



"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks.com

Warning: Keep connected with our support team
for latest updates

Question: 1

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Option

Value

Tier

	▼
Standard	
Premium	

Profile

	▼
Akamai	
Microsoft	

Optimization

	▼
general web delivery	
large file download	
dynamic site acceleration	
video-on-demand media streaming	

Answer:

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: 2

HOTSPOT

You need to secure the Shipping Function app.

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

NOTE: Each correct selection is worth one point.

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:

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.

Reference:

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

Question: 3

You need to secure the Shipping Logic App.

What should you use?

- 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.

Reference:

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

Question: 4

DRAG DROP

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

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

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:

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.

Reference:

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

Question: 5

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: 6

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.

Reference:

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

Question: 7

HOTSPOT

You need to resolve the Shipping web site error.

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

NOTE: Each correct selection is worth one point.

```
<?xml version="1.0" encoding="utf-8"?>
<StorageServiceProperties>
  ...
  <Cors>
    <CorsRule>
      <AllowedHeaders>
        http://*.wideworldimporters.com
        http://test.wideworldimporters.com
        http://test-shippingapi.wideworldimporters.com
        http://www.wideworldimporters.com
      </AllowedHeaders>
      <ExposedHeaders>
      </ExposedHeaders>
      <AllowedMethods>
        GET,PUT
        GET
        POST
        GET,HEAD
      </AllowedMethods>
      <AllowedOrigins>
      </AllowedOrigins>
    </CorsRule>
  </Cors>
</StorageServiceProperties>
```

Answer:

Explanation:

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

```
<Cors>
```

```
<CorsRule>
```

<AllowedOrigins>	http://*.wideworldimporters.com	</AllowedOrigins>
AllowedHeaders	http://test.wideworldimporters.com	AllowedHeaders
ExposedHeaders	http://test-shippingapi.wideworldimporters.com	ExposedHeaders
AllowedMethods	http://www.wideworldimporters.com	AllowedMethods
AllowedOrigins		AllowedOrigins

```
<AllowedMethods>
```

```
</AllowedMethods>
```

GET,PUT
GET
POST
GET,HEAD

```
</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.wideworldimporter3.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

Reference:

<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Access-Control-Allow-Origin>

Question: 8

HOTSPOT

You need to correct the VM issues.

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

NOTE: Each correct selection is worth one point.

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:

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.

Reference:

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

Question: 9

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

az webapp -g shipping-apis-test-rg -n web

--

Answer:

Explanation:

az webapp

cors
config
deployment

add
up
remove

-g shipping-apis-test-rg -n web

slot
allowed-origins
name

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

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/

Reference:

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

Question: 10

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Setting

Value

Plan

▼

Basic
Standard
Premium
Isolated

Instance Count

▼

1
10
20
100

Answer:

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.

Reference:

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

Topic 2, Contoso, Ltd

Case study

Overview

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

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

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

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background

Overview

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

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

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

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

The solution will use eight CPU cores.

Azure Active Directory

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

Requirements

ContentAnalysisService

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

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

Costs

You must minimize costs for all Azure services.

Manual review

To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using React

and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.

High availability

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

Monitoring

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

Security

You have the following security requirements:

Any web service accessible over the Internet must be protected from cross site scripting attacks.

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

Azure Storage access keys must only be stored in memory and must be available only to the service.

All Internal services must only be accessible from internal Virtual Networks (VNETs).

All parts of the system must support inbound and outbound traffic restrictions.

All service calls must be authenticated by using Azure AD.

User agreements

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

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

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

Validation testing

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

significantly deviate from the old version.

Issues

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

Code

ContentUploadService

```
CS01 .tiv^;-n: '2018-10-01'
CS02 ■ : Microsoft.ContainerInstance/containerGroups
CS03 .oca* LoF.: westus
CS04 rams-: contentuploadservice
CS05 properties:
CS06 containers:
CS07 name: service
CS08     properties:
CS09     mage: contoso/contentUploadService:latest
CS10     ports:
CS11     - port: 80
CS12     protocol: TCP
CS13     resources:
CS14     requests:
CS15     cpu: 1.0
CS16     memoryInGB: 1.5
CS17
CS18 ipAddress:
CS19 ip: 10.23.121.112
CS20     ports:
CS21     - port: 80
CS22     protocol: TCP
CS23
CS24
CS25 r.-i- :<: file:
CS26
rd: /subscriptions/96...19/resourceGroups/container/providers/Microsoft.Network/networkprofiles/subnet
AM01 {
AM02   "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
AM03   "appld" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",
AMO 4
AMOS   "createdDateTime" : "2019-12-24T06:01:44Z",
AM06   "logoUrl" : null,
AM07   "logoutUrl" : null,
AM08   "name" : "ContentAnalysisService",
AMO 9
AM10
AM11   "orgRestrictions" : [],
AM12   "parentalControlSettings" : {
AM13   "countriesBlockedForMinors" : [],
AMI4   "legalAgeGroupRule" : "Allow"
AMIS   },
AMI6   "passwordcredentials" : []
AM17 }
```

Question: 11

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,C

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: 12

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: 13

HOTSPOT

You need to implement the bindings for the CheckUserContent function.

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

NOTE: Each correct selection is worth one point.

```
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:

Explanation:

```
public static class CheckUserContent
{
    [FunctionName ("CheckUserContent")]
    public static void Run (
```



```

[QueueTrigger(userContent)]
[BlobTrigger("userContent/{name}")]
[CosmosDBTrigger("content", "userContent")]
[TableTrigger("userContent", "{name}")]

```

▼ string content,

▼ Stream output]

```

QueueTrigger("userContent")
[CosmosDBTrigger("content", "userContent")]
TableTrigger("userContent", "{name}")]
BlobTrigger("userContent/{name}", FileAccess.Write)
}

```

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: 14

DRAG DROP

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

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

NOTE: Each correct selection is worth one point.

Json segments

User

value

role

Application

allowedMemberTypes

allowedAccountTypes

Answer Area

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

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: 15

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

▼

"allowPublicClient":true

"oauth2Permissions":["login"]

"oauth2AllowUrlPathMatching":true

"oauth2AllowIdTokenImplicitFlow":true

▼

"oauth2AllowImplicitFlow": true

"oauth2RequiredPostResponse":true

"preAuthorizedApplications":["SPA"]

"knownClientApplications":["ContentAnalysisService"]

Answer:

Explanation:

```
“allowPublicClient” : true,
“oauth2Permissions” : [“login”],
“oauth2AllowUrlPathMatching” : true,
“oauth2AllowIdTokenImplicitFlow” : true,
“oauth2AllowImplicitFlow” : true,
“oauth2RequiredPostResponse” : true,
“preAuthorizedApplications” : [
  {
    “SPA” : true,
    “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: 17

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 "avgPercentage CPU >8"`
- B. `az monitor metrics alert create --n alert --g ... --scopes... --condition "avgPercentage CPU >800"`
- C. `az monitor metrics alert create --n alert --g ... --scopes... --condition "CPUUsage > 800"`
- D. `az monitor metrics alert create --n alert --g ... --scopes... --condition "CPUUsage > 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: 18

HOTSPOT

You need to ensure that validation testing is triggered per the requirements.

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

NOTE: Each correct selection is worth one point.

