

AZ-203

Number: AZ-203
Passing Score: 800
Time Limit: 120 min
File Version: 1

AZ-203



Website: <https://vceplus.com>
VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>
Facebook: <https://www.facebook.com/VCE.For.All.VN/>
Twitter : https://twitter.com/VCE_Plus

<https://vceplus.com/>

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:

- 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
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 = JObject.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

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**.



<https://vceplus.com/>

- 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:

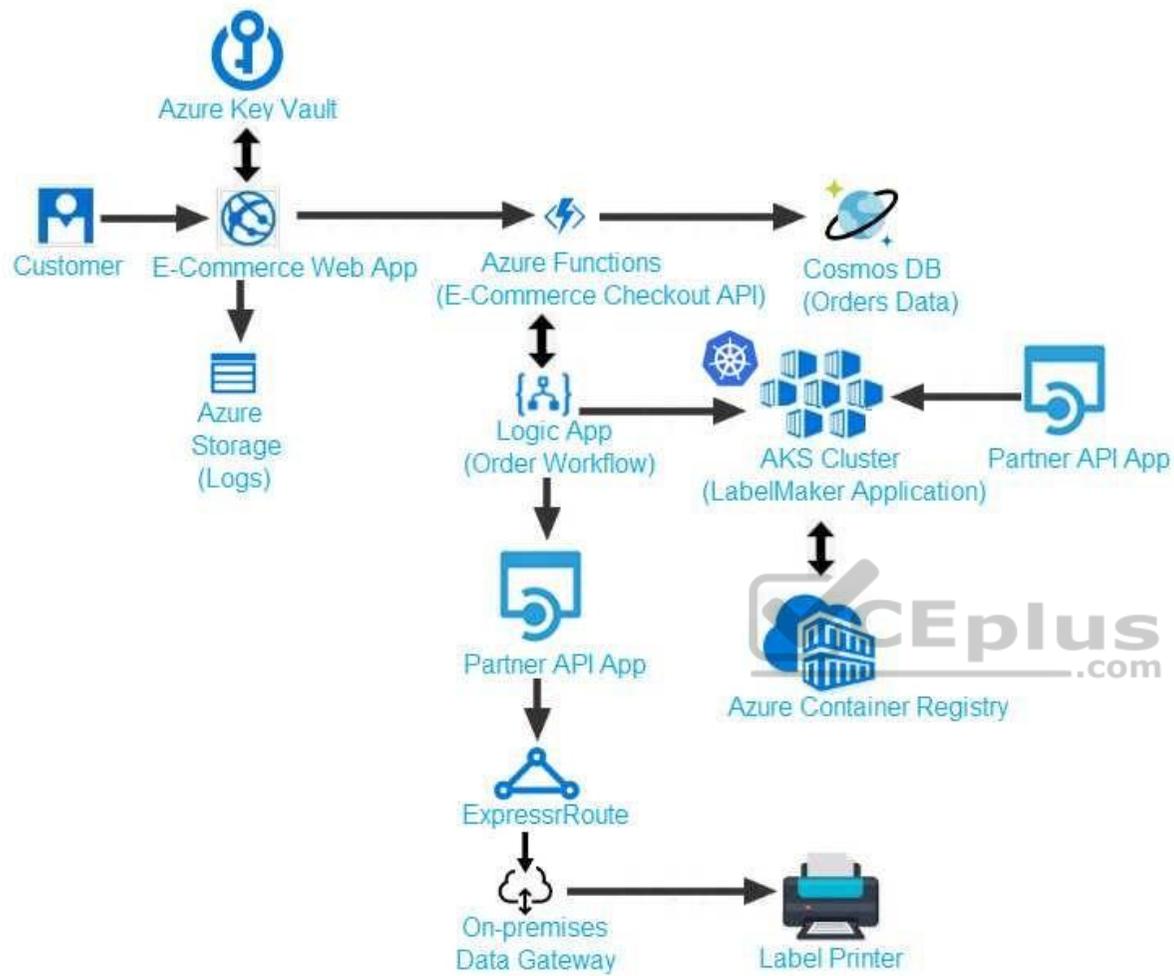
- Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
 - External partners must use their own credentials and authenticate with their organization's identity management solution.
 - External partner logins must be audited monthly 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
```



<https://vceplus.com/>

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

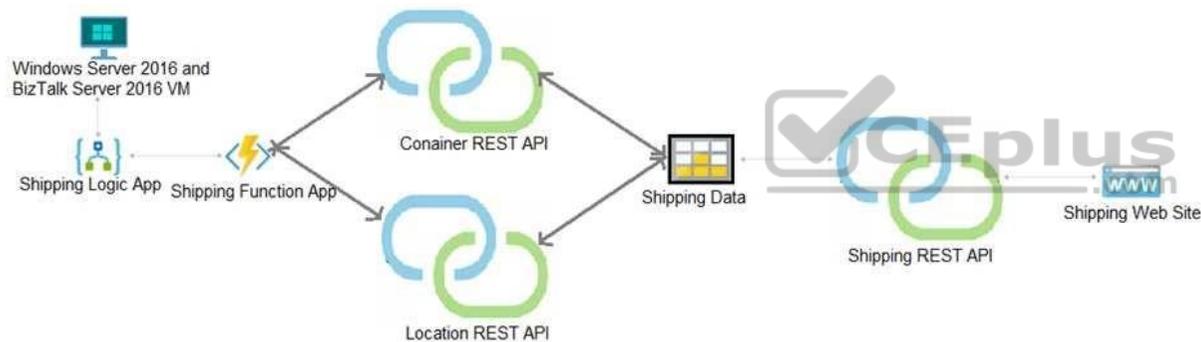
The application uses MongoDB JSON document storage database for all container and transport information.

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 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:

```
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>

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>, CancellationTokens)`
- B. `JobOperations.CreateJob()`
- C. `CloudJob.Enable(IEnumerable<BatchClientBehavior>)`
- D. `JobOperations.EnableJob(String, IEnumerable<BatchClientBehavior>)`

E. CloudJob.CommitAsync(IEnumerable<BatchClientBehavior>, CancellationToken)

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

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

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

What should you do?

- A. In the Azure portal, add a Job to a Batch account.
- B. In a .NET method, call the method: BatchClient.PoolOperations.CreateJob
- C. In Python, implement the class: JobAddParameter
- D. In Azure CLI, run the command: az batch pool create
- E. In a .NET method, call the method: BatchClient.PoolOperations.CreatePool
- F. In Python, implement the class: TaskAddParameter
- G. In the Azure CLI, run the command: az batch account create

Correct Answer: B

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.

Note:

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

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

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

Incorrect Answers:

C, F: To create a Batch pool in Python, the app uses the PoolAddParameter class to set the number of nodes, VM size, and a pool configuration.

E: BatchClient.PoolOperations does not have a CreateJob method.

References:

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

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

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 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 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 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:

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>

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

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     }
PC38     private KeyVaultClient _keyVaultClient = new KeyVaultClient(" . . .");
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
```

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:

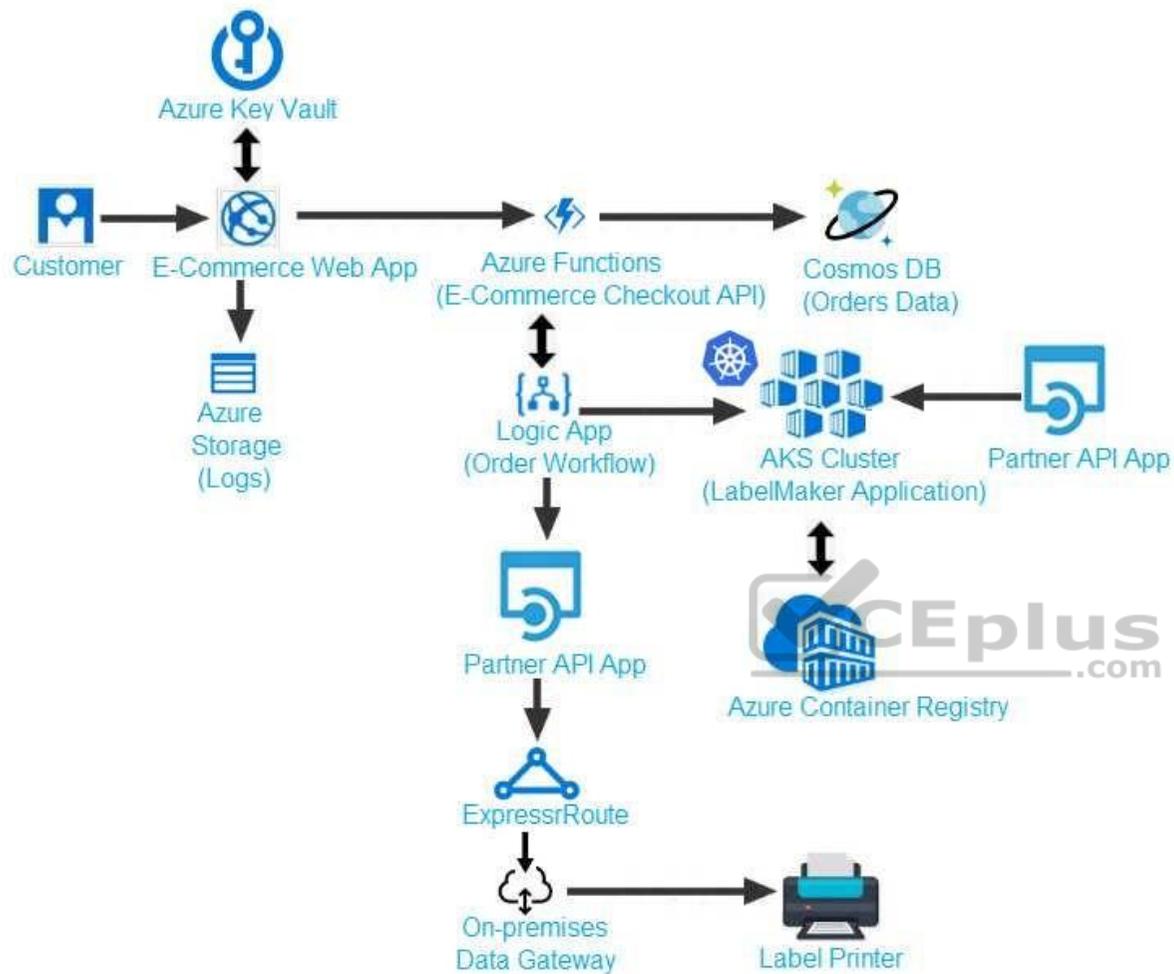
- Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
- External partners must use their own credentials and authenticate with their organization's identity management solution.
- External partner logins must be audited monthly 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
```



