

AZ-204

Number: AZ-204
Passing Score: 800
Time Limit: 120 min
File Version: 1

AZ-204

Developing Solutions for Microsoft Azure



Website: <https://vceplus.com> - <https://vceplus.co>
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 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 questions in 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 section 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 in 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. When you are ready to answer a question, click the **Question** button to return to the question.

Current environment

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk 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 trucking 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:



Requirements

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://test.wideworldimporters.com/' is therefore not allowed access.

QUESTION 1

HOTSPOT

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



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

NOTE: Each correct selection is worth one point.



Hot Area:

Answer Area

Option	Value
Tier	<input type="button" value="▼"/> Standard Premium
Profile	<input type="button" value="▼"/> Akamai Microsoft
Optimization	<input type="button" value="▼"/> general web delivery large file download dynamic site acceleration video-on-demand media streaming

Correct Answer:

Answer Area

Option	Value
Tier	 Standard Premium
Profile	 Akamai Microsoft
Optimization	 general web delivery large file download dynamic site acceleration video-on-demand media streaming

Section: (none)

Explanation:

Explanation/Reference:

Explanation:

Scenario: Shipping website

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

Tier: Standard

Profile: Akamai

Optimization: Dynamic site acceleration

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

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

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

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

QUESTION 2

HOTSPOT

You need to correct the VM issues.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Issue	Tool
Backup and Restore	<div style="border: 1px solid black; padding: 5px;"> Azure Site Recovery Azure Backup Azure Data Box Azure Migrate</div>
Performance	<div style="border: 1px solid black; padding: 5px;"> Azure Network Watcher Azure Traffic Manager ExpressRoute Accelerated Networking</div>

Correct Answer:

Answer Area

Issue	Tool
Backup and Restore	Azure Site Recovery Azure Backup Azure Data Box Azure Migrate
Performance	Azure Network Watcher Azure Traffic Manager ExpressRoute Accelerated Networking

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Azure Backup

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

In-Place restore of disks in IaaS VMs is a feature of Azure Backup.

Performance: Accelerated Networking

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

Box 2: Accelerated networking

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

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

Reference:

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



Develop Azure 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 questions in 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 section 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 in 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. When you are ready to answer a question, click the **Question** button to return to the question.

Background

City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.

Current environment

Architecture overview


The company has a public website located at <http://www.cpndl.com/>. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration. Logic apps are used to orchestrate the data processing while Service Bus and Event Grid handle messaging and events.

The solution uses Application Insights, Azure Monitor, and Azure Key Vault.

Architecture diagram

The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.



User authentication

The following steps detail the user authentication process:

1. The user selects **Sign in** in the website.
2. The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
3. The user signs in.
4. Azure AD redirects the user's session back to the web application. The URL includes an access token.
5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
6. The back-end API validates the access token.

Requirements

Corporate website

- Communications and content must be secured by using SSL.
- Communications must use HTTPS.
- Data must be replicated to a secondary region and three availability zones.
- Data storage costs must be minimized.

Azure Database for PostgreSQL

The database connection string is stored in Azure Key Vault with the following attributes:

- Azure Key Vault name: cpndlkeyvault
- Secret name: PostgreSQLConn
- Id: 80df3e46ffcd4f1cb187f79905e9a1e8

The connection information is updated frequently. The application must always use the latest information to connect to the database.

Azure Service Bus and Azure Event Grid

- Azure Event Grid must use Azure Service Bus for queue-based load leveling.
- Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
- Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security

- All SSL certificates and credentials must be stored in Azure Key Vault.
- File access must restrict access by IP, protocol, and Azure AD rights.
- All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance

Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons. **Issues**

Corporate website

While testing the site, the following error message displays:

CryptographicException: The system cannot find the file specified.

Function app

You perform local testing for the RequestUserApproval function. The following error message displays:

'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'

The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

```
FunctionAppLogs  
| where FunctionName == "RequestUserApproval"
```

Logic app

You test the Logic app in a development environment. The following error message displays:

'400 Bad Request'

Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.

Code

Corporate website

Security.cs:

```
SC01 public class Security  
SC02 {  
SC03 var bytes = System.IO.File.ReadAllBytes("~/var/ssl/private");  
SC04 var cert = new System.Security.Cryptography.X509Certificate2(bytes);  
SC05 var certName = cert.FriendlyName;  
SC06 }
```

Function app

RequestUserApproval.cs:

```
RA01 public static class RequestUserApproval  
RA02 {  
RA03 [FunctionName("RequestUserApproval")]  
RA04 public static async Task<IActionResult> Run(  
RA05 [HttpTrigger(AuthorizationLevel.Function, "get", "post", Route = null)] HttpRequest req,  
ILogger log)  
RA06 {  
RA07     log.LogInformation("RequestUserApproval function processed a request.");  
RA08     ...  
RA09     return ProcessRequest(req)  
RA10     ? (ActionResult)new OkObjectResult($"User approval processed")  
RA11     : new BadRequestObjectResult("Failed to process user approval");  
RA12 }  
RA13 private static bool ProcessRequest(HttpContext req)  
RA14 {  
RA15     ...  
RA16 }  
RA17 }
```

QUESTION 1

You need to correct the RequestUserApproval Function app error.

What should you do?

- A. Update line RA13 to use the `async` keyword and return an `HttpRequest` object value.
- B. Configure the Function app to use an App Service hosting plan. Enable the **Always On** setting of the hosting plan.
- C. Update the function to be stateful by using Durable Functions to process the request payload.
- D. Update the `functionTimeout` property of the `host.json` project file to 15 minutes.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Async operation tracking

The HTTP response mentioned previously is designed to help implement long-running HTTP async APIs with Durable Functions. This pattern is sometimes referred to as the polling consumer pattern.

Both the client and server implementations of this pattern are built into the Durable Functions HTTP APIs.

Function app



You perform local testing for the RequestUserApproval function. The following error message displays:

'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'

The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:

FunctionAppLogs

| where FunctionName == "RequestUserApproval"

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-http-features>

Develop Azure compute solutions

Question Set 3

QUESTION 1

HOTSPOT

You are developing an ASP.NET Core web application. You plan to deploy the application to Azure Web App for Containers.

The application needs to store runtime diagnostic data that must be persisted across application restarts. You have the following code:

```
public void SaveDiagData(string data)
{
    var path = Environment.GetEnvironmentVariable("DIAGDATA");
    File.WriteAllText(Path.Combine(path, "data"), data);
}
```

You need to configure the application settings so that diagnostic data is stored as required.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

App setting	Value
LOCALAPPDATA	true
WEBSITE_LOCALCACHE_ENABLED	
DOTNET_HOSTING_OPTIMIZATION_CACHE	
WEBSITES_ENABLE_APP_SERVICE_STORAGE	
DIAGDATA	

Correct Answer:

Answer Area

App setting	Value
LOCALAPPDATA	true
WEBSITE_LOCALCACHE_ENABLED	
DOTNET_HOSTING_OPTIMIZATION_CACHE	
WEBSITES_ENABLE_APP_SERVICE_STORAGE	
DIAGDATA	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: If WEBSITES_ENABLE_APP_SERVICE_STORAGE

If WEBSITES_ENABLE_APP_SERVICE_STORAGE setting is unspecified or set to true, the /home/ directory will be shared across scale instances, and files written will persist across restarts

Box 2: /home

Reference: <https://docs.microsoft.com/en-us/azure/app-service/containers/app-service-linux-faq>

QUESTION 2

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

All traffic must be secured with SSL. The Azure Application Gateway instance is used by multiple web apps.

You need to configure the Azure Application Gateway for the web app.

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:


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

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

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

In case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway. Therefore, there is no need to add any authentication certificates.





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

QUESTION 3

HOTSPOT

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

- Each instance of the WebJob processes data for a single customer and must run as a singleton instance.
- Each deployment must be tested by using deployment slots prior to serving production data.
- Azure costs must be minimized.

- Azure resources must be located in an isolated network.

You need to configure the App Service plan for the Web App.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

App service plan setting	Value				
Number of VM instances	<input type="button" value="▼"/> <table border="1"><tr><td>2</td></tr><tr><td>4</td></tr><tr><td>8</td></tr><tr><td>16</td></tr></table>	2	4	8	16
2					
4					
8					
16					
Pricing tier	<input type="button" value="▼"/> <table border="1"><tr><td>Isolated</td></tr><tr><td>Standard</td></tr><tr><td>Premium</td></tr><tr><td>Consumption</td></tr></table>	Isolated	Standard	Premium	Consumption
Isolated					
Standard					
Premium					
Consumption					

Correct Answer:

Answer Area

App service plan setting	Value				
Number of VM instances	<input type="button" value="▼"/> <table border="1"><tr><td>2</td></tr><tr><td>4</td></tr><tr><td>8</td></tr><tr><td>16</td></tr></table>	2	4	8	16
2					
4					
8					
16					
Pricing tier	<input type="button" value="▼"/> <table border="1"><tr><td>Isolated</td></tr><tr><td>Standard</td></tr><tr><td>Premium</td></tr><tr><td>Consumption</td></tr></table>	Isolated	Standard	Premium	Consumption
Isolated					
Standard					
Premium					
Consumption					

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Number of VM instances: 4

You are not charged extra for deployment slots.

Pricing tier: Isolated

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

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

QUESTION 4

DRAG DROP You

are a developer for a software as a service (SaaS) company that uses an Azure Function to process orders. The Azure Function currently runs on an Azure Function app that is triggered by an Azure Storage queue.

You are preparing to migrate the Azure Function to Kubernetes using Kubernetes-based Event Driven Autoscaling (KEDA).

You need to configure Kubernetes Custom Resource Definitions (CRD) for the Azure Function.

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

NOTE: Each correct selection is worth one point.

