

Microsoft

AZ-204

**Developing Solutions
for Microsoft Azure**

Version: 4.0

[Total Questions: 142]

Web: www.dumpscollection.com

Email: support@dumpscollection.com



IMPORTANT NOTICE

Feedback

We have developed quality product and state-of-art service to ensure our customers interest. If you have any suggestions, please feel free to contact us at feedback@dumpscollection.com

Support

If you have any questions about our product, please provide the following items:

- 🔗 exam code
- 🔗 screenshot of the question
- 🔗 login id/email

please contact us at support@dumpscollection.com and our technical experts will provide support within 24 hours.

Copyright

The product of each order has its own encryption code, so you should use it independently. Any unauthorized changes will inflict legal punishment. We reserve the right of final explanation for this statement.

Exam Topic Breakdown

Exam Topic	Number of Questions
Topic 1 : Windows Server 2016 virtual machine	10
Topic 2 : ContentAnalysisService	7
Topic 3 : Misc. Questions	125
TOTAL	142



Topic 1, 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 **Next** button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an **All Information** tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the **Question** button to return to the question.

Current environment

Windows Server 2016 virtual machine

The virtual machine (VM) runs BizTalk Server 2016. The VM runs the following workflows:

- ④ Ocean Transport – This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.
- ④ 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:

- ④ Container API – This API provides container information including weight, contents, and other attributes.
- ④ Location API – This API provides location information regarding shipping ports of call and tracking stops.
- ④ Shipping REST API – This API provides shipping information for use and display on the shipping website.

Shipping Data

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


Shipping Web Site

The

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

Proposed solution

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations. You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:



Requirements

Shipping Logic app

The Shipping Logic app must meet the following requirements:

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

Shipping Function app

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

REST APIs

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

- Secure resources to the corporate VNet.
- Allow deployment to a testing location within Azure while not incurring additional costs.
- Automatically scale to double capacity during peak shipping times while not causing application downtime.
- Minimize costs when selecting an Azure payment model.

Shipping data

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

Shipping website

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

Issues

Windows Server 2016 VM

The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.

Shipping website and REST APIs

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

Failed

to load http://test-shippingapi.wideworldimporters.com/: No

'Access-Control-Allow-Origin' header is present on the requested resource. Origin '<http://test.wideworldimporters.com/>' is therefore not allowed access.

Question #:1 - [\(Exam Topic 1\)](#)

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

What should you use?

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

Answer: D

Explanation

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

The gateway supports BizTalk Server 2016.

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

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

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

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

Reference:

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

Question #:2 - (Exam Topic 1)

You need to secure the Shipping Logic App.

What should you use?

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

D. VNet service endpoint

Answer: C

Explanation

Scenario: The Shipping Logic App requires secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.

You can access to Azure Virtual Network resources from Azure Logic Apps by using integration service environments (ISEs).

Sometimes, your logic apps and integration accounts need access to secured resources, such as virtual machines (VMs) and other systems or services, that are inside an Azure virtual network. To set up this access, you can create an integration service environment (ISE) where you can run your logic apps and create your integration accounts.

References:

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

Question #:3 - [\(Exam Topic 1\)](#)

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

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

Actions

Create an integration account in the Azure portal.

Link the custom connector to the Logic App.

Update the Logic App to use the partners, schemas, certificates, maps, and agreements.

Create a custom connector for the Logic App.

Add partners, schemas, certificates, maps, and agreements.

Link the Logic App to the integration account.

Answer Area**Answer:****Yes****No**

The code will successfully insert a player record.

☐☐

The code has a bug and will insert an additional copy of the Game record with a new Id.

☐☒

The code has a bug and will insert the wrong gameId value.

☒☐

There is a valid many-to-many relationship between Players and Games.

☐☒**Explanation**

Create an integration account in the Azure portal.

Link the Logic App to the integration account.

Add partners, schemas, certificates, maps, and agreements.

Create a custom connector for the Logic App.

Step 1: Create an integration account in the Azure portal

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

Step 2: Link the Logic App to the integration account

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

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

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

References:

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

Question #:4 - (Exam Topic 1)

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

Answer:

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

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

Explanation

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

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

Box 1: Function

Box 2: JSON based Token (JWT)

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

Box 3: HTTP

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

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

References:

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

Question #:5 - (Exam Topic 1)

You need to configure Azure App Service to support the REST API requirements.

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

NOTE: Each correct selection is worth one point.

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

Answer:

App service plan setting**Value****Number of VM instances**

2
4
8
16

Pricing tier

Isolated
Standard
Premium
Consumption

Explanation**Setting****Value**

Plan

Basic
Standard
Premium
Isolated

Instance Count

1
10
20
100

Plan: Standard

Standard support auto-scaling

Instance Count: 10

Max instances for standard is 10.

Scenario:

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

- ⦿ Allow deployment to a testing location within Azure while not incurring additional costs.
- ⦿ Automatically scale to double capacity during peak shipping times while not causing application downtime.
- ⦿ Minimize costs when selecting an Azure payment model.

References:

<https://azure.microsoft.com/en-us/pricing/details/app-service/plans/>

Question #:6 - (Exam Topic 1)

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

What should you use?

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

Answer: D

Explanation

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

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

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

References:

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

Question #:7 - (Exam Topic 1)

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

Technology	Value
SSL certificate	Valid root certificate Self-signed certificate
Proxy type	nginx Azure Application Gateway

Explanation

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

Scenario: Shipping website

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

minimizing latency and costs.

Tier: Standard

Profile: Akamai

Optimization: Dynamic site acceleration

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

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

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

Reference:


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

Question #:8 - (Exam Topic 1)

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

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


NOTE: Each correct selection is worth one point.



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://donovanbrown.com/post/How-to-clear-No-Access-Control-Allow-Origin-header-error-with-Azure-App-Service>

Question #:9 - [\(Exam Topic 1\)](#)

You need to resolve the Shipping web site error.

How should you configure the Azure Table Storage service? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
<?xml version="1.0" encoding="utf-8"?>
```

```
<StorageServiceProperties>
```

```
...
```

```
<Cors>
```

```
<CorsRule>
```

```
<
```

AllowedHeaders
ExposedHeaders
AllowedMethods
AllowedOrigins

http://*.wideworldimporters.com
http://test.wideworldimporters.com
http://test-shippingapi.wideworldimporters.com
http://www.wideworldimporters.com

AllowedHeaders
ExposedHeaders
AllowedMethods
AllowedOrigins

```
<AllowedMethods>
```

GET,PUT
GET
POST
GET,HEAD

```
</AllowedMethods>
```

```
...
```

```
</CorsRule>
```

```
</Cors>
```

```
</StorageServiceProperties>
```

Answer:

Item

Value

Powershell command

```
Get-AzureRmRoleDefinition-Name "Reader" | ConvertTo-Json | Out-File C:\SupportRole.json
Get-AzureRmRoleDefinition-Name "Operator" | ConvertTo-Json | Out-File C:\SupportRole.json
Set-AzureRmRoleDefinition-Name "Reader" | Input-File C:\SupportRole.json
Set-AzureRmRoleDefinition Input-File C:\SupportRole.json
```

Actions section

```
"*/read*", "Microsoft.Support/*"
"*/read*"
"*/read*", "Microsoft.Support/*"
"*/read*"
```

Explanation

```

<?xml version="1.0" encoding="utf-8"?>
<StorageServiceProperties>
  ...
  <Cors>
    <CorsRule>
      <AllowedOrigins>
        http://*.wideworldimporters.com
        http://test.wideworldimporters.com
        http://test-shippingapi.wideworldimporters.com
        http://www.wideworldimporters.com
      </AllowedOrigins>
      <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>

Question #:10 - (Exam Topic 1)

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	Azure Site Recovery Azure Backup Azure Data Box Azure Migrate
Performance	Azure Network Watcher Azure Traffic Manager ExpressRoute 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/>



Topic 2, ContentAnalysisService

Case study

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

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

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

To start the case study

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

Requirements

ContentAnalysisService

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

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

Costs

You must minimize costs for all Azure services.

Manual review

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

High availability

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

Monitoring

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

Security

You have the following security requirements:

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

User agreements

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

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

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

Validation testing

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

Issues

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

Code

ContentUploadService

```

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

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

```

Question #:1 - (Exam Topic 2)

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

Answer:

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>

Question #:2 - (Exam Topic 2)

You need to store the user agreements.

Where should you store the agreement after it is completed?

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

Answer: B**Explanation**

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

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

It has the following characteristics:

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

Reference:

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

Question #:3 - (Exam Topic 2)

You need to implement the bindings for the CheckUserContent function.

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

NOTE: Each correct selection is worth one point.


```

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

```

Answer:

Yes

No

The code will successfully insert a player record.

☒
☐

The code has a bug and will insert an additional copy of the Game record with a new Id.

☐
☒

The code has a bug and will insert the wrong gameId value.

☒
☐

There is a valid many-to-many relationship between Players and Games.

☐
☒
Explanation

```

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

```

Box 1: [BlobTrigger(..)]

Box 2: [Blob(..)]

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

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

```

[FunctionName("ResizeImage")]

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

```

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

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

Reference:

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

Question #:4 - (Exam Topic 2)

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

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

NOTE: Each correct selection is worth one point.

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

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

Answer:

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

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

Explanation

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

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

Box 1: "oauth2Permissions": ["login"]

oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

Box 2: "oauth2AllowImplicitFlow":true

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

Implicit Grant flow.

Reference:

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

Question #:5 - (Exam Topic 2)

You need to monitor ContentUploadService according to the requirements.

Which command should you use?

- A. `az monitor metrics alert create -n alert -g ... - -scopes ... - -condition "avg Percentage CPU > 8"`
- B. `az monitor metrics alert create -n alert -g ... - -scopes ... - -condition "avg Percentage CPU > 800"`
- C. `az monitor metrics alert create -n alert -g ... - -scopes ... - -condition "CPU Usage > 800"`
- D. `az monitor metrics alert create -n alert -g ... - -scopes ... - -condition "CPU Usage > 8"`

Answer: B

Explanation

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

Reference:

<https://docs.microsoft.com/sv-se/cli/azure/monitor/metrics/alert>

Question #:6 - (Exam Topic 2)

You need to configure the ContentUploadService deployment.

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

NOTE: Each correct selection is worth one point.

- A. Add the following markup to line CS23:

types: Private

B. Add the following markup to line CS24:

osType: Windows

C. Add the following markup to line CS24:

osType: Linux

D. Add the following markup to line CS23:

types: Public

Answer: A

Explanation

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

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

Reference:

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

Question #:7 - (Exam Topic 2)

You need to ensure that network security policies are met.

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

NOTE: Each correct selection is worth one point.

Technology	Value
SSL certificate	▼ Valid root certificate Self-signed certificate
Proxy type	▼ nginx Azure Application Gateway

Answer:

Technology	Value
SSL certificate	Valid root certificate Self-signed certificate
Proxy type	nginx Azure Application Gateway

Explanation

Technology	Value
SSL certificate	Valid root certificate Self-signed certificate
Proxy type	nginx Azure Application Gateway

Box 1: Valid root certificate

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

Box 2: Azure Application Gateway

Scenario:

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

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

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

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

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

Reference:

<https://docs.microsoft.com/en-us/azure/web-application-firewall/ag/ag-overview>

<https://www.upguard.com/articles/10-tips-for-securing-your-nginx-deployment>

Topic 3, Misc. Questions

Question #:1 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

The screenshot shows an interactive interface with three main panes:

- Components:** A list of components to be dragged: Helm, Draft, Brigade, KubeCtl, Ingress Controller, CoreDNS, and Virtual Kubelet.
- Answer area:** A central area with three requirements:
 - Deploy solution.
 - View cluster and external IP addressing.
 - Implement a single, public IP endpoint that is routed to multiple microservices.
- Component:** A list of three empty boxes labeled "Component" for the final selection.

A mouse cursor is visible over the "Answer area" pane.

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>

Question #:2 - (Exam Topic 3)

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

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

You need to retrieve the information about the current user.

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

NOTE: Each correct selection is worth one point.

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

Answer: A C

Explanation

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

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

References:

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

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

Question #3 - ([Exam Topic 3](#))

You are working for a company that designs mobile applications. They maintain a server where player records are assigned to their different games. The tracking system is new and in development.

The application uses Entity Framework to connect to an Azure Database. The database holds a Player table and Game table.

When adding a player, the code should insert a new player record, and add a relationship between an existing game record and the new player record.

