

Microsoft.AZ-204.v2023-06-14.q167

Exam Code:	AZ-204
Exam Name:	Developing Solutions for Microsoft Azure
Certification Provider:	Microsoft
Free Question Number:	167
Version:	v2023-06-14
# of views:	1724
# of Questions views:	1670
https://www.freepdfdumps.com/Microsoft.AZ-204.v2023-06-14.q167.html	

NEW QUESTION: 1

You develop an image upload service that is exposed using Azure API Management. Images are analyzed after upload for automatic tagging.

Images over 500 KB are processed by a different backend that offers a lower tier of service that costs less money. The lower tier of service is denoted by a header named `x-lsrSe-request`.

Images over 500 KB must never be processed by backends for smaller images and must always be charged the lower price.

You need to implement API Management policies to ensure that images are processed correctly. How should you complete the API Management inbound policy? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
<inbound>
  <base/>
  <set-variable name="imageSize" value="@context.Request.Headers["Content-Length"][0]"/>
  <choose>
    <when condition="@int.Parse(context.Variables.GetValueOrDefault<string>("imageSize"))>512000">
      <set-header name="x-large-request" exists-action="delete"/>
      <value>true</value>
    </set-header>
    </when>
    <otherwise>
      <set-backend-service base-url="{{large-image-host}}"/>
    </set-backend-service>
  </choose>
</inbound>
```

Answer:

Answer Area

```
<inbound>
<base/>
<set-variable name="imageSize" value="@{context.Request.Headers["Content-Length"]}[0]"/>
<choose>
  <when condition="@{int.Parse(context.Variables.GetValueOrDefault<string>("imageSize"))<512000}">
    <set-header name="x-large-request" exists-action="delete">
      <value>true</value>
    </set-header>
  </when>
  <otherwise>
    <set-backend-service base-url="{{large-image-host}}"/>
  </otherwise>
</choose>
</inbound>
```

Explanation

Answer Area

```
<inbound>
<base/>
<set-variable name="imageSize" value="@{context.Request.Headers["Content-Length"]}[0]"/>
<choose>
  <when condition="@{int.Parse(context.Variables.GetValueOrDefault<string>("imageSize"))<512000}">
    <set-header name="x-large-request" exists-action="delete">
      <value>true</value>
    </set-header>
  </when>
  <otherwise>
    <set-backend-service base-url="{{large-image-host}}"/>
  </otherwise>
</choose>
</inbound>
```

NEW QUESTION: 2

You need to update the APIs to resolve the testing error.

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

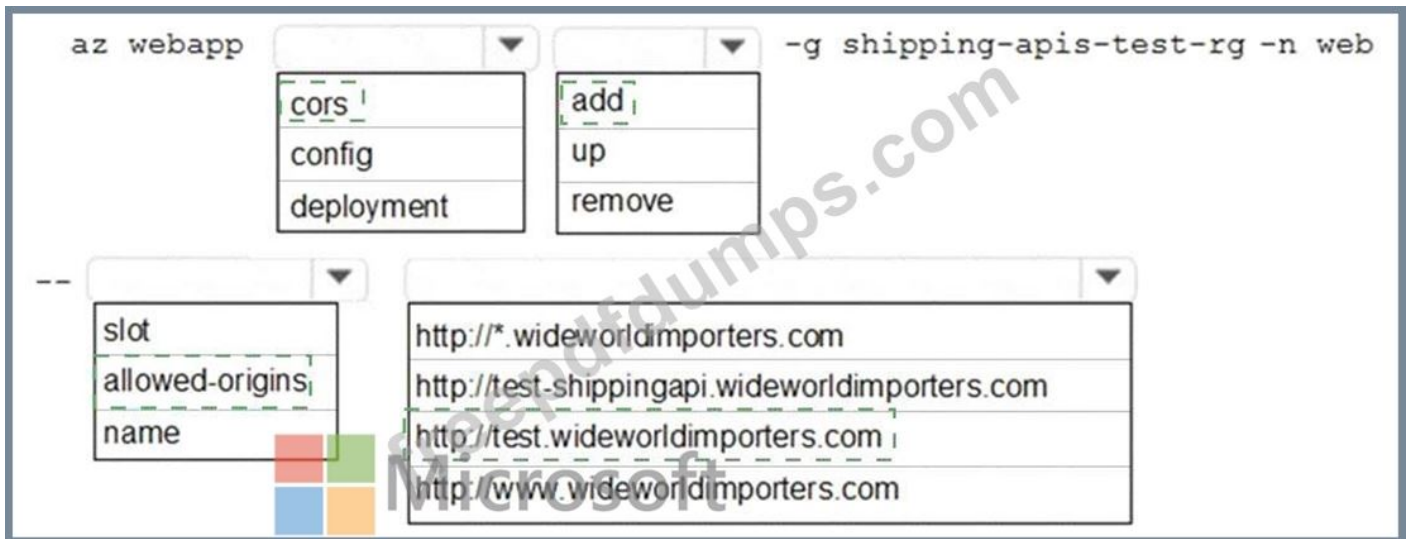
NOTE: Each correct selection is worth one point.

```
az webapp cors add --slot  --allowed-origins  --name  shipping-apis-test-rg -n web
```

<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>

<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>
<input type="text"/>	<input type="text"/>

Answer:



Explanation



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

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

Box 1: cors

Box 2: add

Box 3: allowed-origins

Box

4: http://testwideworldimporters.com/

References:

<http://donovanbr>

[own.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-with-Azure-App-Service](http://donovanbr.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-with-Azure-App-Service)

NEW QUESTION: 3

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

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

You 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

Answer: B (LEAVE A REPLY)

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>

NEW QUESTION: 4

You are developing an ASP.NET Core website that can be used to manage photographs which are stored in Azure Blob Storage containers.

Users of the website authenticate by using their Azure Active Directory (Azure AD) credentials.

You implement role-based access control (RBAC) role permissions on the containers that store photographs.

You assign users to RBAC roles.

You need to configure the website's Azure AD Application so that user's permissions can be used with the Azure Blob containers.

How should you configure the application? To answer, drag the appropriate setting to the correct location.

Each setting can 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.

Settings

client_id
profile
delegated
application
user_impersonation

Answer Area

API	Permission	Type
Azure Storage	Setting	Setting
Microsoft Graph	User.Read	Setting



Answer:

Settings

client_id
profile
delegated
application
user_impersonation

Answer Area



Microsoft

API	Permission	Type
Azure Storage	user_impersonation	delegated
Microsoft Graph	User.Read	delegated

Explanation

API	Permission	Type
Azure Storage	user_impersonation	delegated
Microsoft Graph	User.Read	delegated

Box 1: user_impersonation

Box 2: delegated

Example:

1. Select the API permissions section
2. Click the Add a permission button and then:
Ensure that the My APIs tab is selected
3. In the list of APIs, select the API TodoListService-aspnetcore.
4. In the Delegated permissions section, ensure that the right permissions are checked:
user_impersonation.
5. Select the Add permissions button.

Box 3: delegated

Example

1. Select the API permissions section
2. Click the Add a permission button and then,
Ensure that the Microsoft APIs tab is selected
3. In the Commonly used Microsoft APIs section, click on Microsoft Graph
4. In the Delegated permissions section, ensure that the right permissions are checked:
User.Read. Use the search box if necessary.
5. Select the Add permissions button

Reference:

<https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect->

NEW QUESTION: 5

You are developing a web application that makes calls to the Microsoft Graph API. You register the application in the Azure portal and upload a valid X509 certificate.

You create an appsettings.json file containing the certificate name, client identifier for the application, and the tenant identifier of the Azure active Directory (Azure AD). You create a method named ReadCertificate to return the X509 certificate by name.

You need to implement code that acquires a token by using the certificate.

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.

```
AuthenticationConfig config = AuthenticationConfig.ReadFromJsonFile("appsettings.json");
X509Certificate2 certificate = ReadCertificate(config.CertificateName);
var app = .Create(config.ClientId)


ConfidentialClientApplicationBuilder
GetAccountAsync()
GetAccountsAsync()
ConfidentialClientApplication

.WithCertificate(certificate)
.WithAuthority(new Uri(config.Authority))
.Build();
string[] scopes = new string[] { $"{config.ApiUrl}.default" };
AuthenticationResult result = await app.AcquireTokenForClient().ExecuteAsync();
```

scopes
 app
 config

Answer:

```
AuthenticationConfig config = AuthenticationConfig.ReadFromJsonFile("appsettings.json");
X509Certificate2 certificate = ReadCertificate(config.CertificateName);
var app = .Create(config.ClientId)


ConfidentialClientApplicationBuilder
GetAccountAsync()
GetAccountsAsync()
ConfidentialClientApplication

.WithCertificate(certificate)
.WithAuthority(new Uri(config.Authority))
.Build();
string[] scopes = new string[] { $"{config.ApiUrl}.default" };
AuthenticationResult result = await app.AcquireTokenForClient().ExecuteAsync();
```

scopes
 app
 config

Explanation

Graphical user interface, text, application, email Description automatically generated

```
AuthenticationConfig config = AuthenticationConfig.ReadFromJsonFile("appsettings.json");
X509Certificate2 certificate = ReadCertificate(config.CertificateName);
var app = .Create(config.ClientId)


ConfidentialClientApplicationBuilder
GetAccountAsync()
GetAccountsAsync()
ConfidentialClientApplication

.WithCertificate(certificate)
.WithAuthority(new Uri(config.Authority))
.Build();
string[] scopes = new string[] { $"{config.ApiUrl}.default" };
AuthenticationResult result = await app.AcquireTokenForClient().ExecuteAsync();
```

scopes
 app
 config

<https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-app-configuration?tabs=dotne>

<https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-acquire-token?tabs=dotnet#ac>

NEW QUESTION: 6

You are developing an ASP.NET Core Web API web service. The web service uses Azure Application Insights for all telemetry and dependency tracking. The web service reads and writes data to a database other than Microsoft SQL Server.

You need to ensure that dependency tracking works for calls to the third-party database.

Which two Dependency Telemetry properties should you store in the database? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Telemetry.Context.Operation.Id
- B. Telemetry.Context.Cloud.RoleInstance
- C. Telemetry.Id
- D. Telemetry.ContextSession.Id
- E. Telemetry.Name

Answer: A,C (LEAVE A REPLY)

Explanation

References:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking> Example:

```
public async Task Enqueue(string payload)
```

```
{
```

```
// StartOperation is a helper method that initializes the telemetry item
```

```
// and allows correlation of this operation with its parent and children.
```

```
var operation = telemetryClient.StartOperation<DependencyTelemetry>("enqueue " +  
queueName); operation.Telemetry.Type = "Azure Service Bus"; operation.Telemetry.Data =  
"Enqueue " + queueName; var message = new BrokeredMessage(payload);
```

```
// Service Bus queue allows the property bag to pass along with the message.
```

```
// We will use them to pass our correlation identifiers (and other context)
```

```
// to the consumer.
```

```
message.Properties.Add("ParentId", operation.Telemetry.Id);
```

```
message.Properties.Add("RootId", operation.Telemetry.Context.Operation.Id);
```

```
Reference:  
https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking
```

NEW QUESTION: 7

You are developing an application to store business-critical data in Azure Blob storage. The application must meet the following requirements:

- * Data must not be modified or deleted for a user-specified interval.
- * Data must be protected from overwrites and deletes.

* Data must be written once and allowed to be read many times.

You need to protect the data in the Azure Blob storage account.

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

NOTE: Each correct selection is worth one point.

- A. Enable point-in-time restore for containers in the storage account.
- B. Create an account shared-access signature (SAS).
- C. Enable version-level immutability support for the storage account.
- D. Enable the blob change feed for the storage account.
- E. Create a service shared-access signature (SAS).

Answer: D,E (LEAVE A REPLY)

NEW QUESTION: 8

You are preparing to deploy a medical records application to an Azure virtual machine (VM). The application will be deployed by using a VHD produced by an on-premises build server.

You need to ensure that both the application and related data are encrypted during and after deployment to Azure.

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

Actions

- Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage.
- Run the Azure PowerShell command `Set-AzureRmVMDiskEncryptionExtension`.
- Run the Azure PowerShell command `Set-AzureRmVMOSDisk`.
- Encrypt the on-premises VHD by using BitLocker with a TPM. Upload the VM to Azure Storage.
- Run the Azure PowerShell command `New-AzureRmVM`.

Answer area

Microsoft

freepdfmumps.com

Answer:

Explanation

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

Run the Azure PowerShell command `Set-AzureRmVMOSDisk`.

Run the Azure PowerShell command `Set-AzureRmVMDiskEncryptionExtension`.

Step 1: Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage
 Step 2: Run the Azure PowerShell command `Set-AzureRmVMOSDisk` To use an existing disk instead of creating a new disk you can use the `Set-AzureRmVMOSDisk` command.

Example:

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

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

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

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

Step 3: Run the Azure PowerShell command `Set-AzureRmVMDiskEncryptionExtension` Use the `Set-AzVMDiskEncryptionExtension` cmdlet to enable encryption on a running IaaS virtual machine in Azure.

Incorrect:

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

References:

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

NEW QUESTION: 9

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

Answer: D (LEAVE A REPLY)

Explanation

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

Reference:

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A (LEAVE A REPLY)

Explanation

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

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

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

Reference:

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

NEW QUESTION: 11

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.

Setting	Value				
Authorization level	<div data-bbox="786 775 1518 835">▼</div> <table border="1"><tr><td data-bbox="803 846 971 891">Function</td></tr><tr><td data-bbox="803 913 1031 958">Anonymous</td></tr><tr><td data-bbox="803 981 925 1025">Admin</td></tr></table>	Function	Anonymous	Admin	
Function					
Anonymous					
Admin					
User claims	<div data-bbox="786 1088 1518 1149">▼</div> <table border="1"><tr><td data-bbox="803 1160 1261 1205">JSON Web Token (JWT)</td></tr><tr><td data-bbox="803 1227 1518 1272">Shared Access Signature (SAS) token</td></tr><tr><td data-bbox="803 1294 942 1339">API Key</td></tr></table>	JSON Web Token (JWT)	Shared Access Signature (SAS) token	API Key	
JSON Web Token (JWT)					
Shared Access Signature (SAS) token					
API Key					
Trigger type	<div data-bbox="786 1413 1518 1473">▼</div> <table border="1"><tr><td data-bbox="803 1485 892 1529">blob</td></tr><tr><td data-bbox="803 1552 925 1597">HTTP</td></tr><tr><td data-bbox="803 1619 930 1664">queue</td></tr><tr><td data-bbox="803 1686 908 1731">timer</td></tr></table>	blob	HTTP	queue	timer
blob					
HTTP					
queue					
timer					

Answer:

Setting	Value
Authorization level	<input type="text" value="Function"/> <ul style="list-style-type: none"> Function Anonymous Admin
User claims	<input type="text" value="JSON Web Token (JWT)"/> <ul style="list-style-type: none"> JSON Web Token (JWT) Shared Access Signature (SAS) token API Key
Trigger type	<input type="text" value="HTTP"/> <ul style="list-style-type: none"> blob HTTP queue timer

Explanation



Setting Microsoft

Authorization level

Value

Value
Function
Anonymous
Admin

User claims

Value
JSON Web Token (JWT)
Shared Access Signature (SAS) token
API Key

Trigger type

Value
blob
HTTP
queue
timer

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

Box 1: Function

Box 2: JSON based Token (JWT)

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

Box 3: HTTP

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

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

References:

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

NEW QUESTION: 12

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.

YAML segments

secret

envVar

secretValues

volumes

volumeMounts

environmentVariables

Answer Area

```

YAML segment :
- mountPath: /mnt/secrets
  name: accesskey
YAML segment :
  name: accesskey
YAML segment :
  key: TXkgZmlyc3Qgc2VjcmV0IEZPTwo=
                    
```

Answer:

YAML segments

secret

envVar

secretValues

volumes

volumeMounts

environmentVariables

Answer Area

```

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

Explanation

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

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>
Topic 2, Windows Server 2016 virtual machine

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 button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons

displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the button to return to the question.

Current environment

Windows Server 2016 virtual machine

The 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 tracking 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.

NEW QUESTION: 13

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

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

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

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

NOTE: Each correct selection is worth one point.

```
public class PlayerEntity : TableEntity
```

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

<input type="text"/>	▼
email	
phone	
region	


<input type="text"/>	▼
email	
phone	
region	

```
protected PlayerEntity player;
```

```
async void GetPlayer(string cs,  table, string pk, string rk)
```

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

```
{
    
    TEntity query = TEntity.Retrieve<PlayerEntity>(pk, rk);
    TableOperation query = TableOperation.Retrieve<PlayerEntity>(pk,rk);
    TableResult query = TableQuery.Retrieve<PlayerEntity>(pk,rk);
    TableResultSegment query = TableResult.Retrieve<PlayerEntity>(pk, rk);
```

 Microsoft	▼
TEntity data =await table.ExecuteAsync(query);	
TableOperation data =await table.ExeucteAsync(query);	
TableQuery data =await table.ExecuteAsync(query);	
TableResult data =await table.ExecuteAsync(query);	

```
player=data.Result as PlayerEntity;
}
}
```

Answer:

```

public class PlayerEntity : TableEntity
{
    public PlayerEntity()
    {
    }
    public PlayerEntity(string region, string email)
    {
        PartitionKey =  ;
        RowKey =  ;
        public string Phone { get; set; }
    }
    public class Player
    {
        protected PlayerEntity player;
        async void GetPlayer(string cs,  table, string pk, string rk)
        {
            
            TEntity query = TEntity.Retrieve<PlayerEntity>(pk, rk);
            TOperation query = TOperation.Retrieve<PlayerEntity>(pk, rk);
            TResult query = TResult.Retrieve<PlayerEntity>(pk, rk);
            TSegment query = TSegment.Retrieve<PlayerEntity>(pk, rk);

            TEntity data = await table.ExecuteAsync(query);
            TOperation data = await table.ExeucteAsync(query);
            TResult data = await table.ExecuteAsync(query);
            TSegment data = await table.ExecuteAsync(query);
            player = data.Result as PlayerEntity;
        }
    }
}

```

Explanation

Answer Area

```
public class PlayerEntity : TableEntity
{
    public PlayerEntity()
    {
    }
    public PlayerEntity(string region, string email)
    {
        PartitionKey =  ;
        RowKey =  ;
        public string Phone { get; set; }
    }
    public class Player
    {
        protected PlayerEntity player;
        async void GetPlayer(string cs,  table, string pk, string rk)
    }
}
```

Explanation:

```
{
    TEntity query = TEntity.Retrieve<PlayerEntity>(pk, rk);
    TableOperation query = TableOperation.Retrieve<PlayerEntity>(pk, rk);
    TableResult query = TableQuery.Retrieve<PlayerEntity>(pk, rk);
    TableResultSegment query = TableResult.Retrieve<PlayerEntity>(pk, rk);
    TEntity data = await table.ExecuteAsync(query);
    TableOperation data = await table.ExeucteAsync(query);
    TableQuery data = await table.ExecuteAsync(query);
    TableResult data = await table.ExecuteAsync(query);
    player = data.Result as PlayerEntity;
}
```

Box 1: region

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

Choosing the PartitionKey.

The core of any table's design is based on its scalability, the queries used to access it, and storage operation requirements. The PartitionKey values you choose will dictate how a table will

be partitioned and the type of queries that can be used. Storage operations, in particular inserts, can also affect your choice of PartitionKey values.

Box 2: email

Not phone number some players may not have a phone number.

Box 3: CloudTable

Box 4 : TableOperation query =..

Box 5: TableResult

References:

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

NEW QUESTION: 14

You are developing an Azure Function App. You develop code by using a language that is not supported by the Azure Function App host. The code language supports HTTP primitives.

You must deploy the code to a production Azure Function App environment.

You need to configure the app for deployment.

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

NOTE: Each correct selection is worth one point.

Configuration parameter

Configuration value

Publish

	▼
Code	
Docker Container	

Runtime stack

	▼
Node.js	
Python	
PowerShell Core	
Custom Handler	

Version

	▼
14 LTS	
7.0	
custom	

Answer:

Configuration parameter	Configuration value
Publish	<div style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <div style="padding: 2px;">Code</div> <div style="padding: 2px; border: 1px dashed gray;"> Docker Container </div> </div>
Runtime stack	<div style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <div style="padding: 2px;">Node.js</div> <div style="padding: 2px;">Python</div> <div style="padding: 2px; border: 1px dashed gray;"> PowerShell Core </div> <div style="padding: 2px;">Custom Handler</div> </div>
Version	<div style="border: 1px solid black; padding: 2px;"> <div style="text-align: right; padding-right: 5px;">▼</div> <div style="padding: 2px;">14 LTS</div> <div style="padding: 2px; border: 1px dashed gray;"> 7.0 </div> <div style="padding: 2px;">custom</div> </div>

Explanation

Box 1: Docker container

A custom handler can be deployed to every Azure Functions hosting option. If your handler requires operating system or platform dependencies (such as a language runtime), you may need to use a custom container. You can create and deploy your code to Azure Functions as a custom Docker container.

Box 2: PowerShell core

When creating a function app in Azure for custom handlers, we recommend you select .NET Core as the stack.

A "Custom" stack for custom handlers will be added in the future.

PowerShell Core (PSC) is based on the new .NET Core runtime.

Box 3: 7.0

On Windows: The Azure Az PowerShell module is also supported for use with PowerShell 5.1 on Windows.

On Linux: PowerShell 7.0.6 LTS, PowerShell 7.1.3, or higher is the recommended version of PowerShell for use with the Azure Az PowerShell module on all platforms.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-function-linux-custom-image>

<https://docs.microsoft.com/en-us/powershell/azure/install-az-ps?view=azps-7.1.0>

NEW QUESTION: 15

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

You have the following parameters:

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

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

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

Command segments

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

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

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

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

Answer:

Command segments	Answer Area
<pre>az batch pool create --id mypool --vm-size Standard_A1_v2 --target-dedicated-nodes 2 --image \$image --node-agent-sku-id \$sku</pre>	<pre>az batch pool create --id mypool --vm-size Standard_A1_v2 --target-dedicated-nodes 2 --image \$image --node-agent-sku-id \$sku</pre>
<pre>az batch job create --id myjob --pool-id mypool</pre>	<pre>az batch task create --task-id mytask\$i --job-id myjob --command-line \$script</pre>
<pre>for i in {1..\$numberOfJobs} do</pre>	<pre>az batch job create --id myjob --pool-id mypool</pre>
<pre>az batch task create --task-id mytask\$i --job-id myjob --command-line \$script</pre>	<pre>for i in {1..\$numberOfJobs} do</pre>

Explanation

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

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

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

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

Step 1: az batch pool create

Create a new Linux pool with a virtual machine configuration.

```
az batch pool create \
  --id mypool \
  --vm-size Standard_A1 \
  --target-dedicated 2 \
  --image canonical:ubuntu-server:16.04-LTS \
  --node-agent-sku-id "batch.node.ubuntu 16.04"
```

Step 2: az batch job create

Create a new job to encapsulate the tasks that are added.

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

Step 3: az batch task create

Add tasks to the job. Here the task is a basic shell command.

```
az batch task create \
```

```
--job-id myjob \  
--task-id task1 \  
--command-line "/bin/bash -c 'printenv AZ_BATCH_TASK_WORKING_DIR'"
```

Step 4: for i in {1..\$numberOfJobs} do

References:

<https://docs.microsoft.com/bs-latn-ba/azure/batch/scripts/batch-cli-sample-run-job>

NEW QUESTION: 16

You need to ensure receipt processing occurs correctly.

What should you do?

- A. Use blob properties to prevent concurrency problems
- B. Use blob SnapshotTime to prevent concurrency problems
- C. Use blob metadata to prevent concurrency problems
- D. Use blob leases to prevent concurrency problems

Answer: (SHOW ANSWER)

Explanation

You can create a snapshot of a blob. A snapshot is a read-only version of a blob that's taken at a point in time.

Once a snapshot has been created, it can be read, copied, or deleted, but not modified.

Snapshots provide a way to back up a blob as it appears at a moment in time.

Scenario: Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime.

Once processing is completed, results are stored in Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.

Reference:

<https://docs.microsoft.com/en-us/rest/api/storageservices/creating-a-snapshot-of-a-blob>

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam!

Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 17


You develop a containerized application. You plan to deploy the application to a new Azure Container instance by using a third-party continuous integration and continuous delivery (CI/CD) utility.

The deployment must be unattended and include all application assets. The third-party utility must only be able to push and pull images from the registry. The authentication must be managed by Azure Active Directory (Azure AD). The solution must use the principle of least privilege. You need to ensure that the third-party utility can access the registry. Which authentication options should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Authentication	Option
Registry authentication method	<div style="border: 1px solid #ccc; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid #ccc;">▼</div> <div style="padding: 2px 5px;">Service principal</div> <div style="padding: 2px 5px;">Individual identity</div> <div style="padding: 2px 5px;">Repository-scoped access token</div> <div style="padding: 2px 5px;">Managed identity for Azure resources</div> </div>
RBAC role	<div style="border: 1px solid #ccc; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid #ccc;">▼</div> <div style="padding: 2px 5px;">AcrPull</div> <div style="padding: 2px 5px;">Owner</div> <div style="padding: 2px 5px;">AcrPush</div> <div style="padding: 2px 5px;">Contributor</div> </div>

Answer:

Authentication	Option
Registry authentication method	<div style="border: 1px solid #ccc; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid #ccc;">▼</div> <div style="padding: 2px 5px; border: 1px dashed green;">Service principal</div> <div style="padding: 2px 5px;">Individual identity</div> <div style="padding: 2px 5px;">Repository-scoped access token</div> <div style="padding: 2px 5px;">Managed identity for Azure resources</div> </div>
RBAC role	<div style="border: 1px solid #ccc; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid #ccc;">  ▼ </div> <div style="padding: 2px 5px;">AcrPull</div> <div style="padding: 2px 5px;">Owner</div> <div style="padding: 2px 5px; border: 1px dashed green;">AcrPush</div> <div style="padding: 2px 5px;">Contributor</div> </div>

Explanation

Graphical user interface, text, application, table Description automatically generated



Option

Service principal
Individual identity
Repository-scoped access token
Managed identity for Azure resources

RBAC role

AcrPull
Owner
AcrPush
Contributor

Box 1: Service principal

Applications and container orchestrators can perform unattended, or "headless," authentication by using an Azure Active Directory (Azure AD) service principal.

Box 2: AcrPush

AcrPush provides pull/push permissions only and meets the principle of least privilege.

Reference:

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-authentication?tabs=azure-cli>

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-roles?tabs=azure-cli>

NEW QUESTION: 18

You are developing applications for a company. You plan to host the applications on Azure App Services.

The company has the following requirements:

- * Every five minutes verify that the websites are responsive.
- * Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly.
- * Generate alerts if a website is experiencing issues.
- * If a website fails to load, the system must attempt to reload the site three more times.

You need to implement this process with the least amount of effort.

What should you do?

D18912E1457D5D1DDCCBD40AB3BF70D5D

- A. Create a Selenium web test and configure it to run from your workstation as a scheduled task.
- B. Set up a URL ping test to query the home page.
- C. Create an Azure function to query the home page.
- D. Create a multi-step web test to query the home page.

E. Create a Custom Track Availability Test to query the home page.

Answer: D (LEAVE A REPLY)

Explanation

You can monitor a recorded sequence of URLs and interactions with a website via multi-step web tests.

Reference:

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

NEW QUESTION: 19

You develop and deploy a web app to Azure App Service in a production environment. You scale out the web app to four instances and configure a staging slot to support changes.

You must monitor the web app in the environment to include the following requirements:

- * Increase web app availability by re-routing requests away from instances with error status codes and automatically replace instances if they remain in an error state after one hour.
- * Send web server logs, application logs, standard output and standard error messaging to an Azure Storage blob account.

You need to configure Azure App Service.

Which values should you use? To answer, drag the appropriate configuration value to the correct requirements.

Each configuration value may be used once, more than...

Configuration values	Requirement	Configuration value
Health check	Increase availability	
Diagnostic setting	Send logs	
Deployment slot		
Autoscale rule		
Zone redundancy		

Answer:

Configuration values	Requirement	Configuration value
Health check	Increase availability	Autoscale rule
Diagnostic setting	Send logs	Zone redundancy
Deployment slot		
Autoscale rule		
Zone redundancy		

Explanation



NEW QUESTION: 20

You need to secure the Azure Functions to meet the security requirements.

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

NOTE: Each correct selection is worth one point.

- A. Store the RSA-HSM key in Azure Cosmos DB. Apply the built-in policies for customer-managed keys and allowed locations.
- B. Create a free tier Azure App Configuration instance with a new Azure AD service principal.
- C. Store the RSA-HSM key in Azure Key Vault with soft-delete and purge-protection features enabled.
- D. Store the RSA-HSM key in Azure Blob storage with an Immutability policy applied to the container.
- E. Create a standard tier Azure App Configuration instance with an assigned Azure AD managed identity.

Answer: ([SHOW ANSWER](#))

Explanation

Scenario: All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.

Microsoft Azure Key Vault is a cloud-hosted management service that allows users to encrypt keys and small secrets by using keys that are protected by hardware security modules (HSMs). You need to create a managed identity for your application.