Select and Place:



Answer Area

CRD types	Setting	CRD type
Secret	Azure Function code	<input type="text"/>
Deployment	Polling interval	<input type="text"/>
ScaledObject	Azure Storage connection string	<input type="text"/>
TriggerAuthentication		<input type="text"/>

Correct Answer:

Answer Area

CRD types	Setting	CRD type
Secret	Azure Function code	Deployment
Deployment		ScaledObject
ScaledObject	Polling interval	
TriggerAuthentication	Azure Storage connection string	Secret

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Deployment

To deploy Azure Functions to Kubernetes use the `func kubernetes deploy` command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.

Box 2: ScaledObject

With `--polling-interval`, we can control the interval used by KEDA to check Azure Service Bus Queue for messages. Example of ScaledObject with polling interval apiVersion: `keda.k8s.io/v1alpha1` kind: `ScaledObject` metadata:

```
name: transformer-fn
namespace: tt  labels:
deploymentName: transformer-fn
spec:
  scaleTargetRef:
  deploymentName: transformer-fn
    pollingInterval: 5
minReplicaCount: 0
maxReplicaCount: 100
```

Box 3: Secret

Store connection strings in Kubernetes Secrets.

Example: to create the Secret in our demo Namespace:

```
# create the k8s demo namespace
kubectl create namespace tt

# grab connection string from Azure Service Bus
KEDA_SCALER_CONNECTION_STRING=$(az servicebus queue authorization-rule keys list \
-g $RG_NAME \
--namespace-name $SBN_NAME \
--queue-name inbound \
-n keda-scaler \
--query "primaryConnectionString" \
-o tsv)

# create the kubernetes secret
kubectl create secret generic tt-keda-auth \
--from-literal KedaScaler=$KEDA_SCALER_CONNECTION_STRING \
--namespace tt
```

Reference: <https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/>



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

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

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

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

Solution: Trigger the photo processing from Blob storage events.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You need to catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload

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

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

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

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



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 that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named **Testing** and **Production**. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

Does the solution meet the goal?

A. No

B. Yes

Correct Answer: A

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

Explanation:

Specify custom warm-up.

Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.

```
<system.webServer>
  <applicationInitialization>
    <add initializationPage="/" hostName="[app hostname]" />
    <add initializationPage="/Home/About" hostName="[app hostname]" />
  </applicationInitialization>
</system.webServer>
```

Reference: <https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps>

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 that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named **Testing** and **Production**. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Enable auto swap for the Testing slot. Deploy the app to the Testing slot.

Does the solution meet the goal?

A. No

B. Yes

Correct Answer: B

Section: (none)**Explanation**

Explanation/Reference:

Explanation:

Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.

```
<system.webServer>
  <applicationInitialization>
    <add initializationPage="/" hostName="[app hostname]" />
    <add initializationPage="/Home/About" hostName="[app hostname]" />
  </applicationInitialization>
</system.webServer>
```

Reference: <https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps> **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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named **Testing** and **Production**. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Disable auto swap. Update the app with a method named statuscheck to run the scripts. Re-enable auto swap and deploy the app to the Production slot.

Does the solution meet the goal?

- A. No
- B. Yes

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.

```
<system.webServer>
  <applicationInitialization>
    <add initializationPage="/" hostName="[app hostname]" />
    <add initializationPage="/Home/About" hostName="[app hostname]" />
  </applicationInitialization>
</system.webServer>
```

Reference: <https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps>

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

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

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

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

Solution: Convert the Azure Storage account to a BlockBlobStorage storage account.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

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

Reference:

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

QUESTION 10

DRAG DROP

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

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

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

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

Select and Place:

Azure CLI Commands Answer Area

az group create

az group update

az webapp update

az webapp create

az appservice plan create



Correct Answer:

Azure CLI Commands Answer Area

az group create

az group create

az group update

az appservice plan create

az webapp update

az webapp create

az webapp create



az appservice plan create

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

Step 1: az group create

In the Cloud Shell, create a resource group with the az group create command.

Step 2: az appservice plan create

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

Step 3: az webapp create

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

Reference:

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

QUESTION 11

DRAG DROP Fourth
Coffee has an
ASP.NET Core web
app that runs in
Docker. The app is
mapped to the
[www.fourthcoffee.co
m](http://www.fourthcoffee.com) domain.



Fourth Coffee is migrating this application to Azure.

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

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

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

Select and Place:

Azure CLI Commands

Answer Area

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



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

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

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

Correct Answer:

Azure CLI Commands

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

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

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

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

Answer Area

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



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

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

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

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Step 1: #bin/bash

The appName is used when the webapp-name is created in step 2.

Step 2: az webapp config hostname add

The webapp-name is used when the webapp is created in step 3.

Step 3: az webapp create

Create a web app. In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command.

Step : az webapp config container set

In Create a web app, you specified an image on Docker Hub in the az webapp create command. This is good enough for a public image. To use a private image, you need to configure your Docker account ID and password in your Azure web app.

In the Cloud Shell, follow the az webapp create command with az webapp config container set.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image>

QUESTION 12

DRAG DROP

You are developing a serverless Java application on Azure. You create a new Azure Key Vault to work with secrets from a new Azure Functions application.

The application must meet the following requirements:

- Reference the Azure Key Vault without requiring any changes to the Java code.
- Dynamically add and remove instances of the Azure Functions host based on the number of incoming application events.
- Ensure that instances are perpetually warm to avoid any cold starts.
- Connect to a VNet.
- Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.

You need to grant the Azure Functions application access to the Azure Key Vault.

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

Select and Place:

Actions	Answer Area
Create a user-assigned managed identity for the application.	
Create the Azure Functions app with a Premium plan type.	> <
Create an access policy in Azure Key Vault for the application identity.	↑ ↓
Create an SSL certification in Azure Key Vault for the application identity.	↑ ↓
Create the Azure Functions app with an App Service plan type.	↑ ↓
Create the Azure Functions app with a Consumption plan type.	↑ ↓
Create a system-assigned managed identity for the application.	

Correct Answer:



Actions	Answer Area
Create a user-assigned managed identity for the application.	Create the Azure Functions app with a Consumption plan type.
Create the Azure Functions app with a Premium plan type.	Create a user-assigned managed identity for the application.
Create an access policy in Azure Key Vault for the application identity.	> <
Create an SSL certification in Azure Key Vault for the application identity.	↑ ↓
Create the Azure Functions app with an App Service plan type.	↑ ↓
Create the Azure Functions app with a Consumption plan type.	↑ ↓
Create a system-assigned managed identity for the application.	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Step 1: Create the Azure Functions app with a Consumption plan type.
Use the Consumption plan for serverless.

Step 2: Create a system-assigned managed identity for the application.
Create a system-assigned managed identity for your application.

Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.

Step 3: Create an access policy in Key Vault for the application identity.

Create an access policy in Key Vault for the application identity you created earlier. Enable the "Get" secret permission on this policy. Do not configure the "authorized application" or applicationId settings, as this is not compatible with a managed identity.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references>

QUESTION 13

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

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

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 a virtual machine. Configure a Scale Set to increase the virtual machine instance count when the CPU load is high.
- D. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

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

Incorrect Answers:

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

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

QUESTION 14

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup
az group create --location centralus --name $resourcegroupname
az webapp create --name $webappname --resource-group $resourcegroupname --sku S3
az appservice plan create --name $webappname --resource-group $resourcegroupname --plan $webappname
az webapp deployment slot staging create --name $webappname --resource-group $resourcegroupname --slot staging
az webapp config --name $webappname --resource-group $resourcegroupname \
--slot staging --repo-url $gitrepo --branch master --manual-integration
```

Correct Answer:

Answer Area

```

gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup
az group create --location centralus --name $resourcegroupname
az appservice plan create --name $webappname --resource-group $resourcegroupname --sku S3
az webapp create --name $webappname --resource-group $resourcegroupname --plan $webappname
az webapp deployment slot create --name $webappname --resource-group $resourcegroupname --slot staging
az webapp config --name $webappname --resource-group $resourcegroupname \
--slot staging --repo-url $gitrepo --branch master --manual-integration

```

Section: (none)
Explanation

Explanation/Reference:

Explanation:

Box 1: group # Create a resource group. az group create --location westeurope --name myResourceGroup

Box 2: appservice plan
Create an App Service plan in STANDARD tier (minimum required by deployment slots). az appservice plan create --name \$webappname --resource-group myResourceGroup --sku S1

Box 3: webapp #
Create a web app.
az webapp create --name \$webappname --resource-group myResourceGroup \

```
--plan $webappname
```

Box 4: webapp deployment slot

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

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

Box 5: webapp deployment source

Deploy sample code to "staging" slot from GitHub.

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

Reference:

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

QUESTION 15

HOTSPOT

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

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

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

How should you complete the trigger?

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
function ensureTip() {
    var r = 
    __.value();
    __.readDocument('item');
    getContext().getRequest();
    getContext().getResponse();

    var i = r.getBody();
    
    if (!("tip" in i)) {
        if (request.getValue("tip") == null) {
            if (isNaN(i["tip"]) || i["tip"] == null) {
                if (typeof __.pluck("tip") == 'number') {
                    i["tip"] = 0;
                }
            }
        }
    }

    r.setBody(i);
    r.setValue(i);
    __.upsertDocument(i);
    __.replaceDocument(i);
}
```

Correct Answer:

Answer Area

```
function ensureTip() {
    var r = __.readDocument('item');
    getContext().getRequest();
    getContext().getResponse();

    var i = r.getBody();

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

    r.setBody(i);
    r.setValue(i);
    __.upsertDocument(i);
    __.replaceDocument(i);
}
```

Section: (none)

Explanation

Explanation/Reference:

QUESTION 16

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

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

You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.

The app continues to time out after four minutes. The app must process the blob data.

You need to ensure the app does not time out and processes the blob data.