```
var event = getEvent();
if (event.eventType === 'ImagePushed' || event.eventType === 'RepositoryItem ImageDeployed' || event.eventType === 'RepositoryUpdated' || event.eventType === 'contentanalysiservice' && event.data.target === 'aci' && event.data.target === 'image service' && event.data.target === 'repository' && event.data.target === 'topic' && event.data.target === 'service' && event.data.target === 'repository' && event.data.target === 'imageCollection')
{
    startValidationTesting();
}
```

Answer:

Explanation:

```
var event = getEvent();
if (event.eventType === 'ImagePushed' || event.eventType === 'RepositoryItem Image Deployed RepositoryUpdated' || event.eventType === 'contentanalysiservice' && event.data.target === 'aci' && event.data.target === 'image service' && event.data.target === 'repository' && event.data.target === 'topic' && event.data.target === 'service' && event.data.target === 'repository' && event.data.target === 'imageCollection')
{
    startValidationTesting();
}
```


startValidationTestingO;

Box 1: RepositoryUpdated

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.

Box 2: service

Box 3: imageCollection

Reference:

<https://docs.microsoft.com/en-us/azure/devops/notifications/oob-supported-event-types>

Question: 19

DRAG DROP

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

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

NOTE: Each correct selection is worth one point.

YAML segments

secret

envVar

secretValues

volumes

volumeMounts

environmentVariables

Answer Area

YAML segment :

- mountPath: /mnt/secrets
name: accesskey

YAML segment :

- name: accesskey

YAML segment :

key: TXkgZmlyc3Qgc2VjcmV0IEZPTwo=

Answer:

Explanation:

volumeMounts :

- mountPath: /mnt/secrets name: accesskey

volumes

name: accesskey

secret :

key: TXkgZmlyc3Qgc2VjcmV0IEZPTwo=

Box 1: volumeMounts

Example:

volumeMounts:

- mountPath: /mnt/secrets

name: secretvolume1

volumes:

- name: secretvolume1

secret:

mysecret1: TXkgZmlyc3Qgc2VjcmV0IEZPTwo=

Box 2: volumes

Box 3: secret

Reference:

<https://docs.microsoft.com/en-us/azure/container-instances/container-instances-volume-secret>

Question: 20

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

```
"optionalClaims": [  
  "

|             |
|-------------|
| acct        |
| platf       |
| sid         |
| tenant_ctry |

",  
  "

|           |
|-----------|
| sid       |
| upn       |
| email     |
| enfpolids |

"  
],
```

Answer:

Explanation:

"optionalClaims": [

"

acct
platf
sid
tenant_ctype

"

"

sid
upn
email
enfpolids

"

],

Box 1: sid

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

Scenario: Manual review

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

Box 2: email

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

Question: 21

You need to investigate the http server log output to resolve the issue with the ContentUploadService.

Which command should you use first?

- A. az webapp log
- B. az ams live-output
- C. az monitor activity-log
- D. az container attach

Answer: C

Explanation:

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

"502 bad gateway" and "503 service unavailable" are common errors in your app hosted in Azure App Service.

Microsoft Azure publicizes each time there is a service interruption or performance degradation.

The az monitor activity-log command manages activity logs.

Note: Troubleshooting can be divided into three distinct tasks, in sequential order:

Observe and monitor application behavior

Collect data

Mitigate the issue

Reference:

<https://docs.microsoft.com/en-us/cli/azure/monitor/activity-log>

Question: 22

You need to deploy the CheckUserContent Azure function. The solution must meet the security and cost requirements.

Which hosting model should you use?

- A. Consumption plan
- B. Premium plan
- C. App Service plan

Answer: C

Explanation:

Topic 3, City Power & Light Case study

Overview

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

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

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

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background

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

Current environment

Architecture overview

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

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

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

Architecture diagram

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

User authentication

The following steps detail the user authentication process:

The user selects Sign in in the website.

The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

The user signs in.

Azure AD redirects the user's session back to the web application. The URL includes an access token.

The web application calls an API and includes the access token in the authentication header.

The application ID is sent as the audience ('aud') claim in the access token.

The back-end API validates the access token.

Requirements

Corporate website

Communications and content must be secured by using SSL.

Communications must use HTTPS.

Data must be replicated to a secondary region and three availability zones.

Data storage costs must be minimized.

Azure Database for PostgreSQL

www.atmicnetworks.com

www.atmicnetworks.com

www.atmicnetworks.com

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

Azure Key Vault name: cpandlkeyvault

Secret name: PostgreSQLConn

Id: 80df3e46ffcd4f1cb187f79905e9a1e8

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

Azure Service Bus and Azure Event Grid

Azure Event Grid must use Azure Service Bus for queue-based load leveling.

Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.

Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Security

All SSL certificates and credentials must be stored in Azure Key Vault.

File access must restrict access by IP, protocol, and Azure AD rights.

All user accounts and processes must receive only those privileges which are essential to perform their intended function.

Compliance

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

Issues

Corporate website

While testing the site, the following error message displays:

CryptographicException: The system cannot find the file specified.

Function app

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

displays:

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

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

FunctionAppLogs

| where FunctionName == "RequestUserApproval"

Logic app

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

'400 Bad Request'

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

Code

Corporate website

Security.cs:

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

Function app

RequestUserApproval.cs:

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