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

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)

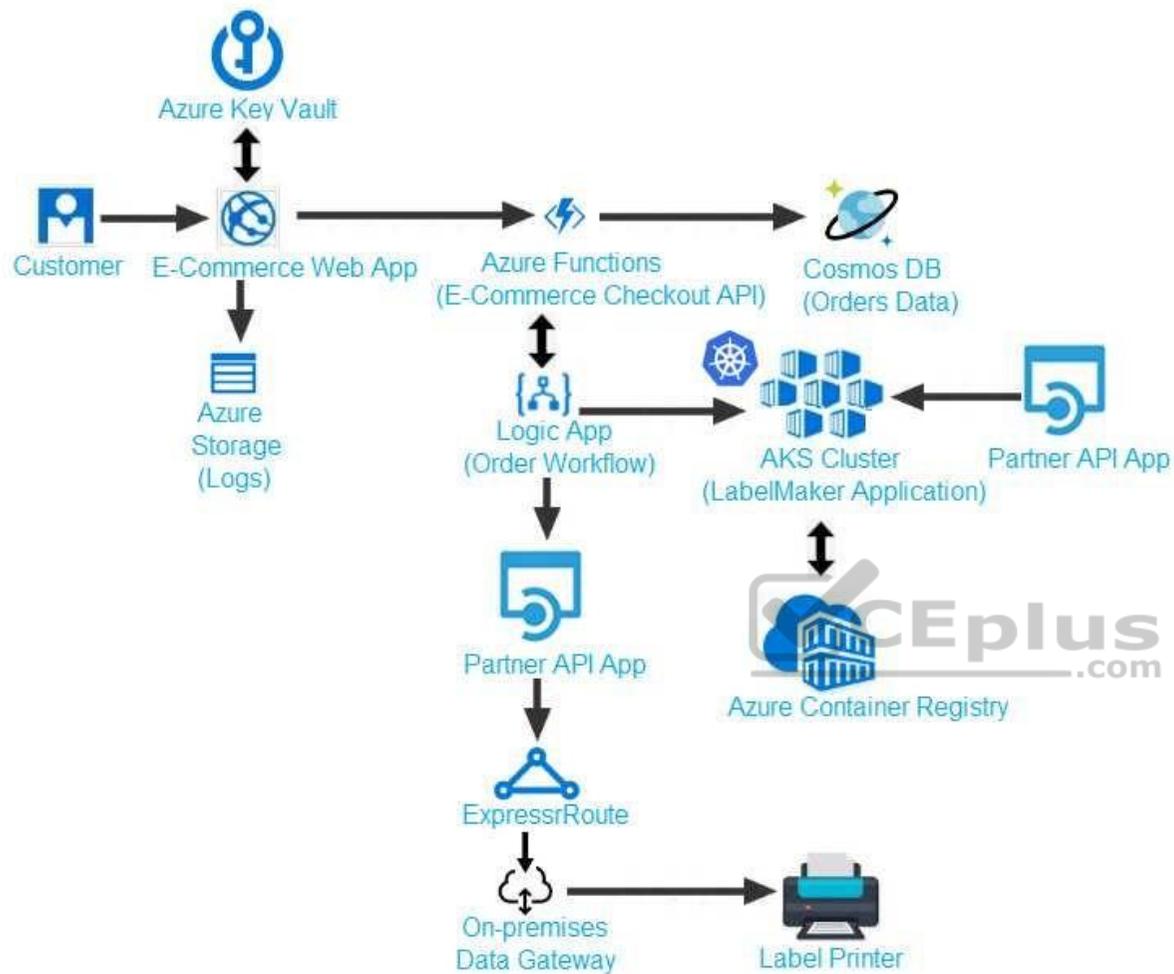
Explanation

Explanation/Reference:

Explanation:

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





Develop Azure Platform as a Service Compute Solutions

Question Set 3

QUESTION 1

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 2

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

The mobile app must meet the following requirements:

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

You need to implement Offline Data Sync.

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

NOTE: Each correct selection is worth one point.

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

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.

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 3

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.

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 4

A company is developing a solution that allows smart refrigerators to send temperature information to a central location.

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?

```
New-AzureRmServiceBusNamespace
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Location fridge-loc
az servicebus namespace create
  - -resource-group fridge-rg
  - -name fridge-ns
  - -location fridge-loc
New-AzureRmResourceGroup
  -Name fridge-rg
  -Location fridge-loc
New-AzureRmServiceBusQueue
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Name fridge-q
  -EnablePartitioning $False
```

A.

B.

C.



D.

```
connectionString=$(az servicebus namespace authorization-rule keys list
  --resource-group fridge-rg
  --fridge-ns fridge-ns
  --name RootManageSharedAccessKey
  --query primaryConnectionString --output tsv)
```

E.

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 5

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

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

- 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 6

You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published.

You must ensure that the website remains available and responsive while minimizing cost.

You need to deploy the website.



What should you do?

- A. Deploy the website to 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.

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

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     }
PC38     private KeyVaultClient _keyVaultClient = new KeyVaultClient(" . . .");
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

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 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>

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.

