the Webhook type property to **GitHub**.
- F. Create an Azure Function using the Generic webhook function template.

Correct Answer: ABD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).

D: A managed identity from Azure Active Directory allows your app to easily access other AAD-protected resources such as Azure Key Vault.

Incorrect Answers:

C: CORS is an HTTP feature that enables a web application running under one domain to access resources in another domain.

References:

<https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity>

QUESTION 2

DRAG DROP

You need to deploy a new version of the LabelMaker application.

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.

Select and Place:

Actions

Answer Area

Restart the cluster.

Create an alias of the image with the a new build number.

Build a new application image by using msbuild.

Create an alias of the image with the fully qualified path to the registry.

Build a new application image by using dockerfile.

Download the image to your local computer.

Log in to the registry and push image.



Correct Answer:

Actions

Answer Area

Restart the cluster.

Create an alias of the image with the a new build number.

Build a new application image by using msbuild.

Create an alias of the image with the fully qualified path to the registry.

Build a new application image by using dockerfile.

Download the image to your local computer.

Log in to the registry and push image.

Build a new application image by using dockerfile.

Create an alias of the image with the fully qualified path to the registry.

Log in to the registry and push image.



Section: (none)
Explanation

Explanation/Reference:

Explanation:

Step 1: Build a new application image by using dockerfile

Step 2: Create an alias if the image with the fully qualified path to the registry

Before you can push the image to a private registry, you've to ensure a proper image name. This can be achieved using the docker tag command. For demonstration purpose, we'll use Docker's hello world image, rename it and push it to ACR.

```
# pulls hello-world from the public docker hub
$ docker pull hello-world
# tag the image in order to be able to push it to a private registry
$ docker tag hello-word <REGISTRY_NAME>/hello-world
# push the image
$ docker push <REGISTRY_NAME>/hello-world
```

Step 3: Log in to the registry and push image

In order to push images to the newly created ACR instance, you need to login to ACR from the Docker CLI. Once logged in, you can push any existing docker image to your ACR instance.

Scenario:

Coho Winery plans to move the application to Azure and continue to support label creation.

LabelMaker app

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

References:

<https://thorsten-hans.com/how-to-use-a-private-azure-container-registry-with-kubernetes-9b86e67b93b6>

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-tutorial-quick-task>

QUESTION 3

You need to provision and deploy the order workflow.

Which three components should you include? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Connections
- B. On-premises Data Gateway
- C. Workflow definition
- D. Resources
- E. Functions

Correct Answer: BCE


Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: The order workflow fails to run upon initial deployment to Azure.



QUESTION 4

HOTSPOT

You need to update the order workflow to address the issue when calling the Printer API App.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
"print_label": {  
  "type": "Http",  
  "inputs": {  
    "method" : "POST",  
    "uri": "https://www.cohowinery.com/printer/printlabel",  
    "retriyPolicy": {  
      "type": "",  


|             |
|-------------|
| default     |
| none        |
| fixed       |
| exponential |

  
    },  
    "interval": "",  


|       |
|-------|
| PT10S |
| PT30S |
| PT60S |
| PT1D  |

  
    "count": ,  


|    |
|----|
| 5  |
| 10 |
| 60 |

  
  },  
},  
}
```

Correct Answer:

Answer Area

```
"print_label": {  
  "type": "Http",  
  "inputs": {  
    "method" : "POST",  
    "uri": "https://www.cohowinery.com/printer/printlabel",  
    "retriyPolicy": {  
      "type": "",  
      "interval": "",  
      "count":   
    }  
  }  
}
```

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Fixed

To specify that the action or trigger waits the specified interval before sending the next request, set the <retry-policy-type> to fixed.

Box 2: PT10S

Box 3: 5

Scenario: Calls to the Printer API App fail periodically due to printer communication timeouts. Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

Incorrect Answers:

Default: If you don't specify a retry policy, the action uses the default policy, which is actually an exponential interval policy that sends up to four retries at exponentially increasing intervals that are scaled by 7.5 seconds. The interval is capped between 5 and 45 seconds.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-exception-handling>

Develop Azure Platform as a Service Compute Solutions

Question Set 3

QUESTION 1

HOTSPOT

You are creating a CLI script that creates an Azure web app and 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.

Hot Area:

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

Correct Answer:

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

Section: (none)

Explanation

Explanation/Reference:

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.

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

References:

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

QUESTION 2

HOTSPOT

You are developing an Azure Web App. You configure TLS mutual authentication for the web app.

You need to validate the client certificate in the web app. To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Hot Area:

Answer Area

Property	Value
Client certificate location	HTTP request header Client cookie HTTP message body URL query string
Encoding type	HTML URL Unicode Base64

Correct Answer:

Answer Area

Property	Value
Client certificate location	HTTP request header Client cookie HTTP message body URL query string
Encoding type	HTML URL Unicode Base64

Section: (none)
Explanation

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation/Reference:

Explanation:

Accessing the client certificate from App Service.

If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be available through the `HttpRequest.ClientCertificate` property. For other application stacks, the client cert will be available in your app through a base64 encoded value in the "X-ARR-ClientCert" request header.

Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.

References:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth>

QUESTION 3

DRAG DROP

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.

Select and Place:

Code segments

```
[Authorize(Policy = "ProviderPartner")]
[Authorize(Role = "SysAdmin")] ]
```

```
[Authorize(Role = "ProviderAdmin")]
[Authorize(Role = "SysAdmin")] ]
```

```
[Authorize(Role = "SysAdmin", "ProviderAdmin")] ]
```

```
[Authorize(Policy = "ProviderPartner",
Role="SysAdmin")] ]
```

Answer Area

```
public class PartnerController : Controller
{
    . . .

    Public ActionResult Manage ()
    {
        . . .
    }
}
```

Correct Answer:**Code segments**

```
[Authorize(Policy = "ProviderPartner")]
[Authorize(Role = "SysAdmin")] ]
```

```
[Authorize(Role = "ProviderAdmin")]
[Authorize(Role = "SysAdmin")] ]
```

```
[Authorize(Role = "SysAdmin", "ProviderAdmin")] ]
```

```
[Authorize(Policy = "ProviderPartner",
Role="SysAdmin")] ]
```

Answer Area

```
[Authorize(Role = "ProviderAdmin")]
[Authorize(Role = "SysAdmin")] ]

public class PartnerController : Controller
{
    . . .

    [Authorize(Policy = "ProviderPartner",
Role="SysAdmin")]
    Public ActionResult Manage ()
    {
        . . .
    }
}
```

Section: (none)**Explanation****Explanation/Reference:**

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.

QUESTION 4**DRAG DROP**

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.

Select and Place:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 5

You are implementing an Azure API app that uses built-in authentication and authorization functionality.

All app actions must be associated with information about the current user.

You need to retrieve the information about the current user.

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

NOTE: Each correct selection is worth one point.

- A. HTTP headers
- B. environment variables
- C. /.auth/me HTTP endpoint
- D. /.auth/login endpoint

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

A: After App Service Authentication has been configured, users trying to access your API are prompted to sign in with their organizational account that belongs to the same Azure AD as the Azure AD application used to secure the API. After signing in, you are able to access the information about the current user through the `HttpContext.Current.User` property.

C: While the server code has access to request headers, client code can access GET /.auth/me to get the same access tokens (

References:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-auth-aad>

<https://docs.microsoft.com/en-us/sharepoint/dev/spfx/web-parts/guidance/connect-to-api-secured-with-aad>

QUESTION 6

HOTSPOT

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.

Hot Area:

Scale rule

Metric source

Storage queue

Service Bus queue

Current resource

Storage queue (classic)

Resource type

Service Bus Namespace

Resource

MessageQueue1103

* Queues

itemqueue

Criteria

* Metric name

Message Count

Active Message Count

1 minute time grain

* Time grain statistic ⓘ

Total

Maximum

Average

Count

* Operator

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

Update

Delete

<https://t.me/learningnets>

Correct Answer:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Scale rule

Metric source

Storage queue

Service Bus queue

Current resource

Storage queue (classic)

Resource type

Service Bus Namespace

Resource

MessageQueue1103

* Queues

itemqueue

Criteria

* Metric name

Message Count

Active Message Count

1 minute time grain

* Time grain statistic ⓘ

Total

Maximum

Average

Count

* Operator

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

Update

Delete

Section: (none)
Explanation

Explanation/Reference:
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

QUESTION 7
HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

gitrepo=https://github.com/Contoso/webapp

webappname=businesswebapp

resourcegroupname=BusinessAppResourceGroup

az	group webapp appservice plan webapp deployment slot webapp deployment source	<pre>create --location centralus - -name \$resourcegroupname create --name \$webappname - -resource-group \$resourcegroupname - -sku S3 create --name \$webappname - -resource-group \$resourcegroupname \ - -plan \$webappname create --name \$webappname - -resource-group \$resourcegroupname \ - -slot staging config - -name \$webappname - -resource-group \$resourcegroupname \ - -slot staging - -repo-url \$gitrepo - -branch master - -manual-integration</pre>
az	group webapp appservice plan webapp deployment slot webapp deployment source	
az	group webapp appservice plan webapp deployment slot webapp deployment source	
az	group webapp appservice plan webapp deployment slot webapp deployment source	
az	group webapp appservice plan webapp deployment slot webapp deployment source	

Correct Answer:

Answer Area

gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup

az	▼ group webapp appservice plan webapp deployment slot webapp deployment source	<pre>create --location centralus - --name \$resourcegroupname create --name \$webappname - --resource-group \$resourcegroupname - --sku S3 create --name \$webappname - --resource-group \$resourcegroupname \ - --plan \$webappname create --name \$webappname - --resource-group \$resourcegroupname \ - --slot staging</pre>
az	▼ group webapp appservice plan webapp deployment slot webapp deployment source	<pre>config - --name \$webappname - --resource-group \$resourcegroupname \ - --slot staging - --repo-url \$gitrepo - --branch master - --manual-integration</pre>
az	▼ group webapp appservice plan webapp deployment slot webapp deployment source	
az	▼ group webapp appservice plan webapp deployment slot webapp deployment source	
az	▼ group webapp appservice plan webapp deployment slot webapp deployment source	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: group

Create a resource group.

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

Box 2: appservice plan

Create an App Service plan in STANDARD tier (minimum required by deployment slots).

```
az appservice plan create --name $webappname --resource-group myResourceGroup --sku S1
```

Box 3: webapp

Create a web app.

```
az webapp create --name $webappname --resource-group myResourceGroup \
--plan $webappname
```

Box 4: webapp deployment slot

#Create a deployment slot with the name "staging".

```
az webapp deployment slot create --name $webappname --resource-group myResourceGroup \
--slot staging
```

Box 5: webapp deployment source

Deploy sample code to "staging" slot from GitHub.

```
az webapp deployment source config --name $webappname --resource-group myResourceGroup \
--slot staging --repo-url $gitrepo --branch master --manual-integration
```

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

References:

<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment>

QUESTION 8

DRAG DROP

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

You need to configure the services.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Services	Scenario	Service
Logic Apps	Process a queue data item.	
WebJobs	Manage all code segments from within the same DevOps environment.	
Flow		

Correct Answer:

Answer Area

Services	Scenario	Service
Logic Apps	Process a queue data item.	WebJobs
WebJobs	Manage all code segments from within the same DevOps environment.	Flow
Flow		

Section: (none)

Explanation

Explanation/Reference:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation:

Box 1: WebJobs

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

Box 2: Flow

Incorrect Answers:

Azure Logic Apps is a cloud service that helps you schedule, automate, and orchestrate tasks, business processes, and workflows when you need to integrate apps, data, systems, and services across enterprises or organizations. Logic Apps simplifies how you design and build scalable solutions for app integration, data integration, system integration, enterprise application integration (EAI), and business-to-business (B2B) communication, whether in the cloud, on premises, or both.

References:

<https://code.msdn.microsoft.com/Processing-Service-Bus-84db27b4>

QUESTION 9

HOTSPOT

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.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)

Explanation

Explanation/Reference:

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/>

QUESTION 10

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?

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

- A. **New-AzureRmServiceBusNamespace**
 -ResourceGroupName fridge-rg
 -NamespaceName fridge-ns
 -Location fridge-loc
- B. **az servicebus namespace create**
 - --resource-group fridge-rg
 - --name fridge-ns
 - --location fridge-loc
- C. **New-AzureRmResourceGroup**
 -Name fridge-rg
 -Location fridge-loc
- D. **New-AzureRmServiceBusQueue**
 -ResourceGroupName fridge-rg
 -NamespaceName fridge-ns
 -Name fridge-q
 -EnablePartitioning \$False

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 11

You are a developer for a SaaS company that offers many web services.

All web services for the company must meet the following requirements:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

- Use API Management to access the services
- Use OpenID Connect for authentication.
- Prevent anonymous usage

A recent security audit found that several web services can be called without any authentication.

Which API Management policy should you implement?

- A. validate-jwt
- B. jsonp
- C. authentication-certificate
- D. check-header

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Add the validate-jwt policy to validate the OAuth token for every incoming request.

Incorrect Answers:

B: The jsonp policy adds JSON with padding (JSONP) support to an operation or an API to allow cross-domain calls from JavaScript browser-based clients. JSONP is a method used in JavaScript programs to request data from a server in a different domain. JSONP bypasses the limitation enforced by most web browsers where access to web pages must be in the same domain.

JSONP - Adds JSON with padding (JSONP) support to an operation or an API to allow cross-domain calls from JavaScript browser-based clients.

References: <https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad>

QUESTION 12

DRAG DROP

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.

Select and Place:

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

Entity

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

List blobs 2

Folder

/items

Show advanced options

Condition

Check LastModified timestamp and whether older than the tier age variable

ticks(items(For_each? [LastModified]))

is less than

ticks(addDaysInMonth(), variables(TierAgeInDays))

Put a message on a queue

Queue Name

processing

Message

Path

Connected to testConnection. [Change connection](#)

Answer area

trigger or action block

Set tier age variable

List blobs

For each

Scan all blobs in this folder

Select an output from previous steps

value

trigger or action block

if true

trigger or action block

Add an action

if false

trigger or action block

Add an action

Add an action

Correct Answer:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Triggers and action blocks

Insert Entity

Table

processing

Entity

Path X

Show advanced options

When there are messages in a queue

Queue Name

processing

Show advanced options

Connected to tableStorageAccountConnection. [Change connection](#)

List blobs 2

Folder

/items

Show advanced options

Answer area

Recurrence

Interval

1

Frequency

Month

Show advanced options

↓

Set tier age variable

↓

List blobs

↓

For each

Scan all blobs in this folder

Select an output from previous steps

value X

Condition

Check LastModified timestamp and whether older than the tier age variable

ticks(items(For_each)?

is less than

ticks(addDaysInMonth(), variables('TierAgeInDays'))

if true

Tier blob

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

Blob path

Path X

Entity

Archive

Add an action

if false

Put a message on a queue

Queue Name

processing

Message

Path X

Connected to testConnection. [Change connection](#)

Add an action

Add an action

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Recurrence

This operation must be performed automatically once a month.

Box 2: Condition

Move blocks to Archive tier after they have not been accessed for 180 days.

Box 3 (if true): Tier Blob

Move blocks to Archive tier after they have not been accessed for 180 days.

Box 4: Put a message in a queue

The path for any item that is not archived must be placed in an existing queue.

References:

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

QUESTION 13

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

You need to deploy the website.

What should you do?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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.

Incorrect Answers:

B: Shared and Free modes do not offer the scaling flexibility of Standard, and they have some important limits.

Shared mode, just as the name states, also uses shared Compute resources, and also has a CPU limit. So, while neither Free nor Shared is likely to be the best choice for your production environment due to these limits.

QUESTION 14

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
```

New-AzWebAppSlot New-AzWebApp New-AzAppServicePlan New-AzResourceGroup	-Name myResourceGroup -Location \$location
New-AzWebAppSlot New-AzWebApp New-AzAppServicePlan New-AzResourceGroup	-Name \$webappname -Location \$location -ResourceGroupName myResourceGroup -Tier Standard
New-AzWebAppSlot New-AzWebApp New-AzAppServicePlan New-AzResourceGroup	-Name \$webappname -Location \$location -AppServicePlan \$webappname -ResourceGroupName myResourceGroup
New-AzWebAppSlot New-AzWebApp New-AzAppServicePlan New-AzResourceGroup	-Name \$webappname -ResourceGroupName myResourceGroup -Slot review

```
@PropertiesObject = @{repoUrl = "$gitrepo";branch = "master";}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup `
-SourceSlotName review -DestinationSlotName production
```

Correct Answer:

Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
```

New-AzWebAppSlot New-AzWebApp New-AzAppServicePlan New-AzResourceGroup	-Name myResourceGroup -Location \$location
New-AzWebAppSlot New-AzWebApp New-AzAppServicePlan New-AzResourceGroup	-Name \$webappname -Location \$location -ResourceGroupName myResourceGroup -Tier Standard
New-AzWebAppSlot New-AzWebApp New-AzAppServicePlan New-AzResourceGroup	-Name \$webappname -Location \$location -AppServicePlan \$webappname -ResourceGroupName myResourceGroup
New-AzWebAppSlot New-AzWebApp New-AzAppServicePlan New-AzResourceGroup	-Name \$webappname -ResourceGroupName myResourceGroup -Slot review

```
@PropertiesObject = @{repoUrl = "$gitrepo";branch = "master";}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup `
-SourceSlotName review -DestinationSlotName production
```

Section: (none) Explanation

Explanation/Reference: Explanation:

The New-AzResourceGroup cmdlet creates an Azure resource group.

The New-AzAppServicePlan cmdlet creates an Azure App Service plan in a given location

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

The New-AzWebApp cmdlet creates an Azure Web App in a given a resource group

The New-AzWebAppSlot cmdlet creates an Azure Web App slot.

References:

<https://docs.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroup?view=azps-2.3.2>

<https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azappserviceplan?view=azps-2.3.2>

<https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebapp?view=azps-2.3.2>

<https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebappslot?view=azps-2.3.2>

QUESTION 15

HOTSPOT

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.

Hot Area:

Answer Area

App service plan setting

Value

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

Correct Answer:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Answer Area

App service plan setting

Value

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

Section: (none)

Explanation

Explanation/Reference:

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/>

QUESTION 16

HOTSPOT

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.

Hot Area:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Answer Area

```
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) |   |

  
}
```

Correct Answer:

Answer Area

```
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)  
  }  
}
```

Section: (none)

Explanation

Explanation/Reference:

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://vceoreteconvert.com/>

<https://t.me/learningnets>

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

QUESTION 17

DRAG DROP

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.

Select and Place:

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



Correct Answer:

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

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

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

```
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
```

Section: (none)

Explanation

Explanation/Reference:

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 config 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>

QUESTION 18

You develop a serverless application that includes Azure Functions by using Visual Studio. These functions connect to data from within the code. You deploy the functions to Azure.

You want to configure tracing for an Azure Function App project.

You need to change configuration settings in the host.json file.

Which tool should you use?

- A. Visual Studio
- B. Azure portal
- C. Azure PowerShell

D. Azure Functions Core Tools (Azure CLI)

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

```
FunctionApp
| - host.json
| - Myfirstfunction
| | - function.json
| | - ...
| - mysecondfunction
| | - function.json
| | - ...
| - SharedCode
| - bin
```

References:

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

QUESTION 19

You are developing a mobile instant messaging app for a company.

The mobile app must meet the following requirements:

- Support offline data sync.
- Update the latest messages during normal sync cycles.

You need to implement Offline Data Sync.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

References:

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

QUESTION 20

DRAG DROP

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.

Select and Place:

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

Correct Answer:

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

Section: (none)

Explanation

Explanation/Reference:

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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.

Incorrect Answers:

Consistent prefix: Updates that are returned contain some prefix of all the updates, with no gaps.

Consistent prefix guarantees that reads never see out-of-order writes.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

QUESTION 21

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.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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>

QUESTION 22

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. **New-AzureRmServiceBusNamespace**
 -ResourceGroupName fridge-rg
 -NamespaceName fridge-ns
 -Location fridge-loc
- B. **az servicebus namespace create**
 - --resource-group fridge-rg
 - --name fridge-ns
 - --location fridge-loc
- C. **New-AzureRmResourceGroup**
 -Name fridge-rg
 -Location fridge-loc
- D. **az servicebus queue create**
 --resource-group fridge-rg
 --namespace-name fridge-ns
 --name fridge-q

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

\$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>

Develop for Azure storage

Testlet 1

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

Background

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web API project that is deployed as an Azure Web App.

Overall architecture

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing

Employees may upload receipts in two ways:

- Uploading using an Azure Files mounted folder
- Uploading using the web application

Data Storage

Receipt and employee information is stored in an Azure SQL database.

Documentation

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for Azure File upload, and instructions on how to configure the mounted folder.

Solution details

Users table

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging

Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled. Application Insights must always contain all log messages.

Requirements

Receipt processing

Concurrent processing of a receipt must be prevented.

Disaster recovery

Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security

- Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
- All certificates and secrets used to secure data must be stored in Azure Key Vault.
- You must adhere to the Least Privilege Principal and provide privileges which are essential to perform the intended function.
- All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI)
- Receipt data must always be encrypted at rest.
- All data must be protected in transit.
- User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment with the remaining parts obscured.
- In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues

Upload format issue

Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal Server error page.

Capacity issue

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Log capacity issue

Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Processing.cs

```
PC01 public static class Processing
PC02 {
PC03     public static class Function
PC04     {
PC05         [FunctionName ("IssueWork")]
PC06         public static async Task Run ([TimerTrigger("0 */5" ***)] TimerInfo timer, ILogger log)
PC07         {
PC08             var container = await GetCloudBlobContainer();
PC09             foreach (var fileItem in await ListFiles())
PC10             {
PC11                 var file = new CloudFile (fileItem.StorageUri.PrimaryUri);
PC12                 var ms = new MemoryStream();
PC13                 await file.DownloadToStreamAsync(ms);
PC14                 var blob = container.GetBlockBlobReference (fileItem.Uri.ToString());
PC15                 await blob.UploadFromStreamAsync(ms);
PC16             }
PC17         }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob (CloudBlockBlob sourceBlob)
PC20     {
PC21         . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25         var cloudBlobClient = new CloudBlobClient (new Uri(" . . ."), await GetCredentials());
PC26
PC27         await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28         return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32         . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36         . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(" . . .");
PC38 }
PC39 }
```

Database.cs

```

DB01 public class Database
DB02 {
DB03     private string ConnectionString =
DB04
DB05     public async Task<object> LoadUserDetails(string userId)
DB06     {
DB07
DB08         return await policy.ExecuteAsync (async () =>
DB09         {
DB10             using (var connection = new SqlConnection (ConnectionString))
DB11             {
DB12                 await connection.OpenAsync();
DB13                 using (var command = new SqlCommand("_", connection))
DB14                 using (var reader = command.ExecuteReader())
DB15                 {
DB16                     -
DB17                 }
DB18             }
DB19         }
DB20     }
DB21 }

```

ReceiptUploader.cs

```

RU01 public class ReceiptUploader
RU02 {
RU03     public async Task UploadFile(string file, byte[] binary)
RU04     {
RU05         var httpClient = new HttpClient();
RU06         var response = await httpClient.PutAsync( "...", new ByteArrayContent(binary));
RU07         while (ShouldRetry (response))
RU08         {
RU09             response = await httpClient.PutAsync ( "...", new ByteArrayContent(binary));
RU10         }
RU11     }
RU12     private bool ShouldRetry(HttpResponseMessage response)
RU13     {
RU14
RU15     }
RU16 }

```

ConfigureSSE.ps1

```

CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "... -AccountName ..."
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "..."
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "..."
CS04 Set-AzureRmKeyVaultAccessPolicy'
CS05 -VaultName $keyVault.VaultName'
CS06 -ObjectId $storageAccount.Identity.PrincipalId'
CS07
CS08
CS09 Set-AzureRmStorageAccount"
CS10 -ResourceGroupName $storageAccount.ResourceGroupName'
CS11 -AccountName $storageAccount.StorageAccountName'
CS12 -EnableEncryptionService File `
CS13 -KeyvaultEncryption'
CS14 -KeyName $key.Name
CS15 -KeyVersion $key.Version'
CS16 -KeyVaultUri $keyVault.VaultUri

```

QUESTION 1

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

DRAG DROP

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

Code segments	Answer Area
<code>MSITokenProvider(".", null)</code>	<code>var tp=new</code>
<code>tp.GetAccessTokenAsync("..")</code>	<code>var t=new TokenCredential(await</code>
<code>AzureServiceTokenProvider()</code>	<code>return new StorageCredentials(t);</code>
<code>StringTokenProvider("storage","msi")</code>	
<code>tp.GetAuthenticationHeaderAsync(CancellationTokens.None)</code>	

Correct Answer:

Code segments	Answer Area
<code>MSITokenProvider(".", null)</code>	<code>var tp=new</code> AzureServiceTokenProvider()
<code>tp.GetAccessTokenAsync("..")</code>	<code>var t=new TokenCredential(await</code> <input ..\")"="" type="text" value="tp.GetAccessTokenAsync(\"/>
<code>AzureServiceTokenProvider()</code>	<code>return new StorageCredentials(t);</code>
<code>StringTokenProvider("storage","msi")</code>	
<code>tp.GetAuthenticationHeaderAsync(CancellationTokens.None)</code>	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: AzureServiceTokenProvider()

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

Acquiring an access token is then quite easy. Example code:

```
private async Task<string> GetAccessTokenAsync()
{
    var tokenProvider = new AzureServiceTokenProvider();
    return await tokenProvider.GetAccessTokenAsync("https://storage.azure.com/");
}
```

References:

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

QUESTION 2

You need to construct the link to the summary report for the email that is sent to users.

What should you do?

- A. Create a SharedAccessBlobPolicy and add it to the containers SharedAccessPolicies. Call GetSharedAccessSignature on the blob and use the resulting link.
- B. Create a SharedAccessAccountPolicy and call GetSharedAccessSignature on storage account and use

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

the resulting link.

- C. Create a SharedAccessBlobPolicy and set the expiry time to two weeks from today. Call GetSharedAccessSignature on the blob and use the resulting link.
- D. Create a SharedAccessBlobPolicy and set the expiry time to two weeks from today. Call GetSharedAccessSignature on the container and use the resulting link.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Create a stored access policy to manage signatures on a container's resources, and then generate the shared access signature on the container, setting the constraints directly on the signature.

Code example: Add a method that generates the shared access signature for the container and returns the signature URI.

```
static string GetContainerSasUri(CloudBlobContainer container)
{
    //Set the expiry time and permissions for the container.
    //In this case no start time is specified, so the shared access signature becomes valid immediately.
    SharedAccessBlobPolicy sasConstraints = new SharedAccessBlobPolicy();
    sasConstraints.SharedAccessExpiryTime = DateTimeOffset.UtcNow.AddHours(24);
    sasConstraints.Permissions = SharedAccessBlobPermissions.List | SharedAccessBlobPermissions.Write;

    //Generate the shared access signature on the container, setting the constraints directly on the signature.
    string sasContainerToken = container.GetSharedAccessSignature(sasConstraints);

    //Return the URI string for the container, including the SAS token.
    return container.Uri + sasContainerToken;
}
```

Incorrect Answers:

C: Call GetSharedAccessSignature on the container, not on the blob.

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-dotnet-shared-access-signature-part-2>

QUESTION 3

HOTSPOT

You need to ensure that security requirements are met.

What value should be used for the ConnectionString field on line DB03 in the Database class? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

`"DataSource=datastore.database.windows.net;Initial Catalog=expense;`

	▼	;
Integrated Security=SSPI		
Trusted_Connection=False		
Network Library=DBMSSOCN		
MultipleActiveResultSets=True		

	▼	;"
Encrypt=True		
Integrated Security=True		
Failover Partner=False		
Named Pipes=True		

Correct Answer:

Answer Area

`"DataSource=datastore.database.windows.net;Initial Catalog=expense;`

	▼	;
Integrated Security=SSPI		
Trusted_Connection=False		
Network Library=DBMSSOCN		
MultipleActiveResultSets=True		

	▼	;"
Encrypt=True		
Integrated Security=True		
Failover Partner=False		
Named Pipes=True		

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Integrated Security=SSPI

Integrated security: For all data source types, connect using the current user account.

For SqlConnection you can use Integrated Security=true; or Integrated Security=SSPI;

Scenario: All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI)

Box 2: Encrypt = True

Scenario: All data must be protected in transit.

References:

<https://docs.microsoft.com/en-us/dotnet/framework/data/adonet/connection-string-syntax>

QUESTION 4

DRAG DROP

You need to ensure disaster recovery requirements are met.

What code should you add at line PC16?

To answer, drag the appropriate code fragments to the correct locations. Each code fragment 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.

Select and Place:

Values	Answer Area
true	<pre>var copyOptions = new CopyOptions {}; var context= new DirectoryTransferContext (source,destination)=>Task.FromResult(true); context. ShouldTransferCallbackAsync =(source, destination) => Task.FromResult(true); await TransferManager.CopyAsync(blob, GetDRBlob(blob), isServiceCopy: false , context:context, options: copyOptions);</pre>
false	
SingleTransferContext	
DirectoryTransferContext	
ShouldTransferCallbackAsync	
ShouldOverwriteCallbackAsync	

Correct Answer:

Values	Answer Area
true	<pre>var copyOptions = new CopyOptions {}; var context= new DirectoryTransferContext (source,destination)=>Task.FromResult(true); context. ShouldTransferCallbackAsync =(source, destination) => Task.FromResult(true); await TransferManager.CopyAsync(blob, GetDRBlob(blob), isServiceCopy: false , context:context, options: copyOptions);</pre>
false	
SingleTransferContext	
DirectoryTransferContext	
ShouldTransferCallbackAsync	
ShouldOverwriteCallbackAsync	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: Disaster recovery. Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Box 1: DirectoryTransferContext

We transfer all files in the directory.

Note: The TransferContext object comes in two forms: SingleTransferContext and DirectoryTransferContext. The former is for transferring a single file and the latter is for transferring a directory of files.

Box 2: ShouldTransferCallbackAsync

The DirectoryTransferContext.ShouldTransferCallbackAsync delegate callback is invoked to tell whether a transfer should be done.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Box 3: False

If you want to use the retry policy in Copy, and want the copy can be resume if break in the middle, you can use SyncCopy (isServiceCopy = false).

Note that if you choose to use service side copy ('isServiceCopy' set to true), Azure (currently) doesn't provide SLA for that. Setting 'isServiceCopy' to false will download the source blob loca

References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-data-movement-library>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.windowsazure.storage.datamovement.directorytransfercontext.shouldtransfercallbackasync?view=azure-dotnet>

Develop for Azure storage

Testlet 2

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

LabelMaker app

Coho Winery produces bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends **data** to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the **LabelMaker** application to include artwork and text for custom label designs.

Requirements

Data

You identify the following requirements for data management and manipulation:

- Order data is stored as nonrelational JSON and must be queried using Structured Query Language (SQL).
- Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Security

You have the following security requirements:

- Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
- External partners must use their own credentials and authenticate with their organization's identity management solution.
- External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.
- Storage of e-commerce application settings must be maintained in Azure Key Vault.
- E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).


- Conditional access policies must be applied at the application level to protect company content
- The LabelMaker applications must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

LabelMaker app

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture



Issues

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communications timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order .json

Relevant portions of the app files are shown below. Line numbers are included for reference only. This JSON file contains a representation of the data for an order that includes a single item.