```
RA13 private static bool ProcessRequest(HttpRequest req)
RAI 4 {
RA15 ...
RA16 }
RA17 }
```

Question: 23

QUESTION 1

You need to correct the RequestUserApproval Function app error.

What should you do?

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

Answer: C

Explanation:

Async operation tracking

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

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

Function app

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

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

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

FunctionAppLogs

| where FunctionName == "RequestUserApproval"

Reference:

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

Question: 24

HOTSPOT

You need to configure the Account Kind, Replication, and Storage tier options for the corporate website's Azure Storage account.

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

NOTE: Each correct selection is worth one point.

Create storage account

Basics Advanced Tags Review + create

Azure Storage is a Microsoft-managed service providing cloud storage that is highly available, secure, durable, scalable, and redundant. Azure Storage includes Azure Blobs (objects), Azure Data Lake Storage Gen2, Azure Files, Azure Queues, and Azure Tables. The cost of your storage account depends on the usage and the options you choose below. [Learn more](#)

PROJECT DETAILS

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

* Subscription

Visual Studio Enterprise

* Resource group

(New) cplcorporatesite

[Create new](#)

INSTANCE DETAILS

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead. [Choose classic deployment model](#)

* Storage account name ⓘ

corporatewebsitecontent

* Location

(US) East US

Performance ⓘ

☒ Standard ☐ Premium

Account kind ⓘ

StorageV2 (general purpose v2)
Storage (general purpose v1)
BlobStorage

Replication ⓘ

Locally-redundant storage (LRS)
Zone-redundant storage (ZRS)
Geo-redundant storage (GRS)
Read-access geo-redundant storage (RA-GRS)
Geo-zone-redundant storage (GZRS)
Read-access geo-zone-redundant storage (RA-GZRS)

Access tier (default) ⓘ

☐ Cool ☐ Hot

Answer:

Explanation:

INSTANCE DETAILS

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead. Choose classic deployment model

* Storage account name 0

corporatewebsitecontent

y

* Location

(US) East US

v

Performance 0

• Standard

Premium

Account kind 0

StorageV2 (general purpose v2)

Storage (general purpose v1)

BlobStorage

Replication 0

v

Locally-redundant storage (LRS)

Zone-redundant storage (ZRS)

Geo-redundant storage (GRS)

Read-access geo-redundant storage (RA-GRS)

Geo-zone-redundant storage (GZRS)

Read-access geo-zone-redundant storage (RA-GZRS)

Access tier (default) 0

Cool

Hot

Account Kind: StorageV2 (general-purpose v2)

Scenario: Azure Storage blob will be used (refer to the exhibit). Data storage costs must be minimized.

General-purpose v2 accounts: Basic storage account type for blobs, files, queues, and tables.

Recommended for most scenarios using Azure Storage.

Reference:

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

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

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

Question: 25

HOTSPOT

You need to retrieve the database connection string.

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

NOTE: Each correct selection is worth one point.

REST API Endpoint:

https://	▼	.vault.azure.net/secrets/	▼	/
cpandlkeyvault PostgreSQLConn 80df3e46ffcd4f1cb187f79905e9a1e8		cpandlkeyvault PostgreSQLConn 80df3e46ffcd4f1cb187f79905e9a1e8		

Variable type to access Azure Key Vault secret values:

▼
Environment
Session
ViewState
QueryString

Answer:

Explanation:

REST API Endpoint:

https://	▼	.vault.azure.net/secrets/	▼	/
cpandlkeyvault PostgreSQLConn 80df3e46ffcd4f1cb187f79905e9a1e8		cpandlkeyvault PostgreSQLConn 80df3e46ffcd4f1cb187f79905e9a1e8		

Variable type to access Azure Key Vault secret values:

▼
Environment
Session
ViewState
QueryString

Azure database connection string retrieve REST API vault.azure.net/secrets/

Box 1: cpandlkeyvault

We specify the key vault, cpandlkeyvault.

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

Azure Key Vault name: cpandlkeyvault

Secret name: PostgreSQLConn

Id: 80df3e46ffcd4f1cb187f79905e9a1e8

Box 2: PostgreSQLConn

We specify the secret, PostgreSQLConn

Example, sample request:

<https://myvault.vault.azure.net//secrets/mysecretname/4387e9f3d6e14c459867679a90fd0f79?api-version=7.1>

Box 3: Querystring

Reference:

<https://docs.microsoft.com/en-us/rest/api/keyvault/getsecret/getsecret>

Question: 26

DRAG DROP

You need to correct the corporate website error.

Which four actions should you recommend be performed 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

Upload the certificate to Azure Key Vault.

Update line SC05 of Security.cs to include error handling and then redeploy the code.

Update line SC03 of Security.cs to include a using statement and then re-deploy the code.


Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting.

Upload the certificate to source control.

Import the certificate to Azure App Service.

Generate a certificate.

Answer Area



Answer:

Explanation:

Generate a certificate.

Upload the certificate to Azure Key Vault.

Import the certificate to Azure App Service.

Update line SC05 of Security.cs to include error handling and then redeploy the code.

Scenario: Corporate website

While testing the site, the following error message displays:

CryptographicException: The system cannot find the file specified.

Step 1: Generate a certificate

Step 2: Upload the certificate to Azure Key Vault

Scenario: All SSL certificates and credentials must be stored in Azure Key Vault.

Step 3: Import the certificate to Azure App Service

Step 4: Update line SC05 of Security.cs to include error handling and then redeploy the code

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/configure-ssl-certificate>

Question: 27
HOTSPOT

You need to configure API Management for authentication.

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

NOTE: Each correct selection is worth one point.

Setting

Value

Policy

	▼
Check HTTP header	
Restrict caller IPs	
Limit call rate by key	
Validate JWT	

Policy section

	▼
Inbound	
Outbound	

Answer:

Explanation:

Setting

Value

Policy

	▼
Check HTTP header	
Restrict caller IPs	
Limit call rate by key	
Validate JWT	

Policy section

	▼
Inbound	
Outbound	

Box 1: Validate JWT

The validate-jwt policy enforces existence and validity of a JWT extracted from either a specified HTTP Header or a specified query parameter.

Scenario: User authentication (see step 5 below)

The following steps detail the user authentication process:

The user selects Sign in in the website.

The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

The user signs in.

Azure AD redirects the user's session back to the web application. The URL includes an access token.

The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.

The back-end API validates the access token.

Box 2: Outbound

Reference:

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

Question: 28

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

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

NOTE: Each correct selection is worth one point.

- A. ID token signature
- B. ID token claims
- C. HTTP response code
- D. Azure AD endpoint URI
- E. Azure AD tenant ID

Answer: A,D

Explanation:

Claims in access tokens

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

Header - Provides information about how to validate the token including information about the type of token and how it was signed.

Payload - Contains all of the important data about the user or app that is attempting to call your service.

Signature - Is the raw material used to validate the token.

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

Scenario: User authentication (see step 5 below)

The following steps detail the user authentication process:

The user selects Sign in in the website.

The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.

The user signs in.

Azure AD redirects the user's session back to the web application. The URL includes an access token.

The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.

The back-end API validates the access token.

Reference:

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

Question: 29

You need to investigate the Azure Function app error message in the development environment.

What should you do?

- A. Connect Live Metrics Stream from Application Insights to the Azure Function app and filter the metrics.
- B. Create a new Azure Log Analytics workspace and instrument the Azure Function app with Application Insights.
- C. Update the Azure Function app with extension methods from Microsoft.Extensions.Logging to log events by using the log instance.
- D. Add a new diagnostic setting to the Azure Function app to send logs to Log Analytics.

Answer: A

Explanation:

Azure Functions offers built-in integration with Azure Application Insights to monitor functions.

The following areas of Application Insights can be helpful when evaluating the behavior, performance, and errors in your functions:

Live Metrics: View metrics data as it's created in near real-time.

Failures

Performance

Metrics

Reference:

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

Question: 30

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

az

	▼
eventgrid	
servicebus	

	▼
event-subscription	
topic	
queue	

 create --source-resource-id \$topicid --name \$name --

endpoint-type

	▼
webhook	
eventhub	
servicebusqueue	

 --endpoint \$endpoint

Answer:

Explanation:

az

	▼
eventgrid	
servicebus	

	▼
event-subscription	
topic	
queue	

 create --source-resource-id \$topicid --name \$name --

endpoint-type

	▼
webhook	
eventhub	
servicebusqueue	

 --endpoint \$endpoint

Box 1: eventgrid

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

Box 2: event-subscription

Box 3: servicebusqueue

Scenario: Azure Service Bus and Azure Event Grid

Azure Event Grid must use Azure Service Bus for queue-based load leveling.

Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.

Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.

Reference:

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

Question: 31

HOTSPOT

You need to correct the Azure Logic app error message.

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

NOTE: Each correct selection is worth one point.

Setting

Value

authentication level

anonymous
function
admin

managed identity

system-assigned
user-assigned

Answer:

Explanation:

Setting

Value

authentication level

anonymous
function
admin

managed identity

system-assigned
user-assigned

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

displays:

'400 Bad Request'

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

Note: If the inbound call's request body doesn't match your schema, the trigger returns an HTTP 400 Bad Request error.

Box 1: function

If you have an Azure function where you want to use the system-assigned identity, first enable authentication for Azure functions.

Box 2: system-assigned

Your logic app or individual connections can use either the system-assigned identity or a single user- assigned identity, which you can share across a group of logic apps, but not both.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/create-managed-service-identity>

Question: 32

HOTSPOT

You need to configure Azure Service Bus to Event Grid integration.

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

NOTE: Each correct selection is worth one point.

Setting

Value

Tier

Basic

Standard

Premium

RBAC role

Owner

Contributor

Azure Service Bus Data Owner

Azure Service Bus Data Receiver

Answer:

Explanation:

Setting

Value

Tier

Basic

Standard

Premium

RBAC role

Owner

Contributor

Azure Service Bus Data Owner

Azure Service Bus Data Receiver

Box 1: Premium

Service Bus can now emit events to Event Grid when there are messages in a queue or a subscription when no receivers are present. You can create Event Grid subscriptions to your Service Bus namespaces, listen to these events, and then react to the events by starting a receiver. With this feature, you can use Service Bus in reactive programming models.

To enable the feature, you need the following items:

A Service Bus Premium namespace with at least one Service Bus queue or a Service Bus topic with at least one subscription.

Contributor access to the Service Bus namespace.

Box 2: Contributor

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-to-event-grid-integration-concept>

Question: 33

HOTSPOT

You need to configure security and compliance for the corporate website files.

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

NOTE: Each correct selection is worth one point.

Action

Setting

Restrict file access

role-based access control (RBAC)

managed identity

shared access signature (SAS) token

connection string

Enable file auditing

access tier

change feed

blob indexer

storage account type

Answer:

Explanation:

Answer:

Action

Setting

Restrict file access

role-based access control (RBAC)

managed identity

shared access signature (SAS) token

connection string

Enable file auditing

V

access tier

change feed

blob indexer

storage account type

Box 1: role-based access control (RBAC)

Azure Storage supports authentication and authorization with Azure AD for the Blob and Queue services via Azure role-based access control (Azure RBAC).

Scenario: File access must restrict access by IP, protocol, and Azure AD rights.

Box 2: change feed

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 file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.

Reference:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-sas-storage-support>

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

Question: 34

You need to ensure that all messages from Azure Event Grid are processed.

What should you use?

- A. Azure Event Grid topic
- B. Azure Service Bus topic
- C. Azure Service Bus queue
- D. Azure Storage queue
- E. Azure Logic App custom connector

Answer: D

Explanation:

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

Your solution needs to receive messages without having to poll the queue. With Service Bus, you can achieve it by using a long-polling receive operation using the TCP-based protocols that Service Bus supports.

Reference:

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

Topic 4, Proseware, Inc

Case study

Overview

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

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

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

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background

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

Requirements

Policy service

You develop and deploy a stateful ASP.NET Core 2.1 web application named Policy service to

an Azure App Service Web App. The application reacts to events from Azure Event Grid and performs policy actions based on those events.

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

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

Policies

Log policy

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

Authentication events

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

PolicyLib

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

- Exclude non-user actions from Application Insights telemetry.

- Provide methods that allow a web service to scale itself.

- Ensure that scaling actions do not disrupt application usage.

Other

Anomaly detection service

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

Health monitoring

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

Issues

Policy loss

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

Performance issue

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

Notification latency

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

App code

EventGridController.cs

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

EventGridController.es