Reference:

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

NEW QUESTION: 21

You develop and deploy an Azure App Service web app that connects to Azure Cache for Redis as a content cache. An resources have been deployed to East US 2 region.

The security team requires the from Azure Cache for Redis:

The number of Redis client connections from an associated IP address.

Redis operations completed on the content cache.

The location (region) in which the Azure Cache for Redis instance was accessed.

The audit information must be captured and analyzed by a security team application deployed to Central US region You need to log information on all client corrections to the cache.

Which configuration values should you use?

Requirement	Configuration value
Store log information.	<ul style="list-style-type: none">Log Analytics workspaceBlob Storage accountData Lake Storage Gen2 Storage accountEvent hub
Enable client connection logging.	<ul style="list-style-type: none">Diagnostic settingManaged identityApp registrationEnvironment variable

Answer:

Requirement	Configuration value
Store log information.	<ul style="list-style-type: none">Log Analytics workspaceBlob Storage accountData Lake Storage Gen2 Storage accountEvent hub
Enable client connection logging.	<ul style="list-style-type: none">Diagnostic settingManaged identityApp registrationEnvironment variable

NEW QUESTION: 22

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

How should you complete the code? To answer, drag the appropriate code segments to the correct locations.

Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Code segments

```

MSITokenProvider("...", null)
tp.GetAccessTokenAsync("...")
AzureServiceTokenProvider()
StringTokenProvider("storage", "msi")
tp.GetAuthenticationHeaderAsync(CancellationToken.None)

```

Answer Area

```

var tp = new [code segment]
var t = new TokenCredential(await [code segment] );
return new StorageCredentials(t);

```

Answer:

Code segments

```

MSITokenProvider("...", null)
tp.GetAccessTokenAsync("...")
AzureServiceTokenProvider()
StringTokenProvider("storage", "msi")
tp.GetAuthenticationHeaderAsync(CancellationToken.None)

```

Answer Area

```

var tp = new AzureServiceTokenProvider()
var t = new TokenCredential(await tp.GetAccessTokenAsync("...") );
return new StorageCredentials(t);

```

Explanation

```

var tp = new AzureServiceTokenProvider()
var t = new TokenCredential(await tp.GetAccessTokenAsync("...") );
return new StorageCredentials(t);

```

Box 1: AzureServiceTokenProvider()

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

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

```

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

```

Reference:

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

Topic 6, Coho Winery

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

Coho Winery has a LabelMaker application that prints labels for wine bottles. The application sends data to several printers. The application consists of five modules that run independently on virtual machines (VMs).

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

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

Requirements. Data

You identify the following requirements for data management and manipulation:

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

Requirements. Security

You have the following security requirements:

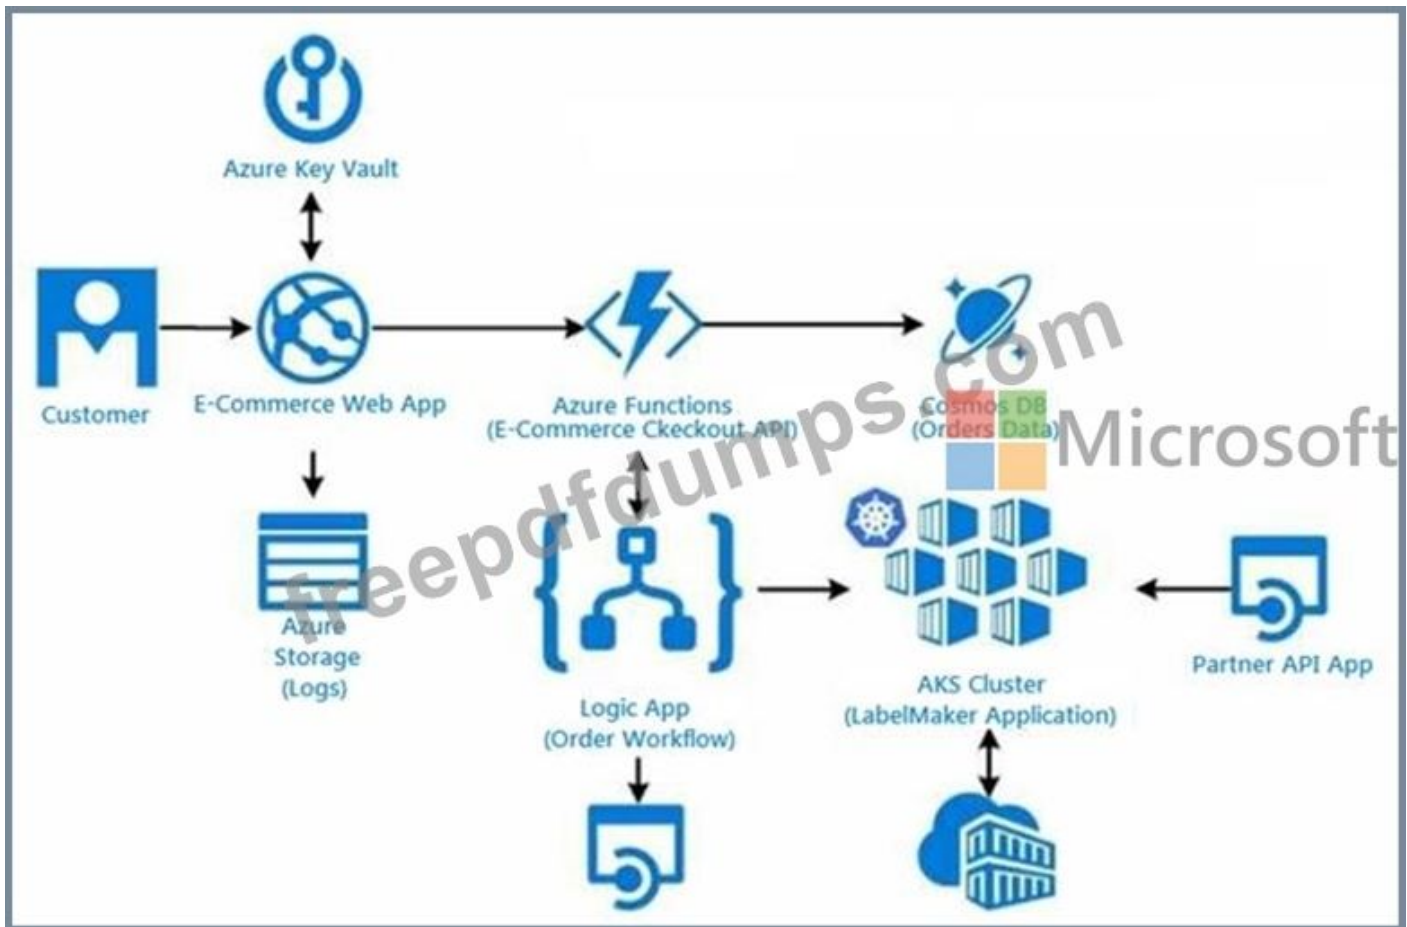
- * Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.
- * External partners must use their own credentials and authenticate with their organization's identity management solution.
- * External partner logins must be audited monthly for application use by a user account administrator to maintain company compliance.
- * Storage of e-commerce application settings must be maintained in Azure Key Vault.
- * E-commerce application sign-ins must be secured by using Azure App Service authentication and Azure Active Directory (AAD).
- * Conditional access policies must be applied at the application level to protect company content.
- * The LabelMaker application must be secured by using an AAD account that has full access to all namespaces of the Azure Kubernetes Service (AKS) cluster.

Requirements. LabelMaker app

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

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

Architecture



Issues

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

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

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

Order.json

Relevant portions of the app files are shown below. Line numbers are included for reference only.

This JSON file contains a representation of the data for an order that includes a single item.

Order.json

```
01 {
02   "id" : 1,
03   "customers" : [
04     {
05       "familyName" : "Doe",
06       "givenName" : "John",
07       "customerid" : 5
08     }
09   ],
10   "line_items" : [
11     {
12       "fulfillable_quantity" : 1,
13       "id" : 6,
14       "price" : "199.99" ,
15       "product_id" : 7513594,
16       "quantity": 1,
17       "requires_shipping" : true ,
18       "sku" : "SFC-342-N" ,
19       "title": "Surface Go" ,
20       "vendor" : "Microsoft" ,
21       "name" : "Surface Go - 8GB" ,
22       "taxable" : true ,
23       "tax_lines" : [
24         {
25           "title" : "State Tax" ,
26           "price" : "3.98" ,
27           "rate" : 0.06
28         }
29       ],
30       "total_discount" : "5.00" ,
31       "discount_allocations" : [
32         {
33           "amount" : "5.00" ,
34           "discount_application_index" : 2
35         }
36       ]
37     }
38   ],
39   "address" : {
40     "state" : "NY" ,
41     "state": "Manhattan" ,
42     "city" : "NYC"
```

```
42     city : ""
43   }
44 }
```

NEW QUESTION: 23

You need to grant access to the retail store location data for the inventory service development effort. What should you use?

- A. Azure AD ID token
- B. Azure AD access token
- C. Shared access signature (SAS) token
- D. Azure AD refresh token
- E. Azure RBAC role

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 24

You need to insert code at line LE03 of LoginEvent.cs to ensure that all authentication events are processed correctly.

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

NOTE: Each correct selection is worth one point.

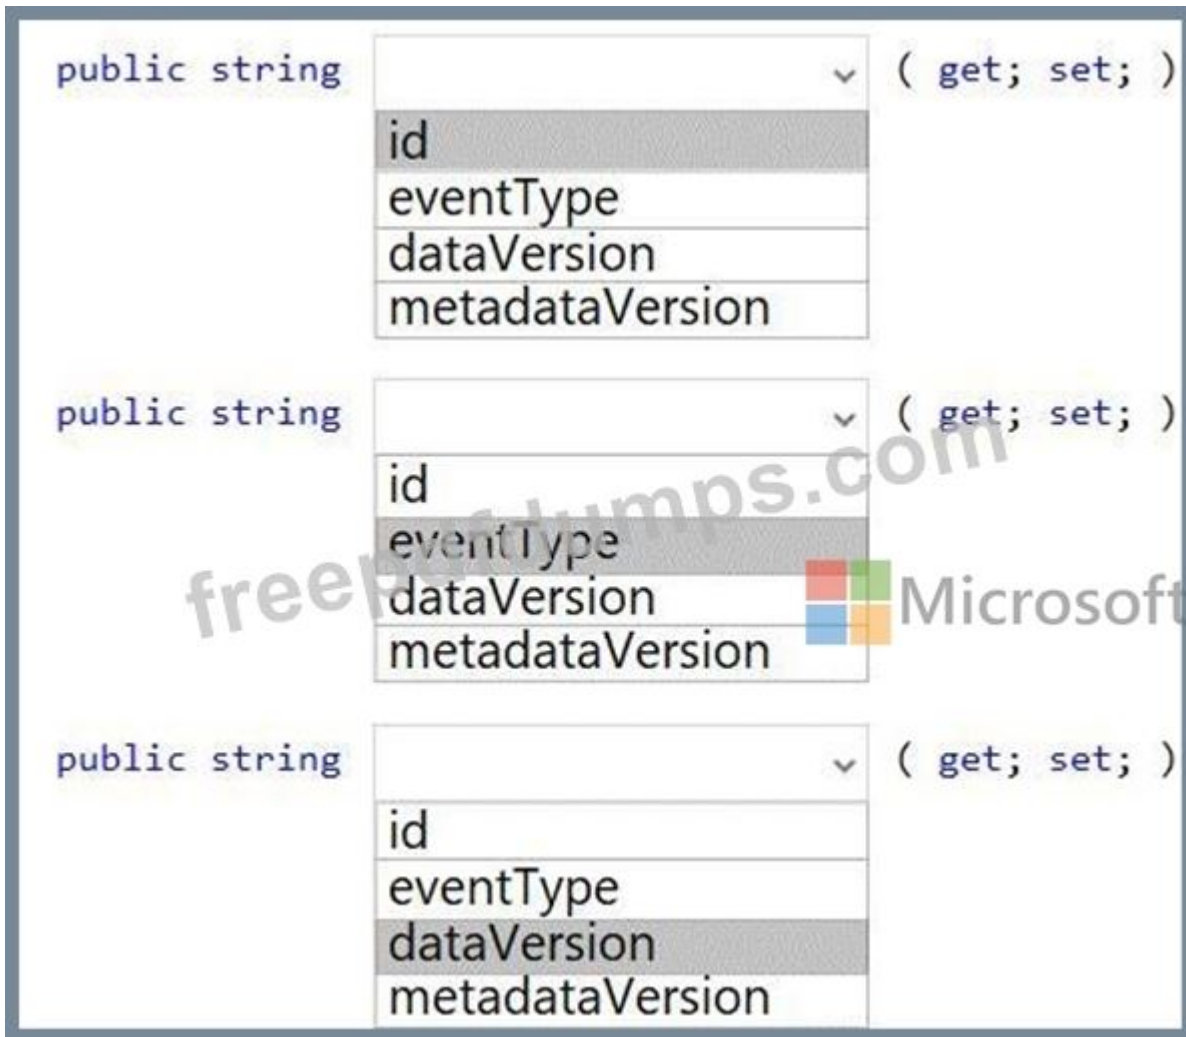
<code>public string</code>	<input type="text"/>	<code>(get; set;)</code>
	id	
	eventType	
	dataVersion	
	metadataVersion	
<code>public string</code>	<input type="text"/>	<code>(get; set;)</code>
	id	
	eventType	
	dataVersion	
	metadataVersion	
<code>public string</code>	<input type="text"/>	<code>(get; set;)</code>
	id	
	eventType	
	dataVersion	
	metadataVersion	

```
public string [ ] ( get; set; )
    id
    eventType
    dataVersion
    metadataVersion

public string [ ] ( get; set; )
    id
    eventType
    dataVersion
    metadataVersion

public string [ ] ( get; set; )
    id
    eventType
    dataVersion
    metadataVersion
```

Explanation



Box 1: id

id is a unique identifier for the event.

Box 2: eventType

eventType is one of the registered event types for this event source.

Box 3: dataVersion

dataVersion is the schema version of the data object. The publisher defines the schema version.

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

The following example shows the properties that are used by all event publishers:

```
[
{
"topic": string,
"subject": string,
"id": string,
"eventType": string,
"eventTime": string,
"data":{
object-unique-to-each-publisher
```

```
},  
"dataVersion": string,  
"metadataVersion": string  
}  
]
```

Reference:

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

NEW QUESTION: 25

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.

Answer: A ([LEAVE A REPLY](#))

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>



NEW QUESTION: 26


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.

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.	



Answer:

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.	

Explanation

 Export a Resource Manager template.

Create a new template deployment.

Modify the template by changing the storage account name and region.

Deploy the template to create a new storage account in the target region.

Use AZCopy to copy the data to the new storage account.

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal#modify-the-te>

NEW QUESTION: 27

You are developing an online game that allows players to vote for their favorite photo that illustrates a word.

The game is built by using Azure Functions and uses durable entities to track the vote count. The voting window is 30 seconds. You must minimize latency.

You need to implement the Azure Function for voting.

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

Answer Area

```
[FunctionName("Vote")]
public static async Task<HttpStatusCode> Run(
    [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
    SignalEntityAsync c,
    [DurableClient] IDurableEntityClient
    [DurableClient] IDurableOrchestrationClient
)
{
    return req.CreateResponse(HttpStatusCode.OK);
}
{
    var eid = new EntityId("pic", id);
    await c.
    return req.Cr
}
```



Answer:

Answer Area

```
[FunctionName("Vote")]
public static async Task<HttpStatusCode> Run(
    [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
    SignalEntityAsync c,
    [DurableClient] IDurableEntityClient
    [DurableClient] IDurableOrchestrationClient
)
{
    return req.CreateResponse(HttpStatusCode.OK);
}
{
    var eid = new EntityId("pic", id);
    await c.
    return req.Cr
}
```




Explanation

D:\mudassar\Untitled.jpg

Answer Area

```
[FunctionName("Vote")]
public static async Task<HttpResponseMessage> Run(
    [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
    SignalEntityAsync c,
    string id)
{
    var eid = new EntityId("pic", id);
    await c. [DurableClient] IDurableEntityClient (eid, "vote");
    return req.CreateResponse(HttpStatusCode.OK);
} }
```



NEW QUESTION: 28

You are building a website that is used to review restaurants. The website will use an Azure CDN to improve performance and add functionality to requests.

You build and deploy a mobile app for Apple iPhones. Whenever a user accesses the website from an iPhone, the user must be redirected to the app store.

You need to implement an Azure CDN rule that ensures that iPhone users are redirected to the app store.

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.

Answer Area

```
"conditions": [ {  
  "name": "IsDevice",  
  "parameters": {  
    "@odata.type": "#Microsoft.Azure.Cdn.Models.  
    "operator": "Equal"  
    "matchValues": [ "  " ]  
  } },  
  {  
    "name": "RequestHeader",  
    "parameters": {  
      "@odata.type": "#Microsoft.Azure.Cdn.Models.  
      "operator": "Contains",  
      "selector": "  "  
      "matchValues": [ "  " ]  
    } }  
  ]  
}
```


dropdown menu: iOS, Mobile, iPhone, Desktop

dropdown menu: DeliveryRulesDeviceConditionParameters, DeliveryRuleCookiesConditionParameters, DeliveryRulePostArgsConditionParameters, DeliveryRuleRequestHeaderConditionParameters

dropdown menu: FROM, PRAGMA, X-POWERED-BY, HTTP_USER_AGENT

dropdown menu: DeliveryRulesDeviceConditionParameters, DeliveryRuleCookiesConditionParameters, DeliveryRulePostArgsConditionParameters, DeliveryRuleRequestHeaderConditionParameters

dropdown menu: iOS, Mobile, iPhone, Desktop



Answer:

Answer Area



```
"conditions": [ {  
  "name": "IsDevice",  
  "parameters": {  
    "@odata.type": "#Microsoft.Azure.Cdn.Models.",  
    "operator": "Equal"  
    "matchValues": [ " " ]  
  } },  
  {  
    "name": "RequestHeader",  
    "parameters": {  
      "@odata.type": "#Microsoft.Azure.Cdn.Models.",  
      "operator": "Contains",  
      "selector": " "  
      "matchValues": [ " " ]  
    } }  
  ]
```

- iOS
- Mobile
- iPhone
- Desktop

- DeliveryRulesDeviceConditionParameters
- DeliveryRuleCookiesConditionParameters
- DeliveryRulePostArgsConditionParameters
- DeliveryRuleRequestHeaderConditionParameters

- FROM
- PRAGMA
- X-POWERED-BY
- HTTP_USER_AGENT

- DeliveryRulesDeviceConditionParameters
- DeliveryRuleCookiesConditionParameters
- DeliveryRulePostArgsConditionParameters
- DeliveryRuleRequestHeaderConditionParameters

- iOS
- Mobile
- iPhone
- Desktop

Explanation

```
"conditions": [ {  
  "name": "IsDevice",  
  "parameters": {  
    "@odata.type": "#Microsoft.Azure.Cdn.Models.",  
    "operator": "Equal"  
    "matchValues": [ " " ]  
  } },  
  {  
    "name": "RequestHeader",  
    "parameters": {  
      "@odata.type": "#Microsoft.Azure.Cdn.Models.",  
      "operator": "Contains",  
      "selector": " "  
      "matchValues": [ " " ]  
    } }  
  ]
```

- iOS
- Mobile
- iPhone
- Desktop

- DeliveryRulesDeviceConditionParameters
- DeliveryRuleCookiesConditionParameters
- DeliveryRulePostArgsConditionParameters
- DeliveryRuleRequestHeaderConditionParameters

- FROM
- PRAGMA
- X-POWERED-BY
- HTTP_USER_AGENT

- DeliveryRulesDeviceConditionParameters
- DeliveryRuleCookiesConditionParameters
- DeliveryRulePostArgsConditionParameters
- DeliveryRuleRequestHeaderConditionParameters

- iOS
- Mobile
- iPhone
- Desktop



Box 1: iOS

Azure AD Conditional Access supports the following device platforms:

- * Android
- * iOS
- * Windows Phone
- * Windows
- * macOS

Box 2: DeliveryRulesDeviceConditionParameters

The DeliveryRulesDeviceCondition defines the IsDevice condition for the delivery rule. parameters defines the parameters for the condition.

Box 3: HTTP_USER_AGENT

Box 4: DeliveryRuleRequestHeaderConditionParameters

DeliveryRuleRequestHeaderCondition defines the RequestHeader condition for the delivery rule. parameters defines the parameters for the condition.

Box 5: iOS

The Require approved client app requirement only supports the iOS and Android for device platform condition.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/concept-conditional-access-condition>

<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/concept-conditional-access-grant>

NEW QUESTION: 29

You develop and deploy the following staticwebapp.config.json file to the app_location value specified in the workflow file of an Azure Static Web app.

```
staticwebapp.config.json
{
  "routes": [
    {
      "route": "/api/**",
      "methods": ["GET"],
      "allowedRoles": ["registeredusers"]
    },
    {
      "route": "/api/**",
      "methods": ["GET", "POST", "PATCH", "DELETE"]
    }
  ]
}
```

Statements	Yes	No
Unauthenticated users are challenged to authenticate with GitHub.	<input type="radio"/>	<input type="radio"/>
A non-existent file in the /images/ folder will generate a 404 response code.	<input type="radio"/>	<input type="radio"/>
HTTP GET method requests from authenticated users in the role named registeredusers are sent to the API folder.	<input type="radio"/>	<input type="radio"/>
Authenticated users that are not in the role named registeredusers and unauthenticated users are served a 401 HTTP error when accessing the API folder.	<input type="radio"/>	<input type="radio"/>

Answer:

Statements	Yes	No
Unauthenticated users are challenged to authenticate with GitHub.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
A non-existent file in the /images/ folder will generate a 404 response code.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
HTTP GET method requests from authenticated users in the role named registerodusers are sent to the API folder.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Authenticated users that are not in the role named registerodusers and unauthenticated users are served a 401 HTTP error when accessing the API folder.	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Explanation

Graphical user interface, text, application, letter, email Description automatically generated

Statements	Yes	No
Unauthenticated users are challenged to authenticate with GitHub.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
A non-existent file in the /images/ folder will generate a 404 response code.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
HTTP GET method requests from authenticated users in the role named registerodusers are sent to the API folder.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Authenticated users that are not in the role named registerodusers and unauthenticated users are served a 401 HTTP error when accessing the API folder.	<input checked="" type="checkbox"/>	<input type="checkbox"/>

NEW QUESTION: 30

You develop an application. You plan to host the application on a set of virtual machines (VMs) in Azure.

You need to configure Azure Monitor to collect logs from the application.

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

Actions	Answer Area
Create a Log Analytics workspace.	
Install agents on the VM and VM scale set to be monitored.	
Send console logs.	
Add a VMInsights solution.	
Create an Application Insights resource.	

Answer:

Actions	Answer Area
Create a Log Analytics workspace.	Create a Log Analytics workspace.
Install agents on the VM and VM scale set to be monitored.	Add a VMInsights solution.
Send console logs.	Install agents on the VM and VM scale set to be monitored.
Add a VMInsights solution.	Create an Application Insights resource.
Create an Application Insights resource.	

Explanation

Graphical user interface, text, application Description automatically generated

Answer Area

Create a Log Analytics workspace.

Add a VMInsights solution.

Install agents on the VM and VM scale set to be monitored.

Create an Application Insights resource.

Step 1: Create a Log Analytics workspace.

First create the workspace.

Step 2: Add a VMInsights solution.

Before a Log Analytics workspace can be used with VM insights, it must have the VMInsights solution installed.

Step 3: Install agents on the VM and VM scale set to be monitored.

Prior to onboarding agents, you must create and configure a workspace. Install or update the Application Insights Agent as an extension for Azure virtual machines and VM scale sets.

Step 4: Create an Application Insights resource

Sign in to the Azure portal, and create an Application Insights resource.

Graphical user interface, application, Word Description automatically generated

Home > New > Application Insights >

Application Insights

Monitor web app performance and usage

Basics Tags Review + create

Create an Application Insights resource to monitor your live web application. With Application Insights, you have full observability into your application across all components and dependencies of your complex distributed architecture. It includes powerful analytics tools to help you diagnose issues and to understand what users actually do with your app. It's designed to help you continuously improve performance and usability. It works for apps on a wide variety of platforms including .NET, Node.js and Java EE, hosted on-premises, hybrid, or any public cloud. [Learn More](#)

PROJECT DETAILS

Select a subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ Visual Studio Enterprise

Resource Group * ⓘ My_Resource_Group

[Create new](#)

INSTANCE DETAILS

Name * ⓘ My_AppInsights_Resource ✓

Region * ⓘ (US) West US 2

Resource Mode * ⓘ Classic **Workspace-based**

WORKSPACE DETAILS

Subscription * ⓘ Visual Studio Enterprise

Log Analytics Workspace * ⓘ my-workspace-name [westus2]

Review + create « Previous Next : Tags »

Once a workspace-based Application Insights resource has been created, configuring monitoring is relatively straightforward.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/vm/vminsights-configure-workspace>

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/create-workspace-resource>

NEW QUESTION: 31

A company is developing a solution that allows smart refrigerators to send temperature information to a central location. You have an existing Service Bus.

The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location. You need to complete the configuration.

Which Azure CLI or PowerShell command should you run?

```

A. az servicebus namespace create
   - --resource-group fridge-rg
   - --name fridge-ns
   - --location fridge-loc

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

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

D. az group create
   --name fridge-rg
   --location fridge-log

```



A. Option A

B. Option B

C. Option C

D. Option D

Answer: B ([LEAVE A REPLY](#))

Explanation

A service bus instance has already been created (Step 2 below). Next is step 3, Create a Service Bus queue.

Note:

Steps:

Step 1: # Create a resource group

```
resourceGroupName="myResourceGroup"
```

```
az group create --name $resourceGroupName --location eastus
```

Step 2: # Create a Service Bus messaging namespace with a unique name

```
namespaceName=myNameSpace$RANDOM
```

```
az servicebus namespace create --resource-group $resourceGroupName --name
```

```
$namespaceName --location eastus Step 3: # Create a Service Bus queue az servicebus queue
```

```
create --resource-group $resourceGroupName --namespace-name $namespaceName
```

```
--name BasicQueue
```

Step 4: # Get the connection string for the namespace

```
connectionString=$(az servicebus namespace authorization-rule keys list --resource-group
```

```
$resourceGroupName --namespace-name $namespaceName --name
```

```
RootManageSharedAccessKey --query primaryConnectionString --output tsv) References:
```

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

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam!
Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 32

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals. You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:

*Share session state across all ASP.NET web applications

*Support controlled, concurrent access to the same session state data for multiple readers and a single writer

*Save full HTTP responses for concurrent requests

You need to store the information.

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B (LEAVE A REPLY)

Explanation

Instead use Azure Cache for Redis.

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

References:

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

NEW QUESTION: 33

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.

Policy segments	Answer Area
<input type="text" value="server"/>	<pre>< Policy segment > <base /> <choose> <when condition = " @ Policy segment .Response.StatusCode == 500 && Policy segment .LastError.Message.Contains " conflict = ")) " > <return-response> < Policy segment > </return-response> </when> <otherwise /> </choose> < Policy segment ></pre>
<input type="text" value="context"/>	
<input type="text" value="on-error"/>	
<input type="text" value="set-status"/>	
<input type="text" value="when-error"/>	
<input type="text" value="override-status"/>	
<input type="text" value=" "/>	

Answer:

Policy segments	Answer Area
<input type="text" value="server"/>	<pre>< on-error > <base /> <choose> <when condition = " @ context .Response.StatusCode == 500 && context .LastError.Message.Contains " conflict = ")) " > <return-response> < set-status > </return-response> </when> <otherwise /> </choose> < on-error ></pre>
<input type="text" value="context"/>	
<input type="text" value="on-error"/>	
<input type="text" value="set-status"/>	
<input type="text" value="when-error"/>	
<input type="text" value="override-status"/>	
<input type="text" value=" "/>	

Explanation

```
< on-error >
<base />
<choose>
  <when condition = " @ context .Response.StatusCode == 500
    && context .LastError.Message.Contains
      " conflict = " )) " >
    <return-response>
      < set-status >
    </return-response>
  </when>
  <otherwise />
</choose>
< on-error >
```

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>

NEW QUESTION: 34

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.

Components

- Helm
- Draft
- Brigade
- KubeCtl
- Ingress Controller
- CoreDNS
- Virtual Kubelet

Answer area

Action

Deploy solution.
View cluster and external IP addressing.
Implement a single, public IP endpoint that is routed to multiple microservices.

Component

- Component
- Component
- Component



Answer:



Components

- Helm
- Draft
- Brigade
- KubeCtl
- Ingress Controller
- CoreDNS
- Virtual Kubelet

Answer area

Action

Deploy solution.
View cluster and external IP addressing.
Implement a single, public IP endpoint that is routed to multiple microservices.

Component

- Helm
- KubeCtl
- Ingress Controller



Explanation

Answer Area

Action

Deploy solution.

View cluster and external IP addressing.

Implement a single, public IP endpoint that is routed to multiple microservices.