```

01 {
02   "id" : 1,
03   "customers" : [
04     {
05       "familyName" : "Doe",
06       "givenName" : "John",
07       "customerid" : 5
08     }
09   ],
10   "line_items" : [
11     {
12       "fulfillable_quantity" : 1,
13       "id" : 6,
14       "price" : "199.99",
15       "product_id" : 7513594,
16       "quantity": 1,
17       "requires_shipping" : true,
18       "sku" : "SFC-342-N" ,
19       "title" : "Surface Go",
20       "vendor" : "Microsoft" ,
21       "name" : "Surface Go - 8GB",
22       "taxable" : true,
23       "tax_lines" : [
24         {
25           "title" : "State Tax",
26           "price" : "3.93",
27           "rate" : 0.06
28         }
29       ],
30       "total_discount" : "5.00"

```

art.com/

```

31   "discount_allocations" : [
32     {

```

QUESTION 1
HOTSPOT

You need to retrieve all order line items sorted alphabetically by the city.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

SELECT li.id AS lineitemid, li.price
FROM

	▼
Orders o	
LineItems li	

JOIN

	▼
li	
o	

 IN

	▼
o.line_items	
li.line_items	
o.address	

ORDER BY

	▼
o.address.city	
li.address.city	
o.city	
li.city	

 ASC

Correct Answer:

Answer Area

SELECT li.id AS lineitemid, li.price
FROM

	▼
Orders o	
LineItems li	

JOIN

	▼
li	
o	

 IN

	▼
o.line_items	
li.line_items	
o.address	

ORDER BY

	▼
o.address.city	
li.address.city	
o.city	
li.city	

 ASC

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Scenario:
Order data is stored as nonrelational JSON and must be queried using Structured Query Language (SQL).
The Order data is stored in a Cosmos database.

QUESTION 2
HOTSPOT

You need to configure Azure Cosmos DB.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Setting	Value										
Consistency Level	<table><tr><td></td><td>▼</td></tr><tr><td colspan="2">Strong</td></tr><tr><td colspan="2">Bounded-staleness</td></tr><tr><td colspan="2">Session</td></tr><tr><td colspan="2">Eventual</td></tr></table>		▼	Strong		Bounded-staleness		Session		Eventual	
	▼										
Strong											
Bounded-staleness											
Session											
Eventual											
API	<table><tr><td></td><td>▼</td></tr><tr><td colspan="2">SQL</td></tr><tr><td colspan="2">MongoDB</td></tr><tr><td colspan="2">Graph</td></tr><tr><td colspan="2">Table</td></tr></table>		▼	SQL		MongoDB		Graph		Table	
	▼										
SQL											
MongoDB											
Graph											
Table											

Correct Answer:

Answer Area

Setting	Value
Consistency Level	Strong Bounded-staleness Session Eventual
API	SQL MongoDB Graph Table

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Strong

When the consistency level is set to strong, the staleness window is equivalent to zero, and the clients are guaranteed to read the latest committed value of the write operation.

Scenario: Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Note: You can choose from five well-defined models on the consistency spectrum. From strongest to weakest, the models are: Strong, Bounded staleness, Session, Consistent prefix, Eventual

Box 2: SQL

Scenario: You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using Structured Query Language (SQL).

Develop for Azure storage

Testlet 3

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment

Windows Server 2016 virtual machine

This virtual machine (VM) runs Biz Talk Server 2016. The VM runs the following workflows:

- Ocean Transport – This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.
- Inland Transport – This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

- Container API – This API provides container information including weight, contents, and other attributes.
- Location API – This API provides location information regarding shipping ports of call and truck stops.
- Shipping REST API – This API provides shipping information for use and display on the shipping website.

Shipping Data

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

Shipping Web Site

The site displays shipping container tracking information and container contents. The site is located at <http://shipping.wideworldimporters.com>


Proposed solution

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:



Shipping Logic App

The Shipping Logic app must meet the following requirements:

- Support the ocean transport and inland transport workflows by using a Logic App.
- Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
- Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
- Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app

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

REST APIs

The REST API's that support the solution must meet the following requirements:

- Secure resources to the corporate VNet.
- 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.

Shipping data

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

Shipping website

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

Issues

Windows Server 2016 VM

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

Shipping website and REST APIs

The following error message displays while you are testing the website:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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.

QUESTION 1

You need to migrate on-premises shipping data to Azure.

What should you use?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Migrate from on-premises or cloud implementations of MongoDB to Azure Cosmos DB with minimal downtime by using Azure Database Migration Service. Perform resilient migrations of MongoDB data at scale and with high reliability.

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

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

References:

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

Develop for Azure storage

Question Set 4

QUESTION 1

You are developing an application that uses Azure Blob storage.

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

You need to process the transaction logs asynchronously.

What should you do?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

References:

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

<https://docs.microsoft.com/en-us/azure/event-grid/event-handlers#azure-functions>

QUESTION 2

HOTSPOT

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 DmModelCreating(DBModelBuilder modelBuilder)
09.         {
10.             modelBuilder.Entity<Player>().HasMany(x => x.Games). WithMany (x => x.Players);
11.         }
12.     }
13.     internal series class dbConfiguration : DbMigrationConfiguration<PlayerDbContext>
14.     {
15.         public dbConfiguration() . {AutomaticMigrationsEnabled = true ; }
16.     {
17.         public class app
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.         }
55.     }

```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

	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.	☐	☐

Correct Answer:

Answer Area

	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.	☐	☒

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Yes

Box 2: No

Box 3: Yes

Box 4: No

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);
}
```

QUESTION 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")`

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 4

HOTSPOT

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.

Hot Area:

Answer Area

```

public class PlayerEntity : TableEntity
{
    public PlayerEntity()
    {
    }
    public PlayerEntity(string region, string email)
    {
        PartitionKey =  ;
        

|        |
|--------|
| email  |
| phone  |
| region |


        RowKey=  ;
        

|        |
|--------|
| email  |
| phone  |
| region |


    }
    public string Phone { get; set; }
}
public class Player
{
    protected PlayerEntity player;
    async void GetPlayer(string cs, 

|                      |
|----------------------|
| <input type="text"/> |
| CloudTable           |
| CloudTableClient     |
| TableEntity          |
| TableEntityAdapter   |

 table, string pk, string rk)
    {
        

|                                                                       |
|-----------------------------------------------------------------------|
| <input type="text"/>                                                  |
| 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); |



|                                                       |
|-------------------------------------------------------|
| <input type="text"/>                                  |
| 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;
    }
}

```

Correct Answer:

Answer Area

```

public class PlayerEntity : TableEntity
{
    public PlayerEntity()
    {
    }
    public PlayerEntity(string region, string email)
    {
        PartitionKey =  ;
        

email
            phone
            region


        RowKey=  ;
        

email
            phone
            region


    }
    public string Phone { get; set; }
}
public class Player
{
    protected PlayerEntity player;
    async void GetPlayer(string cs,  table, string pk, string rk)
    {
        

CloudTable
            CloudTableClient
            TableEntity
            TableEntityAdapter



TEntity 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);



TEntity 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;
    }
}

```

Section: (none)

Explanation

Explanation/Reference:

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

References:

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

QUESTION 5

HOTSPOT

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.

Hot Area:

Answer Area

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

	▼	op = new		▼	() ;
TableOperation			TableOperation		
TableBatchOperaton			TableBatchOperaton		
TableEntity			TableEntity		
TableQuery			TableQuery		

...		▼	(op) ;
table.	ExecuteBatch		
	Execute		
	Insert		
	InsertOrMerge		

Correct Answer:

Answer Area

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

	▼	op = new		▼	() ;
TableOperation			TableOperation		
TableBatchOperaton			TableBatchOperaton		
TableEntity			TableEntity		
TableQuery			TableQuery		

...		▼	(op) ;
table.	ExecuteBatch		
	Execute		
	Insert		
	InsertOrMerge		

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Section: (none)
Explanation

Explanation/Reference:
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>

QUESTION 6

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

You need to capture the data required to implement the Usage Analytics feature of Application Insights.

Which three data values should you capture? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

Correct Answer: ADE
Section: (none)
Explanation

Explanation/Reference:
Explanation:

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

Application Insights manages the ID of a session for you.

References:
<https://github.com/microsoft/ApplicationInsights-Android>

QUESTION 7

DRAG DROP

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

The system has the following requirements for subscriptions:

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

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Filter types	Subscription	Filter type
SQLFilter	FutureOrders	
CorrelationFilter	HighPriorityOrders	
No Filter	InternationalOrders	
	HighQuantityOrders	
	AllOrders	

Correct Answer:

Answer Area

Filter types	Subscription	Filter type
SQLFilter	FutureOrders	SQLFilter
CorrelationFilter	HighPriorityOrders	CorrelationFilter
No Filter	InternationalOrders	SQLFilter
	HighQuantityOrders	SQLFilter
	AllOrders	No Filter

Section: (none)

Explanation

Explanation/Reference:

Explanation:

FutureOrders: SQLFilter

HighPriorityOrders: CorrelationFilter
CorrelationID only

InternationalOrders: SQLFilter
Country NOT USA requires an SQL Filter

HighQuantityOrders: SQLFilter
Need to use relational operators so an SQL Filter is needed.

AllOrders: No Filter

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

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

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

References:

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

QUESTION 8

DRAG DROP

Your company has several websites that use a company logo image. You use Azure Content Delivery

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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.

Select and Place:

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



Correct Answer:

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

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.

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.



Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 9

You develop a solution that uses an Azure SQL Database to store user information for a mobile app.

The app stores sensitive information about users.

You need to hide sensitive information from developers that query the data for the mobile app.

Which three items must you identify when configuring dynamic data masking? Each correct answer presents a part of the solution.

NOTE: Each correct selection is worth one point.

- A. Column
- B. Table
- C. Trigger
- D. Index
- E. Schema

Correct Answer: ABE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In the Dynamic Data Masking configuration page, you may see some database columns that the recommendations engine has flagged for masking. In order to accept the recommendations, just click Add Mask for one or more columns and a mask is created based on the default type for this column. You can change the masking function by clicking on the masking rule and editing the masking field format to a different format of your choice.


Dynamic Data Masking
demo_database

 Save
 Discard
 Add Mask


Downlevel clients require the use of Security Enabled Connection Strings.

Masking Rules

MASK NAME	MASK FUNCTION
You haven't created any masking rules.	

SQL users excluded from masking (administrators are always excluded) ⓘ

SQL users excluded from masking (administrators are always excluded)


Recommended fields to mask

SCHEMA	TABLE	COLUMN	
SalesLT	Customer	FirstName	
SalesLT	Customer	LastName	
SalesLT	Customer	EmailAddress	
SalesLT	Customer	Phone	
SalesLT	CustomerAddress	AddressID	

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started-portal>

QUESTION 10

HOTSPOT

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

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

```
1 CloudTableClient tableClient = account.CreateCloudTableClient();
2 CloudTable table = tableClient.GetTableReference ("people");
3 TableQuery<CustomerEntity> query = new TableQuery < CustomerEntity > ()
4.   .Where (TableQuery.CombineFilters (
5     TableQuery.GenerateFilterCondition (PartitionKey, QueryComparisons.Equal, "Smith"),
6     TableOperators.And , TableQuery.
       GenerateFilterCondition(Email, QueryComparisons.Equal, "ssmith@contoso.com")
7 ));
8 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.

Hot Area:

Answer Area

	Yes	No
The code returns every Record where the surname equals Smith .	☐	☐
The table endpoint <code>https://<mytableendpoint>/People</code> (PartitionKey='Smith',RowKey='Steve') returns the same results as the code.	☐	☐

Correct Answer:

Answer Area

	Yes	No
The code returns every Record where the surname equals Smith .	☐	☒
The table endpoint <code>https://<mytableendpoint>/People</code> (PartitionKey='Smith',RowKey='Steve') returns the same results as the code.	☒	☐

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: No

Only the second row is returned thank to the Email filter condition.

Box 2: Yes

This also returns the second row.

QUESTION 11

A company uses Azure SQL Database to store data for an app. The data includes sensitive information in a column named CustDetails.

You need to implement measures that allow only members of the managers group to see sensitive

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

information.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Include the managers group.
- B. Exclude the managers group.
- C. Exclude the administrators group.
- D. Navigate to the following URL:

```
PUT https://management.azure.com/subscriptions/00000000-1111-2222-3333-444444444444/resourceGroups/rg01/providers/Microsoft.Sql/servers/server01/databases/db01/transparentDataEncryption/current?api-version=2014-04-01
```

- E. Run the following Azure PowerShell command:

```
New-AzureRmSqlDatabaseDataMaskingRule -SchemaName "dbo" -TableName "CustDetails" -ColumnName "CustDetails" -MaskingFunction "Default"
```

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

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

References:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/new-azurermsqldatabasedatamaskingrule?view=azurermps-6.13.0>

QUESTION 12

HOTSPOT

You are developing a ticket reservation system for an airline.

The storage solution for the application must meet the following requirements:

- Ensure at least 99.99% availability and provide low latency.
- Accept reservations even when localized network outages or other unforeseen failures occurs
- Process reservations in the exact sequence as reservation are submitted to minimize overbooking or selling the same seat to multiple travelers.
- Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window.

You provision a resource group named airlineResourceGroup in the Azure South-Central US region.

You need to provision a SQL API Cosmos DB account to support the app.

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

NOTE: Each correct selection is worth one point.

Hot Area:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Answer Area

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

consistencyLevel=

	▼
Strong	
Eventual	
ConsistentPrefix	
BoundedStaleness	

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

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

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

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

```
- -default-consistency-level = $consistencyLevel
```

Correct Answer:

Answer Area

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

consistencyLevel=

	▼
Strong	
Eventual	
ConsistentPrefix	
BoundedStaleness	

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

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

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

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

```
- -default-consistency-level = $consistencyLevel
```

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: 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.

" Process reservations in the exact sequence as reservation are submitted to minimize overbooking or selling the same seat to multiple travelers."

Box 2: Enable-automatic-failover

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

For multi-region Cosmos accounts that are configured with a single-write region, enable automatic-failover by using Azure CLI or Azure portal. After you enable automatic failover, whenever there is a regional disaster, Cosmos DB will automatically failover your account.

" Accept reservations even in network outages or other unforeseen failures"

Box 3: southcentralus

" You provision a resource group named airlineResourceGroup in the Azure South-Central US region."

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/high-availability.md>

QUESTION 13

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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Start-AzureStorageBlobCopy cmdlet starts to copy a blob.

Example 1: Copy a named blob

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

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

References:

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

QUESTION 14

HOTSPOT

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
10 public class PlayerScore : TableEntity
11 {
12     public PlayerScore(string gameId, string playerId, int score, long timePlayed)
13     {
14         PartitionKey = gameId;
15         RowKey = playerId;
16         Score = score;
17         TimePlayed = timePlayed;
18     }
19     public int Score { get; set; }
20     public long TimePlayed { get; set; }
21 }

```

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

```

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

```

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

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

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

```

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

```

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 15

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

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

You need to store the information.

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The worst solution from a performance and scalability standpoint is to use a database backed session state provider. Instead use Azure Cache for Redis.

QUESTION 16

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

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

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

You need to store the information.

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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.

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

[provider](#)

QUESTION 17

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

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

You are developing 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 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.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

References:

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

QUESTION 18

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

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

You are developing 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 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

Correct Answer: B

Section: (none)

Explanation

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation/Reference:

References:

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

QUESTION 19

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

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

You are developing 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 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.

References:

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

QUESTION 20

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

Correct Answer: C

Section: (none)

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation

Explanation/Reference:

Explanation:

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

References:

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

QUESTION 21

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 22

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

- 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

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 23

HOTSPOT

A company develops a series of mobile games. All games use a single leaderboard service.

You have the following requirements:

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

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

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

You develop the following code to save scores in the database. (line numbers are included for reference only.)

```

01 public void SaveScore(string gameId, string playerId, int score, long timePlayed)
02 {
03     CloudStorageAccount storageAccount = CloudStorageAccount.Parse(connectionString);
04     CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
05     CloudTable table = tableClient.GetTableReference("scoreTable");
06     table.CreateIfNotExists();
07     var scoreRecord = new PlayerScore(gameId, playerId, score, timePlayed);
08     TableOperation insertOperation = TableOperation.Insert(scoreRecord);
09     table.Execute(insertOperation);
10 }

```

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

You develop the following code to query the database. (line numbers are included for reference only.)

```
01 CloudTableClient tableClient = account.CreateCloudTableClient();
02 CloudTable table = tableClient.GetTableReference("people");
03 TableQuery<CustomerEntity> query = new TableQuery<CustomerEntity>()
04     .Where(TableQuery.CombineFilters(
05         TableQuery.Generate.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal, "Smith")
06         TableOperators.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal,
07         "ssmith@contoso.com")
08     ));
09 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.

Hot Area:

Answer Area

	Yes	No
SaveScore will work with Cosmos DB.	☐	☐
The save score function will update and replace a record if one already exists with the same playerId and gameId.	☐	☐
Leader board data for the game will be automatically partitioned using gameId	☐	☐
SaveScore will store the values for the gameId and playerId parameters in the database.	☐	☐

Correct Answer:

Answer Area

	Yes	No
SaveScore will work with Cosmos DB.	☒	☐
The save score function will update and replace a record if one already exists with the same playerId and gameId.	☐	☒
Leader board data for the game will be automatically partitioned using gameId	☐	☒
SaveScore will store the values for the gameId and playerId parameters in the database.	☒	☐

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Yes

Code for CosmosDB, example:

// Parse the connection string and return a reference to the storage account.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse(
    CloudConfigurationManager.GetSetting("StorageConnectionString"));
// Create the table client.
CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
// Retrieve a reference to the table.
CloudTable table = tableClient.GetTableReference("people");
// Create the TableOperation object that inserts the customer entity.
TableOperation insertOperation = TableOperation.Insert(customer1);
```

Box 2: No

A new record will always be added as TableOperation.Insert is used, instead of TableOperation.InsertOrReplace.

Box 3: No

No partition key is used.

Box 4: Yes

References:

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

QUESTION 24

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: Enable Application Request Routing (ARR)

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 25

HOTSPOT

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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.

Hot Area:

Answer Area

Application State	Restorability
Data blob	Can be restored Cannot be restored
Snapshot 1	Can be restored Cannot be restored
Snapshot 2	Can be restored Cannot be restored

Correct Answer:

Answer Area

Application State

Restorability

Data blob

	▼
Can be restored	
Cannot be restored	

Snapshot 1

	▼
Can be restored	
Cannot be restored	

Snapshot 2

	▼
Can be restored	
Cannot be restored	

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 26

HOTSPOT

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation

Explanation/Reference:

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>

QUESTION 27

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

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

You are developing 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 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.

References:

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

Implement Azure security

Testlet 1

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

LabelMaker app

Coho Winery produces bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends **data** to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the **LabelMaker** application to include artwork and text for custom label designs.

Requirements

Data

You identify the following requirements for data management and manipulation:

- Order data is stored as nonrelational JSON and must be queried using Structured Query Language (SQL).
- Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Security

You have the following security requirements:

- Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
- External partners must use their own credentials and authenticate with their organization's identity management solution.
- External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.
- Storage of e-commerce application settings must be maintained in Azure Key Vault.
- E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).


- Conditional access policies must be applied at the application level to protect company content
- The LabelMaker applications must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

LabelMaker app

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture



Issues

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communications timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order .json

Relevant portions of the app files are shown below. Line numbers are included for reference only. This JSON file contains a representation of the data for an order that includes a single item.