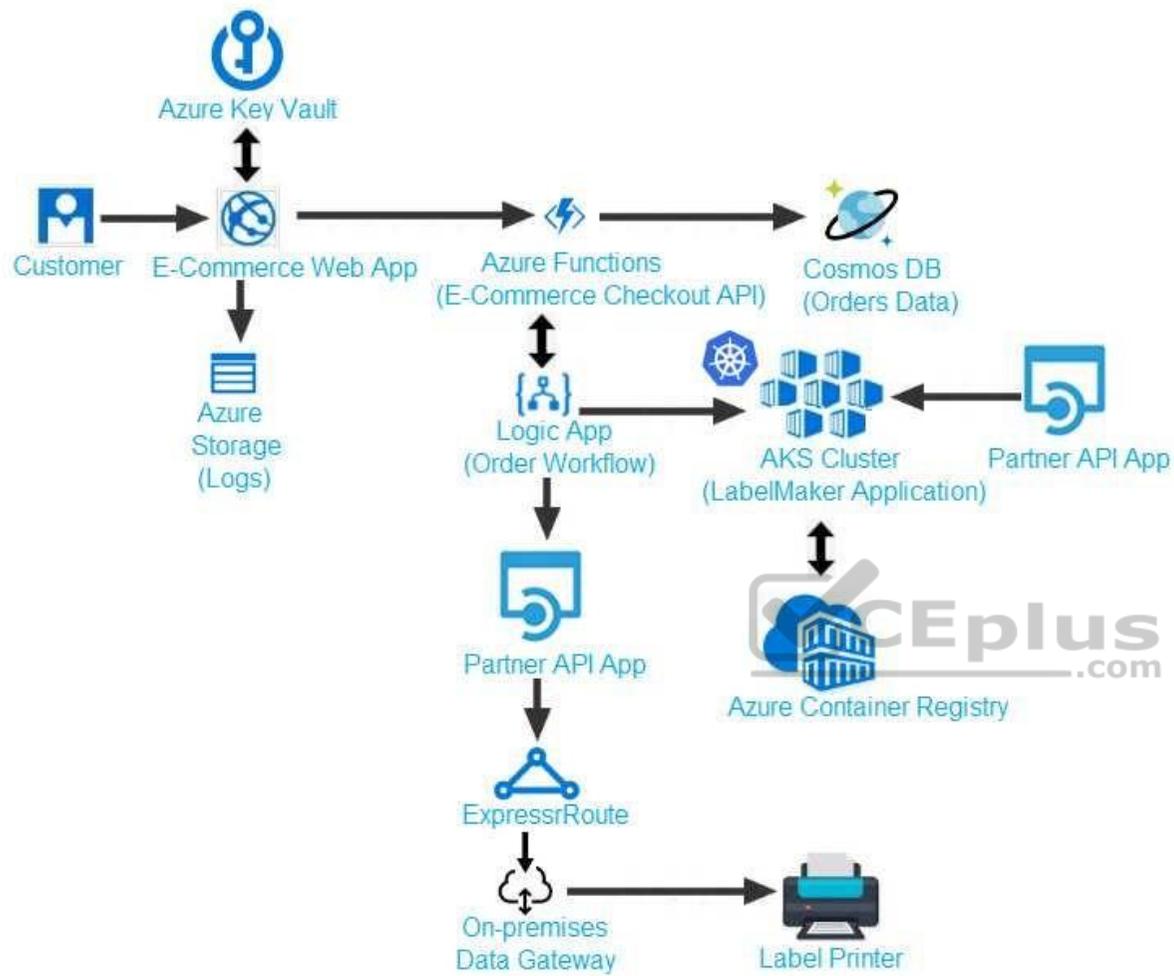
LabelMaker app

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
```



<https://vceplus.com/>

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

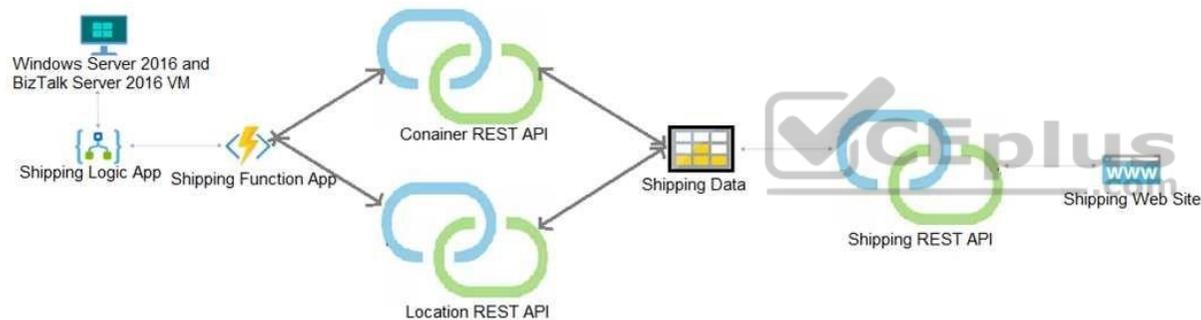
The application uses MongoDB JSON document storage database for all container and transport information.

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 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:

```
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 use Azure Table storage to store customer information for an application. The data contains customer details and is partitioned by last name.

You need to create a query that returns all customers with the last name Smith.

Which code segment should you use?

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

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 2

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 3

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	ADD MASK
SalesLT	Customer	LastName	ADD MASK
SalesLT	Customer	EmailAddress	ADD MASK
SalesLT	Customer	Phone	ADD MASK
SalesLT	CustomerAddress	AddressID	ADD MASK

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

QUESTION 4

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 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/customers  
/transparentDataEncryption/current?api-version=2014-04-01
```

- E. Run the following Azure PowerShell command:

```
New-AzureRmSqlDatabaseDataMaskingRule -SchemaName "dbo" -TableName "customers" ` -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 5

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 6

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

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

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

You need to store the information.

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

Does the solution meet the goal?

- A. Yes
- B. No

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 7

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-provider>

QUESTION 8

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

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

You are developing 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 9

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



Explanation/Reference:

References:

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

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 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 11

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)

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 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. 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 13

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

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

You need to select a partition key.

Which two partition keys can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

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 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. 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:

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:

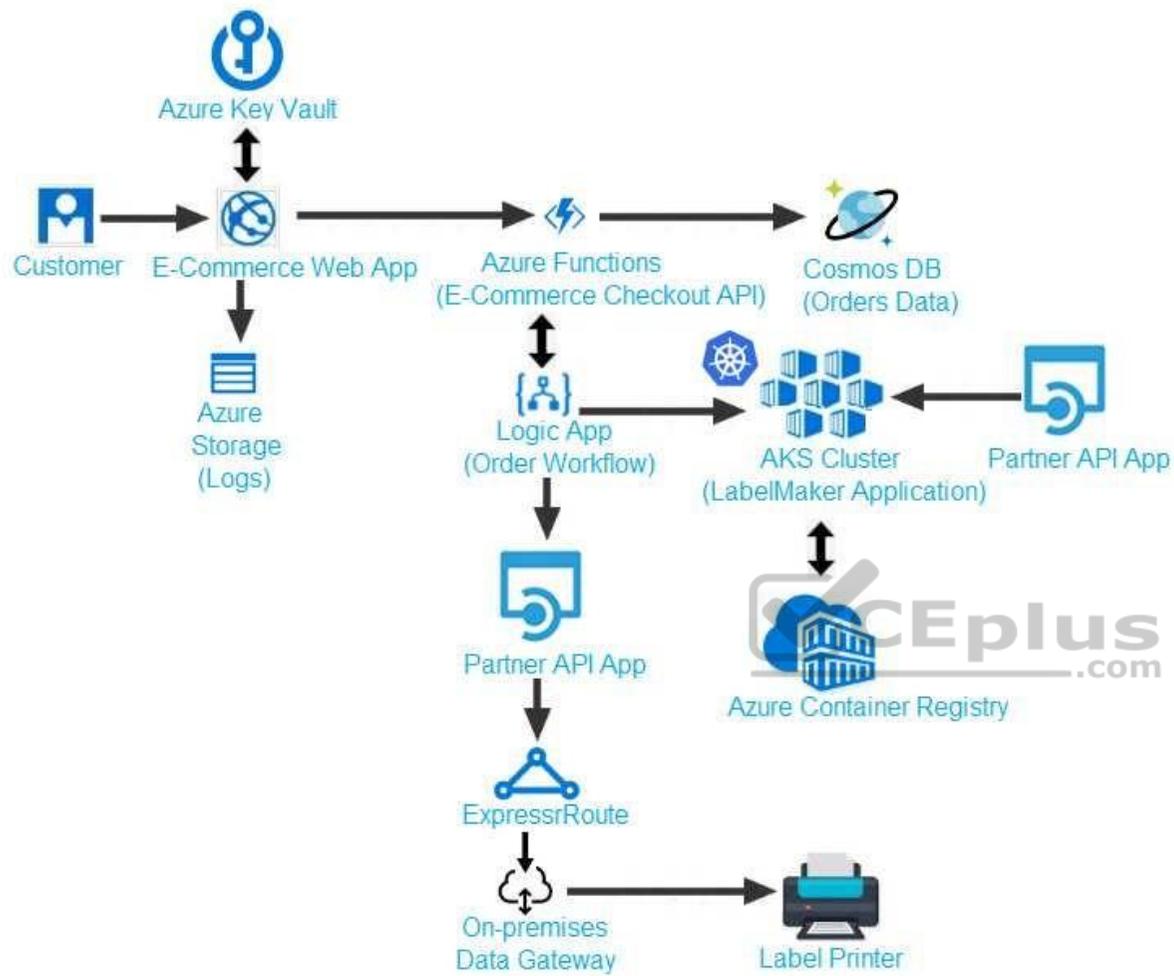
- Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
- External partners must use their own credentials and authenticate with their organization's identity management solution.
- External partner logins must be audited monthly 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
```



QUESTION 1

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.



- 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 2

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 3

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.

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/>

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

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     }
PC38     private KeyVaultClient _keyVaultClient = new KeyVaultClient(" . . .");
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

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 WebAppIdentity 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.

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: 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.

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

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.
- 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 = JObject.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

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

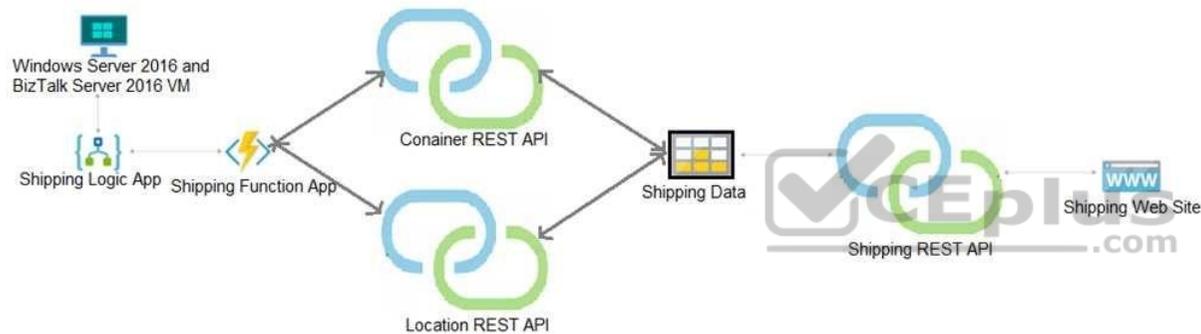
The application uses MongoDB JSON document storage database for all container and transport information.

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 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:

```
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 secure the Shipping Logic App.