```
public class EventGridController : Controller
```

```
{
    public static AsyncLocal<string> EventId ■ new AsyncLocal<string>(); public IActionResult
    Process([FromBody] string events!son) {
        var events = JArray.Parse(events!son);
        foreach (var gevent in events)
        {
            EventId.Value « gevent["id"].ToString();
            if (@event["topic"].ToString().Contains("providers/Microsoft.Storage")) {
                SendToAnomalyDetectionService(@event["data"]["url"].ToString());
            }
        }
        EnsureLogging(^event["subjeef"].ToString());
    }
    return null;
}

private void EnsureLogging(string resource)
{
}

private async Task SendToAnomalyDetectionService(string uri)
(
    var content - GetLogData(uri);
    var scoreRequest ■ new
    {
        Inputs • new Dictionary<string, List<Dictionary<string, string>>>() {
            "input 1",
            new List<Dictionary<string, string>>()
            {
                new Dictionary<string, string>()
                {
                    { "logcontent", content
                }
            }
        }
    }
}
```

LoginEvent.cs

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

LoginEvent.es

```
LE01 public class LoginEvent
```

```

LE02 {
LE03
LE04 public string subject { get; set; }
LE05 public Datetime eventlme { get; set; }
LE06 public Dictionary<string,j string> data { get; set; }
LE07 public string Serialize()
LE08. {
LE09 return IsonConvert.SerializeObject(this);
LE10 }
LE11 }

```

Question: 35

You need to resolve a notification latency issue.

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

NOTE: Each correct selection is worth one point.

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

Answer: A,B

Explanation:

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

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

Anomaly detection service: You have an anomaly detection service that analyzes log information for anomalies. It is implemented as an Azure Machine Learning model. The model is deployed as a web service. If an anomaly is detected, an Azure Function that emails administrators is called by using an HTTP WebHook.

Reference:

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

Question: 36

DRAG DROP

You need to implement the Log policy.

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

NOTE: Each correct selection is worth one point.

Code segment

All
WebHook
EventHub
subjectEndsWith
Microsoft.Storage
subjectBeginsWith
Microsoft.Storage.BlobCreated

Answer Area

```
{
  "name": "newlogs",
  "properties": {
    "topic": "/subscriptions/. . ./providers/Microsoft.EventGrid/topics/. . .",
    "destination": {
      "endpointType": "code segment",
      "filter": {
        "code segment": "code segment",
        "topic": "/blobServices/default/containers/logdrop/"
      },
      "includedEventTypes": [ "code segment" ] },
    "labels": [],
    "eventDeliverySchema": "EventGridSchema"
  }
}
```

Answer:

Explanation:

```
"name": "newlogs", "properties": {
  "topic": "/subscriptions/. . ./providers/Microsoft.EventGrid/topics/. . .",
```

```

"destination": {
  "endpointType" : " WebHook
    },
  "filter": {
    " SUBjeCtBeginsWith
      "/blobServices/default/containers/logdrop/",
    "inciudedEvenrType?": [ " Microsoft.Storage.BlobCreated "
      1 L
    "labels": [], "eventDeliverySchema": "EventGridSchema"
  }
}

```

Box 1:WebHook

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

endpointType: The type of endpoint for the subscription (webhook/HTTP, Event Hub, or queue).

Box 2: SubjectBeginsWith

Box 3: Microsoft.Storage.BlobCreated

Scenario: Log Policy

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

Example subscription schema