The application will call `CreatePlayerWithGame` with the correct `gameId` and the `playerId` to start the process. (Line numbers are included for reference only.)




```
01. namespace ContosoCradt
02. {
03.     public class PlayerDbContext : DbContext
04.     {
05.         public PlayerDbContext() : base ("name-dBConnString") { }
06.         public DbSet<Player> Players { get ; set ; }
07.         public DbSet<Game> Games { get ; set ; }
08.         protected override void OnModelCreating(DBModelBuilder modelBuilder)
09.         {
10.             modelBuilder.Entity<Player>().HasMany(x => x.Games). WithMany (x => x.Players);
11.         }
12.     }
13.     internal class dbConfiguration : DbMigrationConfiguration<PlayerDbContext>
14.     {
15.         public dbConfiguration() . {AutomaticMigrationsEnabled = true ; }
16.     {
17.         public class mp
18.         {
19.             public void CreatePlayerWithGame(int playerId, int gameId) => AddPlayer(playerId, GetGame(gameId));
20.             public Game GetGame(int gameId)
21.             {
22.                 using (var db = new PlayerDbContext())
23.                 {
24.                     return db.Games.FirstOrDefault(x => x.GameId == gameId);
25.                 }
26.             }
27.             public Player AddPlayer (int playerId, Game game)
28.             {
29.                 using (var db = new PlayerDbContext())
30.                 {
31.                     var player = new Player
32.                     {
33.                         PlayerId = playerId,
34.                         Games = new List <Game> {game },
35.                     };
36.                     db.Players.Add(player);
37.                     db.SaveChanges();
38.                     return player;
39.                 }
40.             }
41.         }
42.         public class Player
43.         {
44.             public int PlayerId { get ; set; }
45.             public string PlayerName { get ; set; }
46.             public virtual List<Game> Games { get ; set; }
47.         }
48.         public class Game
49.         {
50.             public int GameId { get ; set ; }
51.             public string Title { get ; set; }
```

```
51. public string Platform { get ; set; }  
52. public virtual List<Player> Players { get ; set; }  
53. }  
54. }
```

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

NOTE: Each correct selection is worth one point.

Yes No

The code will successfully insert a player record.

☐☐

The code has a bug and will insert an additional copy of the Game record with a new Id.

☐☐

The code has a bug and will insert the wrong gameld value.

☐☐

There is a valid many-to-many relationship between Players and Games.

☐☐

Answer:

Yes No

The code will successfully insert a player record.

☒☐

The code has a bug and will insert an additional copy of the Game record with a new Id.

☐☒

The code has a bug and will insert the wrong gameld value.

☒☐

There is a valid many-to-many relationship between Players and Games.

☐☒

Explanation

Yes No

The code will successfully insert a player record.

☐☐

The code has a bug and will insert an additional copy of the Game record with a new Id.

☐☐

The code has a bug and will insert the wrong gameld value.

☐☐

There is a valid many-to-many relationship between Players and Games.

☐☐

Many-to-many relationships without an entity class to represent the join table are not yet supported. However, you can represent a many-to-many relationship by including an entity class for the join table and mapping two separate one-to-many relationships.

```
protected override void OnModelCreating(ModelBuilder modelBuilder)
{
    modelBuilder.Entity<PostTag>()
        .HasKey(t => new { t.PostId, t.TagId });
    modelBuilder.Entity<PostTag>()
        .HasOne(pt => pt.Post)
        .WithMany(p => p.PostTags)
        .HasForeignKey(pt => pt.PostId);
    modelBuilder.Entity<PostTag>()
        .HasOne(pt => pt.Tag)
        .WithMany(t => t.PostTags)
        .HasForeignKey(pt => pt.TagId);
}
```

Question #:4 - [\(Exam Topic 3\)](#)

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

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

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

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

NOTE: Each correct selection is worth one point.

- A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting.
- B. Convert the web app to run in an Azure App service environment (ASE).

- C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application gateway.
- D. In the Azure Application Gateway's HTTP setting, set the value of the Override backend path option to contoso22.azurewebsites.net.

Answer: A D

Explanation

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

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

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

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

Add HTTP setting

saiappgw-appgw

* Protocol

HTTP HTTPS

 Authentication certificates are not required for trusted Azure certificates for end to end ssl to work

* Port 

443 

* Request timeout (seconds)

20

Override backend path 

☒ Use for App service

☒ Use custom probe 

OK

Reference:

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

Question #5 - (Exam Topic 3)

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

- Ⓐ Each instance of the WebJob processes data for a single customer and must run as a singleton instance.
- Ⓑ Each deployment must be tested by using deployment slots prior to serving production data.
- Ⓒ Azure costs must be minimized.
- Ⓓ Azure resources must be located in an isolated network.

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

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

NOTE: Each correct selection is worth one point.

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/>

Question #6 - (Exam Topic 3)

You are developing a data storage solution for a social networking app.

The solution requires a mobile app that stores user information using Azure Table Storage.

You need to develop code that can insert multiple sets of user information.

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

NOTE: Each correct selection is worth one point.


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

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

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

Answer:

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




Explanation

Explanation:

Answer Area

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



Box 1, Box 2: TableBatchOperation

Create the batch operation.

```
TableBatchOperation op = new TableBatchOperation();
```

Box 3: ExecuteBatch

/ Execute the batch operation.

```
table.ExecuteBatch(op);
```

Note: You can insert a batch of entities into a table in one write operation. Some other notes on batch operations:

You can perform updates, deletes, and inserts in the same single batch operation.

A single batch operation can include up to 100 entities.

All entities in a single batch operation must have the same partition key.

While it is possible to perform a query as a batch operation, it must be the only operation in the batch.

References:

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

Question #:7 - (Exam Topic 3)

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

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

You 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: B

Explanation

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

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

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

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

Reference:

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

Question #:8 - (Exam Topic 3)

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

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

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

You need to store the information.

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation

Instead use Azure Cache for Redis.

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

References:

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

Question #:9 - (Exam Topic 3)

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

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

You need to create the custom role.

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

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

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

Question #:10 - [\(Exam Topic 3\)](#)

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

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

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

You need to store the information.

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation

Instead deploy and configure Azure Cache for Redis. Update the web applications.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching#managing-concurrency-in-a-cache>

Question #:11 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

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

Question #:12 - (Exam Topic 3)

You develop and deploy a Java RESTful API to Azure App Service.

You open a browser and navigate to the URL for the API. You receive the following error message:

```
Failed to load http://api.azurewebsites.net:6000/#/api/Products: No 'Access-Control-Allow-Origin' header is present on the requested resource.
Origin 'http://localhost:6000' is therefore not allowed access
```

You need to resolve the error.

What should you do?

- A. Bind an SSL certificate
- B. Enable authentication
- C. Enable CORS
- D. Map a custom domain
- E. Add a CDN

Answer: C

Explanation

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

References:

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

Question #:13 - (Exam Topic 3)

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
☐ Yes	Enable AlwaysOn encryption.	☐
☐ No	Set the column encryption setting to disabled.	☐
	Assign users to the Public fixed database role.	☐
	Store column encryption keys in the system catalog view in the database.	☐

Answer:

Responses	Protection method	Response
☒ Yes	Enable AlwaysOn encryption.	☒ Yes
☒ No	Set the column encryption setting to disabled.	☐ No
	Assign users to the Public fixed database role.	☐ Yes
	Store column encryption keys in the system catalog view in the database.	☒ No

Explanation

Responses	Protection method	Response
☐ Yes	Enable AlwaysOn encryption.	☐ Yes
☐ No	Set the column encryption setting to disabled.	☐ No
	Assign users to the Public fixed database role.	☐ Yes
	Store column encryption keys in the system catalog view in the database.	☐ No

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>

Question #:14 - [\(Exam Topic 3\)](#)

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

Explanation

Change feed support in Azure Blob Storage

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

Reference:

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

Question #:15 - [\(Exam Topic 3\)](#)

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

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

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

on the Production deployment slot.

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation

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

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

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

Reference:

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

Question #:16 - (Exam Topic 3)

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

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

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

Explanation

Use the following method:

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


References:

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

Question #:17 - (Exam Topic 3)


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.)



You use an Azure table as the data source for the import operation. The table contains three records with item inventory data that matches the fields in the Storage data exhibit. These records were imported when the index was created. (Click the Storage Data tab.) When users search with no filter, all three records are displayed.

PartitionKey	RowKey	Timestamp	InStockCount	ItemDescription	ItemName	LocationRow	LocationShelf	SKU
Food	3	2018-08-25T15:47:26.135Z	32	A box of chocolate candy bars	Choco-bar	5	3	123421
Hardware	2	2018-08-25T15:46:08.405Z	2	A bag of bolts	Bolts	1	4	678964
Hardware	1	2018-08-25T15:46:41.405Z	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	☐	☐
You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field.	☐	☐
You can resolve the issue by running the indexer.	☐	☐
You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results	☐	☐

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	☒	☐
You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field.	☐	☒
You can resolve the issue by running the indexer.	☒	☐
You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results	☐	☒

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	☐	☐
You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field.	☐	☐
You can resolve the issue by running the indexer.	☐	☐
You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results	☐	☐

Box 1: Yes

The ItemDescription field is not searchable.

Box 2: No

The ItemDescription field is not searchable, but we would need to recreate the index.

Box 3: Yes

An indexer in Azure Search is a crawler that extracts searchable data and metadata from an external Azure data source and populates an index based on field-to-field mappings between the index and your data source. This approach is sometimes referred to as a 'pull model' because the service pulls data in without you having to write any code that adds data to an index.

Box 4: No

References:

<https://docs.microsoft.com/en-us/azure/search/search-what-is-an-index>

<https://docs.microsoft.com/en-us/azure/search/search-indexer-overview>

Question #:18 - (Exam Topic 3)

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

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

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

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

Azure CLI Commands**Answer Area**`az group create``az group update``az webapp update``az webapp create``az appservice plan create`**Answer:****Azure CLI Commands****Answer Area**`az group create``az group update``az webapp update``az webapp create``az appservice plan create``az group create``az appservice plan create``az webapp create`**Explanation**`az group create``az appservice plan create``az webapp create`

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

Step 1: az group create

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

Step 2: az appservice plan create

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

Step 3: az webapp create

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

References:

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

Question #:19 - (Exam Topic 3)

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

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.

Question #:20 - (Exam Topic 3)

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

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

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

You need to store the information.

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation

The session state provider for Azure Cache for Redis enables you to share session information between different instances of an ASP.NET web application.

The same connection can be used by multiple concurrent threads.

Redis supports both read and write operations.

The output cache provider for Azure Cache for Redis enables you to save the HTTP responses generated by an ASP.NET web application.

Note: Using the Azure portal, you can also configure the eviction policy of the cache, and control access to the cache by adding users to the roles provided. These roles, which define the operations that members can perform, include Owner, Contributor, and Reader. For example, members of the Owner role have complete control over the cache (including security) and its contents, members of the Contributor role can read and write information in the cache, and members of the Reader role can only retrieve data from the cache.

Reference:

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

Question #:21 - ([Exam Topic 3](#))

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

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

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

You create the index in Azure Search.

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

Solution:

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation

Use the following method:

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

References:

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

Question #:22 - [\(Exam Topic 3\)](#)

A company uses Azure SQL Database to store data for an app. The data includes sensitive information.

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

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

NOTE: Each correct selection is worth one point.

- A. Include the managers group.
- B. Exclude the managers group.
- C. Exclude the administrators group.
- D. Navigate to the following URL:
`PUT https://management.azure.com/subscriptions/00000000-1111-2222-3333-444444444444
/resourceGroups/rg01/providers/Microsoft.Sql/servers/server01/databases/customers
/transparentDataEncryption/current?api-version=2014-04-01`
- E. Run the following Azure PowerShell command:
`New-AzureRmSqlDatabaseDataMaskingRule -SchemaName "dbo" -TableName "customers" -
-ColumnName "ssn" -MaskingFunction "Default"`

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

Answer: B E

Explanation

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

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

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

References:

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

Question #23 - (Exam Topic 3)

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

You need to configure the services.

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

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

Box 1: WebJobs

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

Box 2: Flow

Question #:24 - [\(Exam Topic 3\)](#)

You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published. You must ensure that the website remains available and responsive while minimizing cost. You need to deploy the website. What should you do?

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

Answer: C

Explanation

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

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

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

Question #:25 - [\(Exam Topic 3\)](#)

You are developing a .NET Core model-view controller (MVC) application hosted on Azure for a health care system that allows providers access to their information.

You develop the following code:


```

services.AddAuthorization (options =>
{
    options.AddPolicy("ProviderPartner", policy =>
    {
        .policy.AddAuthenticationSchemes("Cookie, Bearer");
        policy.RequireAuthenticatedUser();
        policy.RequireRole("ProviderAdmin", "SysAdmin");
        policy.RequireClaim("editor", "partner");
    });
});

```

You define a role named SysAdmin.

You need to ensure that the application meets the following authorization requirements:

- ⦿ Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.
- ⦿ Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

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

NOTE: Each correct selection is worth one point.

Answer:

```

services.AddAuthorization (options =>
{
    options.AddPolicy("ProviderPartner", policy =>
    {
        .policy.AddAuthenticationSchemes("Cookie, Bearer");
        policy.RequireAuthenticatedUser();
        policy.RequireRole("ProviderAdmin", "SysAdmin");
        policy.RequireClaim("editor", "partner");
    });
});

```

Explanation

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

public class PartnerController : Controller
{
    . . .