```

01 {
02   "id" : 1,
03   "customers" : [
04     {
05       "familyName" : "Doe",
06       "givenName" : "John",
07       "customerid" : 5
08     }
09   ],
10   "line_items" : [
11     {
12       "fulfillable_quantity" : 1,
13       "id" : 6,
14       "price" : "199.99",
15       "product_id" : 7513594,
16       "quantity": 1,
17       "requires_shipping" : true,
18       "sku" : "SFC-342-N" ,
19       "title" : "Surface Go",
20       "vendor" : "Microsoft" ,
21       "name" : "Surface Go - 8GB",
22       "taxable" : true,
23       "tax_lines" : [
24         {
25           "title" : "State Tax",
26           "price" : "3.93",
27           "rate" : 0.06
28         }
29       ],
30       "total_discount" : "5.00"

```

art.com/

```

31   "discount_allocations" : [
32     {

```

QUESTION 1**HOTSPOT**

You need to meet the security requirements for external partners.

Which Azure Active Directory features should you use?

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area**Requirement****Option**

Authentication

	▼
B2C	
B2B	
Self-service signup	
Organizational Units (OU)	

Login Auditing

	▼
Access review	
Risky sign-ins report	
Identity Protection	
Privileged Identity Management	

Correct Answer:

Answer Area

Requirement

Option

Authentication

	▼
B2C	
B2B	
Self-service signup	
Organizational Units (OU)	

Login Auditing

	▼
Access review	
Risky sign-ins report	
Identity Protection	
Privileged Identity Management	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: B2B

Scenario: External partners must use their own credentials and authenticate with their organization's identity management solution.

Azure Active Directory (Azure AD) business-to-business (B2B) collaboration lets you securely share your company's applications and services with guest users from any other organization, while maintaining control over your own corporate data. Work safely and securely with external partners, large or small, even if they don't have Azure AD or an IT department. A simple invitation and redemption process lets partners use their own credentials to access your company's resources. Developers can use Azure AD business-to-business APIs to customize the invitation process or write applications like self-service sign-up portals.

Box 2: Access Review

Scenario: External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.

Azure Active Directory (Azure AD) Access Reviews enable organizations to efficiently manage group memberships, access to enterprise applications, and role assignments.

Administrators can use Azure Active Directory (Azure AD) to create an access review for group members or users assigned to an application. Azure AD automatically sends reviewers an email that prompts them to review access.

References:

<https://docs.microsoft.com/en-us/azure/active-directory/b2b/what-is-b2b>

QUESTION 2

You need to meet the security requirements for the E-Commerce Web App.

Which two steps should you take? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Update the E-Commerce Web App with the service principal's client secret.
- B. Enable Managed Service Identity (MSI) on the E-Commerce Web App.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

- C. Add a policy to the Azure Key Vault to grant access to the E-Commerce Web App.
- D. Create an Azure AD service principal.

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).

A managed identity from Azure Active Directory allows your app to easily access other AAD-protected resources such as Azure Key Vault. T

References:

<https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity>

QUESTION 3

You need to access user claims in the e-commerce web app.

What should you do first?

- A. Using the Azure CLI, enable Cross-origin resource sharing (CORS) from the e-commerce checkout API to the e-commerce web app.
- B. Update the e-commerce web app to read the HTTP request header values.
- C. Assign the Contributor RBAC role to the e-commerce web app by using the Resource Manager create role assignment API.
- D. Write custom code to make a Microsoft Graph API call from the e-commerce web app.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

If you want more information about the user, you'll need to use the Azure AD Graph API.

References: <https://docs.microsoft.com/en-us/azure/architecture/multitenant-identity/claims>

QUESTION 4

You need to meet the LabelMaker application security requirement.

What should you do?

- A. Create a conditional access policy and assign it to the Azure Kubernetes Service cluster.
- B. Place the Azure Active Directory account into an Azure AD group. Create a ClusterRoleBinding and assign it to the group.
- C. Create a RoleBinding and assign it to the Azure AD account.
- D. Create a Microsoft Azure Active Directory service principal and assign it to the Azure Kubernetes Service (AKS) cluster.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: The LabelMaker applications must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Permissions can be granted within a namespace with a RoleBinding, or cluster-wide with a ClusterRoleBinding.

References:

<https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Implement Azure security

Testlet 2

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

Background

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web API project that is deployed as an Azure Web App.

Overall architecture

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing

Employees may upload receipts in two ways:

- Uploading using an Azure Files mounted folder
- Uploading using the web application

Data Storage

Receipt and employee information is stored in an Azure SQL database.

Documentation

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for Azure File upload, and instructions on how to configure the mounted folder.

Solution details

Users table

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging

Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled. Application Insights must always contain all log messages.

Requirements

Receipt processing

Concurrent processing of a receipt must be prevented.

Disaster recovery

Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security

- Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
- All certificates and secrets used to secure data must be stored in Azure Key Vault.
- You must adhere to the Least Privilege Principal and provide privileges which are essential to perform the intended function.
- All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI)
- Receipt data must always be encrypted at rest.
- All data must be protected in transit.
- User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment with the remaining parts obscured.
- In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues

Upload format issue

Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal Server error page.

Capacity issue

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Log capacity issue

Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Processing.cs

```
PC01 public static class Processing
PC02 {
PC03     public static class Function
PC04     {
PC05         [FunctionName ("IssueWork")]
PC06         public static async Task Run ([TimerTrigger("0 */5" ***)] TimerInfo timer, ILogger log)
PC07         {
PC08             var container = await GetCloudBlobContainer();
PC09             foreach (var fileItem in await ListFiles())
PC10             {
PC11                 var file = new CloudFile (fileItem.StorageUri.PrimaryUri);
PC12                 var ms = new MemoryStream();
PC13                 await file.DownloadToStreamAsync(ms);
PC14                 var blob = container.GetBlockBlobReference (fileItem.Uri.ToString());
PC15                 await blob.UploadFromStreamAsync(ms);
PC16             }
PC17         }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob (CloudBlockBlob sourceBlob)
PC20     {
PC21         . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25         var cloudBlobClient = new CloudBlobClient (new Uri(" . . ."), await GetCredentials());
PC26
PC27         await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28         return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32         . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36         . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(" . . .");
PC38 }
PC39 }
```

Database.cs

```

DB01 public class Database
DB02 {
DB03     private string ConnectionString =
DB04
DB05     public async Task<object> LoadUserDetails(string userId)
DB06     {
DB07
DB08     return await policy.ExecuteAsync (async () =>
DB09     {
DB10         using (var connection = new SqlConnection (ConnectionString))
DB11         {
DB12             await connection.OpenAsync();
DB13             using (var command = new SqlCommand("_", connection))
DB14             using (var reader = command.ExecuteReader())
DB15             {
DB16                 -
DB17             }
DB18         }
DB19     }
DB20 }
DB21 }

```

ReceiptUploader.cs

```

RU01 public class ReceiptUploader
RU02 {
RU03     public async Task UploadFile(string file, byte[] binary)
RU04     {
RU05         var httpClient = new HttpClient();
RU06         var response = await httpClient.PutAsync( "...", new ByteArrayContent(binary));
RU07         while (ShouldRetry (response))
RU08         {
RU09             response = await httpClient.PutAsync ( "...", new ByteArrayContent(binary));
RU10         }
RU11     }
RU12     private bool ShouldRetry(HttpResponseMessage response)
RU13     {
RU14
RU15     }
RU16 }

```

ConfigureSSE.ps1

```

CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "... " -AccountName "... "
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "... "
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "... "
CS04 Set-AzureRmKeyVaultAccessPolicy'
CS05 -VaultName $keyVault.VaultName'
CS06 -ObjectId $storageAccount.Identity.PrincipalId'
CS07
CS08
CS09 Set-AzureRmStorageAccount"
CS10 -ResourceGroupName $storageAccount.ResourceGroupName'
CS11 -AccountName $storageAccount.StorageAccountName'
CS12 -EnableEncryptionService File `
CS13 -KeyvaultEncryption'
CS14 -KeyName $key.Name
CS15 -KeyVersion $key.Version'
CS16 -KeyVaultUri $keyVault.VaultUri

```

QUESTION 1

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

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

You need to ensure that the SecurityPin security requirements are met.

Solution: Enable Always Encrypted for the SecurityPin column using a certificate based on a trusted certificate authority. Update the Getting Started document with instructions to ensure that the certificate is installed on user machines.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Enable Always Encrypted is correct, but only the WebApplIdentity service principal should be given access to the certificate.

Scenario: Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.

QUESTION 2

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

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

You need to ensure that the SecurityPin security requirements are met.

Solution: Using the Azure Portal, add Data Masking to the SecurityPin column, and exclude the dbo user. Add a SQL security policy with a filter predicate based on the user identity.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead of DataMasing, enable Always Encrypted for the SecurityPin column.

Scenario: Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.

QUESTION 3

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might

have more than one correct solution, while others might not have a correct solution.

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

You need to ensure that the SecurityPin security requirements are met.

Solution: Enable Always Encrypted for the SecurityPin column using a certificate contained in Azure Key Vault and grant the WebAppIdentity service principal access to the certificate.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.

QUESTION 4

You need to ensure the security policies are met.

What code do you add at line CS07?

- A. `-PermissionsToCertificates create, encrypt, decrypt`
- B. `-PermissionsToKeys wrapkey, unwrapkey, get`
- C. `-PermissionsToCertificates wrapkey, unwrapkey, get`
- D. `-PermissionsToKeys create, encrypt, decrypt`

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: All certificates and secrets used to secure data must be stored in Azure Key Vault.

You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.

The Set-AzureRmKeyVaultAccessPolicy parameter -PermissionsToKeys specifies an array of key operation permissions to grant to a user or service principal. The acceptable values for this parameter: decrypt, encrypt, unwrapKey, wrapKey, verify, sign, get, list, update, create, import, delete, backup, restore, recover, purge

Incorrect Answers:

A, C: The Set-AzureRmKeyVaultAccessPolicy parameter -PermissionsToCertificates specifies an array of certificate permissions to grant to a user or service principal. The acceptable values for this parameter: get, list, delete, create, import, update, managecontacts, getissuers, listissuers, setissuers, deleteissuers, manageissuers, recover, purge, backup, restore

References:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.keyvault/set-azurermkeyvaultaccesspolicy>

QUESTION 5

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

You need to ensure that the SecurityPin security requirements are met.

Solution: Configure the web application to connect to the database using the WebAppIdentity security principal. Using the Azure Portal, add Data Masking to the SecurityPin column and exclude the WebAppIdentity service principal.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead of DataMasing, enable Always Encrypted for the SecurityPin column.

Scenario: Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins. All certificates and secrets used to secure data must be stored in Azure Key Vault.

QUESTION 6

HOTSPOT

You need to ensure that security policies are met.

What code should you add at line PC26?

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

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

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

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

Correct Answer:

Answer Area

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

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

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

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

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: var key = await Resolver.ResolveKeyAsyn(keyBundle,KeyIdentifier.CancellationToken.None);

Box 2: var x = new BlobEncryptionPolicy(key,resolver);

Example:

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Box 3: cloudblobClient. DefaultRequestOptions.EncryptionPolicy = x;

References:

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

QUESTION 7

HOTSPOT

You need to ensure that security requirements are met.

How should you complete the code segment?

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

New-AzureRmSqlDatabaseDataMaskingRule `

-ServerName "datastore" `

-DatabaseName "expense" `

-ResourceGroupName "production" `

-SchemaName "dbo" `

-TableName "Users" `

-ColumnName " 

Userld" `
Supervisorld" `
ExpenseAccount" `
AllowedAmount" `

-MaskingFunction 

Text `
Default `
Number `
NoMasking `

-SuffixSize 

2 `
3 `
4 `
9 `

-ReplacementString " 

XX" `
XXXX" `
XXXXXXX" `
XXXXXXXXXX" `

Correct Answer:

Answer Area

<code>New-AzureRmSqlDatabaseDataMaskingRule</code>	
<code>-ServerName</code>	<code>"datastore"</code>
<code>-DatabaseName</code>	<code>"expense"</code>
<code>-ResourceGroupName</code>	<code>"production"</code>
<code>-SchemaName</code>	<code>"dbo"</code>
<code>-TableName</code>	<code>"Users"</code>
<code>-ColumnName</code>	Userld" Supervisorld" ExpenseAccount" AllowedAmount"
<code>-MaskingFunction</code>	Text ' Default ' Number ' NoMasking '
<code>-SuffixSize</code>	2' 3' 4' 9'
<code>-ReplacementString</code>	XX" XXXX" XXXXXXX" XXXXXXXXXX"

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Box 1: ExpenseAccount

Scenario: User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment. With the remaining parts obscured.

Box 2: Text

If MaskingFunction has a value of Number or Text, you can specify the NumberFrom and NumberTo parameters, for number masking, or the PrefixSize, ReplacementString, and SuffixSize for text masking.

Box 3: 4

-SuffixSize specifies the number of characters at the end of the text that are not masked. Specify this parameter only if you specify a value of Text for the MaskingFunction parameter.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Scenario: Format is 1234-1234-1235

Box 4: xxxx"

Scenario: Format is 1234-1234-1235

References:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/new-azurermsql databasedatamaskingrule?view=azurerm-ps-6.13.0>

QUESTION 8

You need to ensure that the SecurityPin security requirements are met.

- A. Configure the web application to connect to the database using the WebAppIdentity security principal. Using the Azure Portal, add Data Masking to the SecurityPin column and exclude the WebAppIdentity service principal.
- B. Using the Azure Portal, add Data Masking to the SecurityPin column, and exclude the dbo user. Add a SQL security policy with a filter predicate based on the user identity.
- C. Enable Always Encrypted for the SecurityPin column using a certificate based on a trusted certificate authority. Update the Getting Started document with instructions to ensure that the certificate is installed on user machines.
- D. Enable Always Encrypted for the SecurityPin column using a certificate contained in Azure Key Vault and grant the WebAppIdentity service principal access to the certificate.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.

Incorrect Answers:

A, B: Instead of DataMasking, enable Always Encrypted for the SecurityPin column.

C: Enable Always Encrypted is correct, but only the WebAppIdentity service principal should be given access to the certificate.

Implement Azure security

Testlet 3

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

Background

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements

Policy service

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

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

Policies

Log Policy

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named **logdrop**. Logs must remain in the container for 15 days.

Authentication events

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

PolicyLib

You have a shared library named **PolicyLib** that contains functionality common to all ASP.NET Core web services and applications. The **PolicyLib** library must:

- Exclude non-user actions from Application Insights telemetry.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

- Provide methods that allow a web service to scale itself
- Ensure that scaling actions do not disrupt application usage

Other

Anomaly detection service

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring

All web applications and services have health monitoring at the /health service endpoint.

Issues

Policy loss

When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue

When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency

Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

```

EventGridController.cs
EG001 public class EventGridController : Controller
EG002 {
EG003     public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG004     public IActionResult Process([FromBody] string eventsJson
EG005     {
EG006         var events = JArray.Parse(eventsJson);
EG007
EG008         foreach (var @event in events)
EG009         {
EG010             EventId.Value = @event ["id"].ToString();
EG011             if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG012             {
EG013                 SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG014             }
EG015
EG016             {
EG017                 EnsureLogging(@event["subject"].ToString());
EG018             }
EG019         }
EG020         return null;
EG021     }
EG022     private void EnsureLogging(string resource)
EG023     {
EG024         . . .
EG025     }
EG026     private async Task SendToAnomalyDetectionService(string uri)
EG027     {
EG028         var content = GetLogData(uri);
EG029         var scoreRequest = new
EG030         {
EG031             Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG032             {
EG033                 {
EG034                     "input1",
EG035                     new List<Dictionary<string, string>>()
EG036                     {
EG037                         new Dictionary<string, string>()
EG038                         {
EG039                             {
EG040                                 "logcontent", content
EG041                             }
EG042                         }
EG043                     }
EG044                 },
EG045             },
EG046             GlobalParameters = new Dictionary<string, string>() { }
EG047         };
EG048         var result = await (new HttpClient()).PostAsJsonAsync(". . .", scoreRequest);
EG049         var rawModelResult = await result.Content.ReadAsStringAsync();
EG050         var modelResult = JObject.Parse(rawModelResult);
EG051         if (modelResult["notify"].HasValues)
EG052         {
EG053             . . .
EG054         }
EG055     }
EG056     private (string name, string resourceGroup) ParseResourceId(string
resourceId)
EG057     {
EG058         . . .
EG059     }
EG060     private string GetLogData(string uri)
EG061     {
EG062         . . .
EG063     }
EG064     static string BlobStoreAccountSAS(string containerName)
EG065     {
EG066         . . .
EG067     }
EG068 }

```

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```
LE01 public class LoginEvent
LE02 {
LE03
LE04 public string subject { get; set; }
LE05 public DateTime eventTime { get; set; }
LE06 public Dictionary<string, string> data { get; set; }
LE07 public string Serialize()
LE08 {
LE09     return JsonConvert.SerializeObject(this);
LE10 }
LE11 }
```

QUESTION 1

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

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

You need to ensure that authentication events are triggered and processed according to the authentication events policy.