Solution: Use the Durable Function async pattern to process the blob data.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.

Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

QUESTION 17

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

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

You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.

The app continues to time out after four minutes. The app must process the blob data.

You need to ensure the app does not time out and processes the blob data.

Solution: Pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

QUESTION 18

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

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

You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.

The app continues to time out after four minutes. The app must process the blob data.

You need to ensure the app does not time out and processes the blob data.

Solution: Configure the app to use an App Service hosting plan and enable the Always On setting.

Does the solution meet the goal?

- A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.

Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

QUESTION 19

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

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

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

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

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

Solution: Move photo processing to an Azure Function triggered from the blob upload.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

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

Reference:

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

QUESTION 20

You are developing an application that uses Azure Blob storage.

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

You need to process the transaction logs asynchronously.

What should you do?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Change feed support in Azure Blob Storage

The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes. Client applications can read these logs at any time, either in streaming or in batch mode. The change feed enables you to build efficient and scalable solutions that process change events that occur in your Blob Storage account at a low cost.

Reference:

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

QUESTION 21

DRAG DROP

You plan to create a Docker image that runs an ASP.NET Core application named ContosoApp. You have a setup script named setupScript.ps1 and a series of application files including ContosoApp.dll.

You need to create a Dockerfile document that meets the following requirements:

- Call setupScripts.ps1 when the container is built.
- Run ContosoApp.dll when the container starts.

The Dockerfile document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.

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

Select and Place:

Commands	Answer Area
FROM microsoft/aspnetcore:latest	
WORKDIR /apps/ContosoApp	
CMD ["dotnet", "ContosoApp.dll"]	
COPY ./ .	
RUN powershell ./setupScript.ps1	

Correct Answer:

Commands	Answer Area
FROM microsoft/aspnetcore:latest	CMD ["dotnet", "ContosoApp.dll"]
WORKDIR /apps/ContosoApp	FROM microsoft/aspnetcore:latest
CMD ["dotnet", "ContosoApp.dll"]	WORKDIR /apps/ContosoApp
COPY ./ .	COPY ./ .
RUN powershell ./setupScript.ps1	RUN powershell ./setupScript.ps1

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

Explanation:

Box 1: CMD [..]

Cmd starts a new instance of the command interpreter, Cmd.exe.

Syntax: CMD <string>

Specifies the command you want to carry out.

Box 2: FROM microsoft/aspnetcore-build:latest

Box 3: WORKDIR /apps/ContosoApp

Box 4: COPY ./ .

Box 5: RUN powershell ./setupScript.ps1

QUESTION 22

You are developing an Azure Function App that processes images that are uploaded to an Azure Blob container.

Images must be processed as quickly as possible after they are uploaded, and the solution must minimize latency. You create code to process images when the Function App is triggered.

You need to configure the Function App.

What should you do?

- A. Use an App Service plan. Configure the Function App to use an Azure Blob Storage input trigger.
- B. Use a Consumption plan. Configure the Function App to use an Azure Blob Storage trigger.
- C. Use a Consumption plan. Configure the Function App to use a Timer trigger.



- D. Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger.
- E. Use a Consumption plan. Configure the Function App to use an Azure Blob Storage input trigger.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Blob storage trigger starts a function when a new or updated blob is detected. The blob contents are provided as input to the function.
The Consumption plan limits a function app on one virtual machine (VM) to 1.5 GB of memory.

Reference:

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

QUESTION 23

HOTSPOT

You are configuring a new development environment for a Java application.

The environment requires a Virtual Machine Scale Set (VMSS), several storage accounts, and networking components.

The VMSS must not be created until the storage accounts have been successfully created and an associated load balancer and virtual network is configured.

How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
{  
    . . .  
    "resources": [  
        {  
            "apiVersion": "2016-01-01",  
            "type": "Microsoft.Storage/storageAccounts",  
            "name": "[concat('storage', uniqueString(resourceGroup().id))]",  
            "copy": {  
                "index": 1,  
                "dependsOn": [  
                    "storagesetup"  
                ]  
            },  
            "sku": {  
                "name": "Standard_LRS"  
            },  
            "kind": "Storage",  
            "properties": {},  
            "copyIndex": 1,  
            "dependsOn": [  
                "storagesetup"  
            ],  
            "name": "storagesetup",  
            "count": 3  
        },  
        {  
            "apiVersion": "2015-06-15",  
            "type": "Microsoft.Compute/virtualMachines",  
            "name": "[concat('VM', uniqueString(resourceGroup().id))]",  
            "copy": {  
                "index": 1,  
                "dependsOn": [  
                    "loadbalancer",  
                    "virtualNetwork",  
                    "storagesetup"  
                ]  
            },  
            "variables": [  
                "[variables('loadBalancerName')]",  
                "[variables('virtualNetworkName')]",  
                "[variables('storageAccountName')]"  
            ],  
            "dependsOn": [  
                "loadbalancer",  
                "virtualNetwork",  
                "storagesetup"  
            ]  
        }  
    ],  
    "outputs": {}  
}
```

Correct Answer:

Answer Area

```
{  
  . . .  
  "resources": [  
    {  
      "apiVersion": "2016-01-01",  
      "type": "Microsoft.Storage/storageAccounts",  
      "name": "[concat((), 'storage', uniqueString(resourceGroup().id))]",  
      "copy": {},  
      "copyIndex": 1,  
      "priority": null,  
      "dependsOn": null  
      "location": "[resourceGroup().location]",  
      . . .  
      "sku": {  
        "name": "Standard_LRS"  
      },  
      "kind": "Storage",  
      "properties": {},  
      "copy": {},  
      "copyIndex": 1,  
      "priority": null,  
      "dependsOn": null  
      "name": "storagesetup",  
      "count": 3  
    }  
  ],  
  {  
    "apiVersion": "2015-06-15",  
    "type": "Microsoft.Compute/virtualMachines",  
    "name": "[concat('VM', uniqueString(resourceGroup().id))]",  
    "copy": {},  
    "copyIndex": 1,  
    "priority": null,  
    "dependsOn": null  
    "[variables('loadBalancerName')]",  
    "[variables('virtualNetworkName')]",  
    "storagesetup",  
    [],  
    . . .  
  }  
],  
"outputs": {}  
}
```

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

Explanation:

Box 1: copyIndex

Notice that the name of each resource includes the copyIndex() function, which returns the current iteration in the loop. copyIndex() is zero-based.

Box 2: copy

By adding the copy element to the resources section of your template, you can dynamically set the number of resources to deploy.

Box 3: dependsOn

Example:

```
"type": "Microsoft.Compute/virtualMachineScaleSets",
    "apiVersion": "2020-06-01",
    "name": "[variables('namingInfix')]",
"location": "[parameters('location')]",
    "sku": {
        "name": "[parameters('vmSku')]",
        "tier": "Standard",
        "capacity": "[parameters('instanceCount')]"
},
    "dependsOn": [
        "[resourceId('Microsoft.Network/loadBalancers', variables('loadBalancerName'))]",
"[resourceId('Microsoft.Network/virtualNetworks', variables('virtualNetworkName'))]"
],
```

Reference: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/copy-resources> <https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/quick-create-template-windows>

QUESTION 24

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

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

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

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

[config]

command = YOUR COMMAND TO RUN FOR DEPLOYMENT

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

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

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

QUESTION 25

DRAG DROP

You are developing an application to use Azure Blob storage. You have configured Azure Blob storage to include change feeds.

A copy of your storage account must be created in another region. Data must be copied from the current storage account to the new storage account directly between the storage servers.

You need to create a copy of the storage account in another region and copy the data.

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

Select and Place:

Actions	Answer Area
Use AZCopy to copy the data to the new storage account.	
Deploy the template to create a new storage account in the target region.	 
Export a Resource Manager template.	 
Create a new template deployment.	
Modify the template by changing the storage account name and region.	

Correct Answer:

Actions	Answer Area
Use AZCopy to copy the data to the new storage account.	Create a new template deployment.
Deploy the template to create a new storage account in the target region.	Export a Resource Manager template.
Export a Resource Manager template.	 
Create a new template deployment.	Modify the template by changing the storage account name and region.
Modify the template by changing the storage account name and region.	 

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice.

The steps are:

- Export a template.
- Modify the template by adding the target region and storage account name.
- Deploy the template to create the new storage account.
- Configure the new storage account.
- Move data to the new storage account.
- Delete the resources in the source region.

Note: You must enable the change feed on your storage account to begin capturing and recording changes. You can enable and disable changes by using Azure Resource Manager templates on Portal or Powershell.

Reference: <https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move> <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed>

QUESTION 26

DRAG DROP

You are preparing to deploy an Azure virtual machine (VM)-based application.

The VMs that run the application have the following requirements:

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

You need to ensure that the requirements are met.

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

NOTE: Each correct selection is worth one point.

Select and Place:

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

Correct Answer:

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

Section: (none)

Explanation

Explanation/Reference:

Reference:

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

QUESTION 27

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
 -Name myResourceGroup -Location $location
 New-AzWebAppSlot
 New-AzWebApp
 New-AzAppServicePlan
 New-AzResourceGroup
 -Name $webappname -Location $location -ResourceGroupName myResourceGroup -Tier Standard
 New-AzWebAppSlot
 New-AzWebApp
 New-AzAppServicePlan
 New-AzResourceGroup
 -Name $webappname -Location $location -AppServicePlan $webappname -ResourceGroupName myResourceGroup
 New-AzWebAppSlot
 New-AzWebApp
 New-AzAppServicePlan
 New-AzResourceGroup
 -Name $webappname -ResourceGroupName myResourceGroup -Slot review
 New-AzWebAppSlot
 New-AzWebApp
 New-AzAppServicePlan
 New-AzResourceGroup
$PropertiesObject = @{$repoUrl = "$gitrepo",branch = "master"}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup
-SrcSlotName review -DestinationSlotName production
```