Component



Helm

KubeCtl

Ingress Controller

Box 1: Helm

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

Box 2: kubectl

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

Box 3: Ingress Controller

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

Reference:

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

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

NEW QUESTION: 35

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

Answer: A ([LEAVE A REPLY](#))

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>

NEW QUESTION: 36

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

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

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

You create the index in Azure Search.

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

NET SDK.

Solution:

1. Create a SearchServiceClient object to connect to the search index.
2. Create a DataContainer that contains the documents which must be added.
3. Create a DataSource instance and set its Container property to the DataContainer.
4. Set the DataSource property of the SearchServiceClient

Does the solution meet the goal?

A. Yes

B. No

Answer: B (LEAVE A REPLY)

Explanation

Use the following method:

1. Create a SearchIndexClient object to connect to the search index
2. Create an IndexBatch that contains the documents which must be added.
3. Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.

References:

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

NEW QUESTION: 37

You need to access data from the user claim object in the e-commerce web app.

What should you do first?

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

Answer: C (LEAVE A REPLY)

Explanation

Methods to Get User Identity and Claims in a .NET Azure Functions App include:

ClaimsPrincipal from the Request Context

The ClaimsPrincipal object is also available as part of the request context and can be extracted from the HttpRequest.HttpContext.

User Claims from the Request Headers.

App Service passes user claims to the app by using special request headers.

Reference:

<https://levelup.gitconnected.com/four-alternative-methods-to-get-user-identity-and-claims-in-a-net-azurefunction>

NEW QUESTION: 38

You need to add markup at line AM04 to implement the ContentReview 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.

Json segments

- User
- value
- role
- Application
- allowedMemberTypes
- allowedAccountTypes

Answer Area

```
"appRoles" : [
{
  " [ ] ": [
  " [ ] "
],
  "displayName": "ContentReviewer",
  "id": "e1c2ade8-98f8-45fd-aa4a-6d24b512c22a",
  "isEnabled" : true,
  " [ ] " : "ContentReviewer"
}
],
```

Answer:

Json segments

- User
- value
- role
- Application
- allowedMemberTypes
- allowedAccountTypes

Answer Area

```
"appRoles" : [
{
  " allowedMemberTypes " : [
  " User "
],
  "displayName": "ContentReviewer",
  "id": "e1c2ade8-98f8-45fd-aa4a-6d24b512c22a",
  "isEnabled" : true,
  " value " : "ContentReviewer"
}
],
```

Explanation

```

"appRoles" : [
{
  "allowedMemberTypes" : [
  "User"
  ],
  "displayName": "ContentReviewer",
  "id": "e1c2ade8-98f8-45fd-aa4a-6d24b512c22a"
  "isEnabled" : true,
  "value" : "ContentReviewer"
}
],

```

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>

NEW QUESTION: 39

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

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

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

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

Solution:

- 1 Create a SearchIndexClient object to connect to the search index
2. Create an IndexBatch that contains the documents which must be added.
3. Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.

Does the solution meet the goal?

A. Yes

B. No

Answer: A (LEAVE A REPLY)

Explanation

1. The index needs to be populated. To do this, we will need a SearchIndexClient. There are two ways to obtain one: by constructing it, or by calling Indexes.GetClient on the SearchServiceClient. Here we will use the first method.

2. Create the indexBatch with the documents

Something like:

```
var hotels = new Hotel[];
{
new Hotel()
{
HotelId = "3",
BaseRate = 129.99,
Description = "Close to town hall and the river"
}
};
```

...

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

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

Example:

```
try
{
var batch = IndexBatch.Upload(hotels);
}
```

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

References:

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

NEW QUESTION: 40

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.

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

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

Explanation

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

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

NEW QUESTION: 41

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

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

What should you do?

- A. In Python, implement the class: TaskAddParameter
- B. In Python, implement the class: JobAddParameter
- C. In the Azure portal, create a Batch account
- D. In a .NET method, call the method: BatchClient.PoolOperations.CreateJob

Answer: D (LEAVE A REPLY)

Explanation

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

Note:

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

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

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

output to Azure Storage.

NEW QUESTION: 42

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

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

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

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B ([LEAVE A REPLY](#))

Explanation

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

Reference:

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

NEW QUESTION: 43

You are developing an Azure-hosted e-commerce web application. The application will use Azure Cosmos DB to store sales orders. You are using the latest SDK to manage the sales orders in the database.

You create a new Azure Cosmos DB instance. You include a valid endpoint and valid authorization key to an appSettings.json file in the code project.

You are evaluating the following application code: (Line number are included for reference only.)

```

01 using System;
02 using System.Threading.Tasks;
03 using Microsoft.Azure.Cosmos;
04 using Microsoft.Extensions.Configuration;
05 using Newtonsoft.Json;
06 namespace SalesOrders
07 {
08     public class SalesOrder
09     {
10         . . .
11     }
12     internal class ManageSalesOrders
13     {
14         private static async Task GenerateSalesOrders()
15         {
16             IConfigurationRoot configuration = new ConfigurationBuilder().AddJsonFile("appSettings.json").Build();
17             string endpoint = configuration["EndPointUrl"];
18             string authKey = configuration["AuthorizationKey"];
19             using CosmosClient client = new CosmosClient(endpoint, authKey);
20             Database database = null;
21             using (await client.GetDatabase("SalesOrders").DeleteStreamAsync()) { }
22             database = await client.CreateDatabaseIfNotExistsAsync("SalesOrders");
23             Container container1 = await database.CreateContainerAsync(id: "Container1", partitionKeyPath: "/AccountNumber");
24             Container container2 = await database.CreateContainerAsync(id: "Container2", partitionKeyPath: "/AccountNumber");
25             SalesOrder salesOrder1 = new SalesOrder() { AccountNumber = "123456" };
26             await container1.CreateItemAsync(salesOrder1, new PartitionKey(salesOrder1.AccountNumber));
27             SalesOrder salesOrder2 = new SalesOrder() { AccountNumber = "654321" };
28             await container1.CreateItemAsync(salesOrder2, new PartitionKey(salesOrder2.AccountNumber));
29             SalesOrder salesOrder3 = new SalesOrder() { AccountNumber = "109876" };
30             await container2.CreateItemAsync(salesOrder3, new PartitionKey(salesOrder3.AccountNumber));
31             _ = await database.CreateUserAsync("User1");
32             User user1 = database.GetUser("User1");
33             _ = await user1.ReadAsync();
34         }
35     }
36 }

```

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
A database named SalesOrders is created. The database will include two containers.	<input type="radio"/>	<input type="radio"/>
Container1 will contain two items	<input type="radio"/>	<input type="radio"/>
Container2 will contain one item.	<input type="radio"/>	<input type="radio"/>

Answer:

Statements	Yes	No
A database named SalesOrders is created. The database will include two containers.	<input checked="" type="radio"/>	<input type="radio"/>
Container1 will contain two items.	<input checked="" type="radio"/>	<input type="radio"/>
Container2 will contain one item.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation

Graphical user interface, text, application Description automatically generated

Statements	Yes	No
A database named SalesOrders is created. The database will include two containers.	<input type="radio"/>	<input type="radio"/>
Container1 will contain two items.	<input type="radio"/>	<input type="radio"/>
Container2 will contain one item.	<input type="radio"/>	<input type="radio"/>

Box 1: Yes

The createDatabaseIfNotExistsAsync method checks if a database exists, and if it doesn't, create it.

The Database.CreateContainerAsync method creates a container as an asynchronous operation in the Azure Cosmos service.

Box 2: Yes

The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service.

Box 3: Yes

Reference:

[https://docs.microsoft.com/en-](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient.createdatabaseifnotexistsasync)

[us/dotnet/api/microsoft.azure.cosmos.cosmosclient.createdatabaseifnotexistsasync](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient.createdatabaseifnotexistsasync)

[https://docs.microsoft.com/en-](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.database.createcontainerasync)

[us/dotnet/api/microsoft.azure.cosmos.database.createcontainerasync](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.database.createcontainerasync)

<https://docs.microsoft.com/en-us/dotnet/api/azure.cosmos.cosmoscontainer.createitemasync>

NEW QUESTION: 44

A software as a service (SaaS) company provides document management services. The company has a service that consists of several Azure web apps. All Azure web apps run in an Azure App Service Plan named PrimaryASP.

You are developing a new web service by using a web app named ExcelParser. The web app contains a third-party library for processing Microsoft Excel files. The license for the third-party library stipulates that you can only run a single instance of the library.

You need to configure the service.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
Set-AzAppServicePlan `
  -ResourceGroupName $rg `
  -Name "PrimaryASP" `
```

```
NumberOfSites 1
PerSiteScaling $true
TargetWorkerCount = 1
MaxNumberOfWorkers = 1
SiteConfig.NumberOfWorkers = 1
```



Microsoft

```
$app = Get-AzWebApp `
  -ResourceGroupName $rg `
  -Name "ExcelParser"
```

```
$app.
```

```
NumberOfSites 1
PerSiteScaling $true
TargetWorkerCount = 1
MaxNumberOfWorkers = 1
SiteConfig.NumberOfWorkers = 1
```

```
Set-AzWebApp $app
```

Answer:

Answer Area

```
Set-AzAppServicePlan `
  -ResourceGroupName $rg `
  -Name "PrimaryASP" `
```

```
NumberOfSites 1
PerSiteScaling $true
TargetWorkerCount = 1
MaxNumberOfWorkers = 1
SiteConfig.NumberOfWorkers = 1
```

```
$app = Get-AzWebApp `
  -ResourceGroupName $rg `
  -Name "ExcelParser"
```

```
$app
NumberOfSites 1
PerSiteScaling $true
TargetWorkerCount = 1
MaxNumberOfWorkers = 1
SiteConfig.NumberOfWorkers = 1
```

```
Set-AzWebApp $app
```

Explanation

Table Description automatically generated

```

Set-AzAppServicePlan `
  -ResourceGroupName $rg `
  -Name "PrimaryASP" `

```

NumberOfSites 1
PerSiteScaling \$true
TargetWorkerCount = 1
MaxNumberOfWorkers = 1
SiteConfig.NumberOfWorkers = 1

```

$app = Get-AzWebApp `
  -ResourceGroupName $rg `
  -Name "ExcelParser"

```

\$app.

NumberOfSites 1
PerSiteScaling \$true
TargetWorkerCount = 1
MaxNumberOfWorkers = 1
SiteConfig.NumberOfWorkers = 1

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/manage-scale-per-app>

NEW QUESTION: 45

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

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

You need to create the custom role.

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

NOTE: Each correct selection is worth one point.

Item	Value
Powershell command	<pre>Get-AzureRmRoleDefinition-Name "Reader" ConvertTo-Json Out-File C:\SupportRole.json Get-AzureRmRoleDefinition-Name "Operator" ConvertTo-Json Out-File C:\SupportRole.json Set-AzureRmRoleDefinition-Name "Reader" Input-File C:\SupportRole.json Set-AzureRmRoleDefinition Input-File C:\SupportRole.json</pre>
Actions section	<pre>"*/read*", *Microsoft.Support/*" "*/read*" ""* *Microsoft.Support/*" ""*</pre>

Answer:

Item	Value
Powershell command	<pre>Get-AzureRmRoleDefinition-Name "Reader" ConvertTo-Json Out-File C:\SupportRole.json Get-AzureRmRoleDefinition-Name "Operator" ConvertTo-Json Out-File C:\SupportRole.json Set-AzureRmRoleDefinition-Name "Reader" Input-File C:\SupportRole.json Set-AzureRmRoleDefinition Input-File C:\SupportRole.json</pre>
Actions section	<pre>"*/read*", *Microsoft.Support/*" "*/read*" ""* *Microsoft.Support/*" ""*</pre>

Explanation

Item	Value
Powershell command	<pre>Get-AzureRmRoleDefinition-Name "Reader" ConvertTo-Json Out-File C:\SupportRole.json Get-AzureRmRoleDefinition-Name "Operator" ConvertTo-Json Out-File C:\SupportRole.json Set-AzureRmRoleDefinition-Name "Reader" Input-File C:\SupportRole.json Set-AzureRmRoleDefinition Input-File C:\SupportRole.json</pre>
Actions section	<pre>"*/read*", *Microsoft.Support/*" "*/read*" ""* *Microsoft.Support/*" ""*</pre>

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

The Set-AzureRmRoleDefinition cmdlet updates an existing custom role in Azure Role-Based Access Control.

Provide the updated role definition as an input to the command as a JSON file or a PSRoleDefinition object.

The role definition for the updated custom role MUST contain the Id and all other required properties of the role even if they are not updated: DisplayName, Description, Actions, AssignableScope
 Box 2: "*/read*".* Microsoft.Support/*" Microsoft.Support/* Create and manage support tickets

"Microsoft.Support" role definition azure

NEW QUESTION: 46

You develop an ASP.NET Core MVC application. You configure the application to track webpages and custom events.

You need to identify trends in application usage.

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

NOTE: Each correct selection is worth one point.

Answer:

Explanation

Requirement	Feature
Which pages visited by users most often correlate to a product purchase?	Users
How does load time of the product display page affect a user's decision to purchase a product?	Impact
Which events most influence a user's decision to continue to use the application?	Retention
Are there places in the application that users often perform repetitive actions?	User Flows

Box1: Users

Box 2: Impact

One way to think of Impact is as the ultimate tool for settling arguments with someone on your team about how slowness in some aspect of your site is affecting whether users stick around. While users may tolerate a certain amount of slowness, Impact gives you insight into how best to balance optimization and performance to maximize user conversion.

Box 3: Retention

The retention feature in Azure Application Insights helps you analyze how many users return to your app, and how often they perform particular tasks or achieve goals. For example, if you run a game site, you could compare the numbers of users who return to the site after losing a game with the number who return after winning. This knowledge can help you improve both your user experience and your business strategy.

Box 4: User flows

The User Flows tool visualizes how users navigate between the pages and features of your site. It's great for answering questions like:

How do users navigate away from a page on your site?

What do users click on a page on your site?

Where are the places that users churn most from your site?

Are there places where users repeat the same action over and over?

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam! Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 47

You develop Azure Durable Functions to manage vehicle loans.

The loan process includes multiple actions that must be run in a specified order. One of the actions includes a customer credit check process, which may require multiple days to process.

You need to implement Azure Durable Functions for the loan process.

Which Azure Durable Functions type should you use?

- A. activity
- B. entity
- C. orchestrator
- D. client

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 48

You are building an application that stores sensitive customer data in Azure Blob storage. The data must be encrypted with a key that is unique for each customer.

If the encryption key has been corrupted it must not be used for encryption.

You need to ensure that the blob is encrypted.

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

Answer Area

```
from azure.storage.blob import BlobServiceClient

from azure.storage.blob.aio import BlobType x = BlobType(key, verify)
from azure.storage.blob import BlobSasPermissions x = BlobSasPermissions.from_string(key + verify)
from azure.storage.blob import CustomerProvidedEncryptionKey x = CustomerProvidedEncryptionKey(key, verify)
from azure.core.configuration import Configuration x = Configuration(key, verify)

if x.tag == verify:
    creds = ...

if x.maketrans == verify:
    ...

if x.EncryptionKeyHash == verify:
    ...

if x.proxy_policy == verify:
    ...

bsc = BlobServiceClient(credential = creds)
c = bsc.get_blob_client("con", blob)

upload_blob_data(pk=x)
upload_blob_data(bt=x)
upload_blob_data(bsp=x)
upload_blob_data(pk=x)
```

Microsoft

Answer:



```

from azure.storage.blob.aio import BlobType x = BlobType(key, verify)
from azure.storage.blob import BlobSasPermissions x = BlobSasPermissions.from_string(key + verify)
from azure.storage.blob import CustomerProvidedEncryptionKey x = CustomerProvidedEncryptionKey(key, verify)
from azure.core.configuration import Configuration x = Configuration(key, verify)

if x.tag == verify:
    if x.makeitrans == verify:
        if x.EncryptionKeyHash == verify:
            if x.proxy_policy == verify:
                bsc = BlobServiceClient("", credential = creds)
                c = bsc.get_blob_client("con", blob)

c.upload_blob(data, pa=x)
c.upload_blob(data, bt=x)
c.upload_blob(data, bsp=x)
c.upload_blob(data, cpk=x)

```

Explanation

Graphical user interface, application Description automatically generated

```

Answer Area

from azure.storage.blob import BlobServiceClient
from azure.storage.blob import CustomerProvidedEncryptionKey x = CustomerProvidedEncryptionKey(key, verify)

if x.makeitrans == verify:
    bsc = BlobServiceClient("", credential = creds)
    c = bsc.get_blob_client("", blob)

c.upload_blob

```

NEW QUESTION: 49

D18912E1457D5D1DDCBD40AB3BF70D5D

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

Answer: B ([LEAVE A REPLY](#))

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>

NEW QUESTION: 50

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

Answer: A ([LEAVE A REPLY](#))

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>

NEW QUESTION: 51

You develop a gateway solution for a public facing news API.

The news API back end is implemented as a RESTful service and hosted in an Azure App Service instance.

You need to configure back-end authentication for the API Management service instance.

Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values	Answer Area
Azure Resource	Configuration parameter
HTTP(s) endpoint	Target
Basic	Gateway credentials
Client cert	Value

Answer:

Values	Answer Area
Azure Resource	Configuration parameter
HTTP(s) endpoint	Target
Basic	Gateway credentials
Client cert	Value

Explanation

Configuration parameter	Value
Target	Azure Resource
Gateway credentials	Client cert

Box 1: Azure Resource

Box 2: Client cert

API Management allows to secure access to the back-end service of an API using client certificates.

Reference:
<https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-ba>

NEW QUESTION: 52


You are developing an Azure Static Web app that contains training materials for a tool company. Each tool's training material is contained in a static web page that is linked from the tool's publicly available description page.

A user must be authenticated using Azure AD prior to viewing training.


You need to ensure that the user can view training material pages after authentication.

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

NOTE: Each correct selection is worth one point.



Answer:



Explanation

Graphical user interface. application Description automatically generated



NEW QUESTION: 53

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

The mobile app must meet the following requirements:

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

You need to implement Offline Data Sync.

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

NOTE: Each correct selection is worth one point.

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

Answer: B,E ([LEAVE A REPLY](#))

Explanation

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

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

References:

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

NEW QUESTION: 54

You deploy an Azure App Service web app. You create an app registration for the app in Azure Active Directory (Azure AD) and Twitter. the app must authenticate users and must use SSL for all communications.

The app must use Twitter as the identity provider. You need to validate the Azure AD request in the app code.

What should you validate?

- A. HTTP response code
- B. Tenant ID
- C. ID token header
- D. ID token signature

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 55

An organization hosts web apps in Azure. The organization uses Azure Monitor You discover that configuration changes were made to some of the web apps. You need to identify the configuration changes.

Which Azure Monitor log should you review?

- A. AppServiceConsoleLogs
- B. AppServiceApplogs

- C. AppServiceAuditLogs
- D. AppServiceEnvironmentPlatformLogs

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 56

You develop and deploy an Azure Logic App that calls an Azure Function app. The Azure Function App includes an OpenAPI (Swagger) definition and uses an Azure Blob storage account. All resources are secured by using Azure Active Directory (Azure AD).

The Logic App must use Azure Monitor logs to record and store information about runtime data and events.

The logs must be stored in the Azure Blob storage account.

You need to set up Azure Monitor logs and collect diagnostics data for the Azure Logic App.

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

Actions

- Create action groups and alert rules.
- Create a Log Analytics workspace.
- Install the Logic Apps Management solution.
- Add a diagnostic setting to the Azure Function App.
- Create an Azure storage account.
- Add a diagnostic setting to the Azure Logic App.

Answer Area

Navigation icons: left arrow, right arrow, up arrow, down arrow.

Answer:

Actions

- Create action groups and alert rules.
- Create a Log Analytics workspace.
- Install the Logic Apps Management solution.
- Add a diagnostic setting to the Azure Function App.
- Create an Azure storage account.
- Add a diagnostic setting to the Azure Logic App.

Answer Area

- Create a Log Analytics workspace.
- Install the Logic Apps Management solution.
- Add a diagnostic setting to the Azure Logic App.

Navigation icons: left arrow, right arrow, up arrow, down arrow.

Explanation

- Create a Log Analytics workspace.
- Install the Logic Apps Management solution.
- Add a diagnostic setting to the Azure Logic App.

Text Description automatically generated

Step 1: Create a Log Analytics workspace

Before you start, you need a Log Analytics workspace.

Step 2: Install the Logic Apps Management solution

To set up logging for your logic app, you can enable Log Analytics when you create your logic app, or you can install the Logic Apps Management solution in your Log Analytics workspace for existing logic apps.

Step 3: Add a diagnostic setting to the Azure Logic App

Set up Azure Monitor logs

* In the Azure portal, find and select your logic app.

* On your logic app menu, under Monitoring, select Diagnostic settings > Add diagnostic setting.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/monitor-logic-apps-log-analytics>

NEW QUESTION: 57

You are developing an application to store and retrieve data in Azure Blob storage. The application will be hosted in an on-premises virtual machine (VM). The VM is connected to Azure by using a Site-to-Site VPN gateway connection. The application is secured by using Azure Active Directory (Azure AD) credentials.

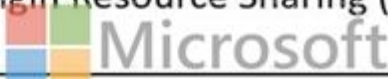
The application must be granted access to the Azure Blob storage account with a start time, expiry time, and read permissions. The Azure Blob storage account access must use the Azure AD credentials of the application to secure data access. Data access must be able to be revoked if the client application security is breached.

You need to secure the application access to Azure Blob storage.

Which security features should you use? To answer select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Component	Security Feature
Application (Client)	<div style="border: 1px solid black; padding: 5px;"><div style="border: 1px solid black; height: 20px; margin-bottom: 5px;"></div><div style="border: 1px solid black; padding: 2px;">Storage Account Access Key</div><div style="border: 1px solid black; padding: 2px;">System-assigned Managed Identity</div><div style="border: 1px solid black; padding: 2px;">Shared access signature (SAS) token</div></div>
Azure Storage (Server)	<div style="border: 1px solid black; padding: 5px;"><div style="border: 1px solid black; height: 20px; margin-bottom: 5px;"></div><div style="border: 1px solid black; padding: 2px;">Stored Access Policy</div><div style="border: 1px solid black; padding: 2px;">User-assigned Managed Identity</div><div style="border: 1px solid black; padding: 2px;">Cross-Origin Resource Sharing (CORS)</div></div>



Answer:

Component	Security Feature								
Application (Client)	<table border="1"><tr><td></td><td>▼</td></tr><tr><td>Storage Account Access Key</td><td></td></tr><tr><td>System-assigned Managed Identity</td><td></td></tr><tr><td>Shared access signature (SAS) token</td><td></td></tr></table>		▼	Storage Account Access Key		System-assigned Managed Identity		Shared access signature (SAS) token	
	▼								
Storage Account Access Key									
System-assigned Managed Identity									
Shared access signature (SAS) token									
Azure Storage (Server)	<table border="1"><tr><td></td><td>▼</td></tr><tr><td>Stored Access Policy</td><td></td></tr><tr><td>User-assigned Managed Identity</td><td></td></tr><tr><td>Cross-Origin Resource Sharing (CORS)</td><td></td></tr></table>		▼	Stored Access Policy		User-assigned Managed Identity		Cross-Origin Resource Sharing (CORS)	
	▼								
Stored Access Policy									
User-assigned Managed Identity									
Cross-Origin Resource Sharing (CORS)									

Explanation

Text. letter Description automatically generated

Component

Security Feature

Application (Client)

	▼
Storage Account Access Key	
System-assigned Managed Identity	
Shared access signature (SAS) token	

Azure Storage (Server)

	▼
Stored Access Policy	
User-assigned Managed Identity	
Cross-Origin Resource Sharing (CORS)	

Box 1: Shared access signature (SAS) token

When your application design requires shared access signatures for access to Blob storage, use Azure AD credentials to create a user delegation SAS when possible for superior security.

Box 2: Stored access policy

Stored access policies give you the option to revoke permissions for a service SAS without having to regenerate the storage account keys.

A shared access signature can take one of the following two forms:

* Service SAS with stored access policy. A stored access policy is defined on a resource container, which can be a blob container, table, queue, or file share. The stored access policy can be used to manage constraints for one or more service shared access signatures. When you associate a service SAS with a stored access policy, the SAS inherits the constraints - the start time, expiry time, and permissions - defined for the stored access policy.

* Ad hoc SAS.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

NEW QUESTION: 58

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

Answer: A ([LEAVE A REPLY](#))

Explanation

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

Reference:

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

NEW QUESTION: 59

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

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

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

Which two policies can you use? Each correct answer presents a complete solution NOTE: Each correct selection is worth one point.

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

D. Certificate Authentication

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 60

You are developing a C++ application that compiles to a native application named process.exe.

The application accepts images as input and returns images in one of the following image formats: GIF, PNG, or JPEG.

You must deploy the application as an Azure Function.

You need to configure the function and host json files.

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

NOTE: Each correct selection is worth one point.

function.json

{

	▼
"type": "http"	
"platform": "gcm"	
"datatype": "stream"	
"path": "process.exe"	

"direction": "out",
"name" : "result"

}

host.json

	▼
"customHandler": { "description": {	
"languageWorker": { "path": {	
"extensions": { "worker": {	
"extensionBundle": {	

"defaultExecutablePath": "process.exe"

},

	▼
"enableForwardingHttpRequest": true	
"enableForwardingHttpRequest": false	

}



Answer:

```
function.json
```

```
{
```

	▼
"type": "http"	
"platform": "gcm"	
"datatype": "stream"	
"path": "process.exe"	

```
  "direction": "out",  
  "name" : "result"
```

```
}
```

```
host.json
```

	▼
"customHandler": { "description": {	
"languageWorker": { "path": {	
"extensions": { "worker": {	
"extensionBundle": {	

```
  "defaultExecutablePath": "process.exe"
```

```
},
```

	▼
"enableForwardingHttpRequest": true	
"enableForwardingHttpRequest": false	

```
}
```



Explanation

```

{
  "type": "http"
  "platform": "gcm"
  "datatype": "stream"
  "path": "process.exe"
}

```

```

  "direction": "out",
  "name" : "result"
}
host.json

```

```

{
  "customHandler": { "description": {
  "languageWorker": { "path": {
  "extensions": { "worker": {
  "extensionBundle": {

```

```

    "defaultExecutablePath": "process.exe"

```

```

},

```

```

{
  "enableForwardingHttpRequest": true
  "enableForwardingHttpRequest": false
}

```

NEW QUESTION: 61

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.

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

Answer:

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

Explanation

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

Box 1: Enterprise Integration Pack

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

Box 2: Code View Editor

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

Box 3: Logical Apps Designer

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

References:

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

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

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

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam!

Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest**

Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, 30%OFF

Special Discount: **Freepdfdumps**)

NEW QUESTION: 62

You are debugging an application that is running on Azure Kubernetes cluster named cluster1.

The cluster uses Azure Monitor for containers to monitor the cluster.

The application has sticky sessions enabled on the ingress controller.

Some customers report a large number of errors in the application over the last 24 hours.

You need to determine on which virtual machines (VMs) the errors are occurring.

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

NOTE: Each correct selection is worth one point.