Solution: Create a new Azure Event Grid subscription for all authentication that delivers messages to an Azure Event Hub. Use the subscription to process signout events.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Use a separate Azure Event Grid topics and subscriptions for sign-in and sign-out events.

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

Implement Azure security

Testlet 4

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment

Windows Server 2016 virtual machine

This virtual machine (VM) runs Biz Talk Server 2016. The VM runs the following workflows:

- Ocean Transport – This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.
- Inland Transport – This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

- Container API – This API provides container information including weight, contents, and other attributes.
- Location API – This API provides location information regarding shipping ports of call and truck stops.
- Shipping REST API – This API provides shipping information for use and display on the shipping website.

Shipping Data

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

Shipping Web Site

The site displays shipping container tracking information and container contents. The site is located at <http://shipping.wideworldimporters.com>


Proposed solution

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:



Shipping Logic App

The Shipping Logic app must meet the following requirements:

- Support the ocean transport and inland transport workflows by using a Logic App.
- Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
- Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
- Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app

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

REST APIs

The REST API's that support the solution must meet the following requirements:

- Secure resources to the corporate VNet.
- 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.

Shipping data

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

Shipping website

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

Issues

Windows Server 2016 VM

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

Shipping website and REST APIs

The following error message displays while you are testing the website:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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.

QUESTION 1

HOTSPOT

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.

Hot Area:

Answer Area

```
<?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>
```

Correct Answer:

Answer Area

```
<?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>
```

Section: (none)

Explanation

Explanation/Reference:

Explanation:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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>

QUESTION 2

HOTSPOT

You need to secure the Shipping Function app.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

Box 1: Function

Box 2: JSON based Token (JWT)

Azure AD uses JSON based tokens (JWTs) that contain claims

Box 3: HTTP

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

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

References:

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

QUESTION 3

You need to secure the Shipping Logic App.

What should you use?

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 4

HOTSPOT

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.

Hot Area:

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

cors	add
config	up
deployment	remove

--

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

Correct Answer:

```
az webapp
```

cors

config

deployment

add

up

remove

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

slot

allowed-origins

name

http://*.wideworldimporters.com

http://test-shippingapi.wideworldimporters.com

http://test.wideworldimporters.com

http://www.wideworldimporters.com

Section: (none)

Explanation

Explanation/Reference:

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>

Implement Azure security

Question Set 5

QUESTION 1

HOTSPOT

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.

You need to create a new custom role by reusing a default role definition and changing the permissions.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Item	Value
Powershell command	▼ Get-AzureRmRoleDefinition-Name"Reader" ConvertTo-Json Out-File C:\SupportRole.json Get-AzureRmRoleDefinition-Name"Operator" ConvertTo-Json Out-File C:\SupportRole.json Set-AzureRmRoleDefinition-Name"Reader" Input-File C:\SupportRole.json Set-AzureRmRoleDefinition Input-File C:\SupportRole.json
Actions section	▼ "*/read", "Microsoft.Support/*" "*/read" "*/", "Microsoft.Support/*" "*/"

Correct Answer:

Answer Area

Item	Value
Powershell command	▼ Get-AzureRmRoleDefinition-Name"Reader" ConvertTo-Json Out-File C:\SupportRole.json Get-AzureRmRoleDefinition-Name"Operator" ConvertTo-Json Out-File C:\SupportRole.json Set-AzureRmRoleDefinition-Name"Reader" Input-File C:\SupportRole.json Set-AzureRmRoleDefinition Input-File C:\SupportRole.json
Actions section	▼ "*/read", "Microsoft.Support/*" "*/read" "*/", "Microsoft.Support/*" "*/"

Section: (none)

Explanation

Explanation/Reference:

Explanation:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Incorrect Answers:

Get-AzureRmRoleDefinition. The Get-AzureRmRoleDefinition command does not have an action section.

First, use the Get-AzureRmRoleDefinition command to retrieve the custom role that you wish to modify. Then, modify the properties that you wish to change. Finally, save the role definition using the Set-AzureRmRoleDefinition command.

References:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/custom-roles-powershell>

QUESTION 2

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

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A: Multi-Factor Authentication comes as part of the following offerings:

- Azure Active Directory Premium licenses - Full featured use of Azure Multi-Factor Authentication Service (Cloud) or Azure Multi-Factor Authentication Server (On-premises).
- Multi-Factor Authentication for Office 365
- Azure Active Directory Global Administrators

C: MFA Enabled by conditional access policy. It is the most flexible means to enable two-step verification for your users. Enabling using conditional access policy only works for Azure MFA in the cloud and is a premium feature of Azure AD.

References:

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

QUESTION 3

DRAG DROP

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

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

- a set-variable policy to store the detected user identity
- a cache-lookup-value policy
- a cache-store-value policy
- a find-and-replace policy to update the response body with the user profile information

To which policy section should you add the policies? To answer, drag the appropriate sections to the correct policies. Each section may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Policy section	Policy	Policy section
Inbound	Set-variable	
Outbound	Cache-lookup-value	
	Cache-store-value	
	Find-and-replace	

Correct Answer:

Answer Area

Policy section	Policy	Policy section
Inbound	Set-variable	Inbound
Outbound	Cache-lookup-value	Inbound
	Cache-store-value	Outbound
	Find-and-replace	Outbound

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Inbound.

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

Example:

<policies>

<inbound>

<!-- How you determine user identity is application dependent -->

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

```
<set-variable  
name="enduserid"  
value="@{context.Request.Headers.GetValueOrDefault("Authorization","").Split(' ')[1].AsJwt()?.Subject}" />  
Etc.
```

Box 2: Inbound

A cache-lookup-value policy

Example:

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

Box 3: Outbound

A cache-store-value policy.

Example:

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

Box 4: Outbound

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

Example:

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

References:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

QUESTION 4

HOTSPOT

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

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

You need to configure Azure Disk Encryption for the VM.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
az provider register -n Microsoft.KeyVault
resourcegroup= "myResourceGroup"
az group create - -name $resourcegroup - -location westus
keyvault_name=myvaultname$RANDOM
```

```
az 

|               |   |
|---------------|---|
|               | ▼ |
| vm            |   |
| keyvault      |   |
| keyvault key  |   |
| vm encryption |   |

 create\
```

```
- -name $keyvault_name \
- -resource-group $resourcegroup\
- -location eastus \
- -enabled-for-disk-encryption True
```

```
az 

|               |   |
|---------------|---|
|               | ▼ |
| vm            |   |
| keyvault      |   |
| keyvault key  |   |
| vm encryption |   |

 create\
```

```
- -vault-name $keyvault_name\
- -name Name1 \
- -protection software
```

```
az 

|               |   |
|---------------|---|
|               | ▼ |
| vm            |   |
| keyvault      |   |
| keyvault key  |   |
| vm encryption |   |

 create\
```

```
- -resource-group $resourcegroup \
- -name Name2 \
- -image Canonical:UbuntuServer:16.04-LTS:latest \
- -admin-username azureuser \
- -generate-ssh-keys \
- -data-disk-sizes-gb 5
```

```
az 

|               |   |
|---------------|---|
|               | ▼ |
| vm            |   |
| keyvault      |   |
| keyvault key  |   |
| vm encryption |   |

 enable\
```

```
- -resource-group $resourcegroup \
- -name Name2 \
- -disk-encryption-keyvault $keyvault_name \
- -key-encryption-key Name1 \
- -volume-type
```

```


|      |   |
|------|---|
|      | ▼ |
| all  |   |
| data |   |
| os   |   |


```

Correct Answer:

Answer Area


```
az provider register -n Microsoft.KeyVault
resourcegroup= "myResourceGroup"
az group create - -name $resourcegroup - -location westus
keyvault_name=myvaultname$RANDOM
```

```
az  create\
vm
keyvault
keyvault key
vm encryption
```


```
- -name $keyvault_name \
- -resource-group $resourcegroup\
- -location eastus \
- -enabled-for-disk-encryption True
```

```
az  create\
vm
keyvault
keyvault key
vm encryption
```

```
- -vault-name $keyvault_name\
- -name Name1 \
- -protection software
```

```
az  create\
vm
keyvault
keyvault key
vm encryption
```

```
- -resource-group $resourcegroup \
- -name Name2 \
- -image Canonical:UbuntuServer:16.04-LTS:latest \
- -admin-username azureuser \
- -generate-ssh-keys \
- -data-disk-sizes-gb 5
```

```
az  enable\
vm
keyvault
keyvault key
vm encryption
```

```
- -resource-group $resourcegroup \
- -name Name2 \
- -disk-encryption-keyvault $keyvault_name \
- -key-encryption-key Name1 \
- -volume-type
```

```

all
data
os
```

Section: (none)

Explanation

Explanation/Reference:

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 \
```

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

```
--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/en-us/azure/virtual-machines/linux/disk-encryption-cli-quickstart>

QUESTION 5

HOTSPOT

You are developing an Azure App Service hosted ASP.NET Core web app to deliver video on-demand streaming media. You enable an Azure Content Delivery Network (CDN) Standard for the web endpoint. Customer videos are downloaded from the web app by using the following example URL:<http://www.contoso.com/content.mp4?quality=1>

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

You need to configure Azure CDN caching rules.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Setting	Action
Caching behavior	Bypass cache Override Set if missing
Cache expiration duration	1 second 1 minute 1 hour 1 day
Query string caching behavior	Ignore query strings Bypass caching for query strings Cache every unique URL

Correct Answer:

Answer Area

Setting	Action
Caching behavior	Bypass cache Override Set if missing
Cache expiration duration	1 second 1 minute 1 hour 1 day
Query string caching behavior	Ignore query strings Bypass caching for query strings Cache every unique URL

Section: (none)
Explanation

Explanation/Reference:
Explanation:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Box 1: Override

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

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

Incorrect:

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

Box 2: 1 hour

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

Box 3: Cache every unique URL

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

Incorrect Answers:

Bypass caching for query strings: In this mode, requests with query strings are not cached at the CDN POP node. The POP node retrieves the asset directly from the origin server and passes it to the requestor with each request.

Ignore query strings: Default mode. In this mode, the CDN point-of-presence (POP) node passes the query strings from the requestor to the origin server on the first request and caches the asset. All subsequent requests for the asset that are served from the POP ignore the query strings until the cached asset expires.

References:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-query-string>

QUESTION 6

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

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

You need to grant the role.

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A: The az role assignment create command creates a new role assignment for a user, group, or service principal.

Example: Create role assignment for an assignee.

az role assignment create --assignee sp_name --role a_role

C: The New-AzureRmRoleAssignment command assigns the specified RBAC role to the specified principal, at the specified scope.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Incorrect Answers:

B, D: Creates a custom role in Azure RBAC.

References:

<https://docs.microsoft.com/en-us/cli/azure/role/assignment?view=azure-cli-latest#az-role-assignment-create>

<https://docs.microsoft.com/en-us/powershell/module/azurerm.resources/new-azurermroleassignment?view=azurerm-6.13.0>

QUESTION 7

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

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

QUESTION 8

DRAG DROP

You maintain an existing Azure SQL Database instance. Management of the database is performed by an external party. All cryptographic keys are stored in an Azure Key Vault.

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Responses	Protection method	Response
☐ Yes	Enable the Always Encrypted feature.	☐
☐ No	Set the column encryption setting to disabled.	☐
	Assign users to the Public fixed database role.	☐
	Store column encryption keys in the system catalog view in the database.	☐

Correct Answer:

Answer Area

Responses	Protection method	Response
☐ Yes	Enable the Always Encrypted feature.	☐ Yes
☐ No	Set the column encryption setting to disabled.	☐ No
	Assign users to the Public fixed database role.	☐ Yes
	Store column encryption keys in the system catalog view in the database.	☐ No

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Yes

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

Box 2: No

Box 3: Yes

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

the two VIEW permissions.

Box 4: No

All cryptographic keys are stored in an Azure Key Vault.

Reference:

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

QUESTION 9

HOTSPOT

You are building a website to access project data related to teams within your organization. The website does not allow anonymous access. Authentication performed using an Azure Active Directory (Azure AD) app named **internal**.

The website has the following authentication requirements:

- Azure AD users must be able to login to the website.
- Personalization of the website must be based on membership in Active Directory groups.

You need to configure the application's manifest to meet the authentication requirements.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
{
  ...
  "appId": "d61126e3-089b-4adb-b721-
d5023213df7d",
   : "All",
  

|                         |
|-------------------------|
| "optionalClaims"        |
| "groupMembershipClaims" |

 : true
  

|                           |
|---------------------------|
| "allowPublicClient"       |
| "oauth2Permissions"       |
| "requiredResourceAccess"  |
| "oauth2AllowImplicitFlow" |


  ...
}
```

Correct Answer:

Answer Area

```
{
  ...
  "appId": "d61126e3-089b-4adb-b721-
d5023213df7d",
  [ ] : "All",
  [
    "optionalClaims"
    "groupMembershipClaims"
  ]
  [ ] : true
  [
    "allowPublicClient"
    "oauth2Permissions"
    "requiredResourceAccess"
    "oauth2AllowImplicitFlow"
  ]
  ...
}
```

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: groupMembershipClaims

Scenario: Personalization of the website must be based on membership in Active Directory groups.

Group claims can also be configured in the Optional Claims section of the Application Manifest. Enable group membership claims by changing the groupMembershipClaim

The valid values are:

"All"

"SecurityGroup"

"DistributionList"

"DirectoryRole"

Box 2: oauth2Permissions

Scenario: Azure AD users must be able to login to the website.

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

Incorrect Answers:

oauth2AllowImplicitFlow. oauth2AllowImplicitFlow specifies whether this web app can request OAuth2.0

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

implicit flow access tokens. The default is false. This flag is used for browser-based apps, like Javascript single-page apps.

References:

<https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-fed-group-claims>

QUESTION 10

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

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

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

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

Solution:

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

- In the website, use the value of the groups claim from the JWT for the user to determine permissions.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To configure Manifest to include Group Claims in Auth Token

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

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

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

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

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

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

References:

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

QUESTION 11

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

questions will not appear in the review screen.

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

References:

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

QUESTION 12

DRAG DROP

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

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

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

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

Settings	Answer Area		
client_id			
delegated			
profile			
application			
user_impersonation			

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

Correct Answer:

Settings	Answer Area		
client_id			
delegated			
profile			
application			
user_impersonation			

API	Permission	Type
Azure Storage	user_impersonation	delegated
Microsoft Graph	User.Read	delegated

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: user_impersonation

Box 2: delegated

Example:

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

Box 3: delegated

Example

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

References:

<https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect-aspnetcore/calling-a-web-api-in-an-aspnet-core-web-application-using-azure-ad/>

QUESTION 13

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

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

References:

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

<https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate>

Monitor, troubleshoot, and optimize Azure solutions

Testlet 1

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

Background

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements

Policy service

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

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

Policies

Log policy

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events

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

PolicyLib

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

- Exclude non-user actions from Application Insights telemetry.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

- Provide methods that allow a web service to scale itself.
- Ensure that scaling actions do not disrupt application usage.

Other

Anomaly detection service

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring

All web applications and services have health monitoring at the /health service endpoint.