    [Authorize(Policy = "ProviderEditor", Role= "SysAdmin")]

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

Box 1:

Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.

Box 2:

Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

Question #:26 - ([Exam Topic 3](#))

You are preparing to deploy an application to an Azure Kubernetes Service (AKS) cluster.

The application must only be available from within the VNet that includes the cluster.

You need to deploy the application.

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

NOTE: Each correct selection is worth one point.

Code segments

Ingress

Service

LoadBalancer

Deployment

ingress.class

azure-load-balancer-internal

Answer Area

apiVersion: v1

kind: Code segment

metadata:

name: web-app

annotations:

service.beta.kubernetes.

Code segment

: "true"

spec:

type: Code segment

ports:

- port: 80

selector:

app: web-app

Answer:

Code segments

Ingress

Service

LoadBalancer

Deployment

ingress.class

azure-load-balancer-internal

Answer Area

apiVersion: v1

kind: Service

metadata:

name: web-app

annotations:

service.beta.kubernetes.

azure-load-balancer-internal

: "true"

spec:

type: LoadBalancer

ports:

- port: 80

selector:

app: web-app

Explanation

```
apiVersion: v1
kind: Service
metadata:
  name: web-app
  annotations:
    service.beta.kubernetes.azure-load-balancer-internal: "true"
spec:
  type: LoadBalancer
  ports:
    - port: 80
  selector:
    app: web-app
```

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

YAML:

```
apiVersion: v1
kind: Service
metadata:
  name: internal-app
  annotations:
    service.beta.kubernetes.io/azure-load-balancer-internal: "true"
spec:
  type: LoadBalancer
  ports:
    - port: 80
  selector:
    app: internal-app
```


References:

<https://docs.microsoft.com/en-us/azure/aks/internal-lb>

Question #27 - (Exam Topic 3)

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

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"]

Answer Area

Answer:

Commands

```
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"]
```

Answer Area

```
WORKDIR /apps/ContosoApp
```

```
COPY ./.
```

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

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

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 ./ContosApp/ /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>

Question #:28 - (Exam Topic 3)

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

Explanation

It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber>

Question #:29 - (Exam Topic 3)

You develop a gateway solution for a public facing news API. The news API back end is implemented as a RESTful service and uses an OpenAPI specification.

You need to ensure that you can access the news API by using an Azure API Management service instance.

Which Azure PowerShell command should you run?

- A. `Import-AzureRmApiManagementApi -Context $ApiMgmtContext -SpecificationFormat "Swagger" -SpecificationPath $SwaggerPath -Path $Path`
- B. `New-AzureRmApiManagementBackend -Context $ApiMgmtContext -Url $Url -Protocol http`
- C. `New-AzureRmApiManagement -ResourceGroupName $ResourceGroup -Name $Name -Location $Location -Organization $Org -AdminEmail $AdminEmail`
- D. `New-AzureRmApiManagementBackendProxy -Url $ApiUrl`

Answer: D

Explanation

`New-AzureRmApiManagementBackendProxy` creates a new Backend Proxy Object which can be piped when creating a new Backend entity.

Example: Create a Backend Proxy In-Memory Object

```
PS C:\>$secpassword = ConvertTo-SecureString "PlainTextPassword" -AsPlainText -Force
```

```
PS C:\>$proxyCreds = New-Object System.Management.Automation.PSCredential ("foo", $secpassword)
```

```
PS C:\>$credential = New-AzureRmApiManagementBackendProxy -Url "http://12.168.1.1:8080" -ProxyCredential $proxyCreds
```

```
PS C:\>$apimContext = New-AzureRmApiManagementContext -ResourceGroupName "Api-Default-WestUS" -ServiceName "contoso"
```

```
PS C:\>$backend = New-AzureRmApiManagementBackend -Context $apimContext -BackendId 123 -Url 'https://contoso.com/awesomeapi' -Protocol http -Title "first backend" -SkipCertificateChainValidation $true -Proxy $credential -Description "backend with proxy server"
```

Creates a Backend Proxy Object and sets up Backend

Question #:30 - (Exam Topic 3)

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

Smith. Which code segment should you use?

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

Answer: C

Explanation

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

Construct the query operation for all customer entities where `PartitionKey="Smith"`.

```
TableQuery<CustomerEntity> query = new  
TableQuery<CustomerEntity>().Where(TableQuery.GenerateFilterCondition("PartitionKey",  
QueryComparisons.Equal, "Smith"));
```

References:

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

Question #:31 - (Exam Topic 3)

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

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

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

You create the index in Azure Search.

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

Solution:

1. Create a `SearchIndexClient` object to connect to the search index.
2. Create a `DataContainer` that contains the documents which must be added.

3. Create a DataSource instance and set its Container property to the DataContainer.
4. Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation

Use the following method:

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

References:

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

Question #32 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

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

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

az  group
webapp
appservice plan
webapp deployment slot
webapp deployment source

az  group
webapp
appservice plan
webapp deployment slot
webapp deployment source

az  group
webapp
appservice plan
webapp deployment slot
webapp deployment source

Answer:

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

```
az create --location centralus - -name $resourcegroupname
create --name $webappname - -resource-group $resourcegroupname
- -sku S3
create --name $webappname - -resource-group $resourcegroupname
\ - -plan $webappname
create --name $webappname - -resource-group $resourcegroupname
\ - -slot staging
config - -name $webappname - -resource-group $resourcegroupname
\ - -slot staging --repo-url
$gitrepo - -branch master - -manual-integration
```

```
az
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
```

```
az
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
```

```
az
group
webapp
appservice plan
webapp deployment slot
webapp deployment source
```

Explanation


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

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

az	▼
group	
webapp	
appservice plan	
webapp deployment slot	
webapp deployment source	

az	▼
group	
webapp	
appservice plan	
webapp deployment slot	
webapp deployment source	

Box 1: group

Create a resource group.

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

Box 2: appservice plan

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

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

Box 3: webapp

Create a web app.

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

Box 4: webapp deployment slot

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

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

Box 5: webapp deployment source

Deploy sample code to "staging" slot from GitHub.

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

References:

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

Question #:33 - (Exam Topic 3)

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

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

You need to complete the source code of the subscription client

What should you do?

- A. `await subscriptionClient.CloseAsync();`
- B. `await subscriptionClient.AddRuleAsync(new RuleDescription(RuleDescription.DefaultRuleName, new`

```
TrueFilter()));
```

- C. `subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);`
- D. `subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName);`

Answer: C

Explanation

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

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

References:

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

Question #:34 - (Exam Topic 3)

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.

all
data
os

all
data
os

Explanation

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

az

	▼
vm	
keyvault	
keyvault key	
vm encryption	

 create\

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

az

	▼
vm	
keyvault	
keyvault key	
vm encryption	

 create\

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

```

az
┌───────────┴───────────▼ create\
vm
keyvault
keyvault key
vm encryption
--resource -group $resourcegroup \
--name Name2
--image Canonical:UbuntuServer:16.04=LTS:latest \
--admin-username azureuser \
--generate-ssh-keys \
--data-disk-sizes-gb 5

```

```

az
┌───────────┴───────────▼ create\
vm
keyvault
keyvault key
vm encryption
--resource-group $resourcegroup \
--name Name2 \
--disk-encryption-keyvault $keyVault_name \
--key-encryption-key Name1 \
--volume-type

```

```

┌───────────┴───────────▼
all
data
os

```

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>

Question #35 - (Exam Topic 3)

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

The solution must receive and store message 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 queue create --resource-group fridge-rg --namespace-name fridge-ns --name fridge-q`
- B. `New-AzureRmResourceGroup -Name fridge-rg -Location fridge-loc`
- C. `New-AzureRmServiceBusNamespace -ResourceGroupName fridge-rg -NamespaceName fridge-loc -Location fridge-loc`
- D. `connectionString=$(az servicebus namespace authorization-rule keys list --resource-group fridge-rg --fridge-ns fridge-ns --query primaryConnectionString -output tsv)`

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation

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

Note:

Steps:

Step 1: # Create a resource group

```
resourceGroupName="myResourceGroup"
```

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

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

```
namespaceName=myNameSpace$RANDOM
```

```
az servicebus namespace create --resource-group $resourceGroupName --name $namespaceName --location eastus
```

Step 3: # Create a Service Bus queue

```
az servicebus queue create --resource-group $resourceGroupName --namespace-name $namespaceName --name BasicQueue
```

Step 4: # Get the connection string for the namespace

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

Reference:

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

Question #:36 - (Exam Topic 3)

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

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

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

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation

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

Reference:

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

Question #:37 - (Exam Topic 3)

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.

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

Answer:

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

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>

Question #:38 - (Exam Topic 3)

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

You are developing and deploying several ASP.NET web applications to Azure App Service. You plan to save session state information and HTML output.

You must use a storage mechanism with the following requirements:

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

You need to store the information.

Solution: Enable Application Request Routing (ARR).

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation

Instead deploy and configure Azure Cache for Redis. Update the web applications.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching#managing-concurrency-in-a-cache>

Question #:39 - (Exam Topic 3)

You develop an Azure web app. You monitor performance of the web app by using Application Insights. You need to ensure the cost for Application Insights does not exceed a preset budget. What should you do?

- A. Implement ingestion sampling using the Azure portal.
- B. Set a daily cap for the Application Insights instance.
- C. Implement adaptive sampling using the Azure portal.
- D. Implement adaptive sampling using the Application Insights SDK.
- E. Implement ingestion sampling using the Application Insights SDK.

Answer: D

Explanation

Sampling is an effective way to reduce charges and stay within your monthly quota.

You can set sampling manually, either in the portal on the Usage and estimated costs page; or in the ASP.NET SDK in the .config file; or in the Java SDK in the ApplicationInsights.xml file, to also reduce the network traffic.

Adaptive sampling is the default for the ASP.NET SDK. Adaptive sampling automatically adjusts to the volume of telemetry that your app sends. It operates automatically in the SDK in your web app so that telemetry traffic on the network is reduced.

References:

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

Question #40 - (Exam Topic 3)

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

Answer:

App setting

▼
LOCALAPPDATA
WEBSITE_LOCALCACHE_ENABLED
DOTNET_HOSTING_OPTIMIZATION_CACHE
WEBSITES_ENABLE_APP_SERVICE_STORAGE
DIAGDATA

Value

true

▼
/home
/local
D:\home
D:\local

Explanation**App setting**

▼
LOCALAPPDATA
WEBSITE_LOCALCACHE_ENABLED
DOTNET_HOSTING_OPTIMIZATION_CACHE
WEBSITES_ENABLE_APP_SERVICE_STORAGE
DIAGDATA

Value

true

▼
/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>

Question #:41 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

JSON segments	Answer Area
orderBy	<pre>{ "automatic": true, "ngMode": "Consistent", "includedPaths": [{ "path": "/*" }], "excludedPaths": [], "": [{ "path": "/name", "order": "descending" }, { "path": "/city", "order": " }] }</pre>
sortOrder	
ascending	
descending	
compositeIndexes	

Answer:

JSON segments

orderBy

sortOrder

ascending

descending

compositeIndexes

Answer Area

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

Explanation

```

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

```

Box 1: compositeIndexes

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

Box 2: descending

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

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

```

{
  "automatic":true,
  "indexingMode":"Consistent",
  "includedPaths":[
    {
      "path": "/*"
    }
  ],

```

```
"excludedPaths":[],  
"compositeIndexes":[  
  [  
    {  
      "path":"/name",  
    },  
    {  
      "path":"/age",  
    }  
  ]  
]
```

Question #:42 - ([Exam Topic 3](#))

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

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

You implement role-based access control (RBAC) role permission on the containers that store photographs. You assign users to RBAC role.

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

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

NOTE: Each correct selection is worth one point.

Settings

client_id

delegated

profile

application

user_impersonation

Answer Area

API

Azure Storage

Microsoft Graph

Permission

Setting

User.Read

Type

Setting

Setting

Answer:

Settings

client_id

delegated

profile

application

user_impersonation

Answer Area

API

Azure Storage

Microsoft Graph

Permission

user_impersonation

User.Read

Type

delegated

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

References:

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

Question #43 - (Exam Topic 3)

You are configuring a development environment for your team. You deploy the latest Visual Studio image from the Azure Marketplace to your Azure subscription.

The development environment requires several software development kits (SDKs) and third-party components to support application development across the organization. You install and customize the deployed virtual machine (VM) for your development team. The customized VM must be saved to allow provisioning of a new team member development environment.

You need to save the customized VM for future provisioning.

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

NOTE: Each correct selection is worth one point.