```
let startTimestamp = 
| 
| ago(1d)
| since(1d)
| totimespan(1d)
| date(now() - 1d)

let ContainerIDs = KubePodInventory
| where ClusterName == "Cluster1"
| 
| top ContainerID
| union ContainerID
| sample ContainerID
| distinct ContainerID

ContainerLog
| 
| fork containerIDs
| where ContainerID in (ContainerIDs)
| restrict ContainerID in (ContainerIDs)
| join ContainerID == ContainerIDs.ContainerID

| where TimeGenerated > startTimestamp
| where LogEntrySource == "stderr"
| 
| project by Computer
| summarize by Computer
| partition count() by Computer
| summarize count() by Computer
```

Answer:

```
let startTimestamp =
```

ago(1d)
since(1d)
totimespan(1d)
date(now() - 1d)

```
let ContainerIDs = KubePodInventory  
| where ClusterName == "Cluster1"
```

```
|
```

top ContainerID
union ContainerID
sample ContainerID
distinct ContainerID

```
;
```

```
ContainerLog
```

```
|
```

fork containerIDs
where ContainerID in (ContainerIDs)
restrict ContainerID in (ContainerIDs)
join ContainerID == ContainerIDs.ContainerID

```
| where TimeGenerated > startTimestamp  
| where LogEntrySource == "stderr"
```

```
|
```

project by Computer
summarize by Computer
partition count() by Computer
summarize count() by Computer

Explanation

let startTimestamp =

	▼
ago(1d)	
since(1d)	
totimespan(1d)	
date(now() - 1d)	

let ContainerIDs = KubePodInventory
| where ClusterName == "Cluster1"

	▼	:
top ContainerID		
union ContainerID		
sample ContainerID		
distinct ContainerID		

ContainerLog

	▼
fork containerIDs	
where ContainerID in (ContainerIDs)	
restrict ContainerID in (ContainerIDs)	
join ContainerID == ContainerIDs.ContainerID	

| where TimeGenerated > startTimestamp
| where LogEntrySource == "stderr"

	▼
project by Computer	
summarize by Computer	
partition count() by Computer	
summarize count() by Computer	

Box 1: ago(1d)

Box 2: distinct containerID

Box 3: where ContainerID in (ContainerIDs)

Box 4: summarize Count by Computer

Summarize: aggregate groups of rows

Use summarize to identify groups of records, according to one or more columns, and apply aggregations to them. The most common use of summarize is count, which returns the number of results in each group.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/get-started-queries>

<https://docs.microsoft.com/en-us/azure/azure-monitor/log-query/query-optimization>

NEW QUESTION: 63

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:

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

You need to implement an Azure Service Bus solution.

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

Actions

- Create a 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.

Answer Area

Answer:

Actions

- 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.

Answer Area

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

Explanation

Create a single Service Bus Namespace.

Create a Service Bus topic for each restaurant for which a driver can receive messages.

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

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>

NEW QUESTION: 64

You need to implement the delivery service telemetry data

How should you configure the solution?

NOTE: Each correct selection is worth one point.



Answer:



Explanation

Graphical user interface, text, application Description automatically generated



NEW QUESTION: 65

You need to authenticate the user to the corporate website as indicated by the architectural diagram.

Which two values should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. ID token signature

B. ID token claims

- C. HTTP response code
- D. Azure AD endpoint URI
- E. Azure AD tenant ID

Answer: ([SHOW ANSWER](#))

Explanation

Claims in access tokens

JWTs (JSON Web Tokens) are split into three pieces:

- * Header - Provides information about how to validate the token including information about the type of token and how it was signed.
- * Payload - Contains all of the important data about the user or app that is attempting to call your service.
- * Signature - Is the raw material used to validate the token.

Your client can get an access token from either the v1.0 endpoint or the v2.0 endpoint using a variety of protocols.

Scenario: User authentication (see step 5 below)

The following steps detail the user authentication process:

- * The user selects Sign in in the website.
- * The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
- * The user signs in.
- * Azure AD redirects the user's session back to the web application. The URL includes an access token.
- * 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.
- * The back-end API validates the access token.

Reference:

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

NEW QUESTION: 66

YOU need to reliably identify the delivery driver profile information.

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

NOTE Each correct selection is worth one point.

Configuration	Value
JSON web token (JWT) type	<input type="text" value="ID"/> <input type="checkbox"/> ID <input type="checkbox"/> Refresh <input type="checkbox"/> Access
Payload claim value	<input type="text" value="oid"/> <input type="checkbox"/> oid <input type="checkbox"/> aud <input type="checkbox"/> idp

Answer:

Configuration	Value
JSON web token (JWT) type	<input checked="" type="checkbox"/> ID <input type="checkbox"/> Refresh <input type="checkbox"/> Access
Payload claim value	<input type="checkbox"/> oid <input type="checkbox"/> aud <input checked="" type="checkbox"/> idp

Explanation

Configuration	Value
JSON web token (JWT) type	ID
Payload claim value	idp

Text, letter Description automatically generated

NEW QUESTION: 67

You need to resolve a notification latency issue.

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

NOTE: Each correct selection is worth one point.

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

Answer: (SHOW ANSWER)

Explanation

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

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

Anomaly detection service: You have an anomaly detection service that analyzes log information

for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Reference:

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

NEW QUESTION: 68

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

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

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B ([LEAVE A REPLY](#))

Explanation

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

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

NEW QUESTION: 69

You have an Azure Cosmos DB instance that uses the Strong consistency level and 10,000 Request Units (RUs) per container. <3eo-replication is enabled.

The instance stores restaurant information including location, menu items, and staff. You currently store information for 1,000 restaurant locations, 500 menu items, and 10,000 staff members. You select the location id as the partition key.

How many logical partitions will be created for the container?

A. 10,000,000

B. 1,100

C. 500

D. 10,000

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 70

you need to reduce read latency for the retail store solution.

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.** Provision an Azure Cosmos DB dedicated gateway, update blob storage to use the new dedicated gateway endpoint.
- B.** Create a new composite index for the store location data queries in Azure Cosmos DB. Modify the queries to support parameterized SQL and update the Azure function app to call the new Queries.
- C.** Provision an Azure Cosmos DB dedicated gateway Update the Azure Function app connection string to use the new dedicated gateway endpoint.
- D.** Configure Azure Cosmos DB consistency to strong consistency Increase the RUs for the container supporting store location data.
- E.** Configure Azure Cosmos DB consistency to session consistency. Cache session tokens in a new Azure Redis cache instance after every write. Update reads to use the session token stored in Azure Redis.

Answer: A,B ([LEAVE A REPLY](#))

NEW QUESTION: 71

You are developing an application to store millions of images in Azure blob storage. The images are uploaded to an Azure blob storage container named companyimages contained in an Azure blob storage account named companymedia. The stored images are uploaded with multiple blob index tags across multiple blobs in the container.


You must find all blobs whose tags match a search expression in the container. The search expression must evaluate an index tag named status with a value of final.

You need to construct the GET method request URL

How should you complete the URI? To answer, drag the appropriate parameters to the correct request URI targets. Each parameter 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.



Answer:



Explanation

D:\mudassar\Untitled.jpg



NEW QUESTION: 72

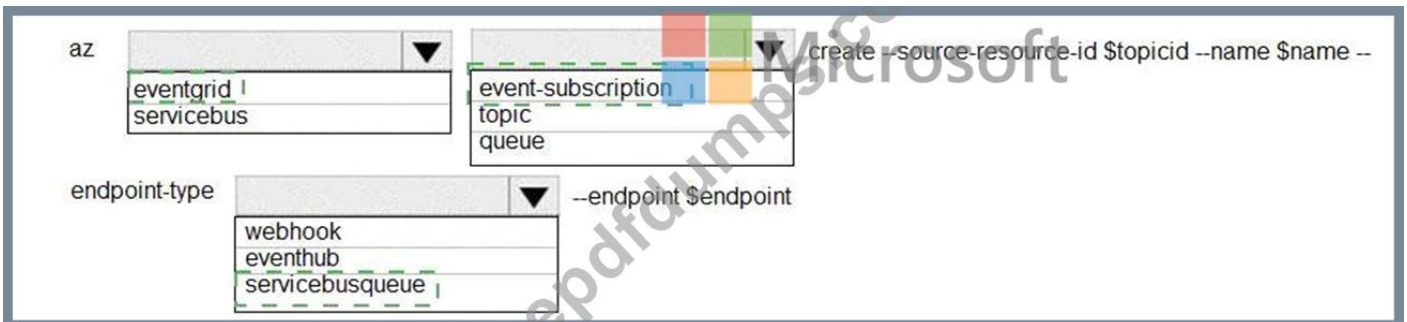
You need to configure the integration for Azure Service Bus and Azure Event Grid.

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

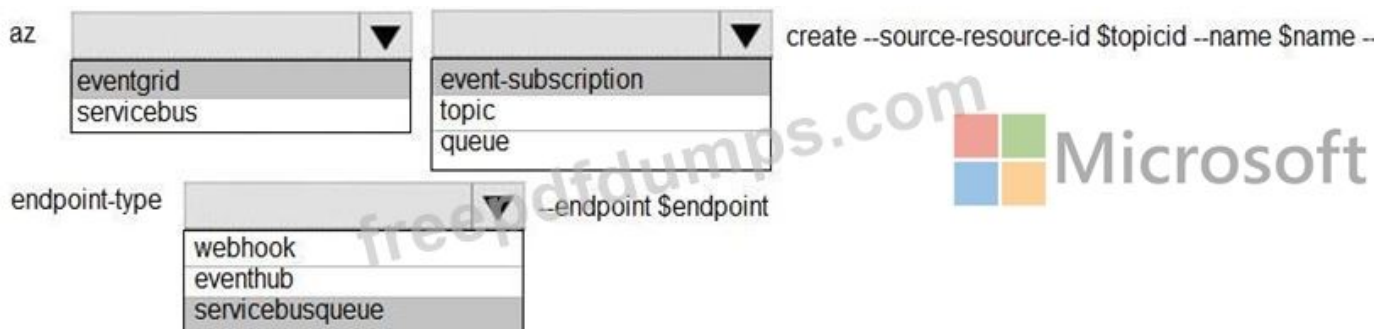
NOTE: Each correct selection is worth one point.



Answer:



Explanation



Box 1: eventgrid

To create event subscription use: az eventgrid event-subscription create Box 2: event-

subscription Box 3: servicebusqueue Scenario: 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.

Reference:

https://docs.microsoft.com/en-us/cli/azure/eventgrid/event-subscription?view=azure-cli-latest#az_eventgrid_eve

NEW QUESTION: 73

You develop and deploy an ASP.NET Core application that connects to an Azure Database for MySQL instance.

Connections to the database appear to drop intermittently and the application code does not handle the connection failure.

You need to handle the transient connection errors in code by implementing retries.

What are three possible ways to achieve this goal? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Disable connection pooling and configure a second Azure Database for MySQL instance.
- B. Set a maximum number of connection attempts to 10 and report an error on subsequent connections.
- C. Wait five seconds before repeating the connection attempt to the database.
- D. Increase connection repeat attempts exponentially up to 120 seconds.
- E. Close the database connection and immediately report an error.

Answer: A,C,E (LEAVE A REPLY)

NEW QUESTION: 74

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

Answer: ([SHOW ANSWER](#))

Reference:

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

NEW QUESTION: 75

You are developing a solution that will use a multi-partitioned Azure Cosmos DB database. You plan to use the latest Azure Cosmos DB SDK for development.

The solution must meet the following requirements:

- * Send insert and update operations to an Azure Blob storage account.
- * Process changes to all partitions immediately.
- * Allow parallelization of change processing.

You need to process the Azure Cosmos DB operations.

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

NOTE: Each correct selection is worth one point.

A. Create an Azure App Service API and implement the change feed estimator of the SDK. Scale the API by using multiple Azure App Service instances.

B. Create a background job in an Azure Kubernetes Service and implement the change feed feature of the SDK.

C. Create an Azure Function to use a trigger for Azure Cosmos DB. Configure the trigger to connect to the container.

D. Create an Azure Function that uses a FeedIterator object that processes the change feed by using the pull model on the container. Use a FeedRange object to parallelize the processing of the change feed across multiple functions.

Answer: ([SHOW ANSWER](#))

Explanation

Azure Functions is the simplest option if you are just getting started using the change feed. Due to its simplicity, it is also the recommended option for most change feed use cases. When you create an Azure Functions trigger for Azure Cosmos DB, you select the container to connect, and the Azure Function gets triggered whenever there is a change in the container. Because Azure Functions uses the change feed processor behind the scenes, it automatically parallelizes change processing across your container's partitions.

Note: You can work with change feed using the following options:

- * Using change feed with Azure Functions
- * Using change feed with change feed processor

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed>

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model>

<https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed#azure-functions>

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model#using-feedrange-for-parallelization>

NEW QUESTION: 76

You are building a website to access project data related to terms within your organization. The website does not allow anonymous access. Authentication performed using an Azure Active Directory (Azure AD) app named internal.

The website has the following authentication requirements:

*Azure AD users must be able to login to the website.

*Personalization of the website must be based on membership in Active Directory groups.

You need to configure the application's manifest to meet the authentication requirements.

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

NOTE: Each correct selection is worth one point.

```
{
  ...
  "appId": "d61126e3-089b-4adb-b721-
d5023213df7d",
  [ ] : "All",
  "optionalClaims"
  "groupMembershipClaims"
  [ ] : true
  "allowPublicClient"
  "oauth2Permissions"
  "requiredResourceAccess"
  "oauth2AllowImplicitFlow"
  ...
}
```

Answer:

```

{
  ...
  "appId": "d61126e3-089b-4adb-b721-
d5023213df7d",
  [redacted] : "All",
  "optionalClaims"
  "groupMembershipClaims"
  [redacted] : true
  "allowPublicClient"
  "oauth2Permissions"
  "requiredResourceAccess"
  "oauth2AllowImplicitFlow"
  ...
}

```

Explanation

Box 1: groupMembershipClaims

Personalization of the website must be based on membership in Active Directory groups.

Group claims can also be configured in the Optional Claims section of the Application Manifest.

Enable group membership claims by changing the groupMembershipClaim The valid values are:

- "All"
- "SecurityGroup"
- "DistributionList"
- "DirectoryRole"

Here we need to mention that we want to get the groups for the users. Hence we need to mention to set the groupMembershipClaims property to All.

Box 2: oauth2AllowImplicitFlow

Azure AD users must be able to login to the website.

oauth2Permissions can only accept collections value like an array, not a boolean.

oauth2AllowImplicitFlow accepts boolean value.

Here from the list of options given, if we want the application to fetch the required tokens , we would need to allow Implicit Flow.

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam!
Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 77

You have an Azure App Services Web App. Azure SQL Database instance. Azure Storage Account and an Azure Redis Cache instance in a resource group.

A developer must be able to publish code to the web app. You must grant the developer the Contributor role to the web app. You need to grant the role.

What two commands can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. New-AzureRmRoleAssignment
- B. az role assignment create
- C. az role definition create
- D. New-AzureRmRoleDefinition

Answer: A,B (LEAVE A REPLY)

Explanation

References:

<https://docs.microsoft.com/en-us/cli/azure/role/assignment?view=azure-cli-latest#az-role-assignment-create>

<https://docs.microsoft.com/en-us/powershell/module/azurerm.resources/new-azurermroleassignment?view=azure>

NEW QUESTION: 78

You need to implement the Log policy.

How should you complete the EnsureLogging method in EventGridController.cs? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

var client = new WebSiteManagementClient(. . .);
var id = ParseResourceID(resource);
var appSettings = new StringDictionary(name: "properties",
    properties: new Dictionary<string, string> {
        {"DIAGNOSTICS_AZUREBLOBCONTAINERSASURL", BlobStoreAccountSAS("
        logs
        logdrop
        ")

        {"DIAGNOSTICS_AZUREBLOBRETENTIONINDAYS", "
        15
        30
        "}

});
client.WebApps.
    UploadLoggingSettings
    UpdateApplicationSetting
    id.resourceGroup,
    id.name, appSettings);

```

Answer:

```

var client = new WebSiteManagementClient(. . .);
var id = ParseResourceID(resource);
var appSettings = new StringDictionary(name: "properties",
    properties: new Dictionary<string, string> {
        {"DIAGNOSTICS_AZUREBLOBCONTAINERSASURL", BlobStoreAccountSAS("
        logs
        logdrop
        ")

        {"DIAGNOSTICS_AZUREBLOBRETENTIONINDAYS", "
        15
        30
        "}

});
client.WebApps.
    UploadLoggingSettings
    UpdateApplicationSetting
    id.resourceGroup,
    id.name, appSettings);

```

Explanation

```

var client = new WebSiteManagementClient(...);
var id = ParseResourceID(resource);
var appSettings = new StringDictionary(name: "properties",
    properties: new Dictionary<string, string> {
        {"DIAGNOSTICS_AZUREBLOBCONTAINERSASURL", BlobStoreAccountSAS("
        {"DIAGNOSTICS_AZUREBLOBRETENTIONINDAYS", "
    });
client.WebApps.
    id.resourceGroup,
    id.name, appSettings);

```

Box 1: logdrop

All log files should be saved to a container named logdrop.

Box 2: 15

Logs must remain in the container for 15 days.

Box 3: UpdateApplicationSettings

All Azure App Service Web Apps must write logs to Azure Blob storage.

Reference:

<https://blog.hompus.nl/2017/05/29/adding-application-logging-blob-to-a-azure-web-app-service-using-powershe>

NEW QUESTION: 79

You develop and deploy a web app to Azure App Service. The Azure App Service uses a Basic plan in a region.

Users report that the web app is responding must capture the complete call stack to help performance issues in code. Call stack data must be correlated across app instances. You must minimize cost and impact to users on the web app.

You need to capture the telemetry.

Which three actions should you perform? Each answer presents part Of the solution NOTE: Each correct selection is worth point

- A. Upgrade the Azure App Service plan to Premium
- B. Enable Profiler.
- C. Enable the Always On setting for the app service.
- D. Enable remote debugging.
- E. Restart all apps in the App Service plan.
- F. Enable Snapshot debugger.
- G. Enable Application Insights site extensions.

Answer: C,E,F ([LEAVE A REPLY](#))

NEW QUESTION: 80

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

The app stores sensitive information about users.

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

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

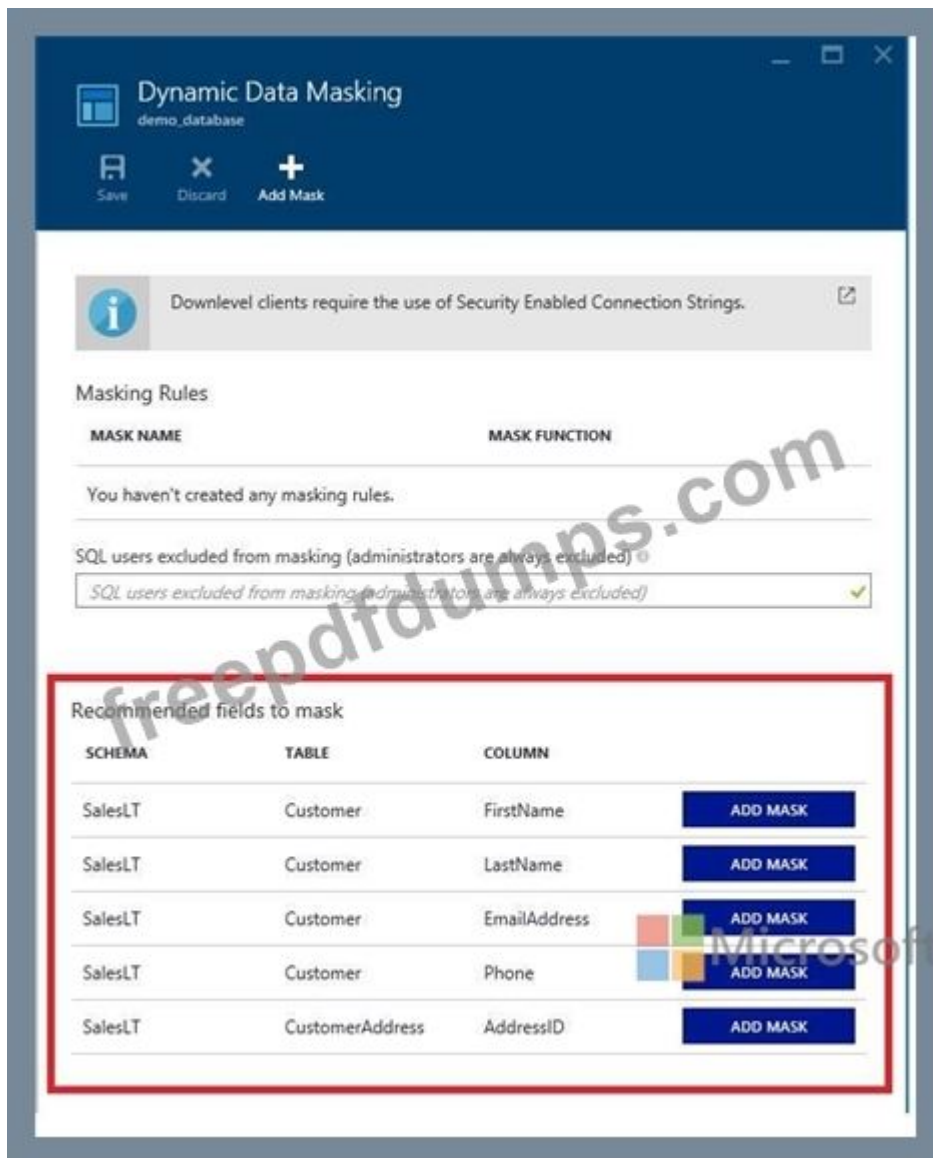
NOTE: Each correct selection is worth one point.

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

Answer: A,B,E ([LEAVE A REPLY](#))

Explanation

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



References:

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

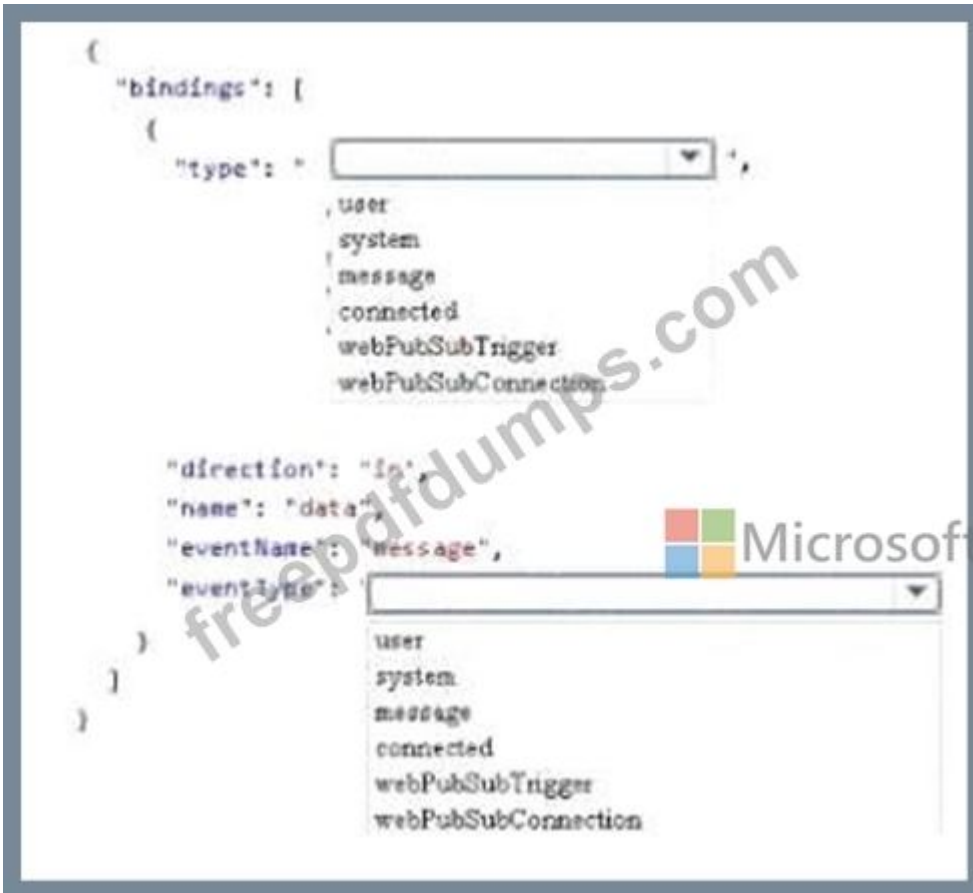
NEW QUESTION: 81

You are developing a service where customers can report news events from a browser using Azure Web PubSub. The service is implemented as an Azure App that the JSON WebSocket subprotocol to receive news events.

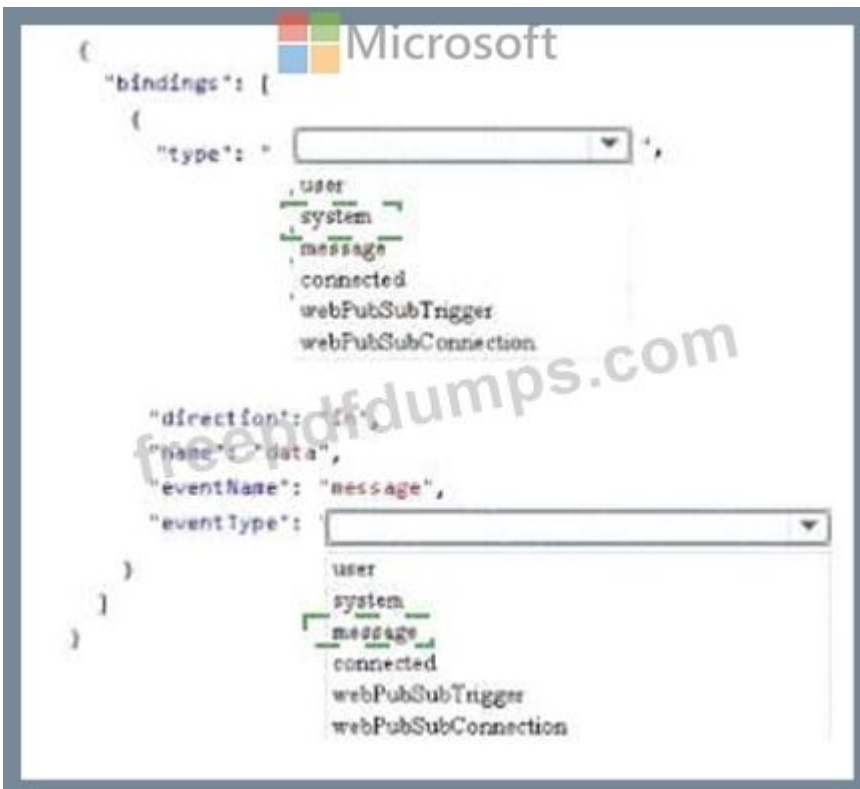
You need to implement the bindings for the Azure Function App.

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

Note: Each Correct Selection is worth one point.



Answer:



Explanation

Graphical user interface, text, application, chat or text message Description automatically generated

```

{
  "bindings": [
    {
      "type": "system",
      "direction": "in",
      "name": "data",
      "eventName": "message",
      "eventType": "message"
    }
  ]
}

```

NEW QUESTION: 82

You are Implementing an Azure solution that uses Azure Cosmos DB and the latest Azure Cosmos DB SDK.

You add a change feed processor to a new container instance.

You attempt to read a batch of 100 documents. The process falls when reading one of the documents. The solution must monitor the progress of the change feed processor instance on the new container as the change feed is read. You must prevent the change feed processor from retrying the entire batch when one document cannot be read.

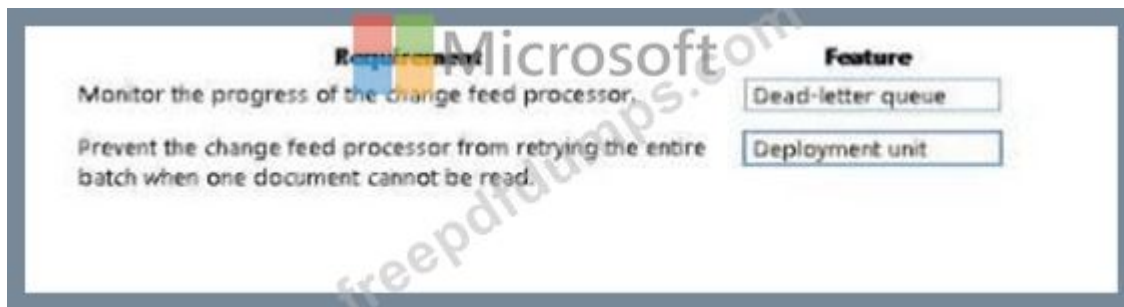
You need to implement the change feed processor to read the documents.

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 bat between panes or scroll to view content Each correct selection is worth one point

Answer:

Explanation

Text, letter Description automatically generated



NEW QUESTION: 83

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 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. Yes
- B. No

Answer: ([SHOW ANSWER](#))

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>

NEW QUESTION: 84

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.

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"/>

Answer:

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"/>

Explanation

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 type="radio"/>	<input type="radio"/>

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-](https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicebus.messaging.queuedescription.lockduration)

[us/dotnet/api/microsoft.servicebus.messaging.queuedescription.lockduration](https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicebus.messaging.queuedescription.lockduration)

NEW QUESTION: 85

You are developing an application to manage shipping information for cargo ships. The application will use Azure Cosmos DB for storage.

The application must run offline when ships are at sea. The application must be connected to Azure when ships are in port.

Which Azure Cosmos DB API should you use for the application?

- A. Core
- B. MongoDB
- C. Cassandra
- D. Gremlin

Answer: C (LEAVE A REPLY)

NEW QUESTION: 86

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

Answer: ([SHOW ANSWER](#))

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>

NEW QUESTION: 87

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.

App service plan setting


Value

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	



Answer:

App service plan setting **Value**

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

Explanation

App service plan setting

Value

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

Number of VM instances: 4

You are not charged extra for deployment slots.

Pricing tier: Isolated

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

References:

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

NEW QUESTION: 88

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

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

*Call setupScript.ps1 when the container is built.

*Run ContosoApp.dll when the container starts.