Issues

Policy loss

When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue

When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency

Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

```

EventGridController.cs
EG01 public class EventGridController : Controller
EG02 {
EG03     public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04     public IActionResult Process([FromBody] string eventsJson)
EG05     {
EG06         var events = JArray.Parse(eventsJson);
EG07
EG08         foreach (var @event in events)
EG09         {
EG10             EventId.Value = @event["id"].ToString();
EG11             if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12             {
EG13                 SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14             }
EG15
EG16             {
EG17                 EnsureLogging(@event["subject"].ToString());
EG18             }
EG19         }
EG20         return null;
EG21     }
EG22     private void EnsureLogging(string resource)
EG23     {
EG24         . . .
EG25     }
EG26     private async Task SendToAnomalyDetectionService(string uri)
EG27     {
EG28         var content = GetLogData(uri)
EG29         var scoreRequest = new
EG30         {
EG31             Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32             {
EG33                 {
EG34                     "input1",
EG35                     new List<Dictionary<string, string>>()
EG36                     {
EG37                         new Dictionary<string, string>()
EG38                         {
EG39                             {
EG40                                 "logcontent", content
EG41                             }
EG42                         }
EG43                     },
EG44                 },
EG45             },
EG46             GlobalParameters = new Dictionary<string, string>() { }
EG47         };
EG48         var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49         var rawModelResult = await result.Content.ReadAsStringAsync();
EG50         var modelResult = JObject.Parse(rawModelResult);
EG51         if (modelResult["notify"].HasValues)
EG52         {
EG53             . . .
EG54         }
EG55     }
EG56     private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57     {
EG58         . . .
EG59     }
EG60     private string GetLogData(string uri)
EG61     {
EG62         . . .
EG63     }
EG64     static string BlobStoreAccountSAS (string containerName)
EG65     {
EG66         . . .
EG67     }
EG68 }

```

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```

LE01 public class LoginEvent
LE02 {
LE03
LE04     public string subject { get ; set ; }
LE05     public DateTime eventTime { get ; set ; }
LE06     public Dictionary<string, string> data { get; set; }
LE07     public string Serialize()
LE08     {
LE09         return JsonConvert.SerializeObject(this);
LE10     }
LE11 }

```

QUESTION 1

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

You need to meet the scaling requirements for Policy Service.

Which information should you cache?

- A. TempData
- B. HttpContext.Items
- C. ViewState
- D. Session state

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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.

Scenario: You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:
Ensure that scaling actions do not disrupt application usage.

References:

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

QUESTION 2

DRAG DROP

You need to implement telemetry for non-user actions.

How should you complete the Filter class? 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.

Select and Place:

Code segments

- `/health`
- `/status`
- `RequestTelemetry`
- `PageViewTelemetry`
- `ITelemetryProcessor`
- `ITelemetryInitializer`

Answer Area

```
public class filter :   
{  
    private readonly  _next;  
    public filter( next)  
    {  
        _next = next;  
    }  
    public void Process(ITelemetry item)  
    {  
        var x = item as ;  
        if (x?.Url.AbsolutePath == "")  
        {  
            return;  
        }  
        _next.Process(item);  
    }  
}
```

Correct Answer:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Code segments

/health

/status

RequestTelemetry

PageViewTelemetry

ITelemetryProcessor

ITelemetryInitializer

Answer Area

```
public class filter : ITelemetryProcessor
{
    private readonly ITelemetryProcessor _next;
    public filter(ITelemetryProcessor next)
    {
        _next = next;
    }

    public void Process(ITelemetry item)
    {
        var x = item as RequestTelemetry;
        if (x?.Url.AbsolutePath == "/health")
        {
            return;
        }
        _next.Process(item);
    }
}
```

Section: (none)**Explanation****Explanation/Reference:**

Explanation:

Scenario: Exclude non-user actions from Application Insights telemetry.

Box 1: ITelemetryProcessor

To create a filter, implement ITelemetryProcessor. This technique gives you more direct control over what is included or excluded from the telemetry stream.

Box 2: ITelemetryProcessor

Box 3: ITelemetryProcessor

Box 4: RequestTelemetry

Box 5: /health

To filter out an item, just terminate the chain.

References:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/api-filtering-sampling>**QUESTION 3**

You need to ensure that the solution can meet the scaling requirements for Policy Service.

Which Azure Application Insights data model should you use?

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

Correct Answer: A**Section: (none)****Explanation****Explanation/Reference:**

Explanation:

Application Insights provides three additional data types for custom telemetry:

<https://vceoreteconvert.com/><https://t.me/learningnets>

Trace - used either directly, or through an adapter to implement diagnostics logging using an instrumentation framework that is familiar to you, such as Log4Net or System.Diagnostics.
Event - typically used to capture user interaction with your service, to analyze usage patterns.
Metric - used to report periodic scalar measurements.

Scenario:

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

References:

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

QUESTION 4

DRAG DROP

You need to implement the Log policy.

How should you complete the Azure Event Grid subscription? 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.

Select and Place:

Code segment

Answer Area

All	<pre>{ "name": "newlogs", "properties": { "topic": "/subscriptions/ . . ./providers/Microsoft.EventGrid/topics/. . .", "destination": { "endpointType": " ", }, "filter": { " ": "/blobServices/default/containers/logdrop/", "includeEventTypes": [" "], }, "labels": [], "eventDeliverySchema": "EventGridSchema" } }</pre>
WebHook	
EventHub	
subjectEndsWith	
Microsoft.Storage	
subjectBeginsWith	
Microsoft.Storage.BlobCreated	

Correct Answer:

Code segment

Answer Area

All	<pre>{ "name": "newlogs", "properties": { "topic": "/subscriptions/ . . ./providers/Microsoft.EventGrid/topics/. . .", "destination": { "endpointType": " WebHook ", }, "filter": { " subjectBeginsWith ": "/blobServices/default/containers/logdrop/", "includeEventTypes": [" Microsoft.Storage.BlobCreated "], }, "labels": [], "eventDeliverySchema": "EventGridSchema" } }</pre>
WebHook	
EventHub	
subjectEndsWith	
Microsoft.Storage	
subjectBeginsWith	
Microsoft.Storage.BlobCreated	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Box 1:WebHook

Scenario: If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

endpointType: The type of endpoint for the subscription (webhook/HTTP, Event Hub, or queue).

Box 2: SubjectBeginsWith

Box 3: Microsoft.Storage.BlobCreated

Scenario: Log Policy

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Example subscription schema

```
{
  "properties": {
    "destination": {
      "endpointType": "webhook",
      "properties": {
        "endpointUrl": "https://example.azurewebsites.net/api/HttpTriggerCSharp1?code=VXbGWce53l48Mt8wuotr0GPmyJ/nDT4hgdFj9DpBiRt38qqnm5OFg=="
      }
    },
    "filter": {
      "includedEventTypes": [ "Microsoft.Storage.BlobCreated", "Microsoft.Storage.BlobDeleted" ],
      "subjectBeginsWith": "blobServices/default/containers/mycontainer/log",
      "subjectEndsWith": ".jpg",
      "isSubjectCaseSensitive ": "true"
    }
  }
}
```

References:

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

QUESTION 5

DRAG DROP

You need to ensure that PolicyLib requirements are met.

How should you complete the code segment? 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.

Select and Place:

Code segments	Answer Area
Process	public class IncludeEventId: []
Initialize	{
telemetry.Sequence	public void []
ITelemetryProcessor	(ITelemetry telemetry)
ITelemetryInitializer	{
Telemetry.Context	[].Properties["EventId"] =
EventGridController.EventId.Value	[]
((EventTelemetry)telemetry).Properties["EventId"]	}

Correct Answer:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Code segments

Process
Initialize
telemetry.Sequence
ITelemetryProcessor
ITelemetryInitializer
Telemetry.Context
EventGridController.EventId.Value
((EventTelemetry) telemetry).Properties["EventId"]

Answer Area

```
public class IncludeEventId: ITelemetryInitializer
{
    public void Initialize
    (ITelemetry telemetry)
    {
        Telemetry.Context.Properties["EventId"] =
        ((EventTelemetry)telemetry).Properties["EventId"]
    }
}
```

Section: (none)

Explanation**Explanation/Reference:**

Explanation:

Scenario: You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

- Exclude non-user actions from Application Insights telemetry.
- Provide methods that allow a web service to scale itself.
- Ensure that scaling actions do not disrupt application usage.

Box 1: ITelemetryInitializer

Use telemetry initializers to define global properties that are sent with all telemetry; and to override selected behavior of the standard telemetry modules.

Box 2: Initialize

Box 3: Telemetry.Context

Box 4: ((EventTelemetry)telemetry).Properties("EventID")

References:

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

Monitor, troubleshoot, and optimize Azure solutions

Testlet 2

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

Background

You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web API project that is deployed as an Azure Web App.

Overall architecture

Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.

Receipt processing

Employees may upload receipts in two ways:

- Uploading using an Azure Files mounted folder
- Uploading using the web application

Data Storage

Receipt and employee information is stored in an Azure SQL database.

Documentation

Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for Azure File upload, and instructions on how to configure the mounted folder.

Solution details

Users table

Column	Description
UserId	unique identifier for and employee
ExpenseAccount	employees expense account number in the format 1234-123-1234
AllowedAmount	limit of allowed expenses before approval is needed
SupervisorId	unique identifier for employee's supervisor
SecurityPin	value used to validate user identity

Web Application

You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.

Processing

Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Logging

Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled. Application Insights must always contain all log messages.

Requirements

Receipt processing

Concurrent processing of a receipt must be prevented.

Disaster recovery

Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.

Security

- Users' SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
- All certificates and secrets used to secure data must be stored in Azure Key Vault.
- You must adhere to the Least Privilege Principal and provide privileges which are essential to perform the intended function.
- All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI)
- Receipt data must always be encrypted at rest.
- All data must be protected in transit.
- User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment with the remaining parts obscured.
- In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues

Upload format issue

Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal Server error page.

Capacity issue

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Log capacity issue

Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.

Processing.cs

```
PC01 public static class Processing
PC02 {
PC03     public static class Function
PC04     {
PC05         [FunctionName ("IssueWork")]
PC06         public static async Task Run ([TimerTrigger("0 */5" ***)] TimerInfo timer, ILogger log)
PC07         {
PC08             var container = await GetCloudBlobContainer();
PC09             foreach (var fileItem in await ListFiles())
PC10             {
PC11                 var file = new CloudFile (fileItem.StorageUri.PrimaryUri);
PC12                 var ms = new MemoryStream();
PC13                 await file.DownloadToStreamAsync(ms);
PC14                 var blob = container.GetBlockBlobReference (fileItem.Uri.ToString());
PC15                 await blob.UploadFromStreamAsync(ms);
PC16             }
PC17         }
PC18     }
PC19     private static CloudBlockBlob GetDRBlob (CloudBlockBlob sourceBlob)
PC20     {
PC21         . . .
PC22     }
PC23     private static async Task<CloudBlobContainer> GetCloudBlobContainer()
PC24     {
PC25         var cloudBlobClient = new CloudBlobClient (new Uri(" . . ."), await GetCredentials());
PC26
PC27         await cloudBlobClient.GetRootContainerReference().CreateIfNotExistsAsync();
PC28         return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32         . . .
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36         . . .
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient(" . . .");
PC38 }
PC39 }
```

Database.cs

```

DB01 public class Database
DB02 {
DB03     private string ConnectionString =
DB04
DB05     public async Task<object> LoadUserDetails(string userId)
DB06     {
DB07
DB08     return await policy.ExecuteAsync (async () =>
DB09     {
DB10         using (var connection = new SqlConnection (ConnectionString))
DB11         {
DB12             await connection.OpenAsync();
DB13             using (var command = new SqlCommand("_", connection))
DB14             using (var reader = command.ExecuteReader())
DB15             {
DB16                 -
DB17             }
DB18         }
DB19     }
DB20 }
DB21 }

```

ReceiptUploader.cs

```

RU01 public class ReceiptUploader
RU02 {
RU03     public async Task UploadFile(string file, byte[] binary)
RU04     {
RU05         var httpClient = new HttpClient();
RU06         var response = await httpClient.PutAsync( "...", new ByteArrayContent(binary));
RU07         while (ShouldRetry (response))
RU08         {
RU09             response = await httpClient.PutAsync ( "...", new ByteArrayContent(binary));
RU10         }
RU11     }
RU12     private bool ShouldRetry(HttpResponseMessage response)
RU13     {
RU14
RU15     }
RU16 }

```

ConfigureSSE.ps1

```

CS01 $storageAccount = Get-AzureRmStorageAccount -ResourceGroupName "... " -AccountName "... "
CS02 $keyVault = Get-AzureRmKeyVault -VaultName "... "
CS03 $key = Get-AzureKeyVaultKey -VaultName $keyVault.VaultName -Name "... "
CS04 Set-AzureRmKeyVaultAccessPolicy'
CS05 -VaultName $keyVault.VaultName'
CS06 -ObjectId $storageAccount.Identity.PrincipalId'
CS07
CS08
CS09 Set-AzureRmStorageAccount"
CS10 -ResourceGroupName $storageAccount.ResourceGroupName'
CS11 -AccountName $storageAccount.StorageAccountName'
CS12 -EnableEncryptionService File `
CS13 -KeyvaultEncryption'
CS14 -KeyName $key.Name
CS15 -KeyVersion $key.Version'
CS16 -KeyVaultUri $keyVault.VaultUri

```

QUESTION 1

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

You need to resolve the log capacity issue.

What should you do?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario, the log capacity issue: Developers report that the number of log message in the trace output for the processor is too high, resulting in lost log messages.

Sampling is a feature in Azure Application Insights. It is the recommended way to reduce telemetry traffic and storage, while preserving a statistically correct analysis of application data. The filter selects items that are related, so that you can navigate between items when you are doing diagnostic investigations. When metric counts are presented to you in the portal, they are renormalized to take account of the sampling, to minimize any effect on the statistics.

Sampling reduces traffic and data costs, and helps you avoid throttling.

References:

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

QUESTION 2

You need to resolve the capacity issue.

What should you do?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

References:

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

QUESTION 3

You need to ensure receipt processing occurs correctly.

What should you do?

- A. Use blob metadata to prevent concurrency problems.
- B. Use blob SnapshotTime to prevent concurrency problems.
- C. Use blob leases to prevent concurrency problems.
- D. Use blob properties to prevent concurrency problems.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can create a snapshot of a blob. A snapshot is a read-only version of a blob that's taken at a point in time. Once a snapshot has been created, it can be read, copied, or deleted, but not modified. Snapshots provide a way to back up a blob as it appears at a moment in time.

Scenario: Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

References:

<https://docs.microsoft.com/en-us/rest/api/storageservices/creating-a-snapshot-of-a-blob>

Monitor, troubleshoot, and optimize Azure solutions

Testlet 3

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

LabelMaker app

Coho Winery produces bottles, and distributes a variety of wines globally. You are a developer implementing highly scalable and resilient applications to support online order processing by using Azure solutions.

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends **data** to several printers. The application consists of five modules that run independently on virtual machines (VMs). Coho Winery plans to move the application to Azure and continue to support label creation.

External partners send data to the **LabelMaker** application to include artwork and text for custom label designs.

Requirements

Data

You identify the following requirements for data management and manipulation:

- Order data is stored as nonrelational JSON and must be queried using Structured Query Language (SQL).
- Changes to the Order data must reflect immediately across all partitions. All reads to the Order data must fetch the most recent writes.

Security

You have the following security requirements:

- Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
- External partners must use their own credentials and authenticate with their organization's identity management solution.
- External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.
- Storage of e-commerce application settings must be maintained in Azure Key Vault.
- E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).


- Conditional access policies must be applied at the application level to protect company content
- The LabelMaker applications must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

LabelMaker app

Azure Monitor Container Health must be used to monitor the performance of workloads that are deployed to Kubernetes environments and hosted on Azure Kubernetes Service (AKS).

You must use Azure Container Registry to publish images that support the AKS deployment.

Architecture



Issues

Calls to the Printer API App fail periodically due to printer communication timeouts.

Printer communications timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.

The order workflow fails to run upon initial deployment to Azure.

Order .json

Relevant portions of the app files are shown below. Line numbers are included for reference only. This JSON file contains a representation of the data for an order that includes a single item.