```

{
  "properties": {
    "destination": {
      "endpointType": "webhook",
      "properties": {
        "endpointUrl":
          "https://example.azurewebsites.net/api/HttpTriggerCSharp1?code=VXbGWce53l48Mt8wuotr0GPmy
            J/nDT4hgdFj9DpBiRt38qqnm5OFg=="
      }
    }
  }
}

```

```
}  
},  
"filter": {  
  "includedEventTypes": [ "Microsoft.Storage.BlobCreated", "Microsoft.Storage.BlobDeleted" ],  
  "subjectBeginsWith": "blobServices/default/containers/mycontainer/log",  
  "subjectEndsWith": ".jpg",  
  "isSubjectCaseSensitive ": "true"  
}  
}  
}
```

Reference:

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

Question: 37

You need to ensure that the solution can meet the scaling requirements for Policy Service.

Which Azure Application Insights data model should you use?

- A. an Application Insights dependency
- B. an Application Insights event
- C. an Application Insights trace
- D. an Application Insights metric

Answer: D

Explanation:

Application Insights provides three additional data types for custom telemetry:

Trace - used either directly, or through an adapter to implement diagnostics logging using an instrumentation

framework that is familiar to you, such as Log4Net or System.Diagnostics.

Event - typically used to capture user interaction with your service, to analyze usage patterns.

Metric - used to report periodic scalar measurements.

Scenario:

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

Reference:

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

Question: 38

DRAG DROP

You need to implement telemetry for non-user actions.

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

NOTE: Each correct selection is worth one point.

Code segments

/health

/status

RequestTelemetry

PageViewTelemetry

ITelemetryProcessor

ITelemetryInitializer

Answer Area

```
public class Filter :   
{  
    private readonly  _next;  
    public (Filter  next)  
    {  
        _next = next;  
    }  
    public void Process(ITelemetry item)  
    {  
        var x = item as  ;  
        if (x?.Url.AbsolutePath == ""  
        {  
            return;  
        }  
        _next.Process(item);  
    }  
}
```

Answer:

Explanation:

```
public class Filter : ITelemetryProcessor
{
    private readonly ITelemetryProcessor _next;
    public Filter(ITelemetryProcessor next)
    {
        _next = next;
    }
    public void Process(ITelemetry item)
    {
        var x = item as RequestTelemetry ;
        if (x?.Url.AbsolutePath == "/health" )
        {
            return;
        }
        _next.Process(item);
    }
}
```

Scenario: Exclude non-user actions from Application Insights telemetry.

Box 1: ITelemetryProcessor

To create a filter, implement ITelemetryProcessor. This technique gives you more direct control over what is included or excluded from the telemetry stream.

Box 2: ITelemetryProcessor

Box 3: ITelemetryProcessor

Box 4: RequestTelemetry

Box 5: /health

To filter out an item, just terminate the chain.

Reference:

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

Question: 39

DRAG DROP

You need to ensure that PolicyLib requirements are met.

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

NOTE: Each correct selection is worth one point.

Code segments	Answer Area
Process	public class IncludeEventId : code segment
Initialize	{
telemetry.Sequence	public void code segment (ITelemetry telemetry)
ITelemetryProcessor	{
ITelemetryInitializer	code segment .Properties["EventId"] =
telemetry.Context	code segment ;
EventGridController.EventId.Value	}
((EventTelemetry)telemetry).Properties["EventId"]	}

Answer:

Explanation:

```
public class IncludeEventId : ITelemetryInitializer
```

```
public void Initialize  
{
```

```
    telemetry.Context
```

```
        ((EventTelemetry)telemetry).Properties["EventId"] =  
    }  
}
```

```
(ITelemetry telemetry)
```

```
.Properties["EventId"] =
```

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

Exclude non-user actions from Application Insights telemetry.

Provide methods that allow a web service to scale itself.

Ensure that scaling actions do not disrupt application usage.

Box 1: ITelemetryInitializer

Use telemetry initializers to define global properties that are sent with all telemetry; and to override selected behavior of the standard telemetry modules.

Box 2: Initialize

Box 3: Telemetry.Context

Box 4: ((EventTelemetry)telemetry).Properties["EventID"]

Reference:

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

Question: 40

DRAG DROP

You need to add code at line EG15 in EventGridController.cs to ensure that the Log policy applies to all services.

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

NOTE: Each correct selection is worth one point.

Code segments

topic
status
eventType
Succeeded
operationName
resourceProvider

Answer Area

```
if {  
    @event[ "data" ][ "code segment" ].ToString() == "code segment"  
    &&  
    @event[ "data" ][ "code segment" ].ToString() == "Microsoft.Web/sites/write"  
}
```

Answer:

Explanation:

```
if {  
    @event[ "data" ][ "status" ].ToString() == "Succeeded"  
    &&  
    @event[ "data" ][ "operationName" ].ToString() == "Microsoft.Web/sites/write"  
}
```

Scenario, Log policy: All Azure App Service Web Apps must write logs to Azure Blob storage.

Box 1: Status

Box 2: Succeeded

Box 3: operationName

Microsoft.Web/sites/write is resource provider operation. It creates a new Web App or updates an existing one.

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/resource-provider-operations>

Question: 41

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

```
( get; set; )  
id  
eventType  
dataVersion  
metadaversion
```

eventType dataVersion
metadaversion

eventType dataVersion
metadata version

Answer:

Explanation:

public string

▼ (get; set;)

id

eventType

dataVersion

metadataVersion

public string

▼ (get; set;)

id

eventType

dataVersion

metadataVersion

public string

▼ (get; set;)

id

eventType

dataVersion

metadataVersion

Box 1: id

id is a unique identifier for the event.

Box 2: eventType

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

Box 3: dataVersion

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

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

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

```
[
{
  "topic": string,

  "subject": string,
  "id": string,
  "eventType": string,
  "eventTime": string,

  "data":{
    object-unique-to-each-publisher
  },

  "dataVersion": string,

  "metadataVersion": string
}
```

Reference:

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

Question: 42

HOTSPOT

You need to implement the Log policy.

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

NOTE: Each correct selection is worth one point.

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

Answer:

Explanation:

```
var client = new WebSiteManagementClient(. . .);
var id = ParseResourceID(resource);
var appSettings = new StringDictionary(name: "properties", properties: new
Dictionary<string, string> {
    {"DIAGNOSTICS_AZUREBLOBCONTAINERSASURL", BlobStoreAccountSAS("
    v ")}},
    {"DIAGNOSTICS_AZUREBLOBRETENTIONINDAYS", " 30 "}
});
client.WebApps.
UploadLoggingSettings
UpdateApplicationSetting
(id.resourceGroup, id.name, appSettings)'
```

Box 1: logdrop

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

Box 2: 15

Logs must remain in the container for 15 days.

Box 3: UpdateApplicationSettings

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

Reference:

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

Topic 5, Litware Inc Case study

Overview

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

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

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

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background

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

Overall architecture

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

Receipt processing

Employees may upload receipts in two ways:

Uploading using an Azure Files mounted folder

Uploading using the web application

Data Storage

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

Documentation

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

Solution details

Users table

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

Web Application

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

Processing

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

Logging

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

Requirements

Receipt processing

Concurrent processing of a receipt must be prevented.

Disaster recovery

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

Security

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

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

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

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

Receipt data must always be encrypted at rest.

All data must be protected in transit.

User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured. In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.

Issues

Upload format issue

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

Capacity issue

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

Log capacity issue

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

Application code

Processing.cs