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


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

Commands	Answer Area
RUN powershell ./setupScript.ps1 CMD ["dotnet", "ContosoApp.dll"]	
EXPOSE ./ContosoApp/ /apps/ContosoApp	
COPY /.	
FROM microsoft/aspnetcore:2.0	
WORKDIR /apps/ContosoApp	
CMD powershell ./setupScript.ps1 ENTRYPOINT ["dotnet", "ContosoApp.dll"]	

 Microsoft

Answer:

Commands	Answer Area
RUN powershell ./setupScript.ps1 CMD ["dotnet", "ContosoApp.dll"]	WORKDIR /apps/ContosoApp
EXPOSE ./ContosoApp/ /apps/ContosoApp	COPY /.
COPY /.	EXPOSE ./ContosoApp/ /apps/ContosoApp
FROM microsoft/aspnetcore:2.0	
WORKDIR /apps/ContosoApp	CMD powershell ./setupScript.ps1 ENTRYPOINT ["dotnet", "ContosoApp.dll"]
CMD powershell ./setupScript.ps1 ENTRYPOINT ["dotnet", "ContosoApp.dll"]	

 Microsoft

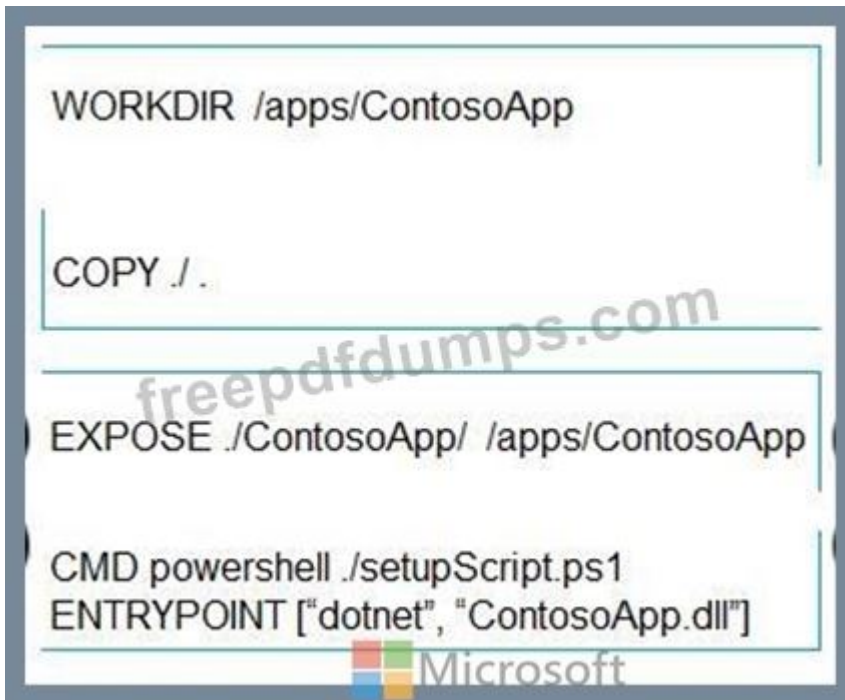
Explanation

```
WORKDIR /apps/ContosoApp

COPY ./

EXPOSE ./ContosoApp/ /apps/ContosoApp

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



Step 1: WORKDIR /apps/ContosoApp

Step 2: COPY ./-

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

Step 3: EXPOSE ./ContosoApp/ /app/ContosoApp

Step 4: CMD powershell ./setupScript.ps1

ENTRYPOINT ["dotnet", "ContosoApp.dll"]

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

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

References:

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

NEW QUESTION: 89

You are building a web application that performs image analysis on user photos and returns metadata containing objects identified. The image is very costly in terms of time and compute resources. You are planning to use Azure Redis Cache so duplicate uploads do not need to be reprocessed.

In case of an Azure data center outage, metadata loss must be kept to a minimum. You need to configure the Azure Redis cache instance.

Which two actions should you perform?

- A. Configure Azure Redis with rob persistence
- B. Configure Azure Redis with AOF persistence
- C. Configure second storage account far persistence.
- D. Set backup frequency to the minimum value.

Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 90

You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway. Data from traffic sensors are stored in Azure Event Hub.

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

Data throughout must be maximized. Latency must be minimized.

You need to implement the Azure Event Hub.

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

NOTE: Each correct selection is worth one point.

Setting



Value

Number of partitions

	▼
3	
4	
6	
12	

Partition Key

	▼
Highway	
Department	
Timestamp	
VM name	

Answer:

Setting

Value

Number of partitions

	▼
3	
4	
6	
12	

Partition Key

	▼
Highway	
Department	
Timestamp	
VM name	

Setting	Value
Number of partitions	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>3</p> <p>4</p> <p style="background-color: #ccc;">6</p> <p>12</p> </div> </div>
Partition Key	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p style="background-color: #ccc;">Highway</p> <p>Department</p> <p>Timestamp</p> <p>VM name</p> </div> </div>

Box 1: 6

The number of partitions is specified at creation and must be between 2 and 32.

There are 6 highways.

Box 2: Highway

References:

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

NEW QUESTION: 91

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.

Answer Area

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



freepdfdumps.com

Answer:

```

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

```

Explanation

```

{
  ...
  "resources": [
    {
      "apiVersion": "2016-01-01",
      "type": "Microsoft.Storage/storageAccounts",
      "name": "[concat(
        ( ), 'storage', uniqueString(resourceGroup().id))]",
      "location": "[resourceGroup().location]",
      ...
      "sku": {
        "name": "Standard_LRS"
      },

```

```

"kind": "Storage",
"properties": {},
"copy": {
  "name": "storagesetup",
  "count": 3
},
},
{
  "apiVersion": "2015-06-15",
  "type": "Microsoft.Compute/virtualMachines",
  "name": "[concat('VM', uniqueString(resourceGroup().id))]",
  "dependsOn": [
    "[variables('loadBalancerName')]",
    "[variables('virtualNetworkName')]",
    "storagesetup",
  ],
}

```

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>

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam!
Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:
https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**
Special Discount: Freepdfdumps)

NEW QUESTION: 92

You are developing an Azure Durable Function to manage an online ordering process.

The process must call an external API to gather product discount information.

You need to implement Azure Durable Function.

Which Azure Durable Function types should you use? Each correct answer presents part of the solution NOTE: Each correct selection is worth one point

- A. Orchestrator
- B. Entity
- C. Activity
- D. Client

Answer: (SHOW ANSWER)

Explanation

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview>

NEW QUESTION: 93

You are developing a solution by using the Azure Event Hubs SDK. You create a standard Azure Event Hub with 16 partitions. You implement eight event processor clients.

You must balance the load dynamically when an event processor client fails. When an event processor client fails, another event processor must continue processing from the exact point at which the failure occurred. All events must be aggregate and upload to an Azure Blob storage account You need to implement event processing recovery for the solution.

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

Each correct selection is worth one point.

Requirement



Ensure that event process clients mark the position within an event sequence.

Feature

- Offset
- Checkpoint
- Namespace
- Capture

Mark the event processor client position within a partition event sequence.

- Offset
- Checkpoint
- Namespace
- Capture

Answer:

Requirement

Ensure that event process clients mark the position within an event sequence.

Feature

- Offset
- Checkpoint
- Namespace
- Capture

Mark the event processor client position within a partition event sequence.

- Offset
- Checkpoint
- Namespace
- Capture



NEW QUESTION: 94

You have an Azure Batch project that processes and converts files and stores the files in Azure storage. You are developing a function to start the batch job.

You add the following parameters to the function.

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

You must ensure that converted files are placed in the container referenced by the outputContainerSasUrl parameter. Files which fail to convert are placed in the container referenced by the failedContainerSasUrl parameter.

You need to ensure the files are correctly processed.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
public List<CloudTasks> StartTasks(List<FileTask> fileTasks, string jobId,
string outputContainerSasUrl, string failedContainerSasUrl)
{
    BatchSharedKeyCredentials sharedKeyCredentials =
        new BatchSharedKeyCredentials(batchAccountUrl, batchAccountName,
batchAccountKey);
    List<CloudTask> tasks = new List<CloudTask>();
    using (BatchClient batchClient = BatchClient.Open(sharedKeyCredentials))
    {
        CloudJob = batchClient.JobOperations.  ();
        job.Id = jobId,
        job.PoolInformation = new PoolInformation { PoolId = poolId };
        job.Commit();
        fileTasks.ForEach((fileTask) =>
        {
            string taskId = $"Task{DateTime.Now.ToFileTimeUtc().ToString()}";
            CloudTask task = new CloudTask (taskId, fileTask.Command);
            List<OutputFile> outputFileList = new List<OutputFile>();
            OutputFileBlobContainerDestination outputContainer =
                new OutputFileBlobContainerDestination (outputContainerSasUrl);
            OutputFileBlobContainerDestination failedContainer =
                new OutputFileBlobContainerDestination (failedContainerSasUrl);
            outputFileList.Add(new OutputFile (fileTask.Output,
            new OutputFileDestination (outputContainer),
            new OutputFileUploadOptions (OutputFileUploadCondition.  ));
            outputFileList.Add(new OutputFile (fileTask.Output,
            new OutputFileDestination (failedContainer),
            new OutputFileUploadOptions (OutputFileUploadCondition,  ));

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

- GetJob
- GetTask
- EnableJob
- CreateJob

- TaskSuccess
- TaskFailure
- TaskCompletion

- TaskSuccess
- TaskFailure
- TaskCompletion

- OutputFiles
- FilesToStage
- ResourceFiles
- StageFiles

Answer:

Answer Area

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

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

            task.
                OutputFiles
                FilesToStage
                ResourceFiles
                StageFiles
            =outputFileList;

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

Explanation

```

CloudJob = batchClient.JobOperations.  ();

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

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

    task  = outputFileList;

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

```

Box 1: CreateJob

Box 2: TaskSuccess

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

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

Box 3: TaskFailure

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

Box 4: OutputFiles

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

References:

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

NEW QUESTION: 95

You plan to deploy a new application to a Linux virtual machine (VM) that is hosted in Azure. The entire VM must be secured at rest by using industry-standard encryption technology to address organizational security and compliance requirements.

You need to configure Azure Disk Encryption for the VM.

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

NOTE: Each correct selection is worth one point.

ANSWER AREA

The screenshot shows a series of Azure CLI commands and their interactive options:

```
az provider register -n Microsoft.KeyVault
resourcegroup="myResourceGroup"
az group create --name $resourcegroup --location westus
keyvault_name=myvaultname$RANDOM
az create \
  vm \
  keyvault \
  keyvault key \
  vm encryption \
  --enabled-for-disk-encryption True
az create \
  vm \
  keyvault \
  keyvault key \
  vm encryption \
  --software
az create \
  vm \
  keyvault \
  keyvault key \
  vm encryption \
  --os UbuntuServer:16.04-LTS:latest \
  --volume-type
az enable \
  vm \
  keyvault \
  keyvault key \
  vm encryption \
  --non-keyvault $keyvault_name \
  --non-key Name1 \
  --volume-type
```

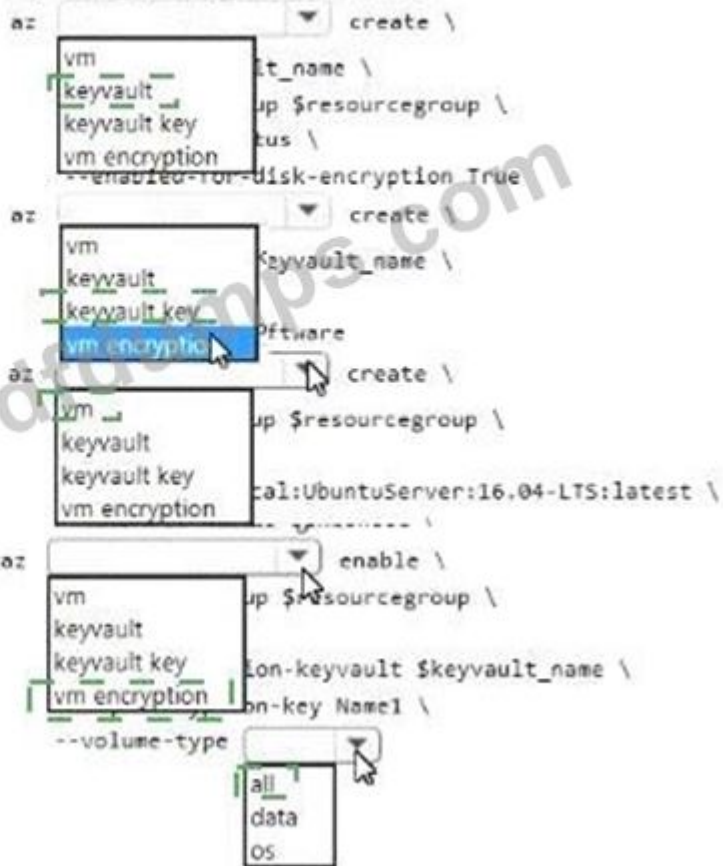
The interactive options shown are:

- vm encryption: --enabled-for-disk-encryption True
- vm encryption: --software
- vm encryption: --os UbuntuServer:16.04-LTS:latest
- vm encryption: --volume-type (options: all, data, OS)

Answer:

```
az provider register -n Microsoft.KeyVault  
resourcegroup="myResourceGroup"  
az group create --name $resourcegroup --location westus  
keyvault_name=myvaultname$RANDOM
```

```
az vm create \n  --name $vm_name \n  --resource-group $resourcegroup \n  --location $location \n  --image $image \n  --enable-secure-boot \n  --enable-disk-encryption True \n  --key-vault-name $keyvault_name \n  --key-name $key_name \n  --os-disk-name $os_disk_name \n  --os-disk-encryption-key-name $key_name \n  --volume-type $volume_type
```



Explanation

```
az provider register -n Microsoft.KeyVault
resourcegroup= "myResourceGroup"
az group create - -name $resourcegroup - -location westus
keyvault name=myvaultname$RANDOM
```

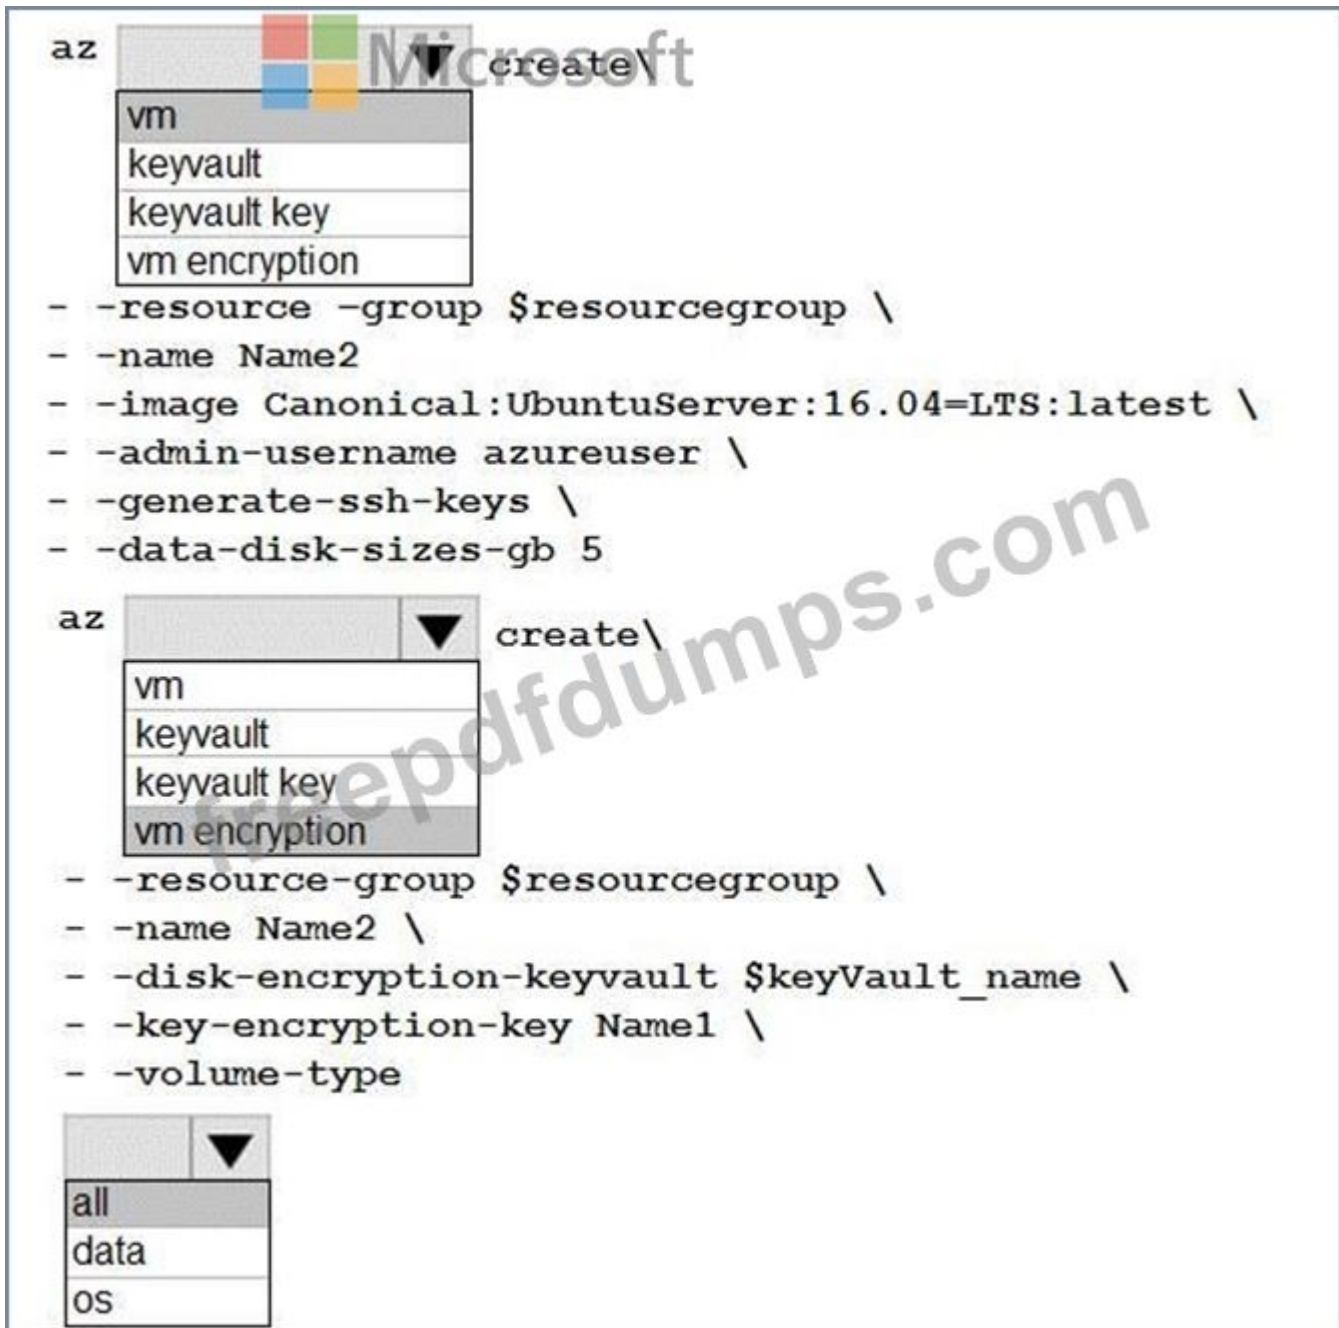
```
az  ▼ create\
vm
keyvault
keyvault key
vm encryption
```

```
- -name $keyvault_name \
- -resource -group $resourcegroup\
- -locstion eastus \
- -enabled for-disk-encryption True
```



```
az  ▼ create\
vm
keyvault
keyvault key
vm encryption
```

```
- -vault-name $keyvault_name\
- -name Name1 \
- -protection software
```



Box 1: keyvault

Create an Azure Key Vault with az keyvault create and enable the Key Vault for use with disk encryption.

Specify a unique Key Vault name for keyvault_name as follows:

keyvault_name=myvaultname\$RANDOM

```
az keyvault create \
```

```
--name $keyvault_name \
```

```
--resource-group $resourcegroup \
```

```
--location eastus \
```

```
--enabled-for-disk-encryption True
```

Box 2: keyvault key

The Azure platform needs to be granted access to request the cryptographic keys when the VM boots to decrypt the virtual disks. Create a cryptographic key in your Key Vault with az keyvault

key create. The following example creates a key named myKey:

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

Box 3: vm

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

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

Box 4: vm encryption

Encrypt your VM with az vm encryption enable:

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

Note: seems to be an error in the question. Should have enable instead of create.

Box 5: all

Encrypt both data and operating system.

References:

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

NEW QUESTION: 96

You are developing a web application that uses the Microsoft identity platform to authenticate users and resources. The web application calls several REST APIs.

The APIs require an access token from the Microsoft identity platform.

You need to request a token.

Which three properties should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Supported account type
- B. Application ID
- C. Redirect URI/URL
- D. Application name
- E. Application secret

Answer: B,C,E (LEAVE A REPLY)

NEW QUESTION: 97

You are developing a solution that uses the Azure Storage Client library for .NET. You have the following code: (Line numbers are included for reference only.)

```
01 CloudBlockBlob src = null;
02 try
03 {
04     src = container.ListBlobs().OfType<CloudBlockBlob>().FirstOrDefault();
05     var id = await src.AcquireLeaseAsync(null);
06     var dst = container.GetBlockBlobReference(src.Name);
07     string cpid = await dst.StartCopyAsync(src);
08     await dst.FetchAttributeAsync();
09     return id;
10 }
11 catch (Exception e)
12 {
13     throw;
14 }
15 finally
16 {
17     if (src != null)
18         await src.FetchAttributesAsync();
19     if (src.Properties.LeaseState != LeaseState.Available)
20         await src.BreakLeaseAsync(new TimeSpan(0));
21 }
```

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

NOTE: Each correct selection is worth one point.

Statement	Yes	No
The code creates an infinite lease	<input type="radio"/>	<input type="radio"/>
The code at line 06 always creates a new blob	<input type="radio"/>	<input type="radio"/>
The finally block releases the lease	<input type="radio"/>	<input type="radio"/>

Answer:

Statement	Yes	No
The code creates an infinite lease	<input checked="" type="radio"/>	<input type="radio"/>
The code at line 06 always creates a new blob	<input type="radio"/>	<input checked="" type="radio"/>
The finally block releases the lease	<input type="radio"/>	<input type="radio"/>

Explanation

Statement	Yes	No
The code creates an infinite lease	<input checked="" type="radio"/>	<input type="radio"/>
The code at line 06 always creates a new blob	<input type="radio"/>	<input checked="" type="radio"/>
The finally block releases the lease	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: Yes

AcquireLeaseAsync does not specify leaseTime.

leaseTime is a TimeSpan representing the span of time for which to acquire the lease, which will be rounded down to seconds. If null, an infinite lease will be acquired. If not null, this must be 15 to 60 seconds.

Box 2: No

The GetBlockBlobReference method just gets a reference to a block blob in this container.

Box 3: Yes

The BreakLeaseAsync method initiates an asynchronous operation that breaks the current lease on this container.

Reference:

[https://docs.microsoft.com/en-](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync)

[us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync)

[https://docs.microsoft.com/en-](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.getblockblobreference)

[us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.getblockblobreferen](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.getblockblobreference)

[https://docs.microsoft.com/en-](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync)

[us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync)

NEW QUESTION: 98

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

Answer: B ([LEAVE A REPLY](#))

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>

NEW QUESTION: 99

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

Answer: B ([LEAVE A REPLY](#))

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>

NEW QUESTION: 100

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.

App setting	Value
<input type="checkbox"/>	true
<input type="checkbox"/> LOCALAPPDATA	<input type="checkbox"/>
<input type="checkbox"/> WEBSITE_LOCALCACHE_ENABLED	/home
<input type="checkbox"/> DOTNET_HOSTING_OPTIMIZATION_CACHE	/local
<input type="checkbox"/> WEBSITES_ENABLE_APP_SERVICE_STORAGE	D:\home
<input type="checkbox"/> DIAGDATA	D:\local

Answer:

App setting	Value
<input checked="" type="checkbox"/>	true
<input type="checkbox"/> LOCALAPPDATA	<input type="checkbox"/>
<input type="checkbox"/> WEBSITE_LOCALCACHE_ENABLED	/home
<input type="checkbox"/> DOTNET_HOSTING_OPTIMIZATION_CACHE	/local
<input type="checkbox"/> WEBSITES_ENABLE_APP_SERVICE_STORAGE	D:\home
<input type="checkbox"/> DIAGDATA	D:\local

Explanation

App setting	Value
LOCALAPPDATA	true
WEBSITE_LOCALCACHE_ENABLED	
DOTNET_HOSTING_OPTIMIZATION_CACHE	
WEBSITES_ENABLE_APP_SERVICE_STORAGE	
DIAGDATA	
	/home
	/local
	D:\home
	D:\local

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>

NEW QUESTION: 101

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

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

You need to select a partition key.

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

NOTE: Each correct selection is worth one point.

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

Answer: A,C (LEAVE A REPLY)

Explanation

You can form a partition key by concatenating multiple property values into a single artificial partitionKey property. These keys are referred to as synthetic keys.

Another possible strategy to distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in this way, you can perform parallel write operations across partitions.

Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands.

The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.

References:

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

NEW QUESTION: 102

You must ensure that the external party cannot access the data in the SSN column of the Person table.

Will each protection method meet the requirement? To answer, drag the appropriate responses to the correct protection methods. Each response may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Responses	Protection method	Response
<input type="checkbox"/> Yes	Enable AlwaysOn encryption.	<input type="checkbox"/>
<input type="checkbox"/> No	Set the column encryption setting to disabled.	<input type="checkbox"/>
	Assign users to the Public fixed database role.	<input type="checkbox"/>
	Store column encryption keys in the system catalog view in the database.	<input type="checkbox"/>

Answer:

Responses	Protection method	Response
<input checked="" type="checkbox"/> Yes	Enable AlwaysOn encryption.	<input checked="" type="checkbox"/> Yes
<input checked="" type="checkbox"/> No	Set the column encryption setting to disabled.	<input checked="" type="checkbox"/> No
	Assign users to the Public fixed database role.	<input checked="" type="checkbox"/> Yes
	Store column encryption keys in the system catalog view in the database.	<input checked="" type="checkbox"/> No

Explanation

Responses	Protection method	Response
<input type="checkbox"/> Yes	Enable AlwaysOn encryption.	<input type="checkbox"/> Yes
<input type="checkbox"/> No	Set the column encryption setting to disabled.	<input type="checkbox"/> No
	Assign users to the Public fixed database role.	<input type="checkbox"/> Yes
	Store column encryption keys in the system catalog view in the database.	<input type="checkbox"/> No

Box 1: Yes

You can configure Always Encrypted for individual database columns containing your sensitive data. When setting up encryption for a column, you specify the information about the encryption algorithm and cryptographic keys used to protect the data in the column.

Box 2: No

Box 3: Yes

In SQL Database, the VIEW permissions are not granted by default to the public fixed database role. This enables certain existing, legacy tools (using older versions of DacFx) to work properly. Consequently, to work with encrypted columns (even if not decrypting them) a database administrator must explicitly grant the two VIEW permissions.

Box 4: No

All cryptographic keys are stored in an Azure Key Vault.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

NEW QUESTION: 103

You develop and deploy a web app to Azure App service. The web app allows users to authenticate by using social identity providers through the Azure B2C service. All user profile information is stored in Azure B2C.

You must update the web app to display common user properties from Azure B2C to include the following information:

- * Email address
- * Job title
- * First name
- * Last name
- * Office Location

You need to implement the user properties in the web app.

Requirement

API to access user properties

Value

- Microsoft Graph
- Azure AD Graph
- Azure Key Vault
- Azure AD entitlement management

Code library to interface to Azure AD B2C

- Microsoft Authentication Library (MSAL)
- Microsoft Azure Key Vault SDK
- Azure Identity library

Answer:

Requirement

API to access user properties

Value

- Microsoft Graph
- Azure AD Graph
- Azure Key Vault
- Azure AD entitlement management

Code library to interface to Azure AD B2C

- Microsoft Authentication Library (MSAL)
- Microsoft Azure Key Vault SDK
- Azure Identity library

Explanation

Requirement

API to access user properties

Value

Azure AD Graph

Code library to interface to Azure AD B2C

Azure Identity library

NEW QUESTION: 104

You are developing an Azure App Service REST API.

The API must be called by an Azure App Service web app. The API must retrieve and update user profile information stored in Azure Active Directory (Azure AD).

You need to configure the API to make the updates.

Which two tools should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Microsoft Graph API

- B. Microsoft Authentication Library (MSAL)
- C. Azure API Management
- D. Microsoft Azure Security Center
- E. Microsoft Azure Key Vault SDK

Answer: A,C (LEAVE A REPLY)

Explanation

A: You can use the Azure AD REST APIs in Microsoft Graph to create unique workflows between Azure AD resources and third-party services.

Enterprise developers use Microsoft Graph to integrate Azure AD identity management and other services to automate administrative workflows, such as employee onboarding (and termination), profile maintenance, license deployment, and more.

C: API Management (APIM) is a way to create consistent and modern API gateways for existing back-end services.

API Management helps organizations publish APIs to external, partner, and internal developers to unlock the potential of their data and services.

Reference:

<https://docs.microsoft.com/en-us/graph/azuread-identity-access-management-concept-overview>

NEW QUESTION: 105

You are developing an application that uses a premium block blob storage account. You are optimizing costs by automating Azure Blob Storage access tiers.

You apply the following policy rules to the storage account. You must determine the implications of applying the rules to the data. (Line numbers are included for reference only.)