What should you use?

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

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>



Implement Azure security

Question Set 5

QUESTION 1

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 2

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.

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-ps-](https://docs.microsoft.com/en-us/powershell/module/azurerm.resources/new-azurermroleassignment?view=azurerm-ps-6.13.0)

[6.13.0](https://docs.microsoft.com/en-us/powershell/module/azurerm.resources/new-azurermroleassignment?view=azurerm-ps-6.13.0)

QUESTION 3

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

Every request to the backend service must include a valid HTTP authorization header.

You need to configure the Azure API Management instance with an authentication policy.

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

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 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 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.

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: 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/>



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.
- 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 = JObject.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             SendToAnomalyDetection Service(@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                         { "logcontent", content
EG40                     }
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

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

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:

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>

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

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     }
PC38     private KeyVaultClient _keyVaultClient = new KeyVaultClient(" . . .");
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

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     }
```

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.

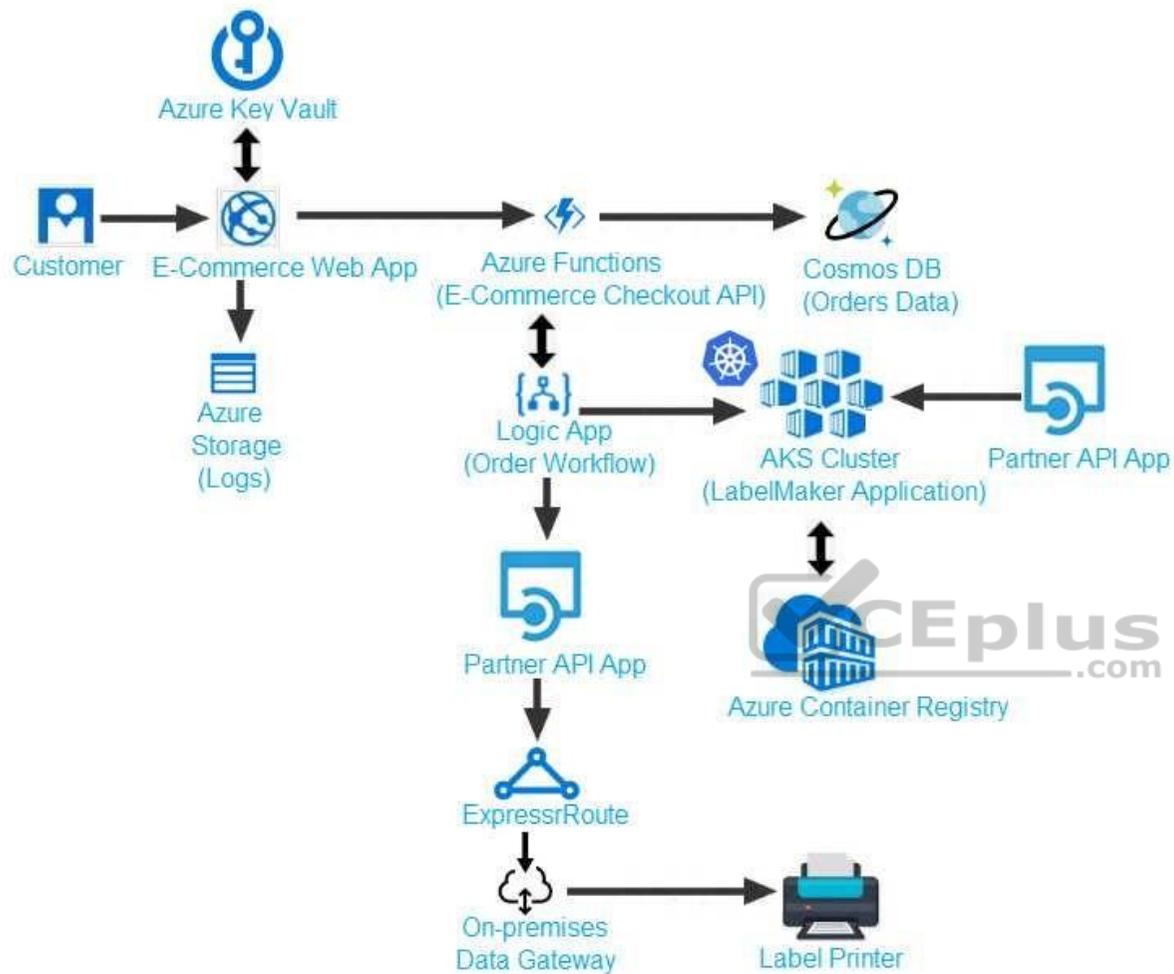
LabelMaker app

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
```