Answer Area	Action	Tool or service
	Generalize the VM.	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	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

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

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>

Question #:44 - (Exam Topic 3)

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";
```

	▼
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)
```

```
{
```

```
...  
}
```

```
...  
}
```

Answer:

Answer Area

```
string notificationHubName = "contoso_hub";
string notificationHubConnection = "connection_string";
```

	▼ hub=
NotificationHubClient	
NotificationHubClientSettings	
NotificationHubJob	
NotificationDetails	

	▼	▼
NotificationHubClient		GetInstallation
NotificationHubClientSettings		CreateClientFromConnectionString
NotificationHubJob		CreateOrUpdateInstallation
NotificationDetails		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


```
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);
```

Box 1: NotificationHubClient

Box 2: NotificationHubClient

Box 3: CreateClientFromConnectionString

// Initialize the Notification Hub

```
NotificationHubClient hub = NotificationHubClient.CreateClientFromConnectionString(listenConnString,
hubName);
```

Box 4: SendWindowsNativeNotificationAsync

Send the push notification.

```
var result = await hub.SendWindowsNativeNotificationAsync(windowsToastPayload);
```

References:

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

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/app-service-mobile/app-service-mobile-windc>

Question #45 - (Exam Topic 3)

Your company is developing an Azure API.

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

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

Which authentication mechanism should you use?

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

Answer: C

Explanation

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

Reference:

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

Question #46 - (Exam Topic 3)

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

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

You 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

Explanation

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

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

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

Reference:

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

Question #:47 - (Exam Topic 3)

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

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation

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

Reference:

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

Question #:48 - (Exam Topic 3)

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

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

You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level. You need to configure authorization.

Solution:

- Create a new Azure AD application's manifest, set value of the groupMembershipClaims option to All.
- In the website, use the value of the groups claim from the JWT for the user to determine permissions.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation

To configure Manifest to include Group Claims in Auth Token

1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:
2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.
3. Locate the “groupMembershipClaims” setting. Set its value to either “SecurityGroup” or “All”. To help you decide which:

“SecurityGroup” - groups claim will contain the identifiers of all security groups of which the user is a member.

“All” - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member

Now your application will include group claims in your manifest and you can use this fact in your code.

References:

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

Question #:49 - (Exam Topic 3)

You are developing a solution for a hospital to support the following use cases:

- The most recent patient status details must be retrieved even if multiple users in different locations have updated the patient record.
- Patient health monitoring data retrieved must be the current version or the prior version.
- After a patient is discharged and all charges have been assessed, the patient billing record contains the final charges.

You provision a Cosmos DB NoSQL database and set the default consistency level for the database account to Strong. You set the value for Indexing Mode to Consistent.

You need to minimize latency and any impact to the availability of the solution. You must override the default consistency level at the query level to meet the required consistency guarantees for the scenarios.

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

NOTE: Each correct selection is worth one point.

Consistency levels**Answer Area**

Strong

Bounded Staleness

Consistent Prefix

Eventual

Return the most recent patient status.

Return health monitoring data that is no less than one version behind.

After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges

Answer:**Consistency levels****Answer Area**

Strong

Bounded Staleness

Consistent Prefix

Eventual

Return the most recent patient status.

Return health monitoring data that is no less than one version behind.

After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges

Strong

Bounded Staleness

Eventual

Explanation

Return the most recent patient status.

Strong

Return health monitoring data that is no less than one version behind.

Bounded Staleness

After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges

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.

Question #:50 - ([Exam Topic 3](#))

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

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 Area



Answer:

Actions

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 Area

Create a new template deployment.

Export a Resource Manager template.

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.

Explanation

Create a new template deployment.

Export a Resource Manager template.

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.

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

The steps are:

- 1. Export a template.

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

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

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move>

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

Question #:51 - (Exam Topic 3)

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. Tetemetry.Context.Cloud.RoleInstance
- C. Telemetry.Id
- D. Telemetry.ContextSession.Id
- E. Telemetry.Name

Answer: A C

Explanation

References:

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

Question #52 - ([Exam Topic 3](#))

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

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

You need to complete the configuration.

Which Azure CLI or PowerShell command should you run?

- A. `az servicebus namespace create`
 - `- --resource-group fridge-rg`
 - `- --name fridge-ns`
 - `- --location fridge-loc`
- B. `az servicebus queue create`
 - `--resource-group fridge-rg`
 - `--namespace-name fridge-ns`
 - `--name fridge-q`
- C. `connectionString=$(az servicebus namespace authorization-rule keys list`
 - `--resource-group fridge-rg`
 - `--fridge-ns fridge-ns`
 - `--name RootManageSharedAccessKey`
 - `--query primaryConnectionString --output tsv)`
- D. `az group create`
 - `--name fridge-rg`
 - `--location fridge-log`

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation

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

Note:

Steps:

Step 1: # Create a resource group

```
resourceGroupName="myResourceGroup"
```

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

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

```
namespaceName=myNameSpace$RANDOM
```

```
az servicebus namespace create --resource-group $resourceGroupName --name $namespaceName --location eastus
```

Step 3: # Create a Service Bus queue

```
az servicebus queue create --resource-group $resourceGroupName --namespace-name $namespaceName --name BasicQueue
```

Step 4: # Get the connection string for the namespace

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

References:

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

Question #:53 - (Exam Topic 3)

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

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

Data throughout must be maximized. Latency must be minimized.

You need to implement the Azure Event Hub.

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

NOTE: Each correct selection is worth one point.

Setting	Value
Number of partitions	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

Explanation

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

Box 1: 6

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

There are 6 highways.

Box 2: Highway

References:

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

Question #:54 - [\(Exam Topic 3\)](#)

ASP.NET Core API app by using C#. The API app will allow users to authenticate by using Twitter and Azure Active Directory (Azure AD).

Users must be authenticated before calling API methods. You must log the user's name for each method call.

You need to configure the API method calls.

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

NOTE: Each correct selection is worth one point.

Code segment**Value**

Attribute

▼
Authorize
AllowAnonymous
AutoValidateAntiforgeryToken

Request Header

▼
X-MS-CLIENT-PRINCIPAL-NAME
Proxy-Authorization
X-Forwarded-For
X-MS-CLIENT-PRINCIPAL-ID

Answer:**Code segment****Value**

Attribute

▼
Authorize
AllowAnonymous
AutoValidateAntiforgeryToken

Request Header

▼
X-MS-CLIENT-PRINCIPAL-NAME
Proxy-Authorization
X-Forwarded-For
X-MS-CLIENT-PRINCIPAL-ID

Explanation

Code segment**Value**

Attribute

Authorize
AllowAnonymous
AutoValidateAntiforgeryToken

Request Header

X-MS-CLIENT-PRINCIPAL-NAME
Proxy-Authorization
X-Forwarded-For
X-MS-CLIENT-PRINCIPAL-ID

+ Explanation:

Box 1: Authorize

Box 2: X-MS-CLIENT-PRINCIPAL-NAME

App Service passes user claims to your application by using special headers. External requests aren't allowed to set these headers, so they are present only if set by App Service. Some example headers include:

X-MS-CLIENT-PRINCIPAL-NAME

X-MS-CLIENT-PRINCIPAL-ID

Here's the set of headers you get from Easy Auth for a Twitter authenticated user:

```
{
  "cookie": "AppServiceAuthSession=Lx43...xHDTA==",
  "x-ms-client-principal-name": "evilSnobu",
  "x-ms-client-principal-id": "35....",
  "x-ms-client-principal-idp": "twitter",
  "x-ms-token-twitter-access-token": "35...Dj",
  "x-ms-token-twitter-access-token-secret": "OK3...Jx",
}
```

References:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-authentication-how-to>

Question #:55 - (Exam Topic 3)

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.	☐	☐
The last message read remains in the queue after the code runs.	☐	☐
The storage queue remains in the storage account after the code runs.	☐	☐

Answer:

Statement**Yes****No**

The code configures the lock duration for the queue.

☐☒

The last message read remains in the queue after the code runs.

☒☐

The storage queue remains in the storage account after the code runs.

☒☐**Explanation****Statement****Yes****No**

The code configures the lock duration for the queue.

☐☒

The last message read remains in the queue after the code runs.

☐☐

The storage queue remains in the storage account after the code runs.

☐☐

Box 1: No

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

Box 2: Yes

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

Box 3: Yes

Reference:

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

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

Question #:56 - (Exam Topic 3)

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

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

You 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

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-vn>

Question #:57 - (Exam Topic 3)

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

You have the following parameters:

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

In which order should you arrange the commands to develop the solution? To answer, move the appropriate

commands from the list of command segments to the answer area and arrange them in the correct order.

Command segments

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

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

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

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

Answer Area

Answer:

Command segments

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

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

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

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

Answer Area

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

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

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

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

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:16.04-LTS \
--node-agent-sku-id "batch.node.ubuntu 16.04"
```

Step 2: az batch job create

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

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

Step 3: az batch task create

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

```
az batch task create \
```

```
--job-id myjob \
```

```
--task-id task1 \
```

```
--command-line "/bin/bash -c 'printenv AZ_BATCH_TASK_WORKING_DIR'"
```

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

References:

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

Question #:58 - ([Exam Topic 3](#))

You

are developing an Azure App Service hosted ASP.NET Core web app to deliver video on-demand streaming media. You enable an Azure Content Delivery Network (CDN) Standard for the web endpoint. Customer videos are downloaded from the web app by using the following example URL.:

<http://www.contoso.com/content.mp4?quality=1>

All media content must expire from the cache after one hour. Customer videos with varying quality must be delivered to the closest regional point of presence (POP) node.

You need to configure Azure CDN caching rules.

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

NOTE: Each correct selection is worth one point.

Setting	Action
Caching behavior	Bypass cache Override Set if missing
Cache expiration duration	1 second 1 minute 1 hour 1 day
Query string caching behavior	Ignore query strings Bypass caching for query strings Cache every unique URL

Answer:

Setting	Action
Caching behavior	Bypass cache Override Set if missing
Cache expiration duration	1 second 1 minute 1 hour 1 day
Query string caching behavior	Ignore query strings Bypass caching for query strings Cache every unique URL

Explanation

Setting	Action
Caching behavior	▼ Bypass cache Override Set if missing
Cache expiration duration	▼ 1 second 1 minute 1 hour 1 day
Query string caching behavior	▼ Ignore query strings Bypass caching for query strings Cache every unique URL

Box 1: Override

Override: Ignore origin-provided cache duration; use the provided cache duration instead. This will not override cache-control: no-cache.

Set if missing: Honor origin-provided cache-directive headers, if they exist; otherwise, use the provided cache duration.

Incorrect:

Bypass cache: Do not cache and ignore origin-provided cache-directive headers.

Box 2: 1 hour

All media content must expire from the cache after one hour.

Box 3: Cache every unique URL

Cache every unique URL: In this mode, each request with a unique URL, including the query string, is treated as a unique asset with its own cache. For example, the response from the origin server for a request for example.ashx?q=test1 is cached at the POP node and returned for subsequent caches with the same query string. A request for example.ashx?q=test2 is cached as a separate asset with its own time-to-live setting.

Reference:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-query-string>

Question #59 - (Exam Topic 3)

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

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

You need to implement a reply trail auditing solution.

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

NOTE: Each correct selection is worth one point.

- A. Assign the value of the hazard message SessionID property to the SequenceNumber property.
- B. Assign the value of the hazard message SequenceNumber property to the DeliveryCount property. C. Assign the value of the hazard message MessageId property to the DeliveryCount property.
- C. Assign the value of the hazard message SessionID property to the ReplyToSessionId property.
- D. Assign the value of the hazard message MessageId property to the SequenceNumber property.
- E. Assign the value of the hazard message MessageId property to the CorrelationId property.

Answer: A B

Question #60 - (Exam Topic 3)

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

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

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

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

Solution: Update the web.config file to include the applicationInitialization configuration element. Specify

custom initialization actions to run the scripts.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation

Specify custom warm-up.

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

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

Reference:

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

Question #:61 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

```

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

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

;
        RowKey= 

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

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

|                    |   |
|--------------------|---|
|                    | ▼ |
| CloudTable         |   |
| CloudTableClient   |   |
| TableEntity        |   |
| TableEntityAdapter |   |

 table, string pk, string rk)
    {
        

|                                                                       |   |
|-----------------------------------------------------------------------|---|
|                                                                       | ▼ |
| TableEntity query =TableEntity.Retrieve<PlayerEntity>(pk, rk);        |   |
| TableOperation query =TableOperation.Retrieve<PlayerEntity>(pk,rk);   |   |
| TableResult query =TableQuery.Retrieve<PlayerEntity>(pk,rk);          |   |
| TableResultSegment query =TableResult.Retrieve<PlayerEntity>(pk, rk); |   |



|                                                       |   |
|-------------------------------------------------------|---|
|                                                       | ▼ |
| TableEntity data =await table.ExecuteAsync(query);    |   |
| TableOperation data =await table.ExeucteAsync(query); |   |
| TableQuery data =await table.ExecuteAsync(query);     |   |
| TableResult data =await table.ExecuteAsync(query);    |   |




        player=data.Result as PlayerEntity;
    }
}



```


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




        
        TableEntity data = await table.ExecuteAsync(query);
        TableOperation data = await table.ExeucteAsync(query);
        TableQuery data = await table.ExecuteAsync(query);
        TableResult data = await table.ExecuteAsync(query);

        player = data.Result as PlayerEntity;
    }
}

```


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:

```

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

    TableEntity data =await table.ExecuteAsync(query);
    TableOperation data =await table.ExeucteAsync(query);
    TableQuery data =await table.ExecuteAsync(query);
    TableResult data =await table.ExecuteAsync(query);

    player=data.Result as PlayerEntity;
}
}

```

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-tables>

Question #:62 - (Exam Topic 3)

You are creating a CLI script that creates an Azure web app related services in Azure App Service. The web app uses the following variables:

Variable name	Value
\$gitrepo	https://github.com/Contos/webapp
&webappname	Webapp1103

You need to automatically deploy code from GitHub to the newly created web app.