```

01 {
02   "rules":
03   {
04     "name": "agingDataRule",
05     "enabled": true,
06     "type": "Microsoft

```

Answer Area		Yes	No
Block blobs prefixed with container1/salesorders or container2/inventory which have not been modified in over 60 days are moved to cool storage. Blobs that have not been modified in 120 days are moved to the archive tier.	<input type="radio"/>	<input type="radio"/>	
Blobs are moved to cool storage if they have not been accessed for 30 days.	<input type="radio"/>	<input type="radio"/>	
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.	<input type="radio"/>	<input type="radio"/>	
All block blobs older than 730 days will be deleted.	<input type="radio"/>	<input type="radio"/>	

Answer:

Answer Area

<p>Block blobs prefixed with container1/salesorders or container2/inventory which have not been modified in over 60 days are moved to cool storage. Blobs that have not been modified in 120 days are moved to the archive tier.</p> <p>Blobs are moved to cool storage if they have not been accessed for 30 days.</p> <p>Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.</p> <p>All block blobs older than 730 days will be deleted.</p>	<p>Yes</p> <input checked="" type="radio"/>	<p>No</p> <input type="radio"/>
---	--	--



Explanation

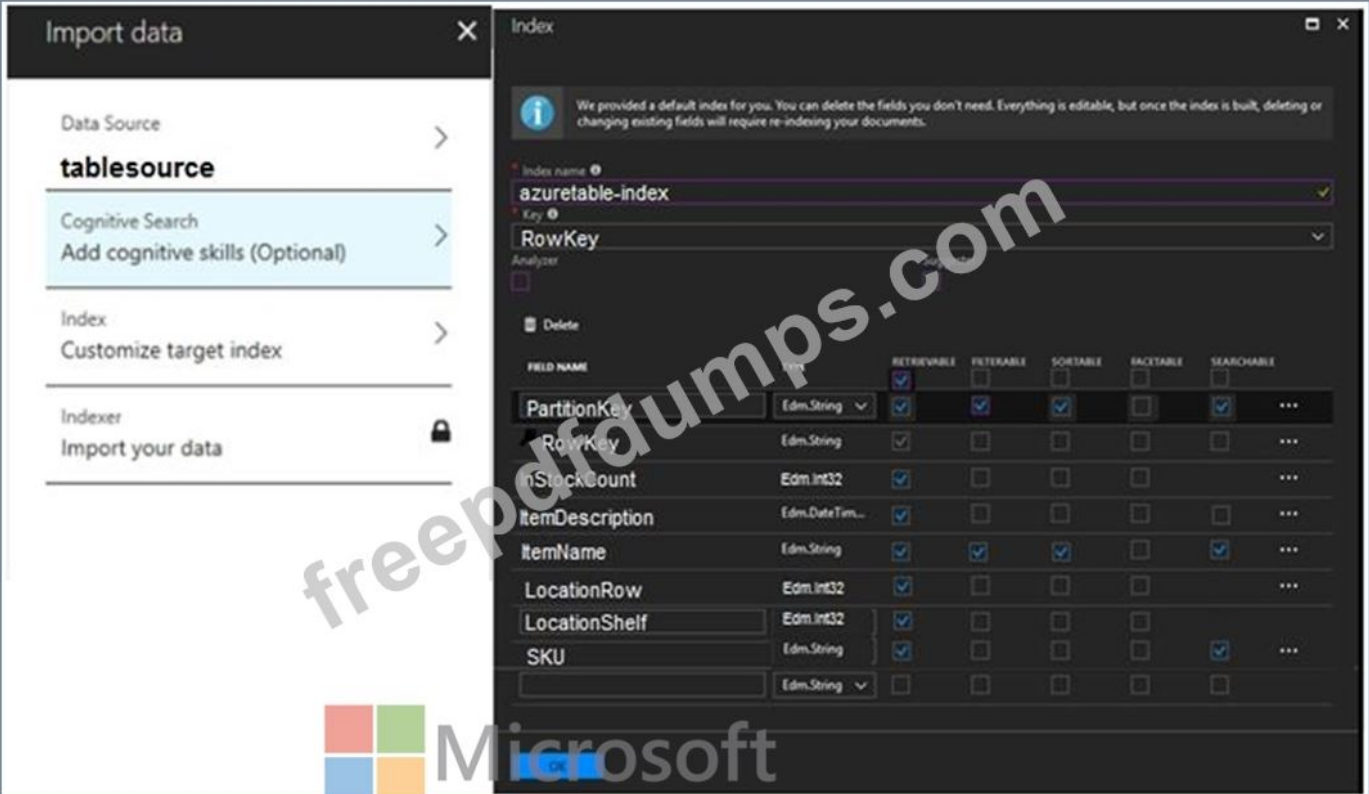
1. Yes
2. Yes
3. Yes
4. No

<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#move-a>


NEW QUESTION: 106

You are validating the configuration of an Azure Search indexer.

The service has been configured with an indexer that uses the Import Data option. The index is configured using options as shown in the Index Configuration exhibit. (Click the Index Configuration tab.)



FIELD NAME	TYPE	RETREIVABLE	PREFERABLE	SORTABLE	FACTABLE	SEARCHABLE
PartitionKey	Edm.String	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
RowKey	Edm.String	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
ItemStockCount	Edm.Int32	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
ItemDescription	Edm.DateTim...	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
ItemName	Edm.String	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
LocationRow	Edm.Int32	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
LocationShelf	Edm.Int32	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
SKU	Edm.String	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>



You use an Azure table as the data source for the import operation. The table contains three records with item inventory data that matches the fields in the Storage data exhibit. These records

were imported when the index was created. (Click the Storage Data tab.) When users search with no filter, all three records are displayed.

Category	RowKey	Timestamp	StockCount	ItemDescription	ItemName	LocationRow	LocationShelf	SKU
Food	3	2018-08-28T15:47:29.130Z	32	A box of chocolate candy bars	Chocolate bar	5	3	123421
Hardware	2	2018-08-28T15:46:08.405Z	2	A bag of bolts	Bolts	1	4	678964
Hardware	1	2018-08-28T15:46:41.402Z	23	A box of nails	Nails	2	1	654365



When users search for items by description, Search explorer returns no records. The Search Explorer exhibit shows the query and results for a test. In the test, a user is trying to search for all items in the table that have a description that contains the word bag. (Click the Search Explorer tab.) You need to resolve the issue.

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

NOTE: Each correct selection is worth one point.

- | | Yes | No |
|--|-----------------------|-----------------------|
| You can resolve the issue by recreating the search index with the same settings for all fields except ItemDescription. Select the SEARCHABLE option for this field | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field. | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by running the indexer. | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results | <input type="radio"/> | <input type="radio"/> |

Answer:

	Yes	No
You can resolve the issue by recreating the search index with the same settings for all fields except ItemDescription. Select the SEARCHABLE option for this field	<input type="radio"/>	<input type="radio"/>
You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field.	<input type="radio"/>	<input checked="" type="radio"/>
You can resolve the issue by running the indexer.	<input checked="" type="radio"/>	<input type="radio"/>
You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results	<input type="radio"/>	<input checked="" type="radio"/>

Explanation

	Yes	No
You can resolve the issue by recreating the search index with the same settings for all fields except ItemDescription. Select the SEARCHABLE option for this field	<input type="radio"/>	<input type="radio"/>
You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field.	<input type="radio"/>	<input checked="" type="radio"/>
You can resolve the issue by running the indexer.	<input checked="" type="radio"/>	<input type="radio"/>
You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results	<input type="radio"/>	<input checked="" type="radio"/>

Box 1: Yes

The ItemDescription field is not searchable.

Box 2: No

The ItemDescription field is not searchable, but we would need to recreate the index.

Box 3: Yes

An indexer in Azure Search is a crawler that extracts searchable data and metadata from an external Azure data source and populates an index based on field-to-field mappings between the index and your data source.

This approach is sometimes referred to as a 'pull model' because the service pulls data in without you having to write any code that adds data to an index.

Box 4: No

References:

<https://docs.microsoft.com/en-us/azure/search/search-what-is-an-index>

<https://docs.microsoft.com/en-us/azure/search/search-indexer-overview>

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam!

Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:

Special Discount: **Freepdfdumps**)

NEW QUESTION: 107

You deploy an API to API Management

You must secure all operations on the API by using a client certificate.

You need to secure access to the backend service of the API by using client certificates.

Which two security features can you use?

- A. Certificate Authority (CA) certificate
- B. Azure AD token
- C. Subscription key
- D. Self-signed certificate
- E. Triple DES (3DES) cipher

Answer: A,D ([LEAVE A REPLY](#))

NEW QUESTION: 108

You are maintaining an existing application that uses an Azure Blob GPV1 Premium storage account. Data older than three months is rarely used.

Data newer than three months must be available immediately. Data older than a year must be saved but does not need to be available immediately.

You need to configure the account to support a lifecycle management rule that moves blob data to archive storage for data not modified in the last year.

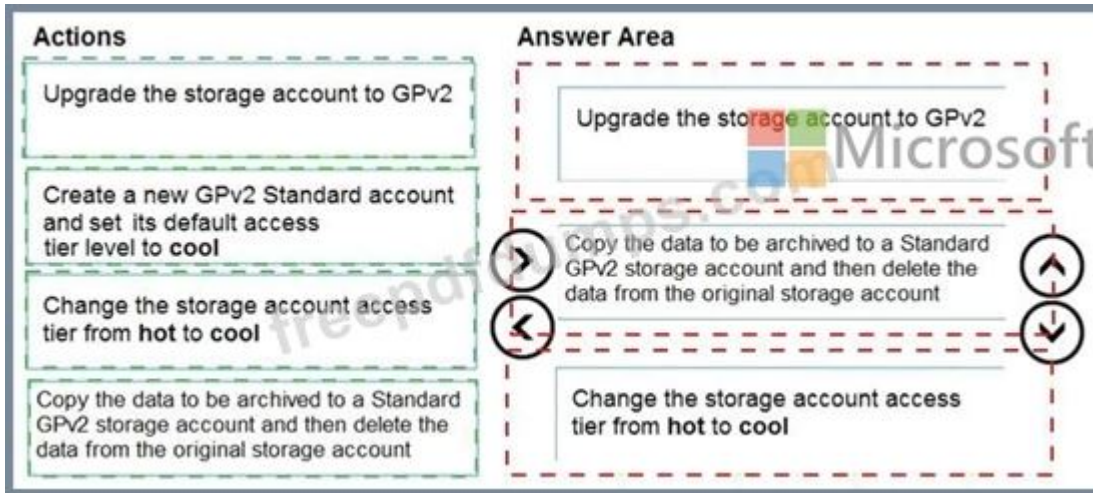
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

The screenshot shows an interactive question interface. On the left, under the heading "Actions", there are four items in a list:

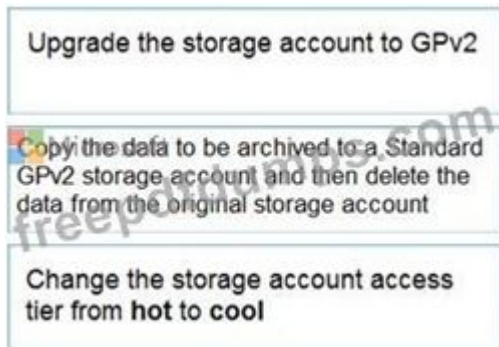
- Upgrade the storage account to GPv2
- Create a new GPv2 Standard account and set its default access tier level to cool
- Change the storage account access tier from hot to cool
- Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account

On the right, under the heading "Answer Area", there are two vertical arrows: an upward arrow and a downward arrow, indicating that the selected actions can be reordered. A watermark "freepdfdumps.com" is visible across the interface, and the Microsoft logo is at the bottom.

Answer:



Explanation



Step 1: Upgrade the storage account to GPv2

Object storage data tiering between hot, cool, and archive is supported in Blob Storage and General Purpose v2 (GPv2) accounts. General Purpose v1 (GPv1) accounts don't support tiering. You can easily convert your existing GPv1 or Blob Storage accounts to GPv2 accounts through the Azure portal.

Step 2: Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account Step 3: Change the storage account access tier from hot to cool Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Only the hot and cool access tiers can be set at the account level. The archive access tier can only be set at the blob level.

Reference:

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

NEW QUESTION: 109

You need to implement the Log policy.

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

NOTE: Each correct selection is worth one point.

Code segment

- All
- WebHook
- EventHub
- subjectEndsWith
- Microsoft.Storage
- subjectBeginsWith
- Microsoft.Storage.BlobCreated

Answer:

Code segment

- All
- WebHook
- EventHub
- subjectEndsWith
- Microsoft.Storage
- subjectBeginsWith
- Microsoft.Storage.BlobCreated

Answer Area

```
{
  "name": "newlogs",
  "properties": {
    "topic": "/subscriptions/.../providers/Microsoft.EventGrid/topics/...",
    "destination": {
      "endpointType": "code segment"
    },
    "filter": {
      "code segment": "/blobServices/default/containers/logdrop/"
    },
    "includedEventTypes": [ "code segment" ]
  },
  "labels": [],
  "eventDeliverySchema": "EventGridSchema"
}
```

Answer Area

```
{
  "name": "newlogs",
  "properties": {
    "topic": "/subscriptions/.../providers/Microsoft.EventGrid/topics/...",
    "destination": {
      "endpointType": "WebHook"
    },
    "filter": {
      "subjectBeginsWith": "/blobServices/default/containers/logdrop/"
    },
    "includedEventTypes": [ "Microsoft.Storage.BlobCreated" ]
  },
  "labels": [],
  "eventDeliverySchema": "EventGridSchema"
}
```

Explanation

```
{
  "name": "newlogs",
  "properties": {
    "topic": "/subscriptions/.../providers/Microsoft.EventGrid/topics/...",
    "destination": {
      "endpointType": "WebHook"
    },
    "filter": {
      "subjectBeginsWith": "/blobServices/default/containers/logdrop/"
    },
    "includedEventTypes": [ "Microsoft.Storage.BlobCreated" ]
  },
  "labels": [],
  "eventDeliverySchema": "EventGridSchema"
}
```

Box 1:WebHook

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

endpointType: The type of endpoint for the subscription (webhook/HTTP, Event Hub, or queue).

Box 2: SubjectBeginsWith

Box 3: Microsoft.Storage.BlobCreated

Scenario: Log Policy

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

Example subscription schema

{

```

"properties": {
"destination": {
"endpointType": "webhook",
"properties": {
"endpointUrl":
"https://example.azurewebsites.net/api/HttpTriggerCSharp1?
code=VXbGWce53l48Mt8wuotr0GPmyJ/nDT4hgd
}
},
"filter": {
"includedEventTypes": [ "Microsoft.Storage.BlobCreated", "Microsoft.Storage.BlobDeleted" ],
"subjectBeginsWith": "blobServices/default/containers/mycontainer/log",
"subjectEndsWith": ".jpg",
"isSubjectCaseSensitive ": "true"
}
}
}

```

Reference:

<https://docs.microsoft.com/en-us/azure/event-grid/subscription-creation-schema>

Topic 5, Litware Inc

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

Overall architecture

Employees upload receipts for the system to process. When processing is complete, the

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

Receipt processing

Employees may upload receipts in two ways:

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

Data Storage

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

Documentation

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

Solution details

Users table

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

Web Application

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

Processing

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

Logging

Azure Application Insights is used for telemetry and logging in both the processor and the web application.

The processor also has TraceWriter logging enabled. Application Insights must always contain all log messages.

Requirements

Receipt processing

Concurrent processing of a receipt must be prevented.

Disaster recovery

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

Security

* User's SecurityPin must be stored in such a way that access to the database does not allow the

viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.

- * All certificates and secrets used to secure data must be stored in Azure Key Vault.
- * You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
- * All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
- * Receipt data must always be encrypted at rest.
- * All data must be protected in transit.
- * User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
- * In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues

Upload format issue

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

Capacity issue

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

Log capacity issue

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

Application code

Processing.cs

```

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

```



Database.cs

```

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

```

ReceiptUploader.cs

```

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

```

ConfigureSSE.ps1

```

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

```

NEW QUESTION: 110

You are working for Contoso, Ltd.

You define an API Policy object by using the following XML markup:

```

<set-variable name= "bodySize" value="@ (context.Request.Headers["Content-Length"] [0])"/>
<choose>
  <when condition= "@ (int.Parse(context.Variables.GetValueOrDefault<string> ("bodySize"))<512000)">
  </when>
  <otherwise>
    <rewrite-uri template= "/put"/>
    <set-backend-service base-url= "http://contoso.com/api/9.1"/>
  </otherwise>
</choose>

```

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

NOTE: Each correct selection is worth one point.

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	<input type="radio"/>	<input type="radio"/>
If the body size is >256k, an error will occur.	<input type="radio"/>	<input type="radio"/>
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	<input type="radio"/>	<input type="radio"/>

Answer:

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	<input type="checkbox"/>	<input type="checkbox"/>
If the body size is >256k, an error will occur.	<input type="checkbox"/>	<input type="checkbox"/>
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	<input type="checkbox"/>	<input type="checkbox"/>

Explanation

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
If the body size is >256k, an error will occur.	<input type="checkbox"/>	<input checked="" type="checkbox"/>
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Box 1: Yes

Use the set-backend-service policy to redirect an incoming request to a different backend than the one specified in the API settings for that operation. Syntax: <set-backend-service base-url="base URL of the backend service" /> Box 2: No The condition is on 512k, not on 256k.

Box 3: No

The set-backend-service policy changes the backend service base URL of the incoming request to the one specified in the policy.

Reference:

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

NEW QUESTION: 111

You develop and deploy an Azure Logic app that calls an Azure Function app. The Azure Function app includes an OpenAPI (Swagger) definition and uses an Azure Blob storage account. All resources are secured by using Azure Active Directory (Azure AD).

The Azure Logic app must securely access the Azure Blob storage account. Azure AD resources must remain if the Azure Logic app is deleted.

You need to secure the Azure Logic app.

What should you do?

- A. Create an Azure AD custom role and assign role-based access controls.
- B. Create an Azure AD custom role and assign the role to the Azure Blob storage account.
- C. Create an Azure Key Vault and issue a client certificate.
- D. Create a user-assigned managed identity and assign role-based access controls.
- E. Create a system-assigned managed identity and issue a client certificate.

Answer: D (LEAVE A REPLY)

Explanation

To give a managed identity access to an Azure resource, you need to add a role to the target resource for that identity.

Note: To easily authenticate access to other resources that are protected by Azure Active Directory (Azure AD) without having to sign in and provide credentials or secrets, your logic app can use a managed identity (formerly known as Managed Service Identity or MSI). Azure manages this identity for you and helps secure your credentials because you don't have to provide or rotate secrets.

If you set up your logic app to use the system-assigned identity or a manually created, user-assigned identity, the function in your logic app can also use that same identity for authentication.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/create-managed-service-identity>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates-for-clients>

NEW QUESTION: 112

You manage a data processing application that receives requests from an Azure Storage queue. You need to manage access to the queue. You have the following requirements:

- * Provide other applications access to the Azure queue.
- * Ensure that you can revoke access to the queue without having to regenerate the storage account keys.
- * Specify access at the queue level and not at the storage account level.

Which type of shared access signature (SAS) should you use?

- A. Service SAS with a stored access policy
- B. Account SAS
- C. User Delegation SAS
- D. Service SAS with ad hoc SAS

Answer: (SHOW ANSWER)

Explanation

A service SAS is secured with the storage account key. A service SAS delegates access to a resource in only one of the Azure Storage services: Blob storage, Queue storage, Table storage, or Azure Files.

Stored access policies give you the option to revoke permissions for a service SAS without having to regenerate the storage account keys.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

NEW QUESTION: 113

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

Answer: ([SHOW ANSWER](#))

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>

NEW QUESTION: 114

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.

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
            (windowsToastPayload);
    ...
}
catch (System.Exception ex)
{
    ...
}
...
```

Answer:

Answer Area

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



Explanation

NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

GetInstallation
CreateClientFromConnectionString
CreateOrUpdateInstallation
PatchInstallation

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

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-manage>

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/app-service-mobile/app-service-mobile-windo>

NEW QUESTION: 115

You are developing an ASP.NET Core time sheet application that runs as an Azure Web App.

Users of the application enter their time sheet information on the first day of every month.

The application uses a third-party web service to validate data.

The application encounters periodic server errors due to errors that result from calling a third-party web server.

Each request to the third-party server has the same chance of failure.

You need to configure an Azure Monitor alert to detect server errors unrelated to the third-party

service. You must minimize false-positive alerts.

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.

```
    "type": "Microsoft.Insights/metricAlerts",
    "properties": {
      "criteria": {
        "odata.type": "...",
        "allOf": [
          {
            "criterionType": "
            DynamicThresholdCriterion
            SingleResourceMultipleMetricCriteria

            "metricName": "
            Http4xx
            Http5xx

            "alertSensitivity": "
            Low
            High

          }
        ]
      }
    }
  }
}
```

Answer:

```
    "type": "Microsoft.Insights/metricAlerts",
    "properties": {
      "criteria": {
        "odata.type": "...",
        "allOf": [
          {
            "criterionType": "
            DynamicThresholdCriterion
            SingleResourceMultipleMetricCriteria

            "metricName": "
            Http4xx
            Http5xx

            "alertSensitivity": "
            Low
            High

          }
        ]
      }
    }
  }
}
```

Explanation

Graphical user interface Description automatically generated

```

"type": "Microsoft.Insights/metricAlerts",
"properties": {
"criteria": {
"odata.type": ". . .",
"allOf": [
{
"criteriaType": "DynamicThresholdCriterion",
"metricName": "Http5xx",
>alertSensitivity": "Low"
}
]
}
}
}
}

```

Box 1: DynamicThresholdCriterion

Box 2: Http5xx

Server errors are in the 5xx range.

Client errors are in the 4xx range

Box 3: Low

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-dynamic-thresholds>

NEW QUESTION: 116

An organization plans to deploy Azure storage services.

You need to configure shared access signature (SAS) for granting access to Azure Storage.

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

Each SAS 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.

SAS types

- Account-level
- Service-level
- User delegation

Answer Area

Requirement

- Delegate access to resources in one or more of the storage services
- Delegate access to a resource in a single storage service
- Secure a resource by using Azure AD credentials

SAS type

Answer:

SAS types	Requirement	SAS type
Account-level	Delegate access to resources in one or more of the storage services	Account-level
Service-level	Delegate access to a resource in a single storage service	Service-level
User delegation	Secure a resource by using Azure AD credentials	User delegation

Explanation

Graphical user interface, text, application, email Description automatically generated

Requirement	SAS type
Delegate access to resources in one or more of the storage services	Account-level
Delegate access to a resource in a single storage service	Service-level
Secure a resource by using Azure AD credentials	User delegation

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

NEW QUESTION: 117

You are developing Azure WebJobs.

You need to recommend a WebJob type for each scenario.

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

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

Box 1: Continuous

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

Box 2: Triggered

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

Box 3: Continuous

Continuous supports remote debugging.

Note:

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

Continuous	Triggered
Starts immediately when the WebJob is created. To keep the job from ending, the program or script typically does its work inside an endless loop. If the job does end, you can restart it.	Starts only when triggered manually or on a schedule.
Runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.	Runs on a single instance that Azure selects for load balancing.
Supports remote debugging.	Doesn't support remote debugging.

References:

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

NEW QUESTION: 118

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); D18912E1457D5D1DDCDBD40AB3BF70D5D`

C. `await subscriptionClient.CloseAsync();`

D.

`subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync,messageHandlerOptions);`

Answer: D (LEAVE A REPLY)

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/>

NEW QUESTION: 119

You need to implement the corporate website.

How should you configure the solution?



Answer:



Explanation

Answer Area

Azure

Configuration

Plan

	▼
Free	
Standard	
Premium	
Isolated	



Service

	▼
App Service Web App	
App Service Static Web App	
Azure Function App	
Azure Blob Storage	

NEW QUESTION: 120

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.

Answer: B (LEAVE A REPLY)

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>

NEW QUESTION: 121

You need to implement a solution to resolve the retail store location data issue.

Which three Azure Blob features should you enable? Each correct answer presents part of the solution.

NOTE Each correct selection is worth one point

- A.** Immutability
- B.** Snapshots
- C.** Versioning
- D.** Soft delete
- E.** Object replication
- F.** Change feed

Answer: C,D,F (LEAVE A REPLY)

Explanation

Scenario: You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.

Before you enable and configure point-in-time restore, enable its prerequisites for the storage account: soft delete, change feed, and blob versioning.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/point-in-time-restore-manage>

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam!
Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 122

You are using Azure Front Door Service.

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

You need to determine the root cause for the issue.

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

NOTE: Each correct selection is worth one point.

Statement	Yes	No
The file MIME type is supported by the service.	<input type="radio"/>	<input type="radio"/>
Edge nodes must be purged of all cache assets.	<input type="radio"/>	<input type="radio"/>
The compression type is supported.	<input type="radio"/>	<input type="radio"/>

Answer:

Statement	Yes	No
The file MIME type is supported by the service.	<input type="radio"/>	<input checked="" type="radio"/>
Edge nodes must be purged of all cache assets.	<input checked="" type="radio"/>	<input type="radio"/>
The compression type is supported.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation

Statement	Yes	No
The file MIME type is supported by the service.	<input type="radio"/>	<input checked="" type="radio"/>
Edge nodes must be purged of all cache assets.	<input checked="" type="radio"/>	<input type="radio"/>
The compression type is supported.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: No

Front Door can dynamically compress content on the edge, resulting in a smaller and faster response to your clients. All files are eligible for compression. However, a file must be of a MIME type that is eligible for compression list.

Box 2: No

Sometimes you may wish to purge cached content from all edge nodes and force them all to retrieve new updated assets. This might be due to updates to your web application, or to quickly update assets that contain incorrect information.

Box 3: Yes

These profiles support the following compression encodings: Gzip (GNU zip), Brotli Reference: <https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching>

NEW QUESTION: 123

You are developing a web application that uses the Microsoft identity platform to authenticate users and resources, The web application calls several REST APIs.

The APIs require an access token from the Microsoft identity platform.

You need to request a token.

Which three properties should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Redirect URI/URL
- B. Application name
- C. Application ID
- D. Supported account type
- E. Application secret

Answer: B,C,E (LEAVE A REPLY)

NEW QUESTION: 124

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 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. Yes

B. No

Answer: A ([LEAVE A REPLY](#))

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>

NEW QUESTION: 125

You develop and deploy a web application to Azure App Service. The application accesses data stored in an Azure Storage account. The account contains several containers with several blobs with large amounts of data.

You deploy all Azure resources to a single region.

You need to move the Azure Storage account to the new region. You must copy all data to the new region.

What should you do first?

A. Export the Azure Storage account Azure Resource Manager template

B. Initiate a storage account failover

C. Configure object replication for all blobs

D. Use the AzCopy command line tool

E. Create a new Azure Storage account in the current region

F. Create a new subscription in the current region

Answer: A (LEAVE A REPLY)

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 and finally, delete the resources in the source region.

To get started, export, and then modify a Resource Manager template.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

NEW QUESTION: 126

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

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

Box 1: No

ExpirationTime - The time that the message expires.

InsertionTime - The time that the message was added to the queue.

Box 2: Yes

maxDequeueCount - The number of times to try processing a message before moving it to the poison queue.

Default value is 5.

Box 3: Yes

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

Box 4: Yes

References:

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

NEW QUESTION: 127

A development team is creating a new REST API. The API will store data in Azure Blob storage. You plan to deploy the API to Azure App Service.

Developers must access the Azure Blob storage account to develop the API for the next two months. The Azure Blob storage account must not be accessible by the developers after the two-month time period.

You need to grant developers access to the Azure Blob storage account.

What should you do?

- A.** Generate a shared access signature (SAS) for the Azure Blob storage account and provide the SAS to all developers.
- B.** Create and apply a new lifecycle management policy to include a last accessed date value. Apply the policy to the Azure Blob storage account.
- C.** Provide all developers with the access key for the Azure Blob storage account. Update the API to include the Coordinated Universal Time (UTC) timestamp for the request header.
- D.** Grant all developers access to the Azure Blob storage account by assigning role-based access control (RBAC) roles.

Answer: A (LEAVE A REPLY)

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

NEW QUESTION: 128

You develop several Azure Grid to include hundreds of event types, such as billing, inventory, and shipping updates.

Events must be sent to a single endpoint for the Azure Functions app to process. The events must be filtered by event type before processing. You must have authorization and authentication control to partition your tenants to receive the event data.

You need to configure Azure Event Grid.

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

NOTE: Each correct selection is worth one point.

The screenshot shows an 'Answer Area' with the Microsoft logo. It contains two requirements and their configuration values:

Requirement	Configuration Value
Third-party system endpoint to send events	system topic
Azure Functions app endpoint to handle filtered events	event domain

Answer:

Answer Area

Requirement

Third-party system endpoint to send events

Azure Functions app endpoint to handle filtered events

Configuration Value

system topic
system topic
custom topic
event domain
event subscription

event domain
system topic
custom topic
event domain
event subscription

Explanation

Answer Area

Requirement

Third-party system endpoint to send events
Azure Functions app endpoint to handle filtered events

Configuration Value

system topic
event domain

NEW QUESTION: 129

You create the following PowerShell script:

```
$source = New-AzScheduledQueryRuleSource -Query 'Heartbeat | where TimeGenerated > ago(1h)' -DataSourceId "contoso"  
$schedule = New-AzScheduledQueryRuleSchedule -FrequencyInMinutes 60 -TimeWindowInMinutes 60  
$triggerCondition = New-AzScheduledQueryRuleTriggerCondition -ThresholdOperator "LessThan" -Threshold 5  
$aznsActionGroup = New-AzScheduledQueryRuleAznsActionGroup -ActionGroup "contoso" -EmailSubject "Custom email subject"  
-CustomWebhookPayload "{ 'alert':'#alertrulename', 'IncludeSearchResults':true }"  
$alertingAction = New-AzScheduledQueryRuleAlertingAction -AznsAction $aznsActionGroup -Severity "3" -Trigger $triggerCondition  
New-AzScheduledQueryRule -ResourceGroupName "contoso" -Location "eastus" -Action $alertingAction -Enabled $true  
-Description "Alert description" -Schedule $schedule -Source $source -Name "Alert Name"
```

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

NOTE: Each correct selection is worth one point.