Correct Answer:

Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
-Name myResourceGroup -Location $location
New-AzWebAppSlot
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup
-Name $webappname -Location $location -ResourceGroupName myResourceGroup -Tier Standard
New-AzWebAppSlot
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup
-Name $webappname -Location $location -AppServicePlan $webappname -ResourceGroupName myResourceGroup
New-AzWebAppSlot
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup
-Name $webappname -ResourceGroupName myResourceGroup -Slot review
New-AzWebAppSlot
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup
$PropertiesObject = @{repoUrl = "$gitrepo";branch = "master"}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup
-SourceSlotName review -DestinationSlotName production
```

Section: (none)
Explanation
Explanation/Reference:

Explanation:

Box 1: New-AzResourceGroup

The New-AzResourceGroup cmdlet creates an Azure resource group.

Box 2: New-AzAppServicePlan

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

Box 3: New-AzWebApp

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

Box 4: New-AzWebAppSlot

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

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



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

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

QUESTION 28

HOTSPOT

You are developing an application that needs access to an Azure virtual machine (VM).

The access lifecycle for the application must be associated with the VM service instance.

You need to enable managed identity for the VM.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
$vm = Get-AzVM -ResourceGroupName "ContosoRG" -Name "ContosoVM"
```

```
Update-AzVM -ResourceGroupName "ContosoRG" -VM $vm
```

-AssignIdentity:
-IdentityId:

\$SystemAssigned
\$UserAssigned

Correct Answer:

Answer Area

```
$vm = Get-AzVM -ResourceGroupName "ContosoRG" -Name "ContosoVM"
```

```
Update-AzVM -ResourceGroupName "ContosoRG" -VM $vm
```

-AssignIdentity:
-IdentityId:

\$SystemAssigned
\$UserAssigned

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: -IdentityType

Enable system-assigned managed identity on an existing Azure VM:

To enable a system-assigned managed identity, use the -IdentityType switch on the Update-AzVM cmdlet (see below).

Box 2: \$SystemAssigned

```
$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM  
Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -IdentityType SystemAssigned
```

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/q-and-a-powershell-windows-vm>

QUESTION 29

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

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

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

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

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

Solution: Create an Azure Function app that uses the Consumption hosting model and that is triggered from the blob upload.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In the Consumption hosting plan, resources are added dynamically as required by your functions.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function>

QUESTION 30

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

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

You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named **Testing** and **Production**. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Update the app with a method named statuscheck to run the scripts. Update the app settings for the app. Set the WEBSITE_SWAP_WARMUP_PING_PATH and WEBSITE_SWAP_WARMUP_PING_STATUSES with a path to the new method and appropriate response codes.

Does the solution meet the goal?

- A. No
- B. Yes

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

Explanation:

These are valid warm-up behavior options, but are not helpful in fixing swap problems.

Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.

```
<system.webServer>
  <applicationInitialization>
    <add initializationPage="/" hostName="[app hostname]" />
    <add initializationPage="/Home/About" hostName="[app hostname]" />
  </applicationInitialization>
</system.webServer>
```

Reference: <https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps> 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 questions in 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 section 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 in 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. When you are ready to answer a question, click the **Question** button to return to the question.

Background

Overview

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

- Messages are sent to ContentUploadService.
- Content is processed by ContentAnalysisService.
- After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements

ContentAnalysisService

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs

You must minimize costs for all Azure services.

Manual review

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes. **High availability**

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security

You have the following security requirements:

- Any web service accessible over the Internet must be protected from cross site scripting attacks.
- All websites and services must use SSL from a valid root certificate authority.
- Azure Storage access keys must only be stored in memory and must be available only to the service.
- All Internal services must only be accessible from internal Virtual Networks (VNets).
- All parts of the system must support inbound and outbound traffic restrictions. ▪
All service calls must be authenticated by using Azure AD.

User agreements

When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.

Information regarding agreements is used by multiple divisions within Contoso, Ltd.

User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing

When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues

Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code

ContentUploadService

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06   containers:
CS07     - name: service
CS08       properties:
CS09         image: contoso/contentUploadService:latest
CS10       ports:
CS11         - port: 80
CS12           protocol: TCP
CS13       resources:
CS14         requests:
CS15           cpu: 1.0
CS16           memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19   ip: 10.23.121.112
CS20   ports:
CS21     - port: 80
CS22       protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```



ApplicationManifest

```
AM01 {  
AM02     "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",  
AM03     "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",  
AM04  
AM05     "createdDateTime" : "2019-12-24T06:01:44Z",  
AM06     "logoUrl" : null,  
AM07     "logoutUrl" : null,  
AM08     "name" : "ContentAnalysisService",  
AM09  
AM10  
AM11     "orgRestrictions" : [],  
AM12     "parentalControlSettings" : {  
AM13         "countriesBlockedForMinors" : [],  
AM14         "legalAgeGroupRule" : "Allow"  
AM15     },  
AM16     "passwordCredentials" : []  
AM17 }
```

QUESTION 1

You need to configure the ContentUploadService deployment.

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

NOTE: Each correct selection is worth one point.

- A. Add the following markup to line CS23:type:

Private

- B. Add the following markup to line CS24:

osType: Windows

- C. Add the following markup to line CS24:

osType: Linux

- D. Add the following markup to line CS23:type:

Public

Correct Answer: A

Section: (none)



Explanation**Explanation/Reference:**

Explanation:

Scenario: All Internal services must only be accessible from Internal Virtual Networks (VNets)

There are three Network Location types – Private, Public and Domain

Reference: <https://devblogs.microsoft.com/powershell/setting-network-location-to-private/>

QUESTION 2

You need to store the user agreements.

Where should you store the agreement after it is completed?

- A. Azure Storage queue
- B. Azure Event Hub
- C. Azure Service Bus topic
- D. Azure Event Grid topic

Correct Answer: B

Section: (none)

Explanation

**Explanation/Reference:**

Explanation:

Azure Event Hub is used for telemetry and distributed data streaming.

This service provides a single solution that enables rapid data retrieval for real-time processing as well as repeated replay of stored raw data. It can capture the streaming data into a file for processing and analysis.

It has the following characteristics:

- low latency
- capable of receiving and processing millions of events per second
- at least once delivery

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

QUESTION 3

HOTSPOT

You need to implement the bindings for the CheckUserContent function.

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

NOTE: Each correct selection is worth one point.

Hot Area:**Answer Area**

```
public static class CheckUserContent
{
    [FunctionName("CheckUserContent")]
    public static void Run(
        [QueueTrigger("userContent")] string content,
        [BlobTrigger("userContent/{name}")] Stream output)
    {
        ...
    }
}
```

Correct Answer:

Answer Area

```
public static class CheckUserContent
{
    [FunctionName("CheckUserContent")]
    public static void Run(
        [QueueTrigger("userContent")] string content,
        [BlobTrigger("userContent/{name}")] Stream output)
    {
        ...
    }
}
```

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

Explanation:

Box 1: [BlobTrigger(..)]

Box 2: [Blob(..)]

Azure Blob storage output binding for Azure Functions. The output binding allows you to modify and delete blob storage data in an Azure Function.

The attribute's constructor takes the path to the blob and a FileAccess parameter indicating read or write, as shown in the following example:

```
[FunctionName("ResizeImage")]
public static void Run(
[BlobTrigger("sample-
images/{name}")] Stream image,
[Blob("sample-images-md/{name}", FileAccess.Write)] Stream imageSmall)
{
...
}
```

Scenario: You must create an Azure Function named CheckUserContent to perform the content checks.

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

Reference:

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



Develop for Azure storage

Question Set 2

QUESTION 1

DRAG DROP

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

```
{  
    "name": "John",  
    "city" : "Seattle"  
}
```

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

```
SELECT*  
FROM People p  
ORDER BY p.name, p.city DESC
```

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

JSON segments

orderBy
sortOrder
ascending
descending
compositeIndexes

Answer Area

```
{  
    "automatic": true,  
    "ngMode": "Consistent",  
    "includedPaths": [  
        {  
            "path": "/**"  
        }  
    ], "excludedPaths": [],  
    "compositeIndexes": [  
        [  
            {  
                "path": "/name", "order": "descending"  
            },  
            {  
                "path": "/city", "order": " "  
            }  
        ]  
    ]  
}
```

Correct Answer:

JSON segments

orderBy
sortOrder
ascending
descending
compositeIndexes

Answer Area

```
{  
    "automatic": true,  
    "ngMode": "Consistent",  
    "includedPaths": [  
        {  
            "path": "/**"  
        }  
    ], "excludedPaths": [],  
    "compositeIndexes": [  
        [  
            {  
                "path": "/name", "order": "descending"  
            },  
            {  
                "path": "/city", "order": "descending"  
            }  
        ]  
    ]  
}
```

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

Explanation:

Box 1: compositeIndexes

You can order by multiple properties. A query that orders by multiple properties requires a composite index.

Box 2: descending

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

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

{

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

QUESTION 2

DRAG DROP

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

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



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

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

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)

Explanation:

Explanation/Reference:

Explanation:

Box 1: Helm

To create the ingress controller, use Helm to install nginx-ingress.