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

NOTE: Each correct selection is worth one point.

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

▼ --name \$webappname --resource-group myResourceGroup --sku FREE

az webapp create
az appservice plan create
az webapp deployment
az group delete

▼ --name \$webappname --resource-group myResourceGroup

az webapp create
az appservice plan create
az webapp deployment
az group delete

▼
- -repo-url \$gitrepo - -branch master - -manual-integration
git clone \$gitrepo
- -plan \$webappname

▼ source config --name \$webappname

az webapp create
az appservice plan create
az webapp deployment
az group delete

- -resource-group myResourceGroup

▼
- -repo-url \$gitrepo - -branch master - -manual-integration
git clone \$gitrepo
- -plan \$webappname

Answer:

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

```
az webapp create --name $webappname --resource-group myResourceGroup --sku FREE
```

```
az appservice plan create  
az webapp deployment  
az group delete
```

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

```
az appservice plan create  
az webapp deployment  
az group delete
```

```
--repo-url $gitrepo --branch master --manual-integration
```

```
git clone $gitrepo
```

```
--plan $webappname
```

```
source config --name $webappname
```

```
az webapp create  
az appservice plan create  
az webapp deployment  
az group delete
```

```
--resource-group myResourceGroup
```

```
--repo-url $gitrepo --branch master --manual-integration
```

```
git clone $gitrepo
```

```
--plan $webappname
```

Explanation

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

```
az webapp create --name $webappname --resource-group myResourceGroup --sku FREE
```

```
az appservice plan create
az webapp deployment
az group delete
```

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

```
az appservice plan create
az webapp deployment
az group delete
```

```
--repo-url $gitrepo --branch master --manual-integration
git clone $gitrepo
--plan $webappname
```

```
source config --name $webappname
```

```
az webapp create
az appservice plan create
az webapp deployment
az group delete
```

```
--resource-group myResourceGroup
```

```
--repo-url $gitrepo --branch master --manual-integration
git clone $gitrepo
--plan $webappname
```

Box 1: az appservice plan create

The azure group creates command successfully returns JSON result. Now we can use resource group to create a azure app service plan

Box 2: az webapp create

Create a new web app..

Box 3: --plan \$webappname

with the serviceplan we created in step 1.

Box 4: az webapp deployment

Continuous Delivery with GitHub. Example:

```
az webapp deployment source config --name firstsamplewebsite1 --resource-group websites --repo-url $gitrepo
```


--branch master --git-token \$token

Box 5: --repo-url \$gitrepo --branch master --manual-integration

Reference:

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

Question #:63 - (Exam Topic 3)

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

You add the following parameters to the function.

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

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

You need to ensure the files are correctly processed.

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

NOTE: Each correct selection is worth one point.

Answer Area

```

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

        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)));
            outputFileList.Add(new OutputFile(fileTask.Output,
                new OutputFileDestination(failedContainer),
                new OutputFileUploadOptions(OutputFileUploadCondition.
                    TaskFailure)));
            task = new CloudTask(taskId, fileTask.Command,
                new OutputFileDestination(failedContainer),
                new OutputFileUploadOptions(OutputFileUploadCondition.
                    TaskCompletion));
            task.Add(task);
        });
    }
}

```

```
});  
}  
return tasks,  
}
```

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(jobId);

        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)));
            outputFileList.Add(new OutputFile(fileTask.Output,
                new OutputFileDestination(failedContainer),
                new OutputFileUploadOptions(OutputFileUploadCondition.
                    TaskFailure)));
            task = new CloudTask(taskId, fileTask.Command,
                outputFileList);
            task.Add(task);
        });
    }
}

```

task = outputFileList;

task.Add(task);

```
});
}
return tasks,
}
```

Explanation

```
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	



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>

Question #:64 - (Exam Topic 3)

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.

NOTE: Each correct selection is worth one point.


```
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}");
    }
}
```

Yes

No

The code will log the time that the order was processed from the queue.

☐☐

When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.

☐☐

When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders.

☐☐

The ProcessOrders function will output the order to an Orders table in Azure Table Storage.

☐☐

Answer:

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

Explanation

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

Box 1: No

ExpirationTime - The time that the message expires.

InsertionTime - The time that the message was added to the queue.

Box 2: Yes

maxDequeueCount - The number of times to try processing a message before moving it to the poison queue.

Default value is 5.

Box 3: Yes

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

Box 4: Yes

References:

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

Question #:65 - (Exam Topic 3)

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

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

```

1 public void GetScore(string playerId, int score, string gameName)
2 {
3     TableQuery<DynamicTableEntity> query = new TableQuery<DynamicTableEntity>().Select(new string[] { "Score" })
        .Where(TableQuery.GenerateFilterConditionForInt("Score", QueryComparisons.GreaterThanOrEqual, 15000)).Take
        (20);
4     EntityResolver<KeyValuePair<string, int?>> resolver =
        (partitionKey, rowKey, ts, props, etag) => new KeyValuePair<string, int?>(rowKey, props["Score"].Int32Value);
5     foreach (var scoreItem in scoreTable.ExecuteQuery(query, resolver, null, null))
6     {
7         Console.WriteLine($"{scoreItem.Key} {scoreItem.Value}");
8     }
9
10 public class PlayerScore : TableEntity
11 {
12     public PlayerScore(string gameId, string playerId, int score, long timePlayed)
13     {
14         PartitionKey = gameId;
15         RowKey = playerId;
16         Score = score;
17         TimePlayed = timePlayed;
18     }
19     public int Score { get; set; }
20     public long TimePlayed { get; set; }
21 }

```

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

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

Box 1: No

Box 2: Yes

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

Example:

```
query.Take(10);
```

Box 3: Yes

Box 4: Yes

References:

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

Question #:66 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

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/>

Question #:67 - (Exam Topic 3)

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

You have the following requirements:

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

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

Yes

No

The code will work with Cosmos DB.

☒☐

The save score function will update and replace a record if one already exists with the same playerId and gameId.

☐☒

The data for the game will be automatically partitioned.

☐☒

This code will store the values for the gameId and playerId parameters in the database.

☒☐

Box 1: Yes

Code for CosmosDB, example:

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

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse(
```

```
CloudConfigurationManager.GetSetting("StorageConnectionString"));
```

```
// Create the table client.
```

```
CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
```

```
// Retrieve a reference to the table.
```

```
CloudTable table = tableClient.GetTableReference("people");
```

```
// Create the TableOperation object that inserts the customer entity.
```

```
TableOperation insertOperation = TableOperation.Insert(customer1);
```

Box 2: No

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

Box 3: No

No partition key is used.

Box 4: Yes

References:

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

Question #:68 - (Exam Topic 3)

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

Explanation

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

Application Insights manages the ID of a session for you.

References:

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

Question #:69 - (Exam Topic 3)


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.



Answer:



Explanation



Box 1: Deployment

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

Box 2: ScaledObject

With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.

Example of ScaledObject with polling interval

apiVersion: keda.k8s.io/v1alpha1

kind: ScaledObject

metadata:

name: transformer-fn

namespace: tt

labels:

deploymentName: transformer-fn

spec:

scaleTargetRef:

deploymentName: transformer-fn

pollingInterval: 5

minReplicaCount: 0

maxReplicaCount: 100

Box 3: Secret

Store connection strings in Kubernetes Secrets.

Example: to create the Secret in our demo Namespace:

create the k8s demo namespace

kubectl create namespace tt

grab connection string from Azure Service Bus

KEDA_SCALER_CONNECTION_STRING=\$(az servicebus queue authorization-rule keys list \

-g \$RG_NAME \

```
--namespace-name $SBN_NAME \  
  
--queue-name inbound \  
  
-n keda-scaler \  
  
--query "primaryConnectionString" \  
  
-o tsv)  
  
# create the kubernetes secret  
  
kubectl create secret generic tt-keda-auth \  
  
--from-literal KedaScaler=$KEDA_SCALER_CONNECTION_STRING \  
  
--namespace tt
```

Reference:

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

Question #:70 - ([Exam Topic 3](#))

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation

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

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

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

Reference:

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

Question #:71 - (Exam Topic 3)

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

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>

Question #:72 - (Exam Topic 3)

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

After you answer a question in this 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

Explanation

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

References:

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

Question #:73 - (Exam Topic 3)

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

Explanation

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

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

Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.

References:

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

Question #:74 - (Exam Topic 3)

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

Explanation

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

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

References:

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

Question #:75 - (Exam Topic 3)

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

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

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

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation

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

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

Reference:

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

Question #:76 - (Exam Topic 3)

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**Answer:****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

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`.

**Explanation**

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

Run the Azure PowerShell command `Set-AzureRmVMDisk`.

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'
```

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

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

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

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

Incorrect:

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

References:

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

Question #:77 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: A D

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>

Question #:78 - (Exam Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the

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

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

You are developing an Azure 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

Explanation

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

Reference:

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

Question #:79 - (Exam Topic 3)

You are developing a project management service by using ASP.NET. The service hosts conversations, files, to-do lists, and a calendar that users can interact with at any time.

The application uses Azure Search for allowing users to search for keywords in the project data.

You need to implement code that creates the object which is used to create indexes in the Azure Search service.

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

NOTE: Each correct selection is worth one point.

- A. SearchService
- B. SearchIndexClient
- C. SearchServiceClient
- D. SearchCredentials

Answer: B C

Explanation

The various client libraries define classes like Index, Field, and Document, as well as operations like Indexes.Create and Documents.Search on the SearchServiceClient and SearchIndexClient classes.

Example:

The sample application we'll be exploring creates a new index named "hotels", populates it with a few documents, then executes some search queries. Here is the main program, showing the overall flow:

/ This sample shows how to delete, create, upload documents and query an index

```
static void Main(string[] args)
{
    IConfigurationBuilder builder = new ConfigurationBuilder().AddJsonFile("appsettings.json");
    IConfigurationRoot configuration = builder.Build();

    SearchServiceClient serviceClient = CreateSearchServiceClient(configuration);

    Console.WriteLine("{0}", "Deleting index...\n");

    DeleteHotelsIndexIfExists(serviceClient);

    Console.WriteLine("{0}", "Creating index...\n");

    CreateHotelsIndex(serviceClient);

    ISearchIndexClient indexClient = serviceClient.Indexes.GetClient("hotels");
```

References:

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

Question #:80 - [\(Exam Topic 3\)](#)

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

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation

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

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

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

Reference:

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

Question #:81 - (Exam Topic 3)

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

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

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

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation

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

References:

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

Question #:82 - (Exam Topic 3)

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

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

You need to implement an Azure Service Bus solution.

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

Actions

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

Answer area**Answer:****Actions**

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

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**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.

C:\Users\wk\Desktop\mudassar\Untitled.jpg

Box 1: Create a single Service Bus Namespace

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

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

Create topics.

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

Topics can have multiple, independent subscriptions.

References:

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

Question #:83 - (Exam Topic 3)

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 Contribute 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

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=azurerm>

Question #:84 - (Exam Topic 3)

You are developing a ticket reservation system for an airline.

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

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Strong
Eventual
ConsistentPrefix
BoundedStaleness

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

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

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

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

```
--default-consistency-level - $consistencylevel
```

Answer:

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

Strong
Eventual
ConsistentPrefix
BoundedStaleness

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

```

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

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

```

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

```
--default-consistency-level - $consistencylevel
```

Explanation

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

Strong
Eventual
ConsistentPrefix
BoundedStaleness

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

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

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

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

```
--default-consistency-level - $consistencylevel
```

Box 1: BoundedStaleness

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. In other words, 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

Reference:

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

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/manage-with-cli.md>

Question #:85 - (Exam Topic 3)

Contoso, Ltd. provides an API to customers by using Azure API Management (APIM). The API authorizes

users with a JWT token.

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

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

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

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

NOTE: Each correct selection is worth one point

Policy section

Inbound

Outbound

Answer Area

Policy

Set-variable

Cache-lookup-value

Cache-store-value

Find-and-replace

Policy section

policy section

policy section

policy section

policy section

Answer:

Policy section

Inbound

Outbound

Answer Area

Policy

Set-variable

Cache-lookup-value

Cache-store-value

Find-and-replace

Policy section

Inbound

Inbound

Outbound

Outbound

Explanation

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

C:\Users\wk\Desktop\mudassar\Untitled.jpg

Box 1: Inbound.