Statements



Yes

No

A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes. Yes No

A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five. Yes No

The log alert is scheduled to run every two hours. Yes No

Answer:

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	<input type="radio"/>	<input checked="" type="radio"/>
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	<input checked="" type="radio"/>	<input type="radio"/>
The log alert is scheduled to run every two hours.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation

Text Description automatically generated

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	<input type="radio"/>	<input checked="" type="radio"/>
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	<input checked="" type="radio"/>	<input type="radio"/>
The log alert is scheduled to run every two hours.	<input type="radio"/>	<input checked="" type="radio"/>

Box 1: No

The AzScheduledQueryRuleSource is Heartbeat, not CPU.

Box 2: Yes

The AzScheduledQueryRuleSource is Heartbeat!

Note: New-AzScheduledQueryRuleTriggerCondition creates an object of type Trigger Condition. This object is to be passed to the command that creates Alerting Action object.

Box 3: No

The schedule is 60 minutes, not two hours.

-FrequencyInMinutes: The alert frequency.

-TimeWindowInMinutes: The alert time window

The New-AzAscheduledQueryRuleSchedule command creates an object of type Schedule. This object is to be passed to the command that creates Log Alert Rule.

Reference:

<https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryrule>

<https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryruletriggercondition>

NEW QUESTION: 130

A company is developing a gaming platform. Users can join teams to play online and see

leaderboards that include player statistics. The solution includes an entity named Team. You plan to implement an Azure Redis Cache instance to improve the efficiency of data operations for entities that rarely change. You need to invalidate the cache when team data is changed. How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
void ClearCachedTeams()
{
   IDatabase cache = Connection.GetDatabase();
   ICache cache = Connection.GetDatabase();

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

  viewBag.nsg += "Team data removed from cache. ";
}
```

Answer:

```
void ClearCachedTeams()
{
   IDatabase cache = Connection.GetDatabase();
   ICache cache = Connection.GetDatabase();

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

  viewBag.nsg += "Team data removed from cache. ";
}
```

Explanation

```

void ClearCachedTeams()
{
    IDatabase cache = Connection.GetDatabase();
    ICache cache = Connection.GetDatabase();

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

    ViewBag.nsg += "Team data removed from cache.";
}

```

Box 1: IDatabase cache = connection.GetDatabase();

Connection refers to a previously configured ConnectionMultiplexer.

Box 2: cache.StringSet("teams", "");

To specify the expiration of an item in the cache, use the TimeSpan parameter of StringSet.

cache.StringSet("key1", "value1", TimeSpan.FromMinutes(90));

References:

<https://azure.microsoft.com/sv-se/blog/lap-around-azure-redis-cache-preview/>

NEW QUESTION: 131

An organization deploys a Blob storage account. Users take multiple snapshots of the blob storage account over time.

You need to delete all snapshots or the blob storage account. You must not delete the blob storage account itself.

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

NOTE: Each correct selection is worth one point.



Answer:



Explanation

Graphical user interface, application Description automatically generated

freepdfdumps.com

NEW QUESTION: 132

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.

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

Answer:

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

Explanation

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

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/>

NEW QUESTION: 133

A company has multiple warehouse. Each warehouse contains IoT temperature devices which deliver temperature data to an Azure Service Bus queue.

You need to send email alerts to facility supervisors immediately if the temperature at a warehouse goes above or below specified threshold temperatures.

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

Actions

- Add a logic app trigger that fires when one or more messages arrive in the queue.
- Add a Recurrence trigger that schedules the app to run every 15 minutes.
- Add an action that sends an email to specified personnel if the temperature is outside of those thresholds.
- Add a trigger that reads IoT temperature data from a Service Bus queue.
- Add a logic app action that fires when one or more messages arrive in the queue.
- Add a condition that compares the temperature against the upper and lower thresholds.
- Create a blank Logic app.
- Add an action that reads IoT temperature data from the Service Bus queue.

Answer Area

Answer:

Actions	Answer Area
Add a logic app trigger that fires when one or more messages arrive in the queue.	Create a blank Logic app.
Add a Recurrence trigger that schedules the app to run every 15 minutes.	Add a logic app trigger that fires when one or more messages arrive in the queue.
Add an action that sends an email to specified personnel if the temperature is outside of those thresholds.	Add a trigger that reads IoT temperature data from a Service Bus queue.
Add a trigger that reads IoT temperature data from a Service Bus queue.	Add a condition that compares the temperature against the upper and lower thresholds.
Add a logic app action that fires when one or more messages arrive in the queue.	Add an action that sends an email to specified personnel if the temperature is outside of those thresholds.
Add a condition that compares the temperature against the upper and lower thresholds.	
Create a blank Logic app.	
Add an action that reads IoT temperature data from the Service Bus queue.	

Explanation

Create a blank Logic app.
Add a logic app action that fires when one or more messages arrive in the queue.
Add an action that reads IoT temperature data from the Service Bus queue.
Add a condition that compares the temperature against the upper and lower thresholds.
Add an action that sends an email to specified personnel if the temperature is outside of those thresholds.

Step 1: Create a blank Logic app.

Create and configure a Logic App.

Step 2: Add a logical app trigger that fires when one or more messages arrive in the queue.

Configure the logic app trigger.

Under Triggers, select When one or more messages arrive in a queue (auto-complete).

Step 3: Add an action that reads IoT temperature data from the Service Bus queue Step 4: Add a condition that compares the temperature against the upper and lower thresholds.

Step 5: Add an action that sends an email to specified personnel if the temperature is outside of

those thresholds Reference:

<https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-monitoring-notifications-with-azure-logic-apps>

NEW QUESTION: 134

You are developing a solution for a hospital to support the following use cases:

*The most recent patient status details must be retrieved even if multiple users in different locations have updated the patient record.

*Patient health monitoring data retrieved must be the current version or the prior version.


*After a patient is discharged and all charges have been assessed, the patient billing record contains the final charges.

You provision a Cosmos DB NoSQL database and set the default consistency level for the database account to Strong. You set the value for Indexing Mode to Consistent.


You need to minimize latency and any impact to the availability of the solution. You must override the default consistency level at the query level to meet the required consistency guarantees for the scenarios.

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

Return the most recent patient status.	<input type="text" value="Strong"/>
Return health monitoring data that is no less than one version behind.	<input type="text" value="Bounded Staleness"/>
After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges	<input type="text" value="Eventual"/>

Box 1: Strong

Strong: Strong consistency offers a linearizability guarantee. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Box 2: Bounded staleness

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is "updates") of an item or by "t" time interval. When you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (t) by which the reads might lag behind the writes

Box 3: Eventual

Eventual: There's no ordering guarantee for reads. In the absence of any further writes, the replicas eventually converge.

NEW QUESTION: 135

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: A,C (LEAVE A REPLY)

Explanation

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

Reference:

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

NEW QUESTION: 136

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

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

You need to implement push notifications.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

▼	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  
(windowsToastPayload);  
...  
}  
catch (System.Exception ex)  
{  
...  
}  
...  
}
```



Answer:

Answer Area

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



	▼
NotificationHubClient	
NotificationHubClientSettings	
NotificationHubJob	
NotificationDetails	

hub=

	▼
NotificationHubClient	
NotificationHubClientSettings	
NotificationHubJob	
NotificationDetails	

	▼
GetInstallation	
CreateClientFromConnectionString	
CreateOrUpdateInstallation	
PatchInstallation	

```
(notificationHubConnection, notificationHubName);  
string windowsToastPayload =  
@"<toast><visual><binding template=""ToastText01""><text id=""1"">" +  
@"New item to view" + @"</text></binding></visual></toast>";  
try  
{  
var result =  
    await hub.

|                                    |   |
|------------------------------------|---|
|                                    | ▼ |
| SendWindowsNativeNotificationAsync |   |
| SubmitNotificationHubJobAsync      |   |
| ScheduleNotificationAsync          |   |
| SendAppleNativeNotificationAsync   |   |

(windowsToastPayload);  
    . . .  
}  
catch (System.Exception ex)  
{  
    . . .  
}  
    . . .
```

Explanation

exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:
https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**
Special Discount: Freepdfdumps)

NEW QUESTION: 137

You develop a serverless application using several Azure Functions. These functions connect to data from within the code.

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

You need to change configuration settings in the hostjson file.

Which tool should you use?

- A. Azure portal
- B. Azure PowerShell
- C. Azure Functions Core Tools (Azure CLI)
- D. Visual Studio

Answer: (SHOW ANSWER)

Explanation

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

References:

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

NEW QUESTION: 138

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: A,D,E (LEAVE A REPLY)

Explanation

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

Application Insights manages the ID of a session for you.

References:

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

NEW QUESTION: 139

You need to configure the Account Kind, Replication, and Storage tier options for the corporate website's Azure Storage account.

How should you complete the configuration? To answer, select the appropriate options in the dialog box in the answer area.

NOTE: Each correct selection is worth one point.

Create storage account

Basics Advanced Tags Review + create

Azure Storage is a Microsoft-managed service providing cloud storage that is highly available, secure, durable, scalable, and redundant. Azure Storage includes Azure Blobs (objects), Azure Data Lake Storage Gen2, Azure Files, Azure Queues, and Azure Tables. The cost of your storage account depends on the usage and the options you choose below. [Learn more](#)



PROJECT DETAILS

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

* Subscription ▼

* Resource group ▼
[Create new](#)

INSTANCE DETAILS

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead. [Choose classic deployment model](#)

* Storage account name ⓘ ✓

* Location ▼

Performance ⓘ Standard Premium

Account kind ⓘ ▼
Storage (general purpose v1)
BlobStorage

Replication ⓘ ▼
Zone-redundant storage (ZRS)
Geo-redundant storage (GRS)
Read-access geo-redundant storage (RA-GRS)
Geo-zone-redundant storage (GZRS)
Read-access geo-zone-redundant storage (RA-GZRS)

Access tier (default) ⓘ Cool Hot

Answer:

Create storage account

Basics Advanced Tags Review + create

Azure Storage is a Microsoft-managed service providing cloud storage that is highly available, secure, durable, scalable, and redundant. Azure Storage includes Azure Blobs (objects), Azure Data Lake Storage Gen2, Azure Files, Azure Queues, and Azure Tables. The cost of your storage account depends on the usage and the options you choose below. [Learn more](#)

PROJECT DETAILS

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

* Subscription

* Resource group

[Create new](#)

INSTANCE DETAILS

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead. [Choose classic deployment model](#)

* Storage account name

* Location


Performance Standard Premium

Account kind

Replication

Access tier (default) Cool Hot

Explanation

INSTANCE DETAILS 

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead. [Choose classic deployment model](#)

* Storage account name ⓘ ✓

* Location ▼

Performance ⓘ Standard Premium

Account kind ⓘ ▼

Replication ⓘ ▼

Access tier (default) ⓘ Cool Hot

Account Kind: StorageV2 (general-purpose v2)

Scenario: Azure Storage blob will be used (refer to the exhibit). Data storage costs must be minimized.

General-purpose v2 accounts: Basic storage account type for blobs, files, queues, and tables. Recommended for most scenarios using Azure Storage.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-overview>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy>

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

NEW QUESTION: 140

You plan to deploy a web app to App Service on Linux. You create an App Service plan. You create and push a custom Docker image that image that contains the web app to Azure Container Registry.

You need to access the console logs generated from inside the container in real-time.

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

NOTE: Each correct selection is worth one point.

az webapp log --name ContosoWeb --resource-group ContosoDevRG

- config
- download
- show
- tail

filesystem

- web-server-logging
- docker-container-logging
- application-logging

az log --name ContosoWeb --resource-group ContosoDevRG

- webapp
- acr
- aks

- config
- download
- show
- tail

Answer:

az webapp log --name ContosoWeb --resource-group ContosoDevRG

- config
- download
- show
- tail

filesystem

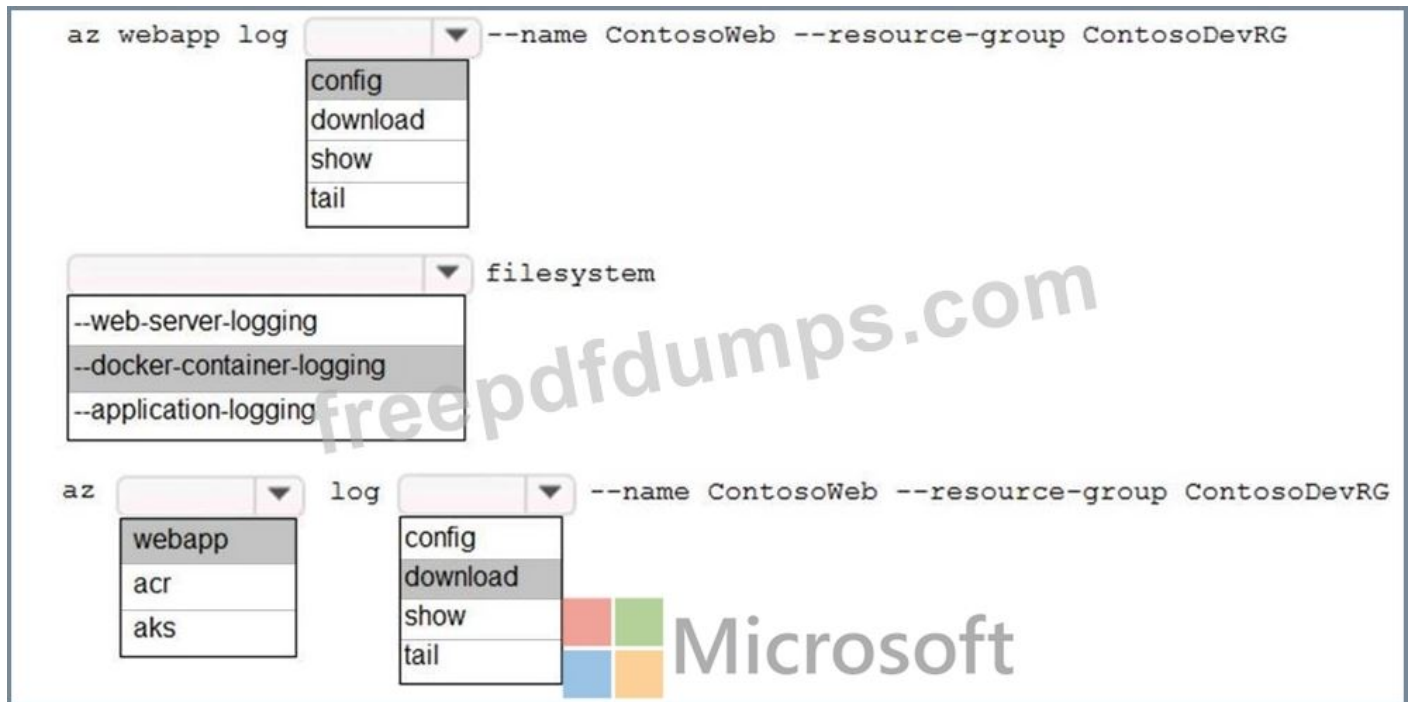
- web-server-logging
- docker-container-logging
- application-logging

az log --name ContosoWeb --resource-group ContosoDevRG

- webapp
- acr
- aks

- config
- download
- show
- tail

Explanation



Box 1: config

To Configure logging for a web app use the command:

```
az webapp log config
```

Box 2: --docker-container-logging

Syntax include:

```
az webapp log config [--docker-container-logging {filesystem, off}]
```

Box 3: webapp

To download a web app's log history as a zip file use the command:

```
az webapp log download
```

Box 4: download

References:

<https://docs.microsoft.com/en-us/cli/azure/webapp/log>

NEW QUESTION: 141

You need to design network connectivity for a subnet in an Azure virtual network. The subnet will contain 30 virtual machines. The virtual machines will establish outbound connections to internet hosts by using the same a pool of four public IP addresses, inbound connections to the virtual machines will be prevented.

What should include in the design?

- A. User Defined Routes
- B. NAT Gateway
- C. Azure Virtual WAN
- D. Azure Private Link

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 142

A company maintains multiple web and mobile applications. Each application uses custom in-house identity providers as well as social identity providers.

You need to implement single sign-on (SSO) for all the applications.

What should you do?

- A. Use Azure Active Directory B2C (Azure AD B2C) with custom policies. Most Voted
- B. Use Azure Active Directory B2B (Azure AD B2B) and enable external collaboration.
- C. Use Azure Active Directory B2C (Azure AD B2C) with user flows.
- D. Use Azure Active Directory B2B (Azure AD B2B).

Answer: A (LEAVE A REPLY)

Explanation

<https://docs.microsoft.com/en-us/azure/active-directory-b2c/custom-policy-reference-ss0>

NEW QUESTION: 143

You need to correct the VM issues.

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

NOTE: Each correct selection is worth one point.

Issue	Tool
Backup and Restore	<input type="checkbox"/> Azure Site Recovery <input type="checkbox"/> Azure Backup <input type="checkbox"/> Azure Data Box <input type="checkbox"/> Azure Migrate
Performance	<input type="checkbox"/> Azure Network Watcher <input type="checkbox"/> Azure Traffic Manager <input type="checkbox"/> ExpressRoute <input type="checkbox"/> Accelerated Networking

Answer:

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

Explanation

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

Backup and Restore: Azure Backup

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

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

Performance: Accelerated Networking

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

Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, greatly improving

its networking performance. This high-performance path bypasses the host from the datapath, reducing latency, jitter, and CPU utilization, for use with the most demanding network workloads on supported VM types.

References:

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

NEW QUESTION: 144

You develop a gateway solution for a public facing news API.

The news API back end is implemented as a RESTful service and hosted in an Azure App Service instance.

You need to configure back-end authentication for the API Management service instance.

Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer:

Explanation

Configuration parameter	Value
Target	Azure Resource
Gateway credentials	Client cert

Box 1: Azure Resource

Box 2: Client cert

API Management allows to secure access to the back-end service of an API using client certificates.

References:

<https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-ba>

NEW QUESTION: 145

You are authoring a set of nested Azure Resource Manager templates to deploy Azure resources. You author an Azure Resource Manager template named mainTemplate.json that contains the following linked templates:

linkedTemplate1.json, linkedTemplate2.json.

You add parameters to a parameters template file named mainTemplate.parameters.json. You save all templates on a local device in the C:\templates\ folder.

You have the following requirements:

- * Store the templates in Azure for later deployment.
- * Enable versioning of the templates.
- * Manage access to the templates by using Azure RBAC

You need to store the templates in Azure.

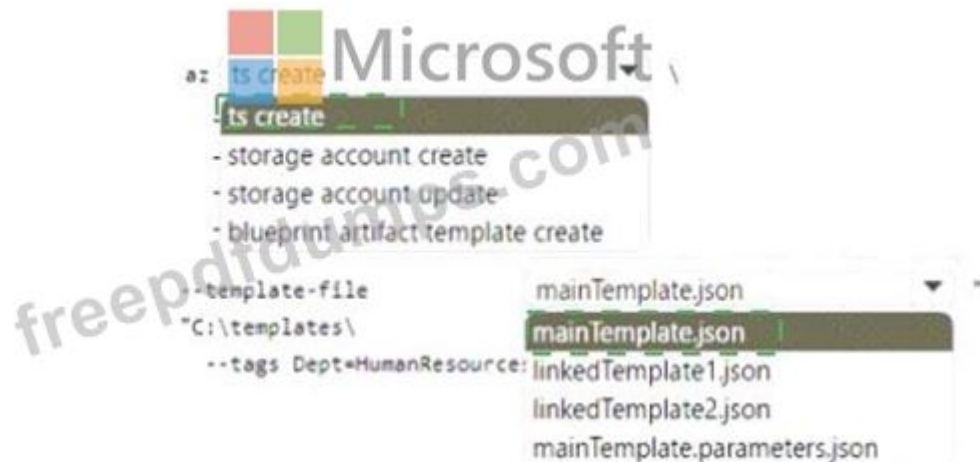
How should you complete the command? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



Answer:

ANSWER AREA



Explanation



NEW QUESTION: 146

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.

Answer: ([SHOW ANSWER](#))

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>

NEW QUESTION: 147

The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location. You need to complete the configuration.

Which Azure CLI or PowerShell command should you run?

A)

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

B)

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

C)

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

D)

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

A. Option B

B. Option A

C. Option D

D. Option C

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 148

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

Answer: B ([LEAVE A REPLY](#))

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>

NEW QUESTION: 149

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

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

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

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

NOTE: Each correct selection is worth one point.

A. Create a file named `.deployment` in the root of the repository that calls a script which generates the static content and deploys the website.

B. Add a PreBuild target in the websites `csproj` project file that runs the static content generation script.

C. Create a file named `run.cmd` in the folder `/run` that calls a script which generates the static content and deploys the website.

D. Add the path to the static content generation tool to `WEBSITE_RUN_FROM_PACKAGE` setting in the `host.json` file.

Answer: ([SHOW ANSWER](#))

Explanation

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

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

[config]

command = YOUR COMMAND TO RUN FOR DEPLOYMENT

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

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

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

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

References:

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

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

NEW QUESTION: 150

You have a single page application (SPA) web application that manages information based on data returned by Microsoft Graph from another company's Azure Active Directory (Azure AD) instance.

Users must be able to authenticate and access Microsoft Graph by using their own company's

Azure AD instance.

You need to configure the application manifest for the app registration.

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

NOTE: Each correct selection is worth one point.

```
{  
  "oauth2AllowImplicitFlow": true,  
  "addIns": [{"  
    "orgRestrictions": "All",  
    "availableToOtherTenants": true,  
    "requiredResourceAccess": [{"  
      "resourceAppId": "00000003-0000-0000-c000-000000000000",  
      "resourceAccess": [{"  
        "id": "24a6cdd6-fab1-4aaf-91b8-3cc8225e90d0",  
        "type": "Scope"}  
      ]  
    }  
  ]}],  
  "signInAudience": "AzureADandPersonalMicrosoftAccount"  
}
```

The image shows a configuration interface for an application manifest. It features several dropdown menus and a list box. The first dropdown menu is for the "oauth2AllowImplicitFlow" property, with options: add, false, spa, true. The second dropdown menu is for the "addIns" property, with options: addIns, orgRestrictions, availableToOtherTenants, requiredResourceAccess. The third dropdown menu is for the "signInAudience" property, with options: All, AzureADMyOrg, AzureADMultipleOrgs, AzureADandPersonalMicrosoftAccount. A watermark "freepdfmups.com" is visible across the image.

Answer:

```

{
  "oauth2AllowImplicitFlow":  ,
  " " : [{"
    addIns
    orgRestrictions
    availableToOtherTenants
    requiredResourceAccess
  }],
  "resourceAppId": "00000003-0000-0000-c000-000000000000",
  "resourceAccess": [
    {
      "id": "24a6cdd6-fab1-4aaf-91b8-3cc8225e90d0",
      "type": "Scope"
    }
  ],
  "signInAudience": 
}

```

add
false
spa
true

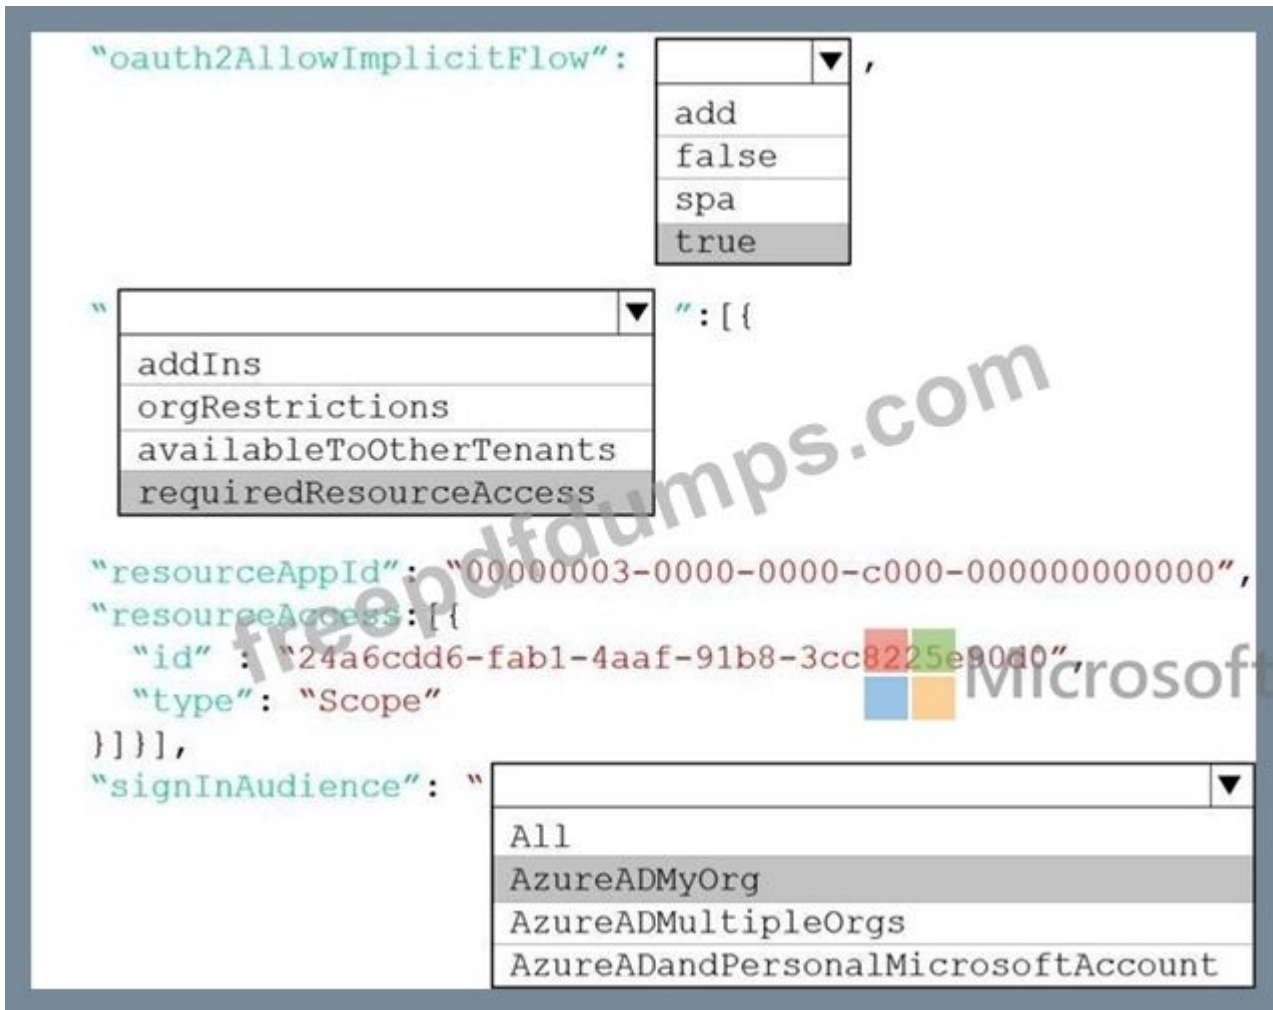
addIns
orgRestrictions
availableToOtherTenants
requiredResourceAccess

Microsoft

All
AzureADMyOrg
AzureADMultipleOrgs
AzureADandPersonalMicrosoftAccount

Explanation

Graphical user interface, text, application, email Description automatically generated



Box 1: true

The `oauth2AllowImplicitFlow` attribute Specifies whether this web app can request OAuth2.0 implicit flow access tokens. The default is false. This flag is used for browser-based apps, like JavaScript single-page apps.

In implicit flow, the app receives tokens directly from the Azure Active Directory (Azure AD) authorize endpoint, without any server-to-server exchange. All authentication logic and session handling is done entirely in the JavaScript client with either a page redirect or a pop-up box.

Box 2: `requiredResourceAccess`

With dynamic consent, `requiredResourceAccess` drives the admin consent experience and the user consent experience for users who are using static consent. However, this parameter doesn't drive the user consent experience for the general case.

`resourceAppId` is the unique identifier for the resource that the app requires access to. This value should be equal to the `appid` declared on the target resource app.

`resourceAccess` is an array that lists the OAuth2.0 permission scopes and app roles that the app requires from the specified resource. Contains the `id` and `type` values of the specified resources.

Example:

```
"requiredResourceAccess": [
{
"resourceAppId": "00000002-0000-0000-c000-000000000000",
"resourceAccess": [
```

```
{  
  "id": "311a71cc-e848-46a1-bdf8-97ff7156d8e6",  
  "type": "Scope"  
}  
]  
}  
],
```

Box 3: AzureADMyOrg

The signInAudience attribute specifies what Microsoft accounts are supported for the current application.