```

01 {
02   "id" : 1,
03   "customers" : [
04     {
05       "familyName" : "Doe",
06       "givenName" : "John",
07       "customerid" : 5
08     }
09   ],
10   "line_items" : [
11     {
12       "fulfillable_quantity" : 1,
13       "id" : 6,
14       "price" : "199.99",
15       "product_id" : 7513594,
16       "quantity" : 1,
17       "requires_shipping" : true,
18       "sku" : "SFC-342-N" ,
19       "title" : "Surface Go",
20       "vendor" : "Microsoft" ,
21       "name" : "Surface Go - 8GB",
22       "taxable" : true,
23       "tax_lines" : [
24         {
25           "title" : "State Tax",
26           "price" : "3.93",
27           "rate" : 0.06
28         }
29       ],
30       "total_discount" : "5.00"

```

art.com/

```

31   "discount_allocations" : [
32     {

```

QUESTION 1

You need to troubleshoot the order workflow.

What should you do? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Review the trigger history.
- B. Review the API connections.
- C. Review the run history.
- D. Review the activity log.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: The order workflow fails to run upon initial deployment to Azure.

Deployment errors arise from conditions that occur during the deployment process. They appear in the activity log.

References:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-audit>

Monitor, troubleshoot, and optimize Azure solutions

Testlet 4

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment

Windows Server 2016 virtual machine

This virtual machine (VM) runs Biz Talk Server 2016. The VM runs the following workflows:

- Ocean Transport – This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.
- Inland Transport – This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

- Container API – This API provides container information including weight, contents, and other attributes.
- Location API – This API provides location information regarding shipping ports of call and truck stops.
- Shipping REST API – This API provides shipping information for use and display on the shipping website.

Shipping Data

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

Shipping Web Site

The site displays shipping container tracking information and container contents. The site is located at <http://shipping.wideworldimporters.com>


Proposed solution

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:



Shipping Logic App

The Shipping Logic app must meet the following requirements:

- Support the ocean transport and inland transport workflows by using a Logic App.
- Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
- Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
- Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app

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

REST APIs

The REST API's that support the solution must meet the following requirements:

- Secure resources to the corporate VNet.
- 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.

Shipping data

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

Shipping website

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

Issues

Windows Server 2016 VM

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

Shipping website and REST APIs

The following error message displays while you are testing the website:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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.

QUESTION 1
HOTSPOT

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.

Hot Area:

Answer Area

Issue	Tool
Backup and Restore	Azure Site Recovery Azure Backup Azure Data Box Azure Migrate
Performance	Azure Network Watcher Azure Traffic Manager ExpressRoute Accelerated Networking

Correct Answer:

Answer Area

Issue	Tool
Backup and Restore	Azure Site Recovery Azure Backup Azure Data Box Azure Migrate
Performance	Azure Network Watcher Azure Traffic Manager ExpressRoute Accelerated Networking

Section: (none)

Explanation

Explanation/Reference:

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/>

QUESTION 2

HOTSPOT

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.

Hot Area:

Answer Area

Option	Value
Tier	Standard Premium
Profile	Akamai Microsoft
Optimization	general web delivery large file download dynamic site acceleration video-on-demand media streaming

Correct Answer:

Answer Area

Option	Value
Tier	Standard Premium
Profile	Akamai Microsoft
Optimization	general web delivery large file download dynamic site acceleration video-on-demand media streaming

Section: (none)

Explanation

Explanation/Reference:

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.

References:

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

Monitor, troubleshoot, and optimize Azure solutions

Question Set 5

QUESTION 1

You are developing an ASP.NET Core Web API web service. The web service uses Azure Application Insights for all telemetry and dependency tracking. The web service reads and writes data to a database other than Microsoft SQL Server.

You need to ensure that dependency tracking works for calls to the third-party database.

Which two Dependency Telemetry properties should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Telemetry.Context.Operation.Id
- B. Telemetry.Name
- C. Telemetry.Context.Cloud.RoleInstance
- D. Telemetry.Context.Session.Id
- E. Telemetry.Id

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

References:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking>

QUESTION 2

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

Correct Answer: DF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

D: CorrelationId: Enables an application to specify a context for the message for the purposes of correlation; for example, reflecting the MessageId of a message that is being replied to.

F: ReplyToSessionId: This value augments the ReplyTo information and specifies which SessionId should be set for the reply when sent to the reply entity.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Incorrect Answers:

A, B: DeliveryCount

Number of deliveries that have been attempted for this message. The count is incremented when a message lock expires, or the message is explicitly abandoned by the receiver. This property is read-only.

C, E: SequenceNumber

The sequence number is a unique 64-bit integer assigned to a message as it is accepted and stored by the broker and functions as its true identifier. For partitioned entities, the topmost 16 bits reflect the partition identifier. Sequence numbers monotonically increase and are gapless. They roll over to 0 when the 48-64 bit range is exhausted. This property is read-only.

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

QUESTION 3

You develop a gateway solution for a public facing news API. The news API back end is implemented as a RESTful service and uses an OpenAPI specification.

You need to ensure that you can access the news API by using an Azure API Management service instance.

Which Azure PowerShell command should you run?

- A. `Import-AzureRmApiManagementApi -Context $ApiMgmtContext -SpecificationFormat "Swagger" -SpecificationPath $SwaggerPath -Path $Path`
- B. `New-AzureRmApiManagementBackend -Context $ApiMgmtContext -Url $Url -Protocol http`
- C. `New-AzureRmApiManagement -ResourceGroupName $ResourceGroup -Name $Name -Location $Location -Organization $Org -AdminEmail $AdminEmail`
- D. `New-AzureRmApiManagementBackendProxy -Url $ApiUrl`

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

New-AzureRmApiManagementBackendProxy creates a new Backend Proxy Object which can be piped when creating a new Backend entity.

Example: Create a Backend Proxy In-Memory Object

```
PS C:\>$secpassword = ConvertTo-SecureString "PlainTextPassword" -AsPlainText -Force
PS C:\>$proxyCreds = New-Object System.Management.Automation.PSCredential ("foo", $secpassword)
PS C:\>$credential = New-AzureRmApiManagementBackendProxy -Url "http://12.168.1.1:8080" -ProxyCredential $proxyCreds
```

```
PS C:\>$apimContext = New-AzureRmApiManagementContext -ResourceGroupName "Api-Default-WestUS" -ServiceName "contoso"
```

```
PS C:\>$backend = New-AzureRmApiManagementBackend -Context $apimContext -BackendId 123 -Url 'https://contoso.com/awesomeapi' -Protocol http -Title "first backend" -SkipCertificateChainValidation $true -Proxy $credential -Description "backend with proxy server"
Creates a Backend Proxy Object and sets up Backend
```

Incorrect Answers:

A: The Import-AzureRmApiManagementApi cmdlet imports an Azure API Management API from a file or a URL in Web Application Description Language (WADL), Web Services Description Language (WSDL), or Swagger format.

B: New-AzureRmApiManagementBackend creates a new backend entity in Api Management.

C: The New-AzureRmApiManagement cmdlet creates an API Management deployment in Azure API Management.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

References:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.apimanagement/new-azurermapimanagementbackendproxy?view=azurerm-6.13.0>

QUESTION 4

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

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 5

DRAW DROP

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

Spikes in traffic have caused increases in page load times.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

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

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

Select and Place:

Actions

Answer Area

- Enable autoscaling on the web app.
- Configure a Scale condition.
- Configure the web app to the Standard App Service tier.
- Configure the web app to the Premium App Service tier.
- Switch to an Azure App Services consumption plan.
- Add a Scale rule.



Correct Answer:

Actions

Answer Area

- Enable autoscaling on the web app.
- Configure a Scale condition.
- Configure the web app to the Standard App Service tier.
- Configure the web app to the Premium App Service tier.
- Switch to an Azure App Services consumption plan.
- Add a Scale rule.



- Configure the web app to the Standard App Service tier.
- Enable autoscaling on the web app.
- Add a Scale rule.
- Configure a Scale condition.

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Step 1: Configure the web app to the Standard App Service Tier
The Standard tier supports auto-scaling, and we should minimize the cost.

Step 2: Enable autoscaling on the web app
First enable autoscale

Step 3: Add a scale rule

Step 4: Add a Scale condidation

References:

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

<https://azure.microsoft.com/en-us/pricing/details/app-service/plans/>

QUESTION 6

DRAG DROP

You are developing an ASP.NET Core Web API web service that uses Azure Application Insights to monitor

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

performance and track events.

You need to enable logging and ensure that log messages can be correlated to events tracked by Application Insights.

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

NOTE: Each correct selection is worth one point.

Select and Place:

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

Correct Answer:

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

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: ApplicationInsightsLoggerOptions

If you want to include the EventId and EventName properties, then add the following to the ConfigureServices method:

```
services
.AddOptions<ApplicationInsightsLoggerOptions>()
.Configure(o => o.IncludeEventId = true);
```

Box 2: IncludeEventId

Box 3: ApplicationServices

In Asp.Net core apps it turns out that trace logs do not show up in Application Insights out of the box. We need to add the following code snippet to our Configure method in Startup.cs:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

```
loggerFactory.AddApplicationInsights(app.ApplicationServices, logLevel);
```

References:

<https://blog.computedcloud.com/enabling-application-insights-trace-logging-in-asp-net-core/>

QUESTION 7

DRAG DROP

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.

Select and Place:

Answer Area

Features	Requirement	Feature
Users	Which pages visited by users most often correlate to a product purchase?	
Funnels	How does load time of the product display page affect a user's decision to purchase a product?	
Impact	Which events most influence a user's decision to continue to use the application?	
Retention		
User Flows	Are there places in the application that users often perform repetitive actions?	

Correct Answer:

Answer Area

Features	Requirement	Feature
Users	Which pages visited by users most often correlate to a product purchase?	Users
Funnels	How does load time of the product display page affect a user's decision to purchase a product?	Impact
Impact	Which events most influence a user's decision to continue to use the application?	Retention
Retention		
User Flows	Are there places in the application that users often perform repetitive actions?	User Flows

Section: (none)

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation

Explanation/Reference:

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?

Incorrect Answers:

Funnel: If your application involves multiple stages, you need to know if most customers are progressing through the entire process, or if they are ending the process at some point. The progression through a series of steps in a web application is known as a funnel. You can use Azure Application Insights Funnels to gain insights into your users, and monitor step-by-step conversion rates.

References:

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

QUESTION 8

HOTSPOT

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.

Hot Area:

Answer Area

```
void ClearCachedTeams()
```

```
{
```

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

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

```
    cache.KeyDelete("Team");
```

```
    cache.StringSet("Team", "");
```

```
    cache.ValueDelete("Team");
```

```
    cache.StringGet("Team", "");
```

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

```
}
```

Correct Answer:

Answer Area

```
void ClearCachedTeams()
```

```
{
```

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

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

```
    cache.KeyDelete("Team");
```

```
    cache.StringSet("Team", "");
```

```
    cache.ValueDelete("Team");
```

```
    cache.StringGet("Team", "");
```

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

```
}
```

Section: (none)
Explanation

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation/Reference:

Explanation:

Box 1: `IDatabase cache = connection.GetDatabase();`
Connection refers to a previously configured `ConnectionMultiplexer`.

Box 2: `cache.StringSet("team",")`

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/>

QUESTION 9**HOTSPOT**

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

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

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

NOTE: Each correct selection is worth one point.

Hot Area:**Answer Area**

`az webapp log` `--name ContosoWeb --resource-group ContosoDevRG`

config
download
show
tail

`filesystem`

--web-server-logging
--docker-container-logging
--application-logging

`az` `log` `--name ContosoWeb --resource-group ContosoDevRG`

webapp
acr
aks

config
download
show
tail

Correct Answer:

Answer Area

az webapp log --name ContosoWeb --resource-group ContosoDevRG

config
download
show
tail

filesystem

--web-server-logging
--docker-container-logging
--application-logging

az log --name ContosoWeb --resource-group ContosoDevRG

webapp
acr
aks

config
download
show
tail

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: config

To Configure logging for a web app use the command:

az webapp log config

Box 2: --docker-container-logging

Syntax include:

az webapp log config [--docker-container-logging {filesystem, off}]

Box 3: webapp

To download a web app's log history as a zip file use the command:

az webapp log download

Box 4: download

References:

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

QUESTION 10

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

Correct Answer: E

Section: (none)

Explanation

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation/Reference:

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.

Incorrect Answers:

B: You can use the daily volume cap to limit the data collected.

To change the daily cap, in the Configure section of your Application Insights resource, in the Usage and estimated costs pane, select Daily Cap.

References:

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

Connect to and Consume Azure Services and Third-party Services

Testlet 1

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a **question**, click the Question button to return to the question.

Background

You are a developer for Proseware, Inc. You are developing an application that applies a set of governance policies for Proseware's internal services, external services, and applications. The application will also provide a shared library for common functionality.

Requirements

Policy service

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

The application must include the Event Grid Event ID field in all Application Insights telemetry.

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

Policies

Log policy

All Azure App Service Web Apps must write logs to Azure Blob storage. All log files should be saved to a container named logdrop. Logs must remain in the container for 15 days.

Authentication events

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

PolicyLib

You have a shared library named PolicyLib that contains functionality common to all ASP.NET Core web services and applications. The PolicyLib library must:

- Exclude non-user actions from Application Insights telemetry.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

- Provide methods that allow a web service to scale itself.
- Ensure that scaling actions do not disrupt application usage.

Other

Anomaly detection service

You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Health monitoring

All web applications and services have health monitoring at the /health service endpoint.

Issues

Policy loss

When you deploy Policy service, policies may not be applied if they were in the process of being applied during the deployment.

Performance issue

When under heavy load, the anomaly detection service undergoes slowdowns and rejects connections.

Notification latency

Users report that anomaly detection emails can sometimes arrive several minutes after an anomaly is detected.

App code

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

```

EventGridController.cs
EG01 public class EventGridController : Controller
EG02 {
EG03     public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04     public IActionResult Process([FromBody] string eventsJson)
EG05     {
EG06         var events = JArray.Parse(eventsJson);
EG07
EG08         foreach (var @event in events)
EG09         {
EG10             EventId.Value = @event["id"].ToString();
EG11             if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12             {
EG13                 SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14             }
EG15
EG16             {
EG17                 EnsureLogging(@event["subject"].ToString());
EG18             }
EG19         }
EG20         return null;
EG21     }
EG22     private void EnsureLogging(string resource)
EG23     {
EG24         . . .
EG25     }
EG26     private async Task SendToAnomalyDetectionService(string uri)
EG27     {
EG28         var content = GetLogData(uri)
EG29         var scoreRequest = new
EG30         {
EG31             Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32             {
EG33                 {
EG34                     "input1",
EG35                     new List<Dictionary<string, string>>()
EG36                     {
EG37                         new Dictionary<string, string>()
EG38                         {
EG39                             {
EG40                                 "logcontent", content
EG41                             }
EG42                         }
EG43                     },
EG44                 },
EG45             },
EG46             GlobalParameters = new Dictionary<string, string>() { }
EG47         };
EG48         var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49         var rawModelResult = await result.Content.ReadAsStringAsync();
EG50         var modelResult = JObject.Parse(rawModelResult);
EG51         if (modelResult ["notify"].HasValues)
EG52         {
EG53             . . .
EG54         }
EG55     }
EG56     private (string name, string resourceGroup) ParseResourceId(string resourceId)
EG57     {
EG58         . . .
EG59     }
EG60     private string GetLogData(string uri)
EG61     {
EG62         . . .
EG63     }
EG64     static string BlobStoreAccountSAS (string containerName)
EG65     {
EG66         . . .
EG67     }
EG68 }

```

Relevant portions of the app files are shown below. Line numbers are included for reference only and include a two-character prefix that denotes the specific file to which they belong.

LoginEvent.cs

```

LE01 public class LoginEvent
LE02 {
LE03
LE04     public string subject { get ; set ; }
LE05     public DateTime eventTime { get ; set ; }
LE06     public Dictionary<string, string> data { get; set; }
LE07     public string Serialize()
LE08     {
LE09         return JsonConvert.SerializeObject(this);
LE10     }
LE11 }

```

QUESTION 1

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

HOTSPOT

You need to implement the Log policy.

How should you complete the EnsureLogging method in EventGridController.cs? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
var client = new WebsiteManagementClient(. . .);
var id = ParseResourceId(resource);
var appSettings = new StringDictionary(name: "properties",
    properties: new Dictionary<string, string> {
        {"DIAGNOSTICS_AZUREBLOBCONTAINERSASURL", BlobStoreAccountsSAS(" ")},

        {"DIAGNOSTICS_AZUREBLOBRETENTIONDAYS", " "}});

client.WebApps. (
    id.resourceGroup,
    id.name, appSettings);
```

Dropdown 1: logs, logdrop

Dropdown 2: 15, 30

Dropdown 3: UploadLoggingSettings, UpdateApplicationSettings

Correct Answer:

Answer Area

```
var client = new WebsiteManagementClient(. . .);
var id = ParseResourceId(resource);
var appSettings = new StringDictionary(name: "properties",
    properties: new Dictionary<string, string> {
        {"DIAGNOSTICS_AZUREBLOBCONTAINERSASURL", BlobStoreAccountsSAS(" ")},

        {"DIAGNOSTICS_AZUREBLOBRETENTIONDAYS", " "}});

client.WebApps. (
    id.resourceGroup,
    id.name, appSettings);
```

Dropdown 1: logs, logdrop (logdrop is highlighted)

Dropdown 2: 15, 30 (15 is highlighted)

Dropdown 3: UploadLoggingSettings, UpdateApplicationSettings (UpdateApplicationSettings is highlighted)

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: logdrop

All log files should be saved to a container named logdrop.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Box 2: 15

Logs must remain in the container for 15 days.