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

Example:

```
<policies>
```

```
<inbound>
```

```
<!-- How you determine user identity is application dependent -->
```

```
<set-variable
```

```
name="enduserid"
```

```
value="@ (context.Request.Headers.GetValueOrDefault("Authorization","").Split(' ')[1].AsJwt()?.Subject)" />
```

Box 2: Inbound

A cache-lookup-value policy

Example:

```
<inbound>
```

```
<base />
```

```
<cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false"
downstream-caching-type="none | private | public" must-revalidate="true | false">
```

```
<vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several
times -->
```

```
</cache-lookup>
```

</inbound>

Box 3: Outbound

A cache-store-value policy.

Example:

<outbound>

<base />

<cache-store duration="3600" />

</outbound>

Box 4: Outbound

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

Example:

<outbound>

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

<find-and-replace

from="\$userprofile\$"

to="@((string)context.Variables["userprofile"])" />

<base />

</outbound>

Reference:

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

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

Question #:86 - (Exam Topic 3)

You are writing code to create and run an Azure Batch job.

You have created a pool of compute nodes.

You need to choose the right class and its method to submit a batch job to the Batch service.

Which method should you use?

- A. JobOperations.CreateJobO
- B. CloudJob.Enable(IEnumerable<BatchClientBehavior>)
- C. CloudJob.CommitAsync(IEnumerable<BatchClientBehavior>, CancellationToken)
- D. JobOperations.EnableJob(String, IEnumerable<BatchClientBehavior>)
- E. JobOperations.EnableJobAsync(String, IEnumerable<BatchClientBehavior>, CancellationToken)

Answer: C

Explanation

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

The Commit method submits the job to the Batch service. Initially the job has no tasks.

```
{  
CloudJob job = batchClient.JobOperations.CreateJob();  
job.Id = JobId;  
job.PoolInformation = new PoolInformation { PoolId = PoolId };  
job.Commit();  
}
```

References:

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

Question #:87 - (Exam Topic 3)

You develop a web app that uses tier D1 app service plan by using the Web Apps feature of Microsoft Azure App Service.

Spikes in traffic have caused increases in page load times.

You need to ensure that the web app automatically scales when CPU load is about 85 percent and minimize costs.

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

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions

Configure the web app to the Premium App Service tier.

Configure the web app to the Standard App Service tier.

Enable autoscaling on the web-app.

Add a Scale rule.

Switch to an Azure App Services consumption plan.

Configure a Scale condition.

Answer Area



Answer:

Actions

Configure the web app to the Premium App Service tier.

Configure the web app to the Standard App Service tier.

Enable autoscaling on the web-app.

Add a Scale rule.

Switch to an Azure App Services consumption plan.

Configure a Scale condition.

Answer Area

Configure the web app to the Premium App Service tier.

Enable autoscaling on the web-app.

Add a Scale rule.

Configure a Scale condition.

Explanation

Configure the web app to the Standard App Service tier.

Enable autoscaling on the web-app.

Add a Scale rule.

Configure a Scale condition.

Step 1: Configure the web app to the Standard App Service Tier

The Standard tier supports auto-scaling, and we should minimize the cost.

Step 2: Enable autoscaling on the web app

First enable autoscale

Step 3: Add a scale rule

Step 4: Add a Scale condition

Reference:

<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-autoscale-get-started>

Question #:88 - (Exam Topic 3)

You are developing an application that use an Azure blob named data to store application data. The application creates blob snapshots to allow application state to be reverted to an earlier state. The Azure storage account has soft deleted enabled.

The system performs the following operations in order:

- The blob is updated
- Snapshot 1 is created.
- Snapshot 2 is created.
- Snapshot 1 is deleted.

A system error then deletes the data blob and all snapshots.

You need to determine which application states can be restored.

What is the restorability of the application data? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Application State**Restorability**

Data blob

	▼
Can be restored	
Cannot be restored	

Snapshot 1

	▼
Can be restored	
Cannot be restored	

Snapshot 2

	▼
Can be restored	
Cannot be restored	

Answer:**Application State****Restorability**

Data blob

	▼
Can be restored	
Cannot be restored	

Snapshot 1

	▼
Can be restored	
Cannot be restored	

Snapshot 2

	▼
Can be restored	
Cannot be restored	

Explanation

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

Box 1: Can be restored

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

Box 2: Cannot be restored

It has been deleted.

Box 3: Can be restored

It has not been deleted.

References:

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

Question #:89 - (Exam Topic 3)

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",
  "optionalClaims": "All",
  "groupMembershipClaims": true
  "allowPublicClient": true
  "oauth2Permissions": true
  "requiredResourceAccess": true
  "oauth2AllowImplicitFlow": true
  ...
}
```

Explanation



Box 1: groupMembershipClaims

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

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

Enable group membership claims by changing the groupMembershipClaim

The valid values are:

"All"

"SecurityGroup"

"DistributionList"

"DirectoryRole"

Box 2: oauth2Permissions

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

oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

Question #:90 - (Exam Topic 3)

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

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

The system has the following requirements for subscriptions

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

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

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

NOTE: Each correct selection is worth one point.

Filter types**Answer Area****Subscription**

FutureOrders

HighPriorityOrders

InternationalOrders

HighQuantityOrders

AllOrders

Filter type**Answer:****Filter types****Answer Area****Subscription**

FutureOrders

HighPriorityOrders

InternationalOrders

HighQuantityOrders

AllOrders

Filter type**Explanation**

Answer Area

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

FutureOrders: SQLFilter

HighPriorityOrders: CorrelationFilter

CorrelationID only

InternationalOrders: SQLFilter

Country NOT USA requires an SQL Filter

HighQuantityOrders: SQLFilter

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

AllOrders: No Filter

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

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

Boolean filters - The TrueFilter and FalseFilter either cause all arriving messages (true) or none of the arriving

messages (false) to be selected for the subscription.

References:

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

Question #:91 - (Exam Topic 3)

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

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

You 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

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-vn>

Question #:92 - (Exam Topic 3)

You are developing an ASP.NET Core Web API web service that uses Azure Application Insights to monitor performance and track events.

You need to enable logging and ensure that log messages can be correlated to events tracked by Application Insights.

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

NOTE: Each correct selection is worth one point.

Code segments

IncludeEventId

ServerFeatures

LoggerFilterOptions

ApplicationServices

ApplicationInsightsLoggerOptions

TrackExceptionsAsExceptionTelemetry

Answer Area

```
public class Startup
{
    ...
    public void ConfigureServices (IServiceCollection services)
    {
        services.AddOptions< >().
        Configure(o => o. = true );
        services.AddMvc();
    }
    public void Configure (IApplicationBuilder app,
        IHostingEnvironment env, ILoggerFactory loggerFactory)
    {
        loggerFactory.AddApplicationInsights(app. ,LogLevel.Trace);
        app.UseMvc();
    }
}
```

Answer:

Code segments

Answer Area

IncludeEventId	<pre> public class Startup { ... public void ConfigureServices (IServiceCollection services) { services.AddOptions< ApplicationInsightsLoggerOptions >(). Configure(o => o. IncludeEventId = true); services.AddMvc(); } public void Configure (IApplicationBuilder app, IHostingEnvironment env, ILoggerFactory loggerFactory) { loggerFactory.AddApplicationInsights(app. ApplicationServices ,LogLevel.Trace); app.UseMvc(); } } </pre>
ServerFeatures	
LoggerFilterOptions	
ApplicationServices	
ApplicationInsightsLoggerOptions	
TrackExceptionsAsExceptionTelemetry	

Explanation

```

public class Startup
{
    ...
    public void ConfigureServices (IServiceCollection services)
    {
        services.AddOptions< ApplicationInsightsLoggerOptions >().
        Configure(o => o. IncludeEventId = true );
        services.AddMvc();
    }
    public void Configure (IApplicationBuilder app,
        IHostingEnvironment env, ILoggerFactory loggerFactory)
    {
        loggerFactory.AddApplicationInsights(app. ApplicationServices ,LogLevel.Trace);
        app.UseMvc();
    }
}

```

Box 1: ApplicationInsightsLoggerOptions

If you want to include the EventId and EventName properties, then add the following to the ConfigureServices method:

services

```
AddOptions<ApplicationInsightsLoggerOptions>()
```

```
Configure(o => o.IncludeEventId = true);
```

Box 2: IncludeEventID

Box 3: ApplicationServices

In Asp.Net core apps it turns out that trace logs do not show up in Application Insights out of the box. We need to add the following code snippet to our Configure method in Startup.cs:

```
loggerFactory.AddApplicationInsights(app.ApplicationServices, logLevel);
```

References:

<https://blog.computedcloud.com/enabling-application-insights-trace-logging-in-asp-net-core/>

Question #:93 - [\(Exam Topic 3\)](#)

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

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

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

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

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

Answer Area**Answer:****Actions**

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

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

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

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

Answer Area

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

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

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

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

**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>

Question #94 - (Exam Topic 3)

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/>

Question #95 - (Exam Topic 3)

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>

Question #:96 - (Exam Topic 3)

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

You need to identify trends in application usage.

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

NOTE: Each correct selection is worth one point.

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?

Question #:97 - ([Exam Topic 3](#))

You are developing Azure WebJobs.

You need to recommend a WebJob type for each scenario.

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

NOTE: Each correct selection is worth one point.

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

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>

Question #:98 - (Exam Topic 3)

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

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

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

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

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

Solution: Provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.

A. Yes

B. No

Answer: A

Explanation

References:

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

Question #:99 - (Exam Topic 3)

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

A rule already exists to scale up the App Service when the average queue length of unprocessed and valid queue messages is greater than 1000.

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

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

NOTE: Each correct selection is worth one point.

Answer Area

Scale rule

Metric source

Storage queue
Service Bus queue
Current resource
Storage queue (classic)

Resource type

Service Bus Namespaces

Resource

MessageQueue1103

* Queues

itemqueue

Criteria

* Metric name

Message Count
Active Message Count

1 minute time grain

* Time grain statistic

Total
Maximum
Average
Count

Greater than
Greater than or equal to
Less than
Less than or equal to

* Threshold

1000

Action

* Operation

Increase count by
Increase count to
Decrease count by
Decrease count to

* Instance count

1

* Cool down (minutes)

5

Answer:



Answer Area

Scale rule

Metric source

Storage queue
Service Bus queue
Current resource
Storage queue (classic)

Resource type

Service Bus Namespaces

Resource

MessageQueue1103

* Queues

itemqueue

Criteria

* Metric name

Message Count
Active Message Count

1 minute time grain

* Time grain statistic

Total
Maximum
Average
Count

Greater than
Greater than or equal to
Less than
Less than or equal to

* Threshold

1000

Action

* Operation

Increase count by
Increase count to
Decrease count by
Decrease count to

* Instance count

1

* Cool down (minutes)

5

Explanation

Answer Area

Scale rule

Metric source

Storage queue
Service Bus queue
Current resource
Storage queue (classic)

Resource type

Service Bus Namespaces

Resource

MessageQueue1103

* Queues

itemqueue

Criteria

* Metric name

Message Count
Active Message Count

1 minute time grain

* Time grain statistic

Total
Maximum
Average
Count

* Operator

Greater than
Greater than or equal to
Less than
Less than or equal to

Threshold
1000

Action
Operation
Increase count by
Increase count to
Decrease count by
Decrease count to

Instance count
1

Cool down (minutes) 5

Box 1: Service bus queue

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

Box 2: ActiveMessage Count

ActiveMessageCount: Messages in the queue or subscription that are in the active state and ready for delivery.

Box 3: Count

Box 4: Less than or equal to

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

Box 5: Decrease count by

Question #:100 - [\(Exam Topic 3\)](#)

A company runs an international travel and bookings management service. The company plans to begin offering restaurant bookings. You must develop a solution that uses Azure Search and meets the following requirements:

- Users must be able to search for restaurants by name, description, location, and cuisine.
- Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

- All words in descriptions must be included in searches.

You need to add annotations to the restaurant class.

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

NOTE: Each correct selection is worth one point.