```

PC01 public static class Processing
PC02 (
PC03 public static class Function
PC04 {
PC05     [Functionflame("IssueWork")]
PC06 public static async Task Run([TimerTrigger("0 */5 * * *")] Tiinerinfo timer, ILogger log)
PC07 {
PC08     var container = await GetCloudBlobContainer();
PC09     foreach (var fileitem in await ListFiles())
PC10     {
PC11         var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12         var ms = new HemoryStream();
PC13         await file.DownloadToStreamAsync(ms);
PC14         var blob = container.Get81ockBlobReference(fileItem.Uri.ToString())j
PC15         await blob.UploadFrofflStreamAsync(ms);
PC16
PC17     }
PC18 }
PC19     private static CloudBlockBlob GetDR81ob(CloudBlockBlob sourceBlob)
PC20 {
PC21 ...
PC22     }
PC23     private static async Task<Cloud81obContainer> GetCloudBlobContainer()
PC24     {
PC25     var cloudBlobClient = new CloudBlobClient(new Uri("..."), await GetCredentials())
PC26
PC27     await cloudBlobClient.GetRootContainerReference().CreatelfNotExistsAsync();
PC28     return cloudBlobClient.GetRootContainerReference();
PC29     }
PC30     private static async Task<StorageCredentials> GetCredentials()
PC31     {
PC32 ...
PC33     }
PC34     private static async Task<List<IListFileItem>> ListFiles()
PC35     {
PC36 ...
PC37     }
PC37     private KeyVaultClient _keyVaultClient = new KeyVaultClient("...");
PC38 }
PC39 }

```

Database.cs

DB01 public class Database

DB02 {

DB03 private string Connectionstring =

0604

DB05 public async 'a *<object> LoadUserDetails(string userid)

OB06 {

DB07

0608 return await policy.ExecuteAsync(async () =>

DB09 {

OB10 using (var connection = new SqlConnection(ConnectionString))

DB11 {

0812 await connection.OpenAsync();

DB13 using (var command = new SqlCommand("-",connection))

DB14 using (var reader = command.ExecuteReader())

DB15 {

DB16

OB17

0618 }

DB19 });

OB20

DB21

ReceiptUploader.cs

RU01 public class ReceiptUploader

RU02 {

RU03 public async Task UploadFile(string file, byte[] binary)

RU04 {

RU05 var httpClient = new HttpClient();

RU06 var response = await httpClient.PutAsync("-", new ByteArrayContent(binary));

RU07 while (ShouldRetry(response))

RU08 {

RU09 response = await httpClient.PutAsync("-", new ByteArrayContent(binary));

RU10 }

RU11 }

RU12 private bool ShouldRetry(HttpResponseMessage response)

RU13 {

RU14

RU15 }

RU16 }

ConfigureSSE.ps1

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

Question: 43

DRAG DROP

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

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

NOTE: Each correct selection is worth one point.

Code segments

```
HSITokenProvider("...", null)
tp.GetAccessTokenAsync("...")
AzureServiceTokenProvider()
StringTokenProvider("storage", "msi")
tp.GetAuthenticationHeaderAsync(CancellationTokens.None)
```

Answer Area

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

Answer:

Explanation:

AzureServiceTokenProvider()

```
var t = new TokenCredential(await tp.GetAccessTokenAsync("..."))
```

```
return new StorageCredentials(t);
```

Box 1: AzureServiceTokenProvider()

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

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

```
private async Task<string> GetAccessTokenAsync()
```

```
{
```

```
var tokenProvider = new AzureServiceTokenProvider();
```

```
return await tokenProvider.GetAccessTokenAsync("https://storage.azure.com/");
```

```
}
```

Reference:

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

Question: 44

DRAG DROP

You need to ensure disaster recovery requirements are met.

What code should you add at line PC16?

To answer, drag the appropriate code fragments to the correct locations. Each code fragment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

true

SingleTransferContext

ShouldTransferCallbackAsync

false

DirectoryTransferContext

ShouldOverwriteCallbackAsync

Answer Area

```
var copyOptions = new CopyOptions { };
var context = new Value { (source, destination) => Task.FromResult(true);
context.Value { (source, destination) => Task.FroaResult(true);
await TransferManager.CopyAsync(blob, GetORBlob(blob), isServiceCopy: Value
, context: context, options:copyOptions);
```

Answer:

Explanation:

```
var copyOptions = new CopyOptions { };
var context = new DirectoryTransferContext = (source, destination) => Task.FromResult(true);
context.ShouldTransferCallbackAsync = (source, destination) => Task.FromResult(true);
await TransferManager.CopyAsync(blob, GetDRBlob(blob), isServiceCopy: False
, context: context, options:copyOptions);
```

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

Box 1: DirectoryTransferContext

We transfer all files in the directory.

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

Box 2: ShouldTransferCallbackAsync

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

Box 3: False

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

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

Reference:

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

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

Question: 45

HOTSPOT

You need to add code at line PC26 of Processing.cs to ensure that security policies are met.

How should you complete the code that you will add at line PC26? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

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

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

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

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

Answer:

Explanation:

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

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

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

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

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

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

Example:

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

```
BlobEncryptionPolicy encryptionPolicy = new BlobEncryptionPolicy(cloudKey1, cachingResolver);
```

```
client.DefaultRequestOptions.EncryptionPolicy = encryptionPolicy;
```

```
Box 3: cloudblobClient. DefaultRequestOptions.EncryptionPolicy = x;
```

Reference:

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

Question: 46

You need to ensure the security policies are met.

What code do you add at line CS07 of ConfigureSSE.ps1?

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

Answer: B

Explanation:

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

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

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

Reference:

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

Question: 47

You need to ensure receipt processing occurs correctly.

What should you do?

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

Answer: D

Explanation:

You can create a snapshot of a blob. A snapshot is a read-only version of a blob that's taken at a point in time. Once a snapshot has been created, it can be read, copied, or deleted, but not modified.

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

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

Reference:

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

Question: 48

You need to resolve the capacity issue.

What should you do?

- A. Convert the trigger on the Azure Function to an Azure Blob storage trigger
- B. Ensure that the consumption plan is configured correctly to allow scaling
- C. Move the Azure Function to a dedicated App Service Plan
- D. Update the loop starting on line PC09 to process items in parallel

Answer: D

Explanation:

If you want to read the files in parallel, you cannot use `forEach`. Each of the `async` callback function calls does return a promise. You can await the array of promises that you'll get with `Promise.all`.

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