Box 3: UpdateApplicationSettings

All Azure App Service Web Apps must write logs to Azure Blob storage.

References:

<https://blog.hompus.nl/2017/05/29/adding-application-logging-blob-to-a-azure-web-app-service-using-powershell/>

QUESTION 2

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 need to ensure that authentication events are triggered and processed according to the authentication events policy.

Solution: Create a new Azure Event Grid topic and add a subscription for the events.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead ensure that signout events have a subject prefix. Create an Azure Event Grid subscription that uses the subjectBeginsWith filter.

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

QUESTION 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 need to ensure that authentication events are triggered and processed according to the authentication events policy.

Solution: Create separate Azure Event Grid topics and subscriptions for sign-in and sign-out events.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation/Reference:

Explanation:

Instead ensure that signout events have a subject prefix. Create an Azure Event Grid subscription that uses the subjectBeginsWith filter.

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

QUESTION 4

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 need to ensure that authentication events are triggered and processed according to the authentication events policy.

Solution: Ensure that signout events have a subject prefix. Create an Azure Event Grid subscription that uses the subjectBeginsWith filter.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

References:

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

QUESTION 5

DRAW DROP

You need to add code at line EG15 in EventGridController.cs to ensure that the Log policy applies to all services.

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.

Select and Place:

Code segments	Answer Area
topic	if (
status	@event["data"][" "] .ToString() == " " "
eventType	&&
Succeeded	@event["data"][" "] .ToString() == "Microsoft.Web/sites/write"
operationName	)
resourceProvider	

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Correct Answer:

Code segments

topic
status
eventType
Succeeded
operationName
resourceProvider

Answer Area

```
if (  
    @event[ "data" ][ " status "].ToString() == " Succeeded "  
    &&  
    @event[ "data" ][ " operationName "].ToString() == "Microsoft.Web/sites/write"  
)
```

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario, Log policy: All Azure App Service Web Apps must write logs to Azure Blob storage.

Box 1: Status

Box 2: Succeeded

Box 3: operationName

Microsoft.Web/sites/write is resource provider operation. It creates a new Web App or updates an existing one.

References:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations>

QUESTION 6

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

`public string` `{ get; set; }`

id
eventType
dataVersion
metadataVersion

`public string` `{ get; set; }`

id
eventType
dataVersion
metadataVersion

`public string` `{ get; set; }`

id
eventType
dataVersion
metadataVersion

Correct Answer:

Answer Area

`public string` `{ get; set; }`

id
eventType
dataVersion
metadataVersion

`public string` `{ get; set; }`

id
eventType
dataVersion
metadataVersion

`public string` `{ get; set; }`

id
eventType
dataVersion
metadataVersion

Section: (none)**Explanation****Explanation/Reference:**

Explanation:

Box 1: id

id is a unique identifier for the event.

Box 2: eventType

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

Box 3: dataVersion

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

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

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

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

References:

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

QUESTION 7

You need to ensure that authentication events are triggered and processed according to the authentication events policy.

What should you do?

- A. Ensure that signout events have a subject prefix. Create an Azure Event Grid subscription that uses the subjectBeginsWith filter.
- B. Create a new Azure Event Grid topic and add a subscription for the events.
- C. Create a new Azure Event Grid subscription for all authentication that delivers messages to an Azure Event Hub. Use the subscription to process signout events.
- D. Create separate Azure Event Grid topics and subscriptions for sign-in and sign-out events.

Correct Answer: A

Section: (none)**Explanation****Explanation/Reference:**

Explanation:

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

References:

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Connect to and Consume Azure Services and Third-party Services

Testlet 2

Case Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

Wide World Importers is moving all their datacenters to Azure. The company has developed several applications and services to support supply chain operations and would like to leverage serverless computing where possible.

Current environment

Windows Server 2016 virtual machine

This virtual machine (VM) runs Biz Talk Server 2016. The VM runs the following workflows:

- Ocean Transport – This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.
- Inland Transport – This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

- Container API – This API provides container information including weight, contents, and other attributes.
- Location API – This API provides location information regarding shipping ports of call and truck stops.
- Shipping REST API – This API provides shipping information for use and display on the shipping website.

Shipping Data

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

Shipping Web Site


The site displays shipping container tracking information and container contents. The site is located at <http://shipping.wideworldimporters.com>

Proposed solution

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:



Shipping Logic App

The Shipping Logic app must meet the following requirements:

- Support the ocean transport and inland transport workflows by using a Logic App.
- Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
- Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
- Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.

Shipping Function app

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

REST APIs

The REST API's that support the solution must meet the following requirements:

- Secure resources to the corporate VNet.
- 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.

Shipping data

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

Shipping website

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

Issues

Windows Server 2016 VM

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

Shipping website and REST APIs

The following error message displays while you are testing the website:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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.

QUESTION 1

DRAG DROP

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?

Select and Place:

Actions	Answer Area
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.	

Correct Answer:

Actions	Answer Area
	Create an integration account in the Azure portal.
Link the custom connector to the Logic App.	Link the Logic App to the integration account.
Update the Logic App to use the partners, schemas, certificates, maps, and agreements.	Add partners, schemas, certificates, maps, and agreements.
	Create a custom connector for the Logic App.

Section: (none)
Explanation

Explanation/Reference:

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>

Connect to and Consume Azure Services and Third-party Services

Question Set 3

QUESTION 1

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 an IndexBatch that contains the documents which must be added.
3. Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

2. Create the indexBatch with the documents

Something like:

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

3. The next step is to populate the newly-created index

Example:

```
var batch = IndexBatch.Upload(hotels);

try
{
    indexClient.Documents.Index(batch);
}
```

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

References:

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

QUESTION 2

DRAG DROP

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.

Select and Place:

Actions

Create a Service Bus subscription for each restaurant for which a driver can receive orders.

Create a single Service Bus topic.

Create a single Service Bus subscription.

Create a Service Bus topic for each restaurant for which a driver can receive messages

Create a single Service Bus Namespace.

Create a Service Bus Namespace for each restaurant for which a driver can receive messages.

Answer area



Correct Answer:

Actions

Create a Service Bus subscription for each restaurant for which a driver can receive orders.

Create a single Service Bus topic.

Create a single Service Bus subscription.

Create a Service Bus topic for each restaurant for which a driver can receive messages

Create a single Service Bus Namespace.

Create a Service Bus Namespace for each restaurant for which a driver can receive messages.

Answer area

Create a single Service Bus Namespace.

Create a Service Bus topic 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.



Section: (none)

Explanation

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation/Reference:

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>

QUESTION 3**HOTSPOT**

A company runs an international travel and bookings management service. The company plans to begin offering restaurant bookings.

You must develop a solution that uses Azure Search and meets the following requirements:

- Users must be able to search for restaurants by name, description, location, and cuisine.
- Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.
- All words in restaurant descriptions must be included in searches.

You need to add annotations to the restaurant class.

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.

Hot Area:

Answer Area

```
[SerializePropertyNameAsCamelCase]
public class Restaurant
{
    [Key, IsFilterable]
    public int RestaurantId { get; set; }
    [IsSearchable, IsFilterable, IsSortable]
    public string Name { get; set; }

    public string location { get; set; }
    public string Phone { get; set; }

    public string Description { get; set; }

    public double Rating { get; set; }

    public List<string> Cuisines { get; set; }

    public bool FamilyFriendly { get; set; }
```

Correct Answer:

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Answer Area

```
[SerializePropertyNameAsCamelCase]
public class Restaurant
{
    [Key, IsFilterable]
    public int RestaurantId { get; set; }
    [IsSearchable, IsFilterable, IsSortable]
    public string Name { get; set; }

    public string location { get; set; }
    public string Phone { get; set; }

    public string Description { get; set; }

    public double Rating { get; set; }

    public List<string> Cuisines { get; set; }

    public bool FamilyFriendly { get; set; }
```

Section: (none)
Explanation

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation/Reference:

Explanation:

Box 1: [IsSearchable.IsFilterable.IsSortable,IsFacetable]

Location

Users must be able to search for restaurants by name, description, location, and cuisine.

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

Box 2: [IsSearchable.IsFilterable.IsSortable,Required]

Description

Users must be able to search for restaurants by name, description, location, and cuisine.

All words in descriptions must be included in searches.

Box 3: [IsFilterable,IsSortable,IsFaceTable]

Rating

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

Box 4: [IsSearchable.IsFilterable,IsFacetable]

Cuisines

Users must be able to search for restaurants by name, description, location, and cuisine.

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

Box 5: [IsFilterable,IsFacetable]

FamilyFriendly

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

References:

<https://www.henkboelman.com/azure-search-the-basics/>

QUESTION 4

DRAG DROP

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

The news API back end is implemented as a RESTful service and hosted in an Azure App Service instance.

You need to configure back-end authentication for the API Management service instance.

Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area**Values**

Azure Resource

HTTP(s) endpoint

Basic

Client cert

Configuration parameter

Target

Gateway credentials

Value

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Correct Answer:

Answer Area

Values	Configuration parameter	Value
Azure Resource	Target	Azure Resource
HTTP(s) endpoint	Gateway credentials	Client cert
Basic		
Client cert		

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Azure Resource

Box 2: Client cert

API Management allows to secure access to the back-end service of an API using client certificates.

References:

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

QUESTION 5

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 ensure that the subscription client processes all messages.

Which code segment should you use?

- 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);`

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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/>

QUESTION 6

DRAG DROP

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.

Select and Place:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)
Explanation

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Explanation/Reference:

Explanation:

QUESTION 7**HOTSPOT**

You are developing a .NET Core MVC application for customers to research hotels. The application will use Azure Search. The application will search the index by using various criteria to locate documents related to hotels. The index will include search fields for rate, a list of amenities, and distance to the nearest airport.

The application must support the following scenarios for specifying search criteria and organizing results:

- Search the index by using regular expressions.
- Organize results by counts for name-value pairs.
- List hotels within a specified distance to an airport and that fall within a specific price range.

You need to configure the SearchParameters class.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Scenario	Property
Search the index by using regular expressions.	QueryType OrderBy SearchMode
Organize results by counts for name-value pairs.	Facets Filter SearchMode
List hotels within a specified distance to an airport and that fall within a specific price range.	Order by Top Filter

Correct Answer:

Answer Area

Scenario	Property
Search the index by using regular expressions.	▼ QueryType OrderBy SearchMode
Organize results by counts for name-value pairs.	▼ Facets Filter SearchMode
List hotels within a specified distance to an airport and that fall within a specific price range.	▼ Order by Top Filter

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: QueryType

The SearchParameters.QueryType Property gets or sets a value that specifies the syntax of the search query. The default is 'simple'. Use 'full' if your query uses the Lucene query syntax.

You can write queries against Azure Search based on the rich Lucene Query Parser syntax for specialized query forms: wildcard, fuzzy search, proximity search, regular expressions are a few examples.

Box 2: Facets

The facets property gets or sets the list of facet expressions to apply to the search query. Each facet expression contains a field name, optionally followed by a comma-separated list of name:value pairs.

Box 3: Filter

The Filter property gets or sets the OData \$filter expression to apply to the search query.

References:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters>

<https://docs.microsoft.com/en-us/azure/search/query-lucene-syntax>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.querytype>

QUESTION 8

HOTSPOT

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

▼ hub=

NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

▼ . ▼

NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

GetInstallation
CreateClientFromConnectionString
CreateOrUpdateInstallation
PatchInstallation

```
(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.▼ (windowsToastPayload);
    SendWindowsNativeNotificationAsync
    SubmitNotificationHubJobAsync
    ScheduleNotificationAsync
    SendAppleNativeNotificationAsync
    . . .
}
catch (System.Exception ex)
{
    . . .
}
. . .
```

Correct Answer:

Answer Area

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

NotificationHubClient hub =
    NotificationHubClientSettings
        NotificationHubJob
        NotificationDetails

NotificationHubClient
    NotificationHubClientSettings
    NotificationHubJob
    NotificationDetails

    .
    GetInstallation
    CreateClientFromConnectionString
    CreateOrUpdateInstallation
    PatchInstallation

(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.
            SendWindowsNativeNotificationAsync
            SubmitNotificationHubJobAsync
            ScheduleNotificationAsync
            SendAppleNativeNotificationAsync
            (windowsToastPayload);

    . . .
}
catch (System.Exception ex)
{
    . . .
}
. . .
```

Section: (none)

Explanation

Explanation/Reference:

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-management>

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/app-service-mobile/app-service-mobile->

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

QUESTION 9
HOTSPOT

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 a specific period of time.

You need to configure Event Grid.

What should you implement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Option	Value
WebHook event delivery	SAS tokens Key authentication Management Access Control
Topic publishing	ValidationCode handshake ValidationURL handshake JWT token

Correct Answer:

Answer Area

Option	Value
WebHook event delivery	SAS tokens Key authentication Management Access Control
Topic publishing	ValidationCode handshake ValidationURL handshake JWT token

Section: (none)

Explanation

Explanation/Reference:

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.

Incorrect Answers:

ValidationURL handshake (manual): In certain cases, you can't access the source code of the endpoint to implement the ValidationCode handshake. For example, if you use a third-party service (like Zapier or IFTTT), you can't programmatically respond with the validation code.

References:

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

QUESTION 10

HOTSPOT

You are validating the configuration of an Azure Search indexer.

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

Index Configuration

Index name: **azuretable-index**

Key: **Rowkey**

FIELD NAME	TYPE	RETRIEVABLE	FILTERABLE	SORTABLE	FACETABLE	SEARCHABLE
PartitionKey	Edm.String	☒	☒	☒	☐	☒
RowKey	Edm.String	☒	☐	☐	☐	☐
InStockCount	Edm.Int32	☒	☐	☐	☐	☐
ItemDescription	Edm.String	☒	☐	☐	☐	☐
ItemName	Edm.String	☒	☒	☒	☐	☒
LocationRow	Edm.Int32	☒	☐	☐	☐	☐
LocationShelf	Edm.Int32	☒	☐	☐	☐	☐
SKU	Edm.String	☒	☐	☐	☐	☒

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

PartitionKey	RowKey	Timestamp	InStockCount	ItemDescription	ItemName	LocationRow	LocationShelf	SKU
Food	3	2018-08-25T 15:47:29.135Z	32	A box of chocolate candy bars	Choco-bar	5	3	123421
Hardware	2	2018-08-25T 15:46:08.409Z	2	A bag of bolts	Bolts	1	4	678564
Hardware	1	2018-08-25T 15:46:41.402Z	23	A box of nails	Nails	2	1	654365

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

Search explorer

Query string: **search=bag** [Search]

Index: **azuretable-index**
API version: **2017-11-11**

Request URL: <https://itemsearch1103search.windows.net/indexes/azuretable-index/docs?api-version=2017-11-11&search=bag>

Results

```

1 {
2   "@odata.context": "https://itemsearch1103search.windows.net/indexes('azuretable-index')/$metadata#docs",
3   "value": []
4 }

```

You need to resolve the issue.

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

NOTE: Each correct selection is worth one point.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Box 1: Yes
The ItemDescription field is not searchable.

Box 2: No
The ItemDescription field is not searchable, but we would need to recreate the index.

Box 3: Yes
An indexer in Azure Search is a crawler that extracts searchable data and metadata from an external Azure data source and populates an index based on field-to-field mappings between the index and your data source. This approach is sometimes referred to as a 'pull model' because the service pulls data in without you having to write any code that adds data to an index.

Box 4: No

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

QUESTION 11
HOTSPOT

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

You are developing an Azure App Service hosted 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.

Hot Area:

Answer Area

Code segment

Value

Attribute

	▼
Authorize	
AllowAnonymous	
AutoValidateAntiforgeryToken	

Request Header

	▼
X-MS-CLIENT-PRINCIPAL-NAME	
Proxy-Authorization	
X-Forwarded-For	
X-MS-CLIENT-PRINCIPAL-ID	

Correct Answer:

Answer Area

Code segment

Value

Attribute

	▼
Authorize	
AllowAnonymous	
AutoValidateAntiforgeryToken	

Request Header

	▼
X-MS-CLIENT-PRINCIPAL-NAME	
Proxy-Authorization	
X-Forwarded-For	
X-MS-CLIENT-PRINCIPAL-ID	

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 12

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.

<https://vceoreteconvert.com/>

<https://t.me/learningnets>

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 13

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 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 DataSources property of the SearchServiceClient.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

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>

QUESTION 14

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 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. Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

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>