Supported values are:

AzureADMyOrg - Users with a Microsoft work or school account in my organization's Azure AD tenant (for example, single tenant)
AzureADMultipleOrgs - Users with a Microsoft work or school account in any organization's Azure AD tenant (for example, multi-tenant)

AzureADandPersonalMicrosoftAccount - Users with a personal Microsoft account, or a work or school account in any organization's Azure AD tenant Reference:

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

<https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-implicit-grant-flow>

NEW QUESTION: 151

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

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

Answer: C ([LEAVE A REPLY](#))

Explanation

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

Construct the query operation for all customer entities where PartitionKey="Smith".

```
TableQuery<CustomerEntity> query = new
```

```
TableQuery<CustomerEntity>().Where(TableQuery.GenerateFilterCondition("PartitionKey",  
QueryComparisons.Equal, "Smith"));
```

 References:

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

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam!
Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 152

You need to configure Azure Service Bus to Event Grid integration.

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

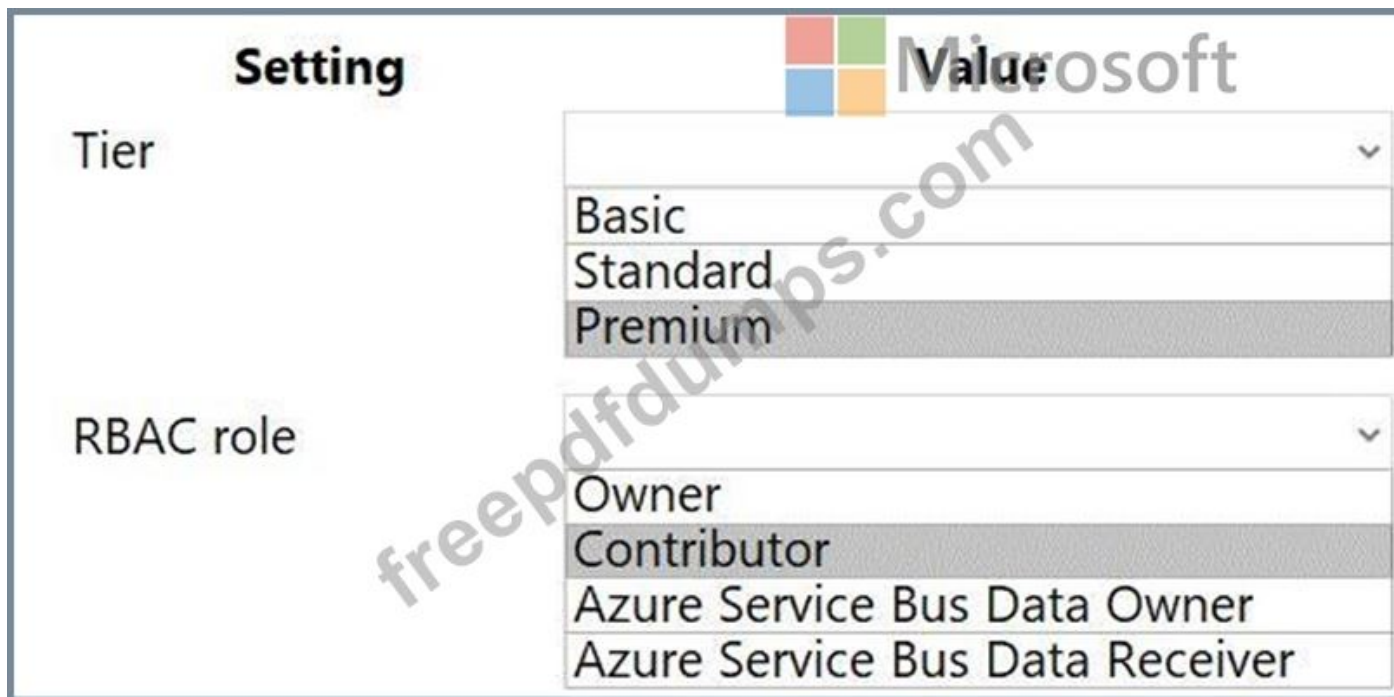
NOTE: Each correct selection is worth one point.

Setting	Value
Tier	<input type="text" value="Basic"/>
RBAC role	<input type="text" value="Contributor"/>

Answer:

Setting	Value
Tier	<input type="text" value="Premium"/>
RBAC role	<input type="text" value="Contributor"/>

Explanation



Box 1: Premium

Service Bus can now emit events to Event Grid when there are messages in a queue or a subscription when no receivers are present. You can create Event Grid subscriptions to your Service Bus namespaces, listen to these events, and then react to the events by starting a receiver. With this feature, you can use Service Bus in reactive programming models.

To enable the feature, you need the following items:

A Service Bus Premium namespace with at least one Service Bus queue or a Service Bus topic with at least one subscription.

Contributor access to the Service Bus namespace.

Box 2: Contributor

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept>

NEW QUESTION: 153

You are developing a user portal for a company.

You need to create a report for the portal that lists information about employees who are subject matter experts for a specific topic. You must ensure that administrators have full control and consent over the data.

Which technology should you use?

- A. Microsoft Graph connectors
- B. Microsoft Graph data connect
- C. Microsoft graph API

Answer: B (LEAVE A REPLY)

NEW QUESTION: 154

You need to resolve the Shipping web site error.

How should you configure the Azure Table Storage service? To answer, select the appropriate

options in the answer area.

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

```

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

```

Box 1: AllowedOrigins

A CORS request will fail if Access-Control-Allow-Origin is missing.

Scenario:

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

```

Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin'
header is present on the requested resource. Origin 'http://testwideworldimporters.com/' is
therefore not allowed access.

```

Box

2: http://test-shippingapi.wideworldimporters.com

Syntax: Access-Control-Allow-Origin: *

Access-Control-Allow-Origin: <origin>

Access-Control-Allow-Origin: null

<origin> Specifies an origin. Only a single origin can be specified.

Box 3: AllowedOrigins

Box 4: POST

The only allowed methods are GET, HEAD, and POST. In this case POST is used.

"<Corsrule>" "allowedmethods" Failed to load no "Access-control-Origin" header is present

References:

<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Access-Control-Allow-Origin>

Topic 3, City Power & Light

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 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 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

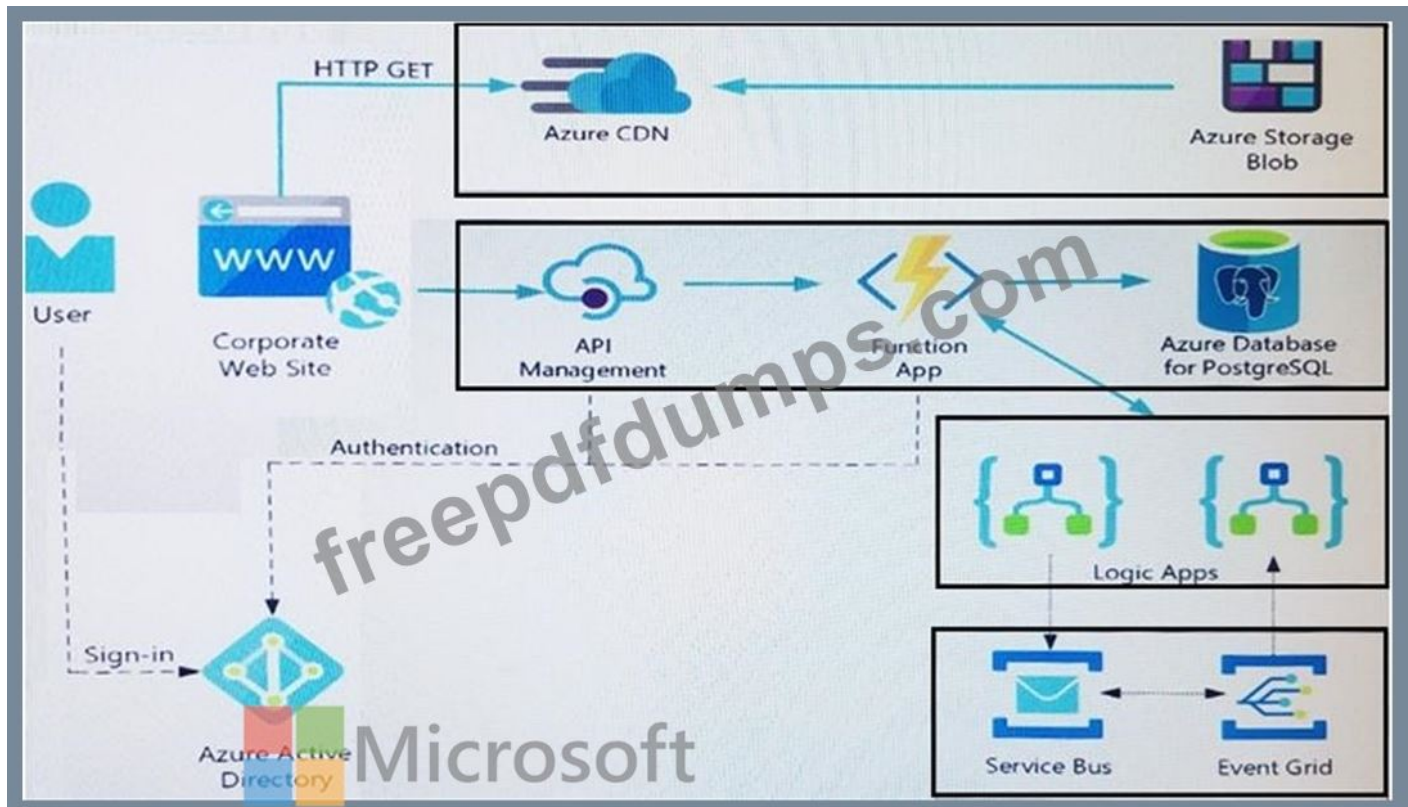
company has a public website located at <http://www.cpandl.com/>. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.

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:

- * The user selects Sign in in the website.
- * The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
- * The user signs in.
- * Azure AD redirects the user's session back to the web application. The URL includes an access token.
- * 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.
- * 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: cpandlkeyvault
- * 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,
RA06 ILogger log)
RA07 {
RA08     log.LogInformation("RequestUserApproval function processed a request.");
RA09     ...
RA10     return ProcessRequest(req)
RA11     ? (ActionResult)new OkObjectResult($"User approval processed")
RA12     : new BadRequestObjectResult("Failed to process user approval");
RA13 }
RA14 private static bool ProcessRequest(HttpRequest req)
RA15 {
RA16     ...
RA17 }

```



NEW QUESTION: 155

You need to audit the retail store sales transactions.

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. Update the retail store location data upload process to include blob index tags. Create an Azure Function to process the blob index tags and filter by store location
- B. Enable blob versioning for the storage account. Use an Azure Function to process a list of the blob versions per day.
- C. Process an Azure Storage blob inventory report by using an Azure Function. Create rule filters on the blob inventory report,
- D. Subscribe to blob storage events by using an Azure Function and Azure Event Grid. Filter the events by store location.
- E. Process the change feed logs of the Azure Blob storage account by using an Azure Function. Specify a time range for the change feed data.

Answer: (SHOW ANSWER)

Explanation

Scenario: Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

"Process the change feed logs of the Azure Blob storage account by using an Azure Function.

Specify a time range for the change feed data": Change feed support is well-suited for scenarios that process data based on objects that have changed. For example, applications can: Store, audit, and analyze changes to your objects, over any period of time, for security, compliance or intelligence for enterprise data management.

"Subscribe to blob storage events by using an Azure Function and Azure Event Grid. Filter the events by store location": Azure Storage events allow applications to react to events, such as the creation and deletion of blobs. It does so without the need for complicated code or expensive and inefficient polling services. The best part is you only pay for what you use.

Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions,

Azure Logic Apps, or even to your own http listener. Event Grid provides reliable event delivery to your applications through rich retry policies and dead-lettering.

Reference:

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

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

NEW QUESTION: 156

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

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

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

Actions	Answer Area
A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.	
Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the files from cache if the TTL has not expired.	
If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.	
The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.	

Answer:

Actions	Answer Area
A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.	A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.
Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the files from cache if the TTL has not expired.	If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.
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.
The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.	Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the files from cache if the TTL has not expired.

Explanation

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

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

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

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

Step 1: A user requests the image..

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

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

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

Step 3: The origin server returns the..

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

An edge server in the POP caches the file and returns the file to the original requestor (Alice).

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

Step 4: Subsequent requests for..

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

If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache.

This process results in a faster, more responsive user experience.

References:

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

NEW QUESTION: 157

You are developing an application that uses Azure Storage to store customer data. The data must only be decrypted by the customer and the customer must be provided a script to rotate keys.

You need to provide a script to rotate keys to the customer.

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

NOTE: Each correct selection is worth one point.

```
$h = $(az keyvault show --hsm-name _ --query "properties.hsmUri")
$x = az keyvault  list-versions --name ""
--vault-name ""
az storage account 
--name _ \
--resource-group 
--resource-group _ \
--encryption-key-name _ |
--encryption-key-version $x
--encryption-key-source 
--encryption-key-vault $|
```

Microsoft

Answer:

```
$h = $(az keyvault show --hsm-name _ --query "properties.hsmUri")
$x = az keyvault  list-versions --name ""
--vault-name ""
az storage account 
--name _ \
--resource-group 
--resource-group _ \
--encryption-key-name _ |
--encryption-key-version $x
--encryption-key-source 
--encryption-key-vault $|
```

Microsoft

NEW QUESTION: 158

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

You implement Azure Service Bus to publish alarms. Each alarm controller uses Azure Service

Bus to receive alarm signals as part of a transaction. Alarm events must be recorded for audit purposes. Each transaction record must include information about the alarm type that was activated.

You need to implement a reply trail auditing solution.

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

NOTE: Each correct selection is worth one point.

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

Answer: A,D (LEAVE A REPLY)

Reference:

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

NEW QUESTION: 159

You are developing an Azure solution to collect inventory data from thousands of stores located around the world. Each store location will send the inventory data hourly to an Azure Blob storage account for processing.

The solution must meet the following requirements:

- * Begin processing when data is saved to Azure Blob storage.
- * Filter data based on store location information.
- * Trigger an Azure Logic App to process the data for output to Azure Cosmos DB.
- * Enable high availability and geographic distribution.
- * Allow 24-hours for retries.
- * Implement an exponential back off data processing.

You need to configure the solution.

What should you implement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Technologies

- Azure Event Hub
- Azure Event Grid
- Azure Service Bus
- Azure Blob Storage
- Azure App Service
- Azure Logic App

Answer Area

Object

- Event Source
- Event Receiver
- Event Handler

Technology

- Technology
- Technology
- Technology

Answer:

The screenshot shows the 'Answer Area' with the following mappings:

- Event Source:** Azure Event Grid
- Event Receiver:** Azure Logic App
- Event Handler:** Azure Service Bus

Explanation

Graphical user interface, application Description automatically generated

Object

Technology

Event Source

Azure Event Grid

Event Receiver

Azure Logic App

Event Handler

Azure Service Bus

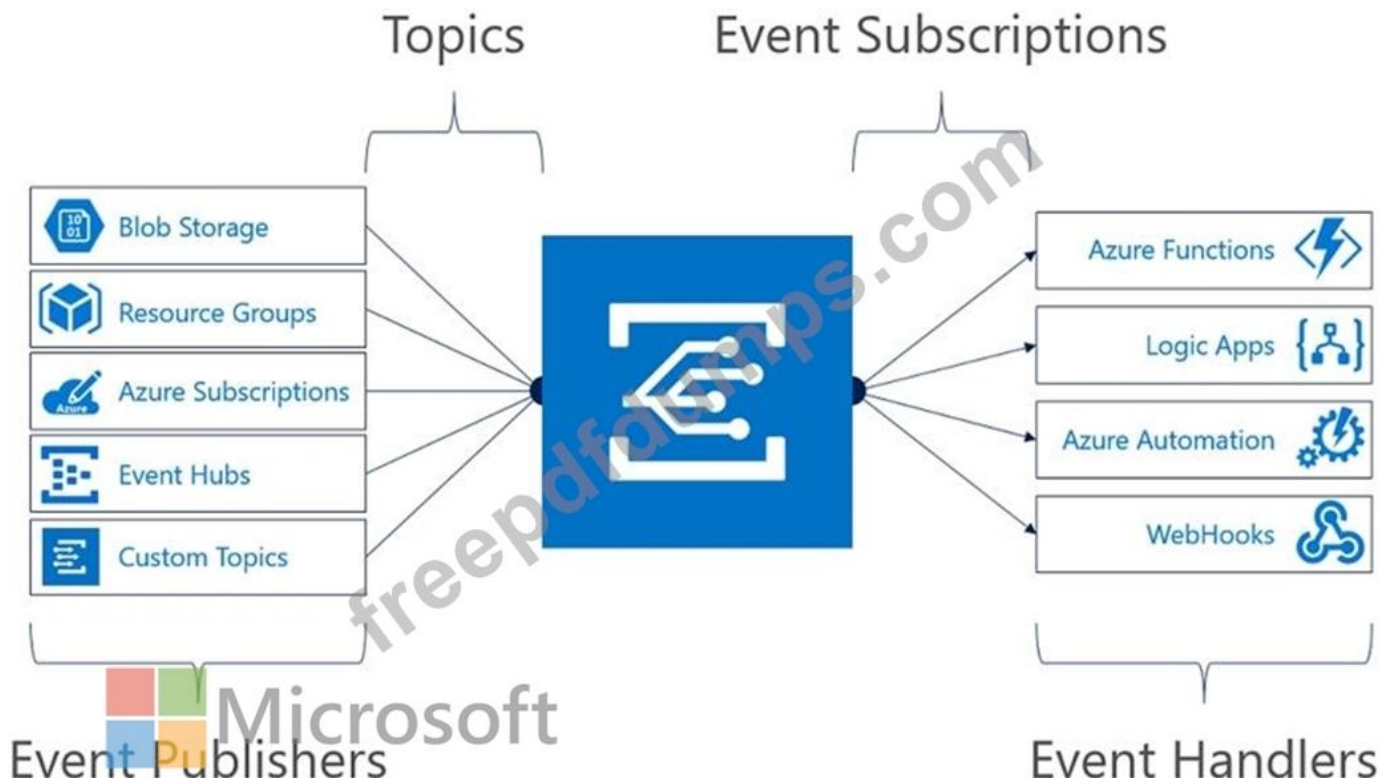
Box 1: Azure Event Grid

Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener. Event Grid provides reliable event delivery to your applications through rich retry policies and dead-lettering.

Box 2: Azure Logic App

Event Grid uses event subscriptions to route event messages to subscribers. This image illustrates the relationship between event publishers, event subscriptions, and event handlers.

Diagram Description automatically generated



Box 3: Azure Service Bus

The Event Grid service doesn't store events. Instead, events are stored in the Event Handlers, including ServiceBus, EventHubs, Storage Queue, WebHook endpoint, or many other supported Azure Services.

Reference:

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

<https://docs.microsoft.com/en-us/java/api/overview/azure/messaging-eventgrid-readme>

NEW QUESTION: 160

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

Answer: A (LEAVE A REPLY)

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>

NEW QUESTION: 161

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 content of the forms must not be compromised.

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

Solution:

Create a Azure Key Vault key named skey.

Encrypt the intake forms using the public key portion of skey.

Store the encrypted data in Azure Blob storage

Does the solution meet the goal?

A. No

B. Yes

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 162

You develop and deploy a Java application to Azure. The application has been instrumented by using the Application Insights SDK.

The telemetry data must be enriched and processed before it is sent to the Application Insights service.

You need to modify the telemetry data.

Which Application Insights SDK 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.

Features

- Sampling
- Telemetry initializer
- Telemetry processor
- Telemetry channel

Answer Area

Requirement

- Reduce the volume of telemetry without affecting statistics.
- Enrich telemetry with additional properties or override an existing one.
- Completely replace or discard a telemetry item.

Feature

Microsoft

Answer:

Features

- Sampling
- Telemetry initializer
- Telemetry processor
- Telemetry channel

Requirement

- Reduce the volume of telemetry without affecting statistics.
- Enrich telemetry with additional properties or override an existing one.
- Completely replace or discard a telemetry item.

Feature

- Sampling
- Telemetry initializer
- Telemetry processor

Answer Area

0.00

Explanation

Features

- Sampling
- Telemetry initializer
- Telemetry processor
- Telemetry channel

Requirement

- Reduce the volume of telemetry without affecting statistics.
- Enrich telemetry with additional properties or override an existing one.
- Completely replace or discard a telemetry item.

Feature

- Sampling
- Telemetry initializer
- Telemetry processor

NEW QUESTION: 163

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

Answer: (SHOW ANSWER)

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.

* In the portal, navigate to Virtual Machines and go to your Windows virtual machine and in the Overview, click Connect.

* Enter in your Username and Password for which you added when you created the Windows VM.

* Now that you have created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.

* 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:

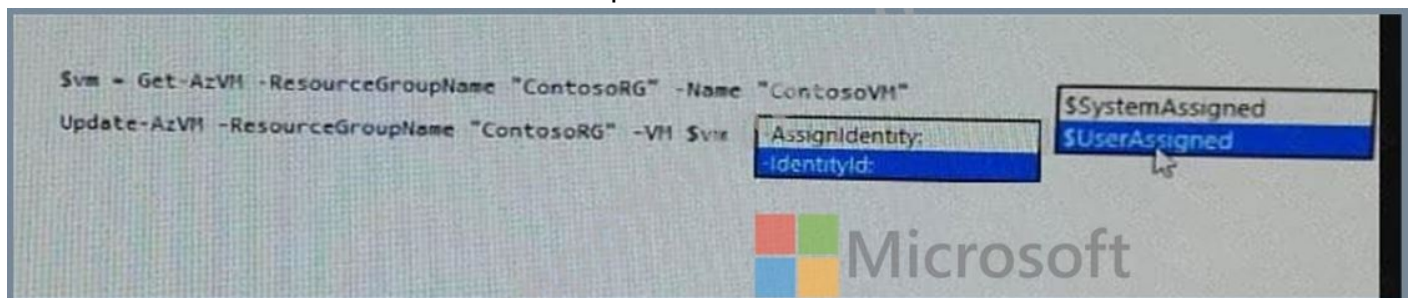
<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm>

NEW QUESTION: 164

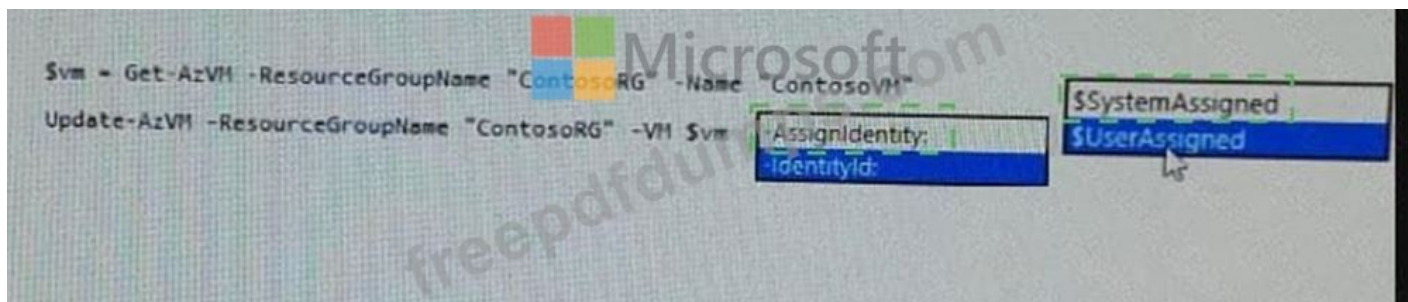
You are developing an application that needs access to an Azure virtual machine (VM). The access lifecycle for the application must be associated with the VM service instance. You need to enable managed identity for the VM.

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.



Answer:



Explanation

```
$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM
```

```
Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -AssignIdentity:
```

```
$SystemAssigned
```

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powers>

NEW QUESTION: 165

You are developing a data storage solution for a social networking app.

The solution requires a mobile app that stores user information using Azure Table Storage.

You need to develop code that can insert multiple sets of user information.

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

NOTE: Each correct selection is worth one point.

```

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

```

```

op = new (op);

```

TableOperation
TableBatchOperaton
TableEntity
TableQuery

```

. . .
table. (op);

```

ExecuteBatch
Execute
Insert
InsertOrMerge

Answer:

```

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

```

```

op = new (op);

```

TableOperation
TableBatchOperaton
TableEntity
TableQuery

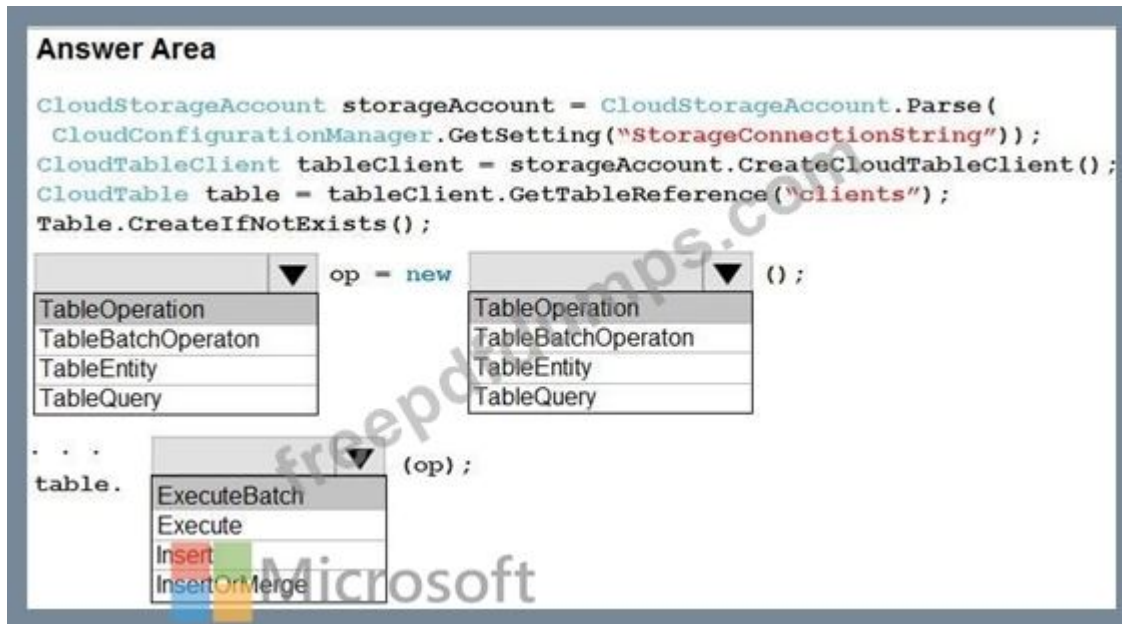
```

. . .
table. (op);

```

ExecuteBatch
Execute
Insert
InsertOrMerge

Explanation:



Box 1, Box 2: TableBatchOperation

Create the batch operation.

```
TableBatchOperation op = new TableBatchOperation();
```

Box 3: ExecuteBatch

/ Execute the batch operation.

```
table.ExecuteBatch(op);
```

Note: You can insert a batch of entities into a table in one write operation. Some other notes on batch operations:

You can perform updates, deletes, and inserts in the same single batch operation.

A single batch operation can include up to 100 entities.

All entities in a single batch operation must have the same partition key.

While it is possible to perform a query as a batch operation, it must be the only operation in the batch.

References:

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

NEW QUESTION: 166

You are configuring a development environment for your team. You deploy the latest Visual Studio image from the Azure Marketplace to your Azure subscription.

The development environment requires several software development kits (SDKs) and third-party components to support application development across the organization. You install and customize the deployed virtual machine (VM) for your development team. The customized VM must be saved to allow provisioning of a new team member development environment.

You need to save the customized VM for future provisioning.


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

NOTE: Each correct selection is worth one point.

ANSWER AREA


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

Answer:

Answer Area  Microsoft

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

Explanation

 Microsoft

Action	Tool or service
Generalize the VM.	<input type="text"/> Azure Power Shell Visual Studio command prompt Azure Migrate Azure Backup
Store images.	<input type="text"/> Azure Blob Storage Visual Data Lake Storage Azure File Storage Azure Table Storage

Box 1: Azure Powershell

Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks.

Before you begin, make sure that you have the latest version of the Azure PowerShell module. You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image.

Box 2: Azure Blob Storage

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource#create-an-image-of-a>

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam! Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 167

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

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

What should you do?

- A. In the Azure portal, create a Batch account.
- B. In a .NET method, call the method: BatchClient.PoolOperations.CreatePool
- C. In Python, implement the class: JobAddParameter
- D. In Python, implement the class: TaskAddParameter

Answer: B (LEAVE A REPLY)

Explanation

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

Valid AZ-204 Dumps shared by Actual4test.com for Helping Passing AZ-204 Exam! Actual4test.com now offer the **newest AZ-204 exam dumps**, the Actual4test.com AZ-204 exam **questions have been updated** and **answers have been corrected** get the **newest**

Actual4test.com AZ-204 dumps with Test Engine here:

https://www.actual4test.com/AZ-204_examcollection.html (478 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)