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

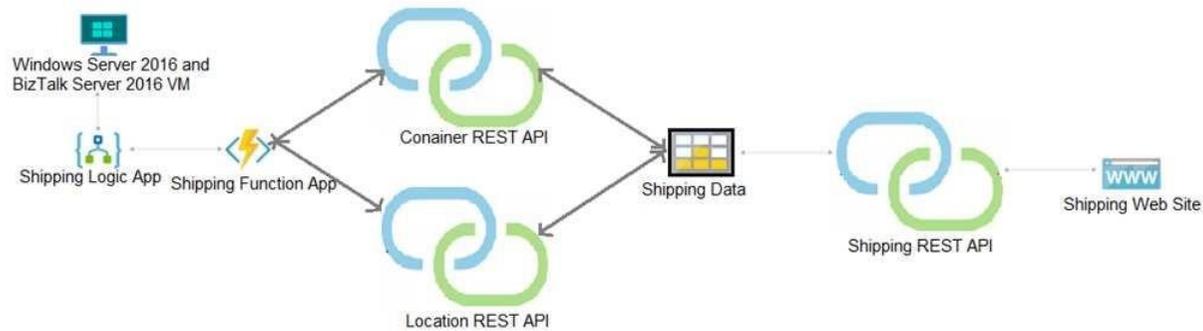
The application uses MongoDB JSON document storage database for all container and transport information.

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 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:

```
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.
```

Monitor, troubleshoot, and optimize Azure solutions

Question Set 5

QUESTION 1

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

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>

QUESTION 2

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 3

You are creating a hazard notification system that has a single signaling server which triggers audio and visual alarms to start and stop.

You implement Azure Service Bus to publish alarms. Each alarm controller uses Azure Service Bus to receive alarm signals as part of a transaction. Alarm events must be recorded for audit purposes. Each transaction record must include information about the alarm type that was activated.

You need to implement a reply trail auditing solution. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Assign the value of the hazard message 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.

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 4

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.

References: <https://docs.microsoft.com/en-us/powershell/module/azurerm.apimanagement/new-azurermapi-management-backend-proxy?view=azurerm-ps-6.13.0>

QUESTION 5

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>

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.
- 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 = JObject.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             SendToAnomalyDetection Service(@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                         { "logcontent", content
EG40                     }
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

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 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 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



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: 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>



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

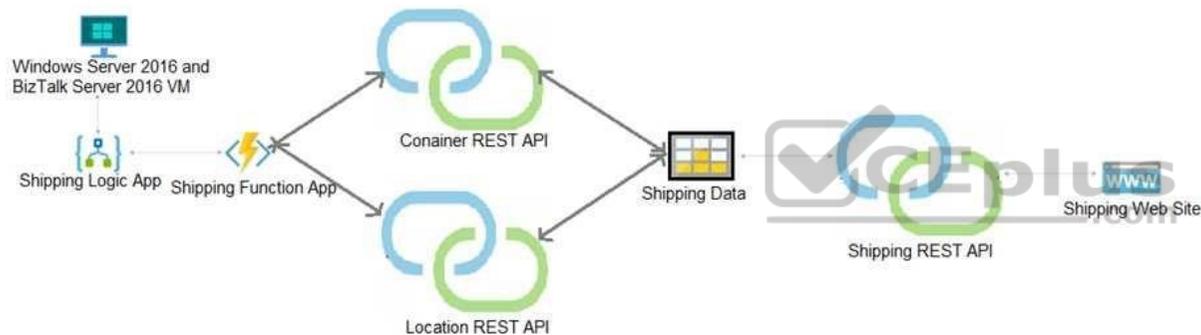
The application uses MongoDB JSON document storage database for all container and transport information.

Shipping Web Site

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

Proposed solution

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:

```
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.
```

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 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>

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.

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 documentsSomething like:

```
var hotels = new Hotel[];
{
new Hotel()
    {
        HotelId = "3",

        BaseRate = 129.99,
        Description = "Close to town hall and the river"
    }
};
...
var batch = IndexBatch.Upload(hotels);
```

3. The next step is to populate the newly-created index

Example:

```
var batch = IndexBatch.Upload(hotels);
```

```
try
{
```

```
indexClient.Documents.Index(batch);  
}
```

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

QUESTION 3

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

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

You need to complete the source code of the subscription client

What should you do?

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

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

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

You create the index in Azure Search.

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

Solution:

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

Does the solution meet the goal?

- A. Yes
- B. No

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 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.

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>



<https://vceplus.com/>