```
PC08 var container = await GetCloudBlobContainer();
PC09     foreach (var fileTea in await ListFiles())
PC10     {
PC11         var file = new CloudFile(fileTea.StorageUri.PrimaryUri);
PC12         var as = new MemoryStream();
PC13         await file.DownloadToStreamAsync(as);
PC14         var blob = container.GetBlockBlobReference(fileTea.Uri.ToString());
PC15         await blob.UploadFromStreamAsync(as);
PC16     }
PC17 }
```

Reference:

<https://stackoverflow.com/questions/37576685/using-async-await-with-a-foreach-loop>

Question: 49

You need to resolve the log capacity issue.

What should you do?

- A. Create an Application Insights Telemetry Filter
- B. Change the minimum log level in the host.json file for the function
- C. Implement Application Insights Sampling
- D. Set a LogCategoryFilter during startup

Answer: C

Explanation:

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

Sampling is a feature in Azure Application Insights. It is the recommended way to reduce telemetry

traffic and storage, while preserving a statistically correct analysis of application data. The filter selects items that are related, so that you can navigate between items when you are doing diagnostic investigations. When metric counts are presented to you in the portal, they are renormalized to take account of the sampling, to minimize any effect on the statistics.

Sampling reduces traffic and data costs, and helps you avoid throttling.

Reference:

Topic 6, Coho Winery

Case study

Overview

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

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

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

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

LabelMaker app

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

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

External partners send data to the LabelMaker application to include artwork and text for

custom label designs.

Requirements. Data

You identify the following requirements for data management and manipulation:

Order data is stored as nonrelational JSON and must be queried using SQL.

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

Requirements. Security

You have the following security requirements:

Users of Coho Winery applications must be able to provide access to documents, resources, and applications to external partners.

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

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

Storage of e-commerce application settings must be maintained in Azure Key Vault.

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

Conditional access policies must be applied at the application level to protect company content.


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

Requirements. LabelMaker app

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

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

Architecture



Issues

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

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

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

Order.json

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

only.

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

Order.json

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


Question: 50

DRAG DROP

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

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

Log in to the registry and push image.

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

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

Download the image to your local computer.

Build a new application image by using dockerfile.

Answer Area



Answer:

Explanation:

Build a new application image by using dockerfile.

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

Log in to the registry and push image.

Step 1: Build a new application image by using dockerfile

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

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