Box 2: kubectl

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

Box 3: Ingress Controller

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

Incorrect Answers:

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

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

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

QUESTION 3

DRAG DROP

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

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

The system has the following requirements for subscriptions:

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

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

Filter types

- SQLFilter
- CorrelationFilter
- No Filter

Answer Area
Subscription

- FutureOrders
- HighPriorityOrders
- InternationalOrders
- HighQuantityOrders
- AllOrders

Filter type

- filter type

Correct Answer:

Filter types

- SQLFilter
- CorrelationFilter
- No Filter

Answer Area
Subscription

- FutureOrders
- HighPriorityOrders
- InternationalOrders
- HighQuantityOrders
- AllOrders

Filter type

- SQLFilter
- CorrelationFilter
- SQLFilter
- SQLFilter
- No Filter

Section: (none)

Explanation

Explanation/Reference:

Explanation:

FutureOrders: SQLFilter

HighPriorityOrders: CorrelationFilter CorrelationID only

InternationalOrders: SQLFilter
Country NOT USA requires an SQL Filter

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

AllOrders: No Filter

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

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

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

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

QUESTION 4

DRAG DROP

Your company has several websites that use a company logo image. You use Azure Content Delivery Network (CDN) to store the static image.

You need to determine the correct process of how the CDN and the Point of Presence (POP) server will distribute the image and list the items in the correct order.

In which order do the actions occur? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

Actions Answer Area

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



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

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

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

Correct Answer:

Actions Answer Area

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



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

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

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

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

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

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

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

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Step 1: A user requests the image..

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

Step 2: If no edge servers in the POP have the..

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

Step 3: The origin server returns the..

The origin server returns the file to an edge server in the POP.

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

Step 4: Subsequent requests for..

Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.

If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.

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



QUESTION 5

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

E. a hash suffix appended to a property value

Correct Answer: DE

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.

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

QUESTION 6

You are building a website that uses Azure Blob storage for data storage. You configure Azure Blob storage lifecycle to move all blobs to the archive tier after 30 days.

Customers have requested a service-level agreement (SLA) for viewing data older than 30 days.

You need to document the minimum SLA for data recovery.

Which SLA should you use?

- A. at least two days
- B. between one and 15 hours
- C. at least one day
- D. between zero and 60 minutes

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The archive access tier has the lowest storage cost. But it has higher data retrieval costs compared to the hot and cool tiers. Data in the archive tier can take several hours to retrieve depending on the priority of the rehydration. For small objects, a high priority rehydrate may retrieve the object from archive in under 1 hour.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal>

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 that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

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

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

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

Solution: Provision an Azure Event Grid. Configure 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:

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

QUESTION 8

You develop Azure solutions.

A .NET application needs to receive a message each time an Azure virtual machine finishes processing data. The messages must NOT persist after being processed by the receiving application.

You need to implement the .NET object that will receive the messages.

Which object should you use?

- A. QueueClient
- B. SubscriptionClient
- C. TopicClient
- D. CloudQueueClient

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A queue allows processing of a message by a single consumer. Need a CloudQueueClient to access the Azure VM.

Incorrect Answers:

B, C: In contrast to queues, topics and subscriptions provide a one-to-many form of communication in a publish and subscribe pattern. It's useful for scaling to large numbers of recipients.

Reference: <https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-queues-topics-subscriptions>

QUESTION 9

You develop Azure solutions.

You must connect to a No-SQL globally-distributed database by using the .NET API.

You need to create an object to configure and execute requests in the database.

Which code segment should you use?

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

C. new CosmosClient(EndpointUri, PrimaryKey);

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Example:

```
// Create a new instance of the Cosmos Client  
this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)
```

//ADD THIS PART TO YOUR CODE

```
await this.CreateDatabaseAsync();
```

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

QUESTION 10

You have an existing Azure storage account that stores large volumes of data across multiple containers.

You need to copy all data from the existing storage account to a new storage account. The copy process must meet the following requirements:

- Automate data movement.
- Minimize user input required to perform the operation.
- Ensure that the data movement process is recoverable.

What should you use?

- A. AzCopy
- B. Azure Storage Explorer
- C. Azure portal
- D. .NET Storage Client Library

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You can copy blobs, directories, and containers between storage accounts by using the AzCopy v10 command-line utility.

The copy operation is synchronous so when the command returns, that indicates that all files have been copied.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy>

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 questions in 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 section 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 in 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. When you are ready to answer a question, click the **Question** button to return to the question.

Background**Overview**

You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.

You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:

- Messages are sent to ContentUploadService.

- Content is processed by ContentAnalysisService.
- After processing is complete, the content is posted to the social network or a rejection message is posted in its place.

The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.

The solution will use eight CPU cores.

Azure Active Directory

Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.

Requirements

ContentAnalysisService

The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.

You must create an Azure Function named CheckUserContent to perform the content checks.

Costs

You must minimize costs for all Azure services.



Manual review

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes. **High availability**

All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.

Monitoring

An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.

Security

You have the following security requirements:

- Any web service accessible over the Internet must be protected from cross site scripting attacks.
- All websites and services must use SSL from a valid root certificate authority.
- Azure Storage access keys must only be stored in memory and must be available only to the service.

- All Internal services must only be accessible from internal Virtual Networks (VNets).
- All parts of the system must support inbound and outbound traffic restrictions.
- All service calls must be authenticated by using Azure AD.

User agreements

When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.

Information regarding agreements is used by multiple divisions within Contoso, Ltd.

User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.

Validation testing

When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Issues

Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.

Code

ContentUploadService

```
CS01 apiVersion: '2018-10-01'
CS02 type: Microsoft.ContainerInstance/containerGroups
CS03 location: westus
CS04 name: contentUploadService
CS05 properties:
CS06   containers:
CS07     - name: service
CS08       properties:
CS09         image: contoso/contentUploadService:latest
CS10       ports:
CS11         - port: 80
CS12           protocol: TCP
CS13       resources:
CS14         requests:
CS15           cpu: 1.0
CS16           memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19   ip: 10.23.121.112
CS20   ports:
CS21     - port: 80
CS22       protocol: TCP
CS23
CS24
CS25 networkProfile:
CS26
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest

```
AM01 {
AM02   "id" : "2b079f03-9b06-2d44-98bb-e9182901fcbb",
AM03   "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AM04
AM05   "createdDateTime" : "2019-12-24T06:01:44Z",
AM06   "logoUrl" : null,
AM07   "logoutUrl" : null,
AM08   "name" : "ContentAnalysisService",
AM09
AM10
AM11   "orgRestrictions" : [],
AM12   "parentalControlSettings" : {
AM13     "countriesBlockedForMinors" : [],
AM14     "legalAgeGroupRule" : "Allow"
AM15   },
AM16   "passwordCredentials" : []
AM17 }
```



QUESTION 1**DRAG DROP**

You need to add markup at line AM04 to implement the ContentReviewer role.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Json segments	Answer Area
User	"appRoles": [
value	{ "": [
role	"": "
Application],
allowedMemberTypes	"displayName": "ContentReviewer",
allowedAccountTypes	"id": "e1c2ade8-98f8-45fd-aa4a-6d24b512c22a", "isEnabled": true, "": "ContentReviewer"

Correct Answer:

Json segments	Answer Area
User	"appRoles": [
value	{ "allowedMemberTypes": [
role	"User"]
Application	,
allowedMemberTypes	"displayName": "ContentReviewer",
allowedAccountTypes	"id": "e1c2ade8-98f8-45fd-aa4a-6d24b512c22a",
	"isEnabled": true,
	" value " : "ContentReviewer"
	}
],

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: allowedMemberTypes allowedMemberTypes specifies whether this app role definition can be assigned to users and groups by setting to "User", or to other applications (that are accessing this application in daemon service scenarios) by setting to "Application", or to both.

Note: The following example shows the appRoles that you can assign to users.

```
"appId": "8763f1c4-f988-489c-a51e-158e9ef97d6a",
"appRoles": [
{
  "allowedMemberTypes": [
    "User"
  ],
  "displayName": "Writer",
  "id": "d1c2ade8-98f8-45fd-aa4a-6d06b947c66f",
  "isEnabled": true,
  "description": "Writers Have the ability to create tasks.",
  "value": "Writer"
},
],
"availableToOtherTenants": false,
```



Box 2: User

Scenario: In order to review content a user must be part of a ContentReviewer role.

Box 3: value value specifies the value which will be included in the roles claim in authentication and access tokens.

Reference: <https://docs.microsoft.com/en-us/graph/api/resources/approle>

QUESTION 2

HOTSPOT

You need to add code at line AM09 to ensure that users can review content using ContentAnalysisService.

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

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

```
"allowPublicClient":true  
"oauth2Permissions": ["login"]  
"oauth2AllowUrlPathMatching":true  
"oauth2AllowIdTokenImplicitFlow":true
```

```
"oauth2AllowImplicitFlow": true  
"oauth2RequiredPostResponse":true  
"preAuthorizedApplications":["SPA"]  
"knownClientApplications":["ContentAnalysisService"]
```

Correct Answer:

Answer Area

```
"allowPublicClient":true
"oauth2Permissions": ["login"]
"oauth2AllowUrlPathMatching":true
"oauth2AllowIdTokenImplicitFlow":true
```