```
[SerializePropertyNameAsCamelCase]
```

```
public class Restaurant
```

```
{
```

```
    [Key, IsFilterable]
```

```
    public int RestaurantId { get; set; }
```

```
    [IsSearchable, IsFilterable, IsSortable]
```

```
    public string Name { get; set; }
```

```
    [IsSearchable.IsFilterable.IsSortable, IsFacetable]
```

```
    [IsFilterable.IsFacetable, Required]
```

```
    [IsSearchable]
```

```
    [IsSearchable, Required]
```

```
    public string location { get; set; }
```

```
    public string Phone { get; set; }
```

```
    [Required]
```

```
    [IsSearchable]
```

```
    [IsFilterable, IsFacetable, Required]
```

```
    [IsFilterable, IsFacetable, IsSortable]
```

```
    public string Description { get; set; }
```

```
    [IsFilterable, IsSortable, IsSearchable]
```

```
    [IsFilterable, IsSortable, IsFacetable]
```

```
    [IsFilterable, IsSortable, Key]
```

```
    [IsFilterable, IsSortable, IsSearchable, Required]
```

```
    public double Rating { get; set; }
```

```
    [IsSearchable, IsFilterable, IsFacetable]
```

```
    [IsFilterable, IsSortable, Key]
```

```
    [IsFilterable, IsSortable, IsSearchable]
```

```
    [IsFilterable, IsSortable, Key, Required]
```

```
    public List<string> Cuisines { get; set; }
```

```
    [IsFilterable, IsSortable, Key, Required]
```

```
    [IsSearchable, IsSortable, IsFacetable]
```

```
    [IsFilterable, IsSortable, IsSearchable]
```

```
    [IsFilterable, IsSortable, Key, Required]
```

```
    [IsFilterable, IsFacetable]
```

```
    public bool FamilyFriendly { get; set; }
```

Answer:

```
[SerializePropertyNameAsCamelCase]
```

```
public class Restaurant
```

```
{
```

```
    [Key, IsFilterable]
```

```
    public int RestaurantId { get; set; }
```

```
    [IsSearchable, IsFilterable, IsSortable]
```

```
    public string Name { get; set; }
```

```
    [IsSearchable, IsFilterable, IsSortable, IsFacetable]
```

```
    [IsFilterable, IsFacetable, Required]
```

```
    [IsSearchable]
```

```
    [IsSearchable, Required]
```

```
    public string location { get; set; }
```

```
    public string Phone { get; set; }
```

```
    [Required]
```

```
    [IsSearchable]
```

```
    [IsFilterable, IsFacetable, Required]
```

```
    [IsFilterable, IsFacetable, IsSortable]
```

```
    public string Description { get; set; }
```

```
    [IsFilterable, IsSortable, IsSearchable]
```

```
    [IsFilterable, IsSortable, IsFacetable]
```

```
    [IsFilterable, IsSortable, Key]
```

```
    [IsFilterable, IsSortable, IsSearchable, Required]
```

```
    public double Rating { get; set; }
```

```
    [IsSearchable, IsFilterable, IsFacetable]
```

```
    [IsFilterable, IsSortable, Key]
```

```
    [IsFilterable, IsSortable, IsSearchable]
```

```
    [IsFilterable, IsSortable, Key, Required]
```

```
    public List<string> Cuisines { get; set; }
```

```
    [IsFilterable, IsSortable, Key, Required]
```

```
    [IsSearchable, IsSortable, IsFacetable]
```

```
    [IsFilterable, IsSortable, Key, IsSearchable]
```

```
    [IsFilterable, IsFacetable]
```

```
    public bool FamilyFriendly { get; set; }
```


Explanation

Answer Area

```
[SerializePropertyNameAsCamelCase]
public class Restaurant
{
    [Key, IsFilterable]
    public int RestaurantId { get; set; }
    [IsSearchable, IsFilterable, IsSortable]
    public string Name { get; set; }

    public string location { get; set; }
    public string Phone { get; set; }


    public string Description { get; set; }
```

Dropdown 1 options:

- [IsSearchable.IsFilterable.IsSortable, IsFacetable]
- [IsFilterable.IsFacetable, Required]
- [IsSearchable]
- [IsSearchable, Required]

Dropdown 2 options:

- [Required]
- [IsSearchable]
- [IsFilterable, IsFacetable, Required]
- [IsFilterable, IsFacetable, IsSortable]



Box 1: [IsSearchable.IsFilterable.IsSortable,IsFacetable]

Location

Users must be able to search for restaurants by name, description, location, and cuisine.

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

Box 2: [IsSearchable.IsFilterable.IsSortable,Required]

Description

Users must be able to search for restaurants by name, description, location, and cuisine.

All words in descriptions must be included in searches.

Box 3: [IsFilterable,IsSortable,IsFaceTable]

Rating

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

Box 4: [IsSearchable.IsFilterable,IsFacetable]

Cuisines

Users must be able to search for restaurants by name, description, location, and cuisine.

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

Box 5: [IsFilterable,IsFacetable]

FamilyFriendly

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.

References:

<https://www.henkboelman.com/azure-search-the-basics/>

Question #:101 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: A C

Question #:102 - (Exam Topic 3)

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

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

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

How should you complete the trigger?

NOTE: Each correct selection is worth one point.

```
function ensureTip() {  
  var r =  
    _value();  
    _readDocument('item');  
    getContext().getRequest();  
    getContext().getResponse();  
  var i = r.getBody();  
  if (!("tip" in i)) {  
    if (request.getValue("tip") === null){  
      if (isNaN(i)["tip"] || i["tip"] === null) {  
        if (typeof _pluck("tip") == 'number') {  
          i["tip"] = 0;  
        }  
      }  
    }  
    r.setBody(i);  
    r.setValue(i);  
    _upsertDocument(i);  
    _replaceDocument(i);  
  }  
}
```

Answer:

```
function ensureTip() {  
  var r =  
    _value();  
    _readDocument('item');  
    getContext().getRequest();  
    getContext().getResponse();  
  
  var i = r.getBody();  
  
  if (!("tip" in i)) {  
    if (request.getValue("tip") === null){  
      if (isNaN(i)["tip"] || i["tip"] === null) {  
        if (typeof _pluck("tip") === 'number') {  
          i["tip"] = 0;  
        }  
      }  
    }  
  
    r.setBody(i);  
    r.setValue(i);  
    _upsertDocument(i);  
    _replaceDocument(i);  
  }  
}
```

Explanation

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

  var i = r.getBody();

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

    r.setBody(i);
    r.setValue(i);
    _upsertDocument(i);
    _replaceDocument(i)
  }
}
```

Box 1: getContext().getRequest();

Box 2: if(isNaN(i)["tip"] ..

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

isNaN() – Stands for “is Not a Number”, if variable is not a number, it return true, else return false.

typeof – If variable is a number, it will returns a string named “number”.

Box 3:r.setBody(i);

// update the item that will be created

References:

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

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

Question #:103 - (Exam Topic 3)

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: A

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>

Question #:104 - (Exam Topic 3)

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

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

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

NOTE: Each correct selection is worth one point.

Triggers and Action Blocks

Insert Entity

Table

processing

Entity

Path X

Show advanced options

Tier blob

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

Blob path

Path X

Blob Tier

Archive

When there are messages in a queue

Queue Name

processing

Show advanced options

Connected to tableStorageAccountConnection. [Change connection.](#)

Recurrence

Interval

1

Frequency

Month

Show advanced options

Answer Area

Set tier age variable

Set tier age variable

For each

Scan all blobs in this folder

Select an output from previous steps

value X

When there are messages in a queue

Queue Name

processing

Show advanced options

Connected to tableStorageAccountConnection. [Change connection.](#)

If true



if false



Add an action



Add an action



Add an action

Answer:



Triggers and Action Blocks

Insert Entity

Table

processing

Entity

Path X

Show advanced options

Tier blob

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

Blob path

Path X

Blob Tier

Archive

When there are messages in a queue

Queue Name

processing

Show advanced options

Connected to tableStorageAccountConnection. [Change connection.](#)

Recurrence

Interval

1

Frequency

Month

Show advanced options

Answer Area

Recurrence

Interval

1

Frequency

Month

Show advanced options

Set tier age variable

Set tier age variable

For each

Scan all blobs in this folder

Select an output from previous steps

value X

When there are messages in a queue

Queue Name

processing

Show advanced options

Connected to tableStorageAccountConnection. [Change connection.](#)

If true

Recurrence

Interval

1

Frequency

Month

Show advanced options

if false

When there are messages in a queue

Queue Name

processing

Show advanced options

Connected to tableStorageAccountConnection. [Change connection.](#)

Add an action

Add an action

Add an action

Explanation

Answer Area

Recurrence

* Interval: 1

* Frequency: Month ✓

Show advanced options ✓



{x} Set tier age variable



10 01 Set tier age variable



For each

Scan all blobs in this folder



* Select an output from previous steps

When there are messages in a queue

* Queue Name

Show advanced options

Connected to tableStorageAccountConnection. [Change connection.](#)

 If true

Recurrence

* Interval * Frequency

Show advanced options

 if false

Box 1: Recurrence

Box 2: Insert Entity

Box 3 (if true): Tier Blob

Box 4: (if false):

Leave blank.

References:

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

Question #:105 - (Exam Topic 3)

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>

Question #:106 - (Exam Topic 3)

You are developing an application. You have an Azure user account that has access to two subscriptions.

You need to retrieve a storage account key secret from Azure Key Vault.

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

Powershell commands

Answer Area

```
$secretvalue = ConvertTo-SecureString  
$storAcctkey -AsPlainText  
-Force  
Set-AzKeyVaultSecret -VaultName  
$vaultName -Name $secretName  
-SecretValue $secretvalue
```

```
Get-AzStorageAccountKey -  
ResourceGroupName $resGroup -Name  
$storAcct
```

```
Set-AzContext -SubscriptionId  
$subscriptionID
```

```
Get-AzKeyVaultSecret -VaultName  
$vaultName
```

```
Get-AzSubscription
```



Answer:

Powershell commands

```
$secretvalue = ConvertTo-SecureString  
$storAcctkey -AsPlainText  
-Force  
Set-AzKeyVaultSecret -VaultName  
$vaultName -Name $secretName  
-SecretValue $secretvalue
```

```
Get-AzStorageAccountKey -  
ResourceGroupName $resGroup -Name  
$storAcct
```

```
Set-AzContext -SubscriptionId  
$subscriptionID
```

```
Get-AzKeyVaultSecret -VaultName  
$vaultName
```

```
Get-AzSubscription
```

Answer Area

```
Get-AzSubscription
```

```
Set-AzContext -SubscriptionId  
$subscriptionID
```

```
Get-AzStorageAccountKey -  
ResourceGroupName $resGroup -Name  
$storAcct
```

```
$secretvalue = ConvertTo-SecureString  
$storAcctkey -AsPlainText  
-Force  
Set-AzKeyVaultSecret -VaultName  
$vaultName -Name $secretName  
-SecretValue $secretvalue
```

```
Get-AzKeyVaultSecret -VaultName  
$vaultName
```

Explanation


```
Get-AzSubscription
```

```
Set-AzContext -SubscriptionId  
$subscriptionID
```

```
Get-AzStorageAccountKey -  
ResourceGroupName $resGroup -Name  
$storAcct
```

```
$secretvalue = ConvertTo-SecureString  
$storAcctkey -AsPlainText  
-Force  
Set-AzKeyVaultSecret -VaultName  
$vaultName -Name $secretName  
-SecretValue $secretvalue
```

```
Get-AzKeyVaultSecret -VaultName  
$vaultName
```



Step 1: Get-AzSubscription

If you have multiple subscriptions, you might have to specify the one that was used to create your key vault. Enter the following to see the subscriptions for your account:

```
Get-AzSubscription
```

Step 2: Set-AzContext -SubscriptionId

To specify the subscription that's associated with the key vault you'll be logging, enter:

```
Set-AzContext -SubscriptionId <subscriptionID>
```

Step 3: Get-AzStorageAccountKey

You must get that storage account key.

```
Step 4: $secretvalue = ConvertTo-SecureString <storageAccountKey> -AsPlainText -Force
```

```
Set-AzKeyVaultSecret -VaultName <vaultName> -Name <secretName> -SecretValue $secretvalue
```

After retrieving your secret (in this case, your storage account key), you must convert that key to a secure string, and then create a secret with that value in your key vault.

Step 5: Get-AzKeyVaultSecret

Next, get the URI for the secret you created. You'll need this URI in a later step to call the key vault and retrieve your secret. Run the following PowerShell command and make note of the ID value, which is the secret's URI:

```
Get-AzKeyVaultSecret -VaultName <vaultName>
```

Reference:

<https://docs.microsoft.com/bs-latn-ba/Azure/key-vault/key-vault-key-rotation-log-monitoring>

Question #:107 - (Exam Topic 3)

You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage

containers named Container1 and Container2.

Uploading of videos occurs on an irregular basis.

You need to copy specific blobs from Container1 to Container2 in real time when specific requirements are met, excluding backup blob copies.

What should you do?

- A. Download the blob to a virtual machine and then upload the blob to Container2.
- B. Run the Azure PowerShell command Start-AzureStorageBlobCopy.
- C. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API.
- D. Use AzCopy with the Snapshot switch blobs to Container2.

Answer: B

Explanation

The Start-AzureStorageBlobCopy cmdlet starts to copy a blob.

Example 1: Copy a named blob

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

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

References:

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

Question #108 - (Exam Topic 3)

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

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

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

Explanation

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

2. Create the indexBatch with the documents

Something like:

```
var hotels = new Hotel[];  
  
{
```

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

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

Example:

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

References:

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

Question #:109 - (Exam Topic 3)

A company is developing a mobile app for field service employees using Azure App Service Mobile Apps as the backend.

The company's network connectivity varies throughout the day. The solution must support offline use and synchronize changes in the background when the app is online app.