```
"oauth2AllowImplicitFlow": true
"oauth2RequiredPostResponse":true
"preAuthorizedApplications":["SPA"]
"knownClientApplications":["ContentAnalysisService"]
```

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

Explanation:

Box 1: "oauth2Permissions": ["login"] oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

Box 2: "oauth2AllowImplicitFlow":true

For applications (Angular, Ember.js, React.js, and so on), Microsoft identity platform supports the OAuth 2.0 Implicit Grant flow.

Reference: <https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest>

QUESTION 3**HOTSPOT**

You need to ensure that network security policies are met.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Technology	Value
SSL certificate	<input type="checkbox"/> Valid root certificate <input type="checkbox"/> Self-signed certificate
Proxy type	<input type="checkbox"/> nginx <input type="checkbox"/> Azure Application Gateway

Correct Answer:

Answer Area

Technology	Value
SSL certificate	<div style="border: 1px solid #ccc; padding: 5px; display: inline-block;">▼</div> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 5px; background-color: #e0f2e0; display: inline-block;">Valid root certificate</div> <div style="border: 1px solid #ccc; padding: 5px; background-color: #fff; display: inline-block;">Self-signed certificate</div>
Proxy type	<div style="border: 1px solid #ccc; padding: 5px; display: inline-block;">▼</div> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 5px; background-color: #e0f2e0; display: inline-block;">nginx</div> <div style="border: 1px solid #ccc; padding: 5px; background-color: #e0f2e0; display: inline-block;">Azure Application Gateway</div>

Section: (none)
Explanation:

Explanation/Reference:

Explanation:

Box 1: Valid root certificate

Scenario: All websites and services must use SSL from a valid root certificate authority.

Box 2: Azure Application Gateway

Scenario:

- Any web service accessible over the Internet must be protected from cross site scripting attacks.
- All Internal services must only be accessible from Internal Virtual Networks (VNets)
- All parts of the system must support inbound and outbound traffic restrictions.

Azure Web Application Firewall (WAF) on Azure Application Gateway provides centralized protection of your web applications from common exploits and vulnerabilities. Web applications are increasingly targeted by malicious attacks that exploit commonly known vulnerabilities. SQL injection and cross-site scripting are among the most common attacks.

Application Gateway supports autoscaling, SSL offloading, and end-to-end SSL, a web application firewall (WAF), cookie-based session affinity, URL path-based routing, multisite hosting, redirection, rewrite HTTP headers and other features.

Note: Both Nginx and Azure Application Gateway act as a reverse proxy with Layer 7 load-balancing features plus a WAF to ensure strong protection against common web vulnerabilities and exploits.

You can modify Nginx web server configuration/SSL for X-XSS protection. This helps to prevent cross-site scripting exploits by forcing the injection of HTTP headers with X-XSS protection.

Reference:

<https://docs.microsoft.com/en-us/azure/web-application-firewall/ag/ag-overview> <https://www.upguard.com/articles/10-tips-for-securing-your-nginx-deployment>

QUESTION 4

DRAG DROP

You need to add YAML markup at line CS17 to ensure that the ContentUploadService can access Azure Storage access keys.

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

NOTE: Each correct selection is worth one point.

Select and Place:

YAML segments	Answer Area
secret	<input type="text"/> YAML segment : - mountPath: /mnt/secrets name: accesskey
envVar	<input type="text"/> YAML segment : - name: accesskey
secretValues	<input type="text"/> YAML segment : key: TXkgZmlyc3Qgc2VjcmV0IEZPTwo=
volumes	
volumeMounts	
environmentVariables	

Correct Answer:**YAML segments**

envVar

secretValues

environmentVariables

Answer Area

```
volumeMounts :  
  - mountPath: /mnt/secrets  
    name: accesskey  
volumes :  
  - name: accesskey  
secret :  
  key: TXkgZmlyc3Qgc2VjcmV0IEZPTwo=
```

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

Explanation:

Box 1: volumeMounts

Example:

volumeMounts:

```
- mountPath: /mnt/secrets    name:  
secretvolume1 volumes: - name: secretvolume1  
secret: mysecret1:  
TXkgZmlyc3Qgc2VjcmV0IEZPTwo=
```

Box 2: volumes

Box 3: secret

Reference: <https://docs.microsoft.com/en-us/azure/container-instances/container-instances-volume-secret>**QUESTION 5****HOTSPOT**

You need to add code at line AM10 of the application manifest to ensure that the requirement for manually reviewing content can be met.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

"optionalClaims": [

"	acct	platf	sid	tenant_ctry
---	------	-------	-----	-------------

"	sid	upn	email	enfpolids
---	-----	-----	-------	-----------

],

Correct Answer:

Answer Area

"optionalClaims": [

"	acct	,
	platf	
	sid	
	tenant_ctry	

"	sid	,
	upn	
	email	
	enfpolids	



Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: sid

Sid: Session ID, used for per-session user sign-out. Personal and Azure AD accounts.

Scenario: Manual review

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role.

Box 2: email

Scenario: All completed reviews must include the reviewer's email address for auditing purposes.

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 questions in 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 section 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 in 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. When you are ready to answer a question, click the **Question** button to return to the question.

Current environment

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk 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 trucking 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:



Requirements

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://test.wideworldimporters.com/`' is therefore not allowed access.

QUESTION 1

HOTSPOT

You need to secure the Shipping Function app.

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

NOTE: Each correct selection is worth one point.


Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: (none)

Explanation:

Explanation/Reference:

Explanation:

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

Box 1: Function

Box 2: JSON based Token (JWT)

Azure AD uses JSON based tokens (JWTs) that contain claims

Box 3: HTTP

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

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

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

QUESTION 2

You need to secure the Shipping Logic App.

What should you use?

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

Correct Answer: B

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.

Reference:

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

Implement Azure security

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

You develop Azure solutions.

You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.

You need to obtain an Azure Resource Manager access token.

Solution: Run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

Explanation:

Get an access token using the VM's system-assigned managed identity and use it to call Azure Resource Manager
You will need to use PowerShell in this portion.

1. In the portal, navigate to Virtual Machines and go to your Windows virtual machine and in the Overview, click Connect.
2. Enter in your Username and Password for which you added when you created the Windows VM.
3. Now that you have created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.
4. Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.

Example:

```
$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}
```

Reference: <https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

QUESTION 2

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 when a new video is uploaded.

What should you do?

- A. Copy blobs to Container2 by using the **Put Blob** operation of the Blob Service REST API
- B. Create an **Event Grid** topic that uses the `Start-AzureStorageBlobCopy` cmdlet
- C. Use `AzCopy` with the `Snapshot` switch to copy blobs to Container2
- D. Download the blob to a virtual machine and then upload the blob 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.

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

QUESTION 3

You are developing an ASP.NET Core website that uses Azure FrontDoor. The website is used to build custom weather data sets for researchers. Data sets are downloaded by users as Comma Separated Value (CSV) files. The data is refreshed every 10 hours.

Specific files must be purged from the FrontDoor cache based upon Response Header values.

You need to purge individual assets from the Front Door cache.

Which type of cache purge should you use?

- A. single path
- B. wildcard
- C. root domain

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

These formats are supported in the lists of paths to purge:

- Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example, /pictures/strasbourg.png;
- Wildcard purge: Asterisk (*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /*, for example, /pictures/*.
- Root domain purge: Purge the root of the endpoint with "/" in the path.

Reference: <https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching>



QUESTION 4

Your company is developing an Azure API.

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

- All API calls must be secure.
- Callers to the API must not send credentials to the API.

Which authentication mechanism should you use?

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

Correct Answer: C

Section: (none)

Explanation**Explanation/Reference:**

Explanation:

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

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

Incorrect Answers:

A: 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.

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

QUESTION 6

DRAG DROP

Contoso, Ltd. provides an API to customers by using Azure API Management (APIM). The API authorizes users with a JWT token.

You must implement response caching for the APIM gateway. The caching mechanism must detect the user ID of the client that accesses data for a given location and cache the response for that user ID.

You need to add the following policies to the policies file:

- a set-variable policy to store the detected user identity
- a cache-lookup-value policy
- a cache-store-value policy
- a find-and-replace policy to update the response body with the user profile information

To which policy section should you add the policies? To answer, drag the appropriate sections to the correct policies. Each section may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Policy section	Policy	Policy section
	Set-variable	
Inbound	Cache-lookup-value	
Outbound	Cache-store-value	
	Find-and-replace	

Correct Answer:



Policy section	Policy	Policy section
	Set-variable	Inbound
Inbound	Cache-lookup-value	Inbound
Outbound	Cache-store-value	Outbound
	Find-and-replace	Outbound

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

Explanation:

Box 1: Inbound.

A set-variable policy to store the detected user identity.

Example:

```
<policies>
  <inbound>
    <!-- How you determine user identity is application dependent -->
    <set-variable      name="enduserid"
value="@{context.Request.Headers.GetValueOrDefault("Authorization","",).Split(' ')[1].AsJwt()?.Subject)" />
```

Box 2: Inbound

A cache-lookup-value policy

Example:

```
<inbound>
  <base />
  <cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false" downstream-caching-type="none | private | public" must-revalidate="true | false">
    <vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several times --
  </cache-lookup>  </inbound>
```

Box 3: Outbound

A cache-store-value policy.

Example:

```
<outbound>
  <base />
  <cache-store duration="3600" />
</outbound>
```

Box 4: Outbound

A find-and-replace policy to update the response body with the user profile information.

Example:

```
<outbound>
  <!-- Update response body with user profile-->
```

```
<find-and-replace
from="$UserProfile$"
to="@((string)context.Variables["UserProfile"])" />

<base />
</outbound>
```

Reference: <https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies> <https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

QUESTION 7

DRAG DROP

You develop a web application.

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

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

Select and Place:



Actions	Answer Area
Select Manifest from the middle-tier service registration.	 
In Enterprise Applications, select New application .	
Add a Cryptographic key.	 
Create a new application and provide the name, account type, and redirect URI.	
Select the Azure AD instance.	
Use an access token to access the secure resource.	
In App Registrations, select New registration .	



Correct Answer:



Actions	Answer Area
Select Manifest from the middle-tier service registration.	In App Registrations, select New registration .
In Enterprise Applications, select New application .	Select the Azure AD instance.
Add a Cryptographic key.	Create a new application and provide the name, account type, and redirect URI.
Create a new application and provide the name, account type, and redirect URI.	Create a new application and provide the name, account type, and redirect URI.
Select the Azure AD instance.	Select the Azure AD instance.
Use an access token to access the secure resource.	Use an access token to access the secure resource.
In App Registrations, select New registration .	In App Registrations, select New registration .

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Register a new application using the Azure portal

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



QUESTION 8

You have a new Azure subscription. You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (Azure AD) for authentication.

You need to implement multifactor authentication for the website.

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

B: 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.

C: 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 (Onpremises).
- Multi-Factor Authentication for Office 365
- Azure Active Directory Global Administrators

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

QUESTION 9

You are developing a Java application that uses Cassandra to store key and value data. You plan to use a new Azure Cosmos DB resource and the Cassandra API in the application. You create an Azure Active Directory (Azure AD) group named **Cosmos DB Creators** to enable provisioning of Azure Cosmos accounts, databases, and containers.

The Azure AD group must not be able to access the keys that are required to access the data.

You need to restrict access to the Azure AD group.

Which role-based access control should you use?

- A. DocumentDB Accounts Contributor
- B. Cosmos Backup Operator
- C. Cosmos DB Operator
- D. Cosmos DB Account Reader

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Azure Cosmos DB now provides a new RBAC role, Cosmos DB Operator. This new role lets you provision Azure Cosmos accounts, databases, and containers, but can't access the keys that are required to access the data. This role is intended for use in scenarios where the ability to grant access to Azure Active Directory service principals to manage deployment operations for Cosmos DB is needed, including the account, database, and containers.

Reference:

<https://azure.microsoft.com/en-us/updates/azure-cosmos-db-operator-role-for-role-based-access-control-rbac-is-now-available/>

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

You are developing 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.

Reference:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

QUESTION 11

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

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

You are developing 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. In the 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.

Reference:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

QUESTION 12

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

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

You are developing 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. In the application’s manifest, define application roles that match the required permission levels for the application.
- Assign the appropriate Azure AD group to each role. In the website, use the value of the roles claim from the JWT for the user to determine permissions.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

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.

Reference:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

QUESTION 13

You have an application that includes an Azure Web app and several Azure Function apps. Application secrets including connection strings and certificates are stored in Azure Key Vault.

Secrets must not be stored in the application or application runtime environment. Changes to Azure Active Directory (Azure AD) must be minimized.

You need to design the approach to loading application secrets.

What should you do?

- A. Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity.
- B. Create a single Azure AD Service Principal with permission to access Key Vault and use a client secret from within the App Services to access Key Vault.
- C. Create a system assigned Managed Identity in each App Service with permission to access Key Vault.
- D. Create an Azure AD Service Principal with Permissions to access Key Vault for each App Service and use a certificate from within the App Services to access Key Vault.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Use Key Vault references for App Service and Azure Functions.

Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references>

QUESTION 14

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

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

You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms.

If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised.

You need to store the intake forms according to the requirements.

Solution:

1. Create an Azure Key Vault key named **skey**.
2. Encrypt the intake forms using the public key portion of **skey**.
3. Store the encrypted data in Azure Blob storage.

Does the solution meet the goal?

- A. Yes
B. No



Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:

QUESTION 15

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

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

You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms.

If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised.

You need to store the intake forms according to the requirements.

Solution:

1. Create an Azure Cosmos DB database with Storage Service Encryption enabled.
2. Store the intake forms in the Azure Cosmos DB database.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

QUESTION 16

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

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

You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms.

If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised.

You need to store the intake forms according to the requirements.

Solution: Store the intake forms as Azure Key Vault secrets.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

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

Explanation:

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

QUESTION 17

Your company is developing an Azure API hosted in Azure.

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

- All API calls must be authenticated.
- Callers to the API must not send credentials to the API.

Which authentication mechanism should you use?

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

**Correct Answer: C****Section: (none)****Explanation****Explanation/Reference:**

Explanation:

Azure Active Directory Managed Service Identity (MSI) gives your code an automatically managed identity for authenticating to Azure services, so that you can keep credentials out of your code.

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

Incorrect Answers:

A: Use the authentication-basic policy to authenticate with a backend service using Basic authentication. This policy effectively sets the HTTP Authorization header to the value corresponding to the credentials provided in the policy. B: Anonymous is no authentication at all.

D: Your code needs credentials to authenticate to cloud services, but you want to limit the visibility of those credentials as much as possible. Ideally, they never appear on a developer's workstation or get checked-in to source control. Azure Key Vault can store credentials securely so they aren't in your code, but to retrieve them you need to authenticate to Azure Key Vault. To authenticate to Key Vault, you need a credential! A classic bootstrap problem.

Reference:

<https://azure.microsoft.com/en-us/blog/keep-credentials-out-of-code-introducing-azure-ad-managed-service-identity/>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies>

QUESTION 18

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

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

You develop Azure solutions.

You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.

You need to obtain an Azure Resource Manager access token.

Solution: Use an X.509 certificate to authenticate the VM with Azure Resource Manager.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation



Explanation/Reference:

Explanation:

Instead run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

QUESTION 19

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

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

You develop Azure solutions.

You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.

You need to obtain an Azure Resource Manager access token.

Solution: Use the Reader role-based access control (RBAC) role to authenticate the VM with Azure Resource Manager.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation



Explanation/Reference:

Explanation:

Instead run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

QUESTION 20

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

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

You are developing 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 and use Integrated Windows Authentication in the website.
- In the website, query Microsoft Graph API to load the group to which the user is a member.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Microsoft Graph is a RESTful web API that enables you to access Microsoft Cloud service resources.

Instead in the 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.

Reference: <https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

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 questions in 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 section 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 in 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. When you are ready to answer a question, click the **Question** button to return to the question.

Current environment

Windows Server 2016 virtual machine

This virtual machine (VM) runs BizTalk 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 trucking 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:



Requirements

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://test.wideworldimporters.com/' is therefore not allowed access.

QUESTION 1

DRAG DROP

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

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

Select and Place:

Actions	Answer Area
Link the Logic App to the integration account.	
Add partners, schemas, certificates, maps, and agreements.	
Update the Logic App to use the partners, schemas, certificates, maps, and agreements.	
Create a custom connector for the Logic App.	
Link the custom connector to the Logic App.	
Create an integration account in the Azure portal.	

**Correct Answer:**

Actions	Answer Area
Link the Logic App to the integration account.	Create an integration account in the Azure portal.
Add partners, schemas, certificates, maps, and agreements.	Link the Logic App to the integration account.
Update the Logic App to use the partners, schemas, certificates, maps, and agreements.	 Add partners, schemas, certificates, maps, and agreements.
Create a custom connector for the Logic App.	 Create a custom connector for the Logic App.
Link the custom connector to the Logic App.	
Create an integration account in the Azure portal.	

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

Explanation:

Step 1: Create an integration account in the Azure portal

You can define custom metadata for artifacts in integration accounts and get that metadata during runtime for your logic app to use. For example, you can provide metadata for artifacts, such as partners, agreements, schemas, and maps - all store metadata using key-value pairs.

Step 2: Link the Logic App to the integration account

A logic app that's linked to the integration account and artifact metadata you want to use.

Step 3: Add partners, schemas, certificates, maps, and agreements

Step 4: Create a custom connector for the Logic App.

Reference: <https://docs.microsoft.com/bs-latn-ba/azure/logic-apps/logic-apps-enterprise-integration-metadata>

QUESTION 2

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

What should you use?

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



Correct Answer: C

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.

Reference:

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



Connect to and consume Azure services and third-party services

Question Set 2

QUESTION 1

DRAG DROP

You have an application that provides weather forecasting data to external partners. You use Azure API Management to publish APIs.

You must change the behavior of the API to meet the following requirements:



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

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

NOTE: Each correct selection is worth one point.

Select and Place:

Policy types

- Inbound
- Outbound
- Backend

Answer Area
Requirement

Rewrite the request URL to match to the format expected by the web service.

Policy type

- policy type
- policy type
- policy type

Remove formatting text from responses.

Forward the user ID that is associated with the subscription key for the original request to the back-end service.

Correct Answer:
Policy types

- Inbound
- Outbound
- Backend

Answer Area
Requirement

Rewrite the request URL to match to the format expected by the web service.

Policy type

- Outbound
- Inbound
- Backend

Remove formatting text from responses.

Forward the user ID that is associated with the subscription key for the original request to the back-end service.

Section: (none)
Explanation
Explanation/Reference:
QUESTION 2

You are developing an e-commerce solution that uses a microservice architecture.

You need to design a communication backplane for communicating transactional messages between various parts of the solution. Messages must be communicated in first-in-first-out (FIFO) order.

What should you use?

- A. Azure Storage Queue
- B. Azure Event Hub

- C. Azure Service Bus
- D. Azure Event Grid

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

As a solution architect/developer, you should consider using Service Bus queues when:

- Your solution requires the queue to provide a guaranteed first-in-first-out (FIFO) ordered delivery.

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-azure-and-service-bus-queues-compared-contrasted>

QUESTION 3

DRAG DROP

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


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

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

NOTE: Each correct selection is worth one point.

Select and Place:

Triggers and action blocks




Answer area