You need to implement the solution.

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

NOTE: Each correct selection is worth one point.

```
var client = new MobileServiceClient("MOBILE_APP_URL");
var store = new MobileServiceSQLiteStore
(Constants.OfflineDbPath);
store.DefineTable<ToDoItem>();
await client.SyncContext.IntitalizeAsync(store);
```

var todoTable = client.GetSyncTable<ToDoItem>();
var todoTable = client.GetTable<ToDoItem>();
var todoTable = client.SyncTable;
var todoTable = client.Table;

```
await client.SyncContext.PushAsync();
```

await todoTable.PullAsync("allToDoItems",todoTable.CreateQuery());
await todoTable.UpdateAsync();
todoTable.PullAsync("allToDoItems", todoTable.CreateQuery());
todoTable.UpdateAsync();

Answer:

```
var client = new MobileServiceClient("MOBILE_APP_URL");
var store = new MobileServiceSQLiteStore
(Constants.OfflineDbPath);
store.DefineTable<ToDoItem>();
await client.SyncContext.IntitalizeAsync(store);
```

var todoTable = client.GetSyncTable<ToDoItem>();
var todoTable = client.GetTable<ToDoItem>();
var todoTable = client.SyncTable;
var todoTable = client.Table;

```
await client.SyncContext.PushAsync();
```

await todoTable.PullAsync("allToDoItems",todoTable.CreateQuery());
await todoTable.UpdateAsync();
todoTable.PullAsync("allToDoItems", todoTable.CreateQuery());
todoTable.UpdateAsync();

Explanation

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

```
var todoTable = client.GetSyncTable<TodoItem>();
var todoTable = client.GetTable<TodoItem>();
var todoTable = client.SyncTable;
var todoTable = client.Table;
```

```
await client.SyncContext.PushAsync();
```

```
await todoTable.PullAsync("allTodoItems",todoTable.CreateQuery());
await todoTable.UpdateAsync();
todoTable.PullAsync("allTodoItems", todoTable.CreateQuery());
todoTable.UpdateAsync();
```

Box 1: `var todoTable = client.GetSyncTable<TodoItem>()`

To setup offline access, when connecting to your mobile service, use the method `GetSyncTable` instead of `GetTable` (example):

```
IMobileServiceSyncTable todoTable = App.MobileService.GetSyncTable(); /
```

Box 2: `await todoTable.PullAsync("allTodoItems",todo.Table.CreateQuery());`

Your app should now use `IMobileServiceSyncTable` (instead of `IMobileServiceTable`) for CRUD operations. This will save changes to the local database and also keep a log of the changes. When the app is ready to synchronize its changes with the Mobile Service, use the methods `PushAsync` and `PullAsync` (example):

```
await App.MobileService.SyncContext.PushAsync();
```

```
await todoTable.PullAsync();
```

References:

<https://azure.microsoft.com/es-es/blog/offline-sync-for-mobile-services/>

Question #:110 - (Exam Topic 3)

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.	☐	☐
If the body size is >256k, an error will occur.	☐	☐
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	☐	☐

Answer:

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	☒	☐
If the body size is >256k, an error will occur.	☐	☒
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	☐	☒

Explanation

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	☒	☐
If the body size is >256k, an error will occur.	☐	☒
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	☐	☒

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>

Question #:111 - (Exam Topic 3)

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**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.

Question #:112 - (Exam Topic 3)

You are developing a .NET Core MVC application for customers to research hotels. The application will use Azure Search. The application will search the index by using various criteria to locate documents related to hotels. The index will include search fields for rate, a list of amenities, and distance to the nearest airport.

The application must support the following scenarios for specifying search criteria and organizing results:



- Search the index by using regular expressions.
- Organize results by counts for name-value pairs.
- List hotels within a specified distance to an airport and that fall within a specific price range.

You need to configure the `SearchParameters` class.

Which properties should you configure? To answer, select the appropriate options in the answer area.

NOTE Each correct selection is worth one point.

Scenario	Property
Search the index by using regular expressions.	☒ QueryType ☐ OrderBy ☐ SearchMode
Organize results by counts for name-value pairs.	☒ Facets ☐ Filter ☐ SearchMode
List hotels within a specified distance to an airport and that fall within a specific price range.	☐ Order by ☒ Top ☐ Filter

Answer:**Explanation****Box 1: QueryType**

The SearchParameters.QueryType Property gets or sets a value that specifies the syntax of the search query. The default is 'simple'. Use 'full' if your query uses the Lucene query syntax.

You can write queries against Azure Search based on the rich Lucene Query Parser syntax for specialized query forms: wildcard, fuzzy search, proximity search, regular expressions are a few examples.

Box 2: Facets

The facets property gets or sets the list of facet expressions to apply to the search query. Each facet expression contains a field name, optionally followed by a comma-separated list of name:value pairs.

Box 3: Filter

The Filter property gets or sets the OData \$filter expression to apply to the search query.

References:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters>

<https://docs.microsoft.com/en-us/azure/search/query-lucene-syntax>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.querytype>

Question #:113 - (Exam Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.


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

New-AzWebAppSlot	-Name myResourceGroup -Location \$location
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

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

New-AzWebAppSlot	-Name \$webappname -Location \$location -AppServicePlan \$webappname -ResourceGroupName myResourceGroup
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

New-AzWebAppSlot	-Name \$webappname -ResourceGroupName myResourceGroup -Slot review
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

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

Answer:


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

New-AzWebAppSlot -Name myResourceGroup -Location $location
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

New-AzWebAppSlot -Name $webappname -Location $location -ResourceGroupName myResourceGroup -Tier Standard
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

New-AzWebAppSlot -Name $webappname -Location $location -AppServicePlan $webappname -ResourceGroupName myResourceGroup
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

New-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup -Slot review
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

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

Explanation

```

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

New-AzWebAppSlot -Name myResourceGroup -Location $location
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

New-AzWebAppSlot -Name $webappname -Location $location -ResourceGroupName myResourceGroup -Tier Standard
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

New-AzWebAppSlot -Name $webappname -Location $location -AppServicePlan $webappname -ResourceGroupName myResourceGroup
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

New-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup -Slot review
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

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

```

The New-AzResourceGroup cmdlet creates an Azure resource group.

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

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

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

References:

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

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

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

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

Question #:114 - [\(Exam Topic 3\)](#)

You develop Azure solutions.

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

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

Which code segment should you use?

- A. `new Container(EndpointUri, PrimaryKey);`
- B. `new Database(Endpoint, PrimaryKey);`
- C. `new CosmosClient(EndpointUri, PrimaryKey);`

Answer: C

Explanation

Example:

```
// Create a new instance of the Cosmos Client
```

```
this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)
```

```
//ADD THIS PART TO YOUR CODE
```

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

Reference:

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

Question #:115 - (Exam Topic 3)

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

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

You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level. You need to configure authorization.

Solution: Configure the Azure Web App for the website to allow only authenticated requests and require Azure AD log on.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation

Instead in the Azure AD application's manifest, set value of the groupMembershipClaims option to All.

References:

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

Question #:116 - (Exam Topic 3)

You are creating an app that uses Event Grid to connect with other services. Your app's event data will be sent to a serverless function that checks compliance. This function is maintained by your company.

You write a new event subscription at the scope of your resource. The event must be invalidated after 3 specific period of time. You need to configure Event Grid to ensure security.

What should you implement? To answer, select the appropriate options in [he answer area.

NOTE: Each correct selection is worth one point

Authentication	Type
WebHook event delivery	SAS tokens Key authentication JWT token
Topic publishing	ValidationCode handshake ValidationURL handshake Management Access Control

Answer:

Authentication	Type
WebHook event delivery	SAS tokens Key authentication JWT token
Topic publishing	ValidationCode handshake ValidationURL handshake Management Access Control

Explanation

Authentication	Type
WebHook event delivery	SAS tokens Key authentication JWT token
Topic publishing	ValidationCode handshake ValidationURL handshake Management Access Control

Box 1: SAS tokens

Custom topics use either Shared Access Signature (SAS) or key authentication. Microsoft recommends SAS, but key authentication provides simple programming, and is compatible with many existing webhook publishers.

In this case we need the expiration time provided by SAS tokens.

Box 2: ValidationCode handshake

Event Grid supports two ways of validating the subscription: ValidationCode handshake (programmatic) and ValidationURL handshake (manual).

If you control the source code for your endpoint, this method is recommended.

Question #:117 - ([Exam Topic 3](#))

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.	☐	☐
Edge nodes must be purged of all cache assets.	☐	☐
The compression type is supported.	☐	☐

Answer:

Statement**Yes****No**

The file MIME type is supported by the service.

☐☒

Edge nodes must be purged of all cache assets.

☒☐

The compression type is supported.

☒☐**Explanation****Statement****Yes****No**

The file MIME type is supported by the service.

☐☒

Edge nodes must be purged of all cache assets.

☒☐

The compression type is supported.

☒☐

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>

Question #118 - (Exam Topic 3)

You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (AAD) for authentication. You need to implement multifactor authentication for the website.

What should you do? Each correct answer presents part of the solution.

NOTE; Each correct selection is worth one point.

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

Answer: A E

Explanation

References:

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

Question #119 - (Exam Topic 3)

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

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

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

You create the index in Azure Search.

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

Solution:

1. Create a SearchIndexClient object to connect to the search index.

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Question #:120 - ([Exam Topic 3](#))

Fourth Coffee has an ASP.NET Core web app that runs in Docker. The app is mapped to the www.fourthcoffee.com domain.

Fourth Coffee is migrating this application to Azure.

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

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

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

Answer area

Azure CLI commands

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

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

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

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

Answer:

Azure CLI commands

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

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

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

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

Answer area

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

⏮

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

⏭

⏮

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

⏭

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

Explanation

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

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

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

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

Step 1: #bin/bash

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

Step 2: az webapp config hostname add

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

Step 3: az webapp create

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

Step : az webapp config container set

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

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

References:

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

Question #121 - (Exam Topic 3)

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

The app stores sensitive information about users.

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

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

NOTE: Each correct selection is worth one point.

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

Answer: A B E

Explanation

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

Dynamic Data Masking
demo_database

Save Discard Add Mask

Downlevel clients require the use of Security Enabled Connection Strings.

Masking Rules

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

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

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

Recommended fields to mask

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

References:

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

Question #:122 - (Exam Topic 3)

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

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

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

Command segments

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

Answer Area

Diagram showing the command segments and the answer area. The answer area is empty, with arrows indicating the correct order of commands: az group create, az aks create, kubectl apply, and az aks get-credentials.

Answer:

Command segments

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

Answer Area

Diagram showing the command segments and the answer area. The answer area contains the following commands in order: az group create, az aks create, kubectl apply, and az aks get-credentials. The answer area is highlighted with a red dashed border.

Explanation

```
az group create
```

```
az aks create
```

```
kubectl apply
```

```
az aks get-credentials
```

Step 1: az group create

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

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

```
az group create --name myAKSCluster --location eastus
```

Step 2 : az aks create

Use the az aks create command to create an AKS cluster.

Step 3: kubectl apply

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

Step 4: az aks get-credentials

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

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

References:

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

Question #:123 - (Exam Topic 3)

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

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

You 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: A

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-vn>

Question #:124 - (Exam Topic 3)

You develop a web application.

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

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

Actions

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

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

Add a Cryptographic key.

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

Select the Azure AD instance.

Use an access token to access the secure resource.

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

Answer Area



Answer:

Actions

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

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

Add a Cryptographic key.

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

Select the Azure AD instance.

Use an access token to access the secure resource.

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

Answer Area

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

Select the Azure AD instance.

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

Explanation

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

Select the Azure AD instance.

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

Register a new application using the Azure portal

- 🕒 Sign in to the Azure portal using either a work or school account or a personal Microsoft account.

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

Question #:125 - ([Exam Topic 3](#))

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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

Explanation

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

Accessing the client certificate from App Service.

If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be

available through the `HttpRequest.ClientCertificate` property. For other application stacks, the client cert will be available in your app through a base64 encoded value in the "X-ARR-ClientCert" request header. Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.

References:

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




About dumpscollection.com

dumpscollection.com was founded in 2007. We provide latest & high quality IT / Business Certification Training Exam Questions, Study Guides, Practice Tests.

We help you pass any IT / Business Certification Exams with 100% Pass Guaranteed or Full Refund. Especially Cisco, CompTIA, Citrix, EMC, HP, Oracle, VMware, Juniper, Check Point, LPI, Nortel, EXIN and so on.

View list of all certification exams: [All vendors](#)



We prepare state-of-the art practice tests for certification exams. You can reach us at any of the email addresses listed below.

- 🔗 Sales: sales@dumpscollection.com
- 🔗 Feedback: feedback@dumpscollection.com
- 🔗 Support: support@dumpscollection.com
- 🔗 Skype ID: crack4sure@gmail.com

Any problems about IT certification or our products, You can write us back and we will get back to you within 24 hours.

15% Discount Coupon Code:

DC15disc