```

graph TD
    A[trigger or action block] --> B[Set tier age variable]
    B --> C[List blobs]
    C --> D[For each]
    D --> E[trigger or action block]
    D --> F[if true]
    F --> G[trigger or action block]
    D --> H[if false]
    H --> I[trigger or action block]
    G --> J[Add an action]
    H --> K[Add an action]
    I --> L[Add an action]
  
```

Correct Answer:



Section: (none)

Explanation


Explanation/Reference:

Explanation:

Box 1: Reoccurrence..

To regularly run tasks, processes, or jobs on specific schedule, you can start your logic app workflow with the built-in Recurrence - Schedule trigger. You can set a date and time as well as a time zone for starting the workflow and a recurrence for repeating that workflow.

Set the interval and frequency for the recurrence. In this example, set these properties to run your workflow every week.



Box 2: Condition..

To run specific actions in your logic app only after passing a specified condition, add a conditional statement. This control structure compares the data in your workflow against specific values or fields. You can then specify different actions that run based on whether or not the data meets the condition.

Box 3: Put a message on a queue

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

Note: Under If true and If false, add the steps to perform based on whether the condition is met.

Box 4: ..tier it to Cool or Archive

tier. Archive item. Box 5: List blobs

2

Reference: <https://docs.microsoft.com/en-us/azure/connectors/connectors-native-recurrence>

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-control-flow-loops>

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-control-flow-conditional-statement>

QUESTION 4

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

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

You are developing an Azure Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.

You have the following requirements:

- Queue size must not grow larger than 80 gigabytes (GB).
- Use first-in-first-out (FIFO) ordering of messages.

Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Service Bus Queue from the mobile application. Create an Azure Function App that uses an Azure Service Bus Queue trigger.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

Explanation:

You can create a function that is triggered when messages are submitted to an Azure Storage queue.

Reference: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function>

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

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

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

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

Solution: Provision an Azure Notification Hub. Register all devices with the hub.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

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 that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

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 that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

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

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

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

Solution: Provision an Azure Event Grid. Configure event filtering to evaluate the device identifier.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

Reference:

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

QUESTION 8

DRAG DROP

You manage several existing Logic Apps.

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

**Answer Area**

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

Correct Answer:

Answer Area

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

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

Explanation:

Box 1: Enterprise Integration Pack

For business-to-business (B2B) solutions and seamless communication between organizations, you can build automated scalable enterprise integration workflows by using the Enterprise Integration Pack (EIP) with Azure Logic Apps.

Box 2: Code View Editor

Edit JSON - Azure portal

1. Sign in to the Azure portal.
2. From the left menu, choose All services. In the search box, find "logic apps", and then from the results, select your logic app.
3. On your logic app's menu, under Development Tools, select Logic App Code View.
4. The Code View editor opens and shows your logic app definition in JSON format. Box 3: Logic Apps Designer

Reference:

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

QUESTION 9

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?

az servicebus queue create A.
--resource-group fridge-rg
--namespace-name fridge-ns
--name fridge-q

New-AzureRmResourceGroup
-Name fridge-rg
-Location fridge-loc

B.

```
az servicebus namespace create  
--resource-group fridge-rg  
--name fridge-ns  
--location fridge-loc  
  
connectionString-$)az serviceBus namespace authorization-rule keys list  
--resource-group fridge-rg  
--fridge-ns fridge-ns  
--query primaryConnectionString -output tsv)
```

C.

D.

Correct Answer: A

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)
```

Reference:

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

QUESTION 10

HOTSPOT

You are developing an application that uses Azure Storage Queues.

You have the following code:

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse  
(CloudConfigurationManager.GetSetting("StorageConnectionString"));  
CloudQueueClient queueClient = storageAccount.CreateCloudQueueClient();  
  
CloudQueue queue = queueClient.GetQueueReference("appqueue");  
await queue.CreateIfNotExistsAsync();  
  
CloudQueueMessage peekedMessage = await queue.PeekMessageAsync();  
if (peekedMessage != null)  
{  
    Console.WriteLine("The peeked message is: {0}", peekedMessage.AsString);  
}  
CloudQueueMessage message = await queue.GetMessageAsync() ;
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statement	Yes	No
The code configures the lock duration for the queue.	<input type="radio"/>	<input type="radio"/>
The last message read remains in the queue after the code runs.	<input type="radio"/>	<input type="radio"/>
The storage queue remains in the storage account after the code runs.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Statement	Yes	No
The code configures the lock duration for the queue.	<input type="radio"/>	<input checked="" type="radio"/>
The last message read remains in the queue after the code runs.	<input checked="" type="radio"/>	<input type="radio"/>
The storage queue remains in the storage account after the code runs.	<input checked="" type="radio"/>	<input type="radio"/>

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: No

The QueueDescription.LockDuration property gets or sets the duration of a peek lock; that is, the amount of time that the message is locked for other receivers. The maximum value for LockDuration is 5 minutes; the default value is 1 minute.

Box 2: Yes

You can peek at the message in the front of a queue without removing it from the queue by calling the PeekMessage method.

Box 3: Yes

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

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicebus.messaging.queuedescription.lockduration>

QUESTION 11

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?

```
az group create
  --name fridge-rg
  --location fridge-loc

New-AzureRmServiceBusNamespace
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Location fridge-loc

New-AzureRmServiceBusQueue
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Name fridge-q
  -EnablePartitioning $False

az servicebus namespace create
  --resource-group fridge-rg
  --name fridge-rg
  --location fridge-loc
```

A.



B.

C.

D.

Correct Answer: C

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)
```

Reference:

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

QUESTION 12

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

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

You are developing an Azure Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.

You have the following requirements:

- Queue size must not grow larger than 80 gigabytes (GB).
- Use first-in-first-out (FIFO) ordering of messages.

Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application. Create an Azure Function App that uses an Azure Storage Queue trigger.



Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Create an Azure Function App that uses an Azure Service Bus Queue trigger.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function>

QUESTION 13

DRAG DROP

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

1. A driver selects the restaurants for which they will deliver orders.
2. Orders are sent to all available drivers in an area.
3. Only orders for the selected restaurants will appear for the driver.
4. The first driver to accept an order removes it from the list of available orders.

You need to implement an Azure Service Bus solution.

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

Select and Place:

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

Correct Answer:

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

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Create a single Service Bus Namespace

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

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

Create topics.

Box 3: Create a Service Bus subscription for each restaurant for which a driver can receive orders.

Topics can have multiple, independent subscriptions.

Reference:

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

QUESTION 14

HOTSPOT

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

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

You need to implement push notifications.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
string notificationHubName = "contoso_hub";
string notificationHubConnection = "connection_string";
    hub =
NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

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

Correct Answer:

Answer Area

```
string notificationHubName = "contoso_hub";
string notificationHubConnection = "connection_string";
hub =
NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

GetInstallation
CreateClientFromConnectionString
CreateOrUpdateInstallation
PatchInstallation

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

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

Explanation:

Box 1: NotificationHubClient

Box 2: NotificationHubClient

Box 3: CreateClientFromConnectionString

```
// Initialize the Notification Hub
NotificationHubClient hub = NotificationHubClient.CreateClientFromConnectionString(listenConnString, hubName);
```

Box 4: SendWindowsNativeNotificationAsync Send the push notification.
var result
= await hub.SendWindowsNativeNotificationAsync(windowsToastPayload);

Reference: <https://docs.microsoft.com/en-us/azure/notification-hubs/notification-hubs-push-notification-registration-management>

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/app-service-mobile/app-service-mobile-windows-store-dotnet-get-started-push.md>

QUESTION 15

You are developing an Azure messaging solution.

You need to ensure that the solution meets the following requirements:

- Provide transactional support.
- Provide duplicate detection.
- Store the messages for an unlimited period of time.

Which two technologies will meet the requirements? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Azure Service Bus Topic
- B. Azure Service Bus Queue
- C. Azure Storage Queue
- D. Azure Event Hub



Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Azure Service Bus Queue and Topic has duplicate detection.

Enabling duplicate detection helps keep track of the application-controlled MessageId of all messages sent into a queue or topic during a specified time window.

Incorrect Answers:

C: There is just no mechanism that can query a Storage queue and find out if a message with the same contents is already there or was there before.

D: Azure Event Hub does not have duplicate detection

Reference:

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

QUESTION 16

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

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

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: AC

Section: (none)

Explanation



Explanation/Reference:

Explanation:

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

Reference:

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

QUESTION 17

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

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

You need to ensure that the subscription client processes all messages.

Which code segment should you use?

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

Correct Answer: D

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);
```

Reference:

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

QUESTION 18

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

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

You are developing an Azure Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.

You have the following requirements:

- Queue size must not grow larger than 80 gigabytes (GB).
- Use first-in-first-out (FIFO) ordering of messages. ▪ Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application. Create an Azure VM that is triggered from Azure Storage Queue events.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Don't use a VM, instead create an Azure Function App that uses an Azure Service Bus Queue trigger.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function>

QUESTION 19

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

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

You are developing an Azure Service application that processes queue data when it receives a message from a mobile application. Messages may not be sent to the service consistently.

You have the following requirements:

- Queue size must not grow larger than 80 gigabytes (GB).
- Use first-in-first-out (FIFO) ordering of messages. ▪

Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Service Bus Queue from the mobile application. Create an Azure Windows VM that is triggered from Azure Service Bus Queue.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Don't use a VM, instead create an Azure Function App that uses an Azure Service Bus Queue trigger.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function>

QUESTION 20

DRAG DROP

You are developing a REST web service. Customers will access the service by using an Azure API Management instance.

The web service does not correctly handle conflicts. Instead of returning an HTTP status code of 409, the service returns a status code of 500. The body of the status message contains only the word conflict.

You need to ensure that conflicts produce the correct response.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Policy segments	Answer Area
server	< Policy segment >
context	<base />
on-error	<choose>
set-status	<when condition = " @ [Policy segment].Response.StatusCode == 500
when-error	&& [Policy segment].LastError.Message.Contains
override-status	" conflict = "))" >

Correct Answer:

Policy segments	Answer Area
server	< [on-error] >
context	<base />
on-error	<choose>
set-status	<when condition = " @ [context].Response.StatusCode == 500 && [context].LastError.Message.Contains " conflict = ")) " >
when-error	<return-response>
override-status	< [set-status] >

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

Explanation:

Box 1: on-error

Policies in Azure API Management are divided into inbound, backend, outbound, and on-error.

If there is no on-error section, callers will receive 400 or 500 HTTP response messages if an error condition occurs.

Box 2: context

Box 3: context

Box 4: set-status

The return-response policy aborts pipeline execution and returns either a default or custom response to the caller. Default response is 200 OK with no body.

Custom response can be specified via a context variable or policy statements.

Syntax:

```
<return-response response-variable-name="existing context variable">  
  <set-header/>  
  <set-body/>  
  <set-status/>  
</return-response>
```

Box 5: on-error



Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-error-handling-policies> <https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies>