```
# pulls hello-world from the public docker hub
```

```
$ docker pull hello-world
```

```
# tag the image in order to be able to push it to a private registry
```

```
$ docker tag hello-word <REGISTRY_NAME>/hello-world
```

```
# push the image
```

```
$ docker push <REGISTRY_NAME>/hello-world
```

Step 3: Log in to the registry and push image

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

Scenario:

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

LabelMaker app

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

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

Reference:

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

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

Question: 51

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

What should you do first?

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

Answer: C

Explanation:

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

ClaimsPrincipal from the Request Context

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

User Claims from the Request Headers.

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

Reference:

<https://levelup.gitconnected.com/four-alternative-methods-to-get-user-identity-and-claims-in-a-net-azurefunctions-app-df98c40424bb>

Question: 52

HOTSPOT

You need to configure Azure Cosmos DB.

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

NOTE: Each correct selection is worth one point.

Setting

Value

Consistency Level

	▼
Strong	
Bounded-staleness	
Session	
Eventual	

API

	▼
SQL	
MongoDB	
Graph	
Table	

Answer:

Explanation:

Setting

Value

Consistency Level

	▼
Strong	
Bounded-staleness	
Session	
Eventual	

API

	▼
SQL	
MongoDB	
Graph	
Table	

Box 1: Strong

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

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

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

Box 2: SQL

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

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

Question: 53
HOTSPOT

You need to retrieve all order line items from Order.json and sort the data alphabetically by the city.

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

NOTE: Each correct selection is worth one point.

SELECT li.id AS lineitemid, li.price

FROM

Orders o

LineItems li

JOIN

li

o

IN

o.line_items

li.line_items

o.address

ORDER BY

o.address.city

li.address.city

o.city

li.city

ASC

Answer:

Explanation:

SELECT li.id AS lineitemid, li.price

FROM

Orders o
LineItems li

JOIN

IN

li
o

o.line_items
li.line_items
o.address

ORDER BY

ASC

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

Box 1: orders o

Scenario: Order data is stored as nonrelational JSON and must be queried using SQL.

Box 2:li

Box 3: o.line_items

Box 4: o.city

The city field is in Order, not in the 2s.

Topic 7, VanArsdel. Ltd

Case study

Overview

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.

Background

VanArsdel. Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

Current environment

Requirements

The application components must meet the following requirements:

Corporate website

- Secure the website by using SSL
- Minimize costs for data storage and hosting.
- Implement native GitHub workflows for continuous integration and continuous deployment (CI/CO).
- Distribute the website content globally for local use.
- Implement monitoring by using Application Insights and availability web tests

including

SSL certificate validity and custom header value verification.

- The website must have 99.95 percent uptime.

Corporate website

The company provides a public website located at <http://www.vanaisdeltd.com>. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure functions.

Retail store locations

- Azure Functions must process data immediately when data is uploaded to Blob storage.

Azure Functions must update Azure Cosmos DB by using native SQL language queries.

- Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

Delivery services

- Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.
- Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

Inventory services

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

Security

- All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
- Authentication and authorization must use Azure AD and services must use managed identities where possible.

Retail Store Locations

- You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
- Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Question: 54
HOTSPOT

You need to implement the retail store location Azure Function.

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

NOTE: Each correct selection is worth one point.

Answer Area

Configuration	Value
Binding	Blob storage Azure Cosmos DB Event Grid HTTP
Binding Direction	Input Output
Trigger	Blob storage Azure Cosmos DB Event Grid HTTP

Answer:

Explanation:

Configuration

Value

Binding

	▼
Blob storage	
Azure Cosmos DB	
Event Grid	
HTTP	

Binding Direction

	▼
Input	
Output	

Trigger

	▼
Blob storage	
Azure Cosmos DB	
Event Grid	
HTTP	

Scenario: Retail store locations: Azure Functions must process data immediately when data is uploaded to Blob storage.

Box 1: HTTP

Binding configuration example: https://<storage_account_name>.blob.core.windows.net

Box 2: Input

Read blob storage data in a function: Input binding

Box 3: Blob storage

The Blob storage trigger starts a function when a new or updated blob is detected.

Azure Functions integrates with Azure Storage via triggers and bindings. Integrating with Blob storage allows you to build functions that react to changes in blob data as well as read and write values.

Reference:

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

Question: 55

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

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

NOTE: Each correct selection is worth one point.

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

Answer: C,E

Explanation:

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

Microsoft Azure Key Vault is a cloud-hosted management service that allows users to encrypt keys and small secrets by using keys that are protected by hardware security modules (HSMs).

You need to create a managed identity for your application.

Reference:

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

Question: 56

You need to audit the retail store sales transactions.

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

NOTE: Each correct selection is worth one point.

- A. Update the retail store location data upload process to include blob index tags. Create an Azure Function to process the blob index tags and filter by store location
- B. Enable blob versioning for the storage account. Use an Azure Function to process a list of the blob versions per day.
- C. Process an Azure Storage blob inventory report by using an Azure Function. Create rule filters on the blob inventory report,
- D. Subscribe to blob storage events by using an Azure Function and Azure Event Grid. Filter the events by store location.
- E. Process the change feed logs of the Azure Blob storage account by using an Azure Function. Specify a time range for the change feed data.

Answer: D,E

Explanation:

Scenario: Audit store sale transaction information nightly to validate data, process sales financials, and

reconcile inventory.

"Process the change feed logs of the Azure Blob storage account by using an Azure Function. Specify a time range for the change feed data": Change feed support is well-suited for scenarios that process data based on objects that have changed. For example, applications can:

Store, audit, and analyze changes to your objects, over any period of time, for security, compliance or intelligence for enterprise data management.

"Subscribe to blob storage events by using an Azure Function and Azure Event Grid. Filter the events by store location": Azure Storage events allow applications to react to events, such as the creation and deletion of blobs. It does so without the need for complicated code or expensive and inefficient polling services. The best part is you only pay for what you use.

Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener. Event Grid provides reliable event delivery to your applications through rich retry policies and dead-lettering.

Reference:

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

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

Question: 57

You need to implement a solution to resolve the retail store location data issue.

Which three Azure Blob features should you enable? Each correct answer presents part of the solution.

NOTE Each correct selection is worth one point

- A. Immutability
- B. Snapshots
- C. Versioning

D. Soft delete

E. Object replication

F. Change feed

Answer: C,D,F

Explanation:

Scenario: You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.

Before you enable and configure point-in-time restore, enable its prerequisites for the storage account: soft delete, change feed, and blob versioning.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/point-in-time-restore-manage>

Question: 58

HOTSPOT

YOU need to reliably identify the delivery driver profile information.

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

NOTE Each correct selection is worth one point.

Cuirfigwation

15 l-l ..1'0 10^ < WT) typt

Value



P vnatr r ^" n VMIF

md
*ud
»dp

Answer:

Explanation:

Value

IS N.»eni0ke (Wljtypr
P4 r'&.K ■ J TI . tJP

ID
l dji

Question: 59

HOTSPOT

You need to implement event routing for retail store location data.

Which configuration should you use?

Event diU Configuratkhtt

sa H'kt

Amit El:b Suit** AilbE EvuJ 0»d

AUlttJrIrvXiISii Amr EvPGI Hub

Reiver

Anu*Ev»:J Grid Anu- tvirtHUB Aiurt Sefflcs &w .-.Ju Jr BLbSrCdcJc

Hwidlff

Amrs rail'll. nAip ^py ■ ır Ajj Amir K#tf Grit AriMb Slih J[,ı tf ■

Answer:

Explanation:

Evert data Configuration

Sodf®

MUfrEvdlt '.III.

RtiTinn

AUNT Evort Hrt

HvtoUtr

Ira# l&pr ■ ■ pc

Question: 60

HOTSPOT

You need to implement the delivery service telemetry data

How should you configure the solution?

NOTE: Each correct selection is worth one point.

AiM^Uiikt CE

V1^

API

1

Cui* (3Qi ■ Grsndm

Tabi* li ^?JE'

1

^ "Hr* nn l <7

h**i ii

Muck liccnic^Hc

Tthick p4cki<* r^vcy

Vehicle kc&Otl cOOiiin<kc

Answer:

Explanation:

Azure Cosmos DB	Value
API	Core (SQL)
Partition Key	Vehicle package capacity

Question: 61

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

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

NOTE: Each correct selection is worth one point.

- A. Create a new composite index for the store location data queries in Azure Cosmos DB. Modify the queries to support parameterized SQL and update the Azure function app to call the new Queries.
- B. Configure Azure Cosmos DB consistency to strong consistency Increase the RUs for the container supporting store location data.
- C. Provision an Azure Cosmos OB dedicated gateway, update blob storage to use the new dedicated gateway endpoint.

- D. Configure Azure Cosmos DB consistency to session consistency. Cache session tokens in a new Azure Redis cache instance after every write. Update reads to use the session token stored in Azure Redis.
- E. Provision an Azure Cosmos DB dedicated gateway Update the Azure Function app connection string to use the new dedicated gateway endpoint.

Answer: C,D

Explanation:

Question: 62

HOTSPOT

You need to implement the corporate website.

How should you configure the solution?

Azure Configuration

PUn

Free
Standard
Premium

Icodiatd

Service

App 3«we Web Ay:
Ap} -Jwc* St-ii^ WfbApp
Aten Fuj.^nr \pj .-.2.3c
Slix Sitetge

Answer:

Explanation:

Answer Area

Azure

Configuration

Plan

	▼
Free	
Standard	
Premium	
Isolated	

Service

	▼
App Service Web App	
App Service Static Web App	
Azure Function App	
Azure Blob Storage	

Question: 63

You need to test the availability of the corporate website.

Which two test types can you use?

- A. Custom testing using the TrackAvailability API method
- B. Standard

C. URL Ping

D. Multi-step

Answer: A,B

Explanation:

Question: 64

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

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

NOTE: Each correct selection is worth one point.

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

Answer: B,C

Explanation:

Question: 65

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

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

E. Azure AD refresh token

Answer: C

Explanation:

Topic 8, Munson's Pickles and Preserves Farm

Overview

Munson's Pickles and Preserves Farm is an agricultural cooperative corporation based in Washington, US, with farms located across the United States. The company supports agricultural production resources by distributing seeds, fertilizers, chemicals, fuel, and farm machinery to the farms.

The company is migrating all applications from an on-premises datacenter to Microsoft Azure. Applications support distributors, farmers, and internal company staff.

Corporate website

- The company hosts a public website located at <http://www.munsonspicklesandpreservesfarm.com>. The site supports farmers and distributors who request agricultural production resources.

Farms

- The company created a new customer tenant in the Microsoft Entra admin center to support authentication and authorization for applications.

Distributors

- Distributors integrate their applications with data that is accessible by using APIs hosted at <http://www.munsonspicklesandpreservesfarm.com/api> to receive and update resource data. The application components must meet the following requirements:

Corporate website

- The site must be migrated to Azure App Service.
- Costs must be minimized when hosting in Azure.
- Applications must automatically scale independent of the compute resources.
- All code changes must be validated by internal staff before release to production.
- File transfer speeds must improve, and webpage-load performance must increase.
- All site settings must be centrally stored, secured without using secrets, and encrypted at rest and in transit.
- A queue-based load leveling pattern must be implemented by using Azure Service Bus queues to support high volumes of website agricultural production resource requests.

Farms

- Farmers must authenticate to applications by using Microsoft Entra ID.

Distributors

- The company must track a custom telemetry value with each API call and monitor

performance of all APIs.

- API telemetry values must be charted to evaluate variations and trends for resource data. **Internal staff**
- App and API updates must be validated before release to production.
- Staff must be able to select a link to direct them back to the production app when validating an app or API update.
- Staff profile photos and email must be displayed on the website once they authenticate to applications by using their Microsoft Entra ID.

Security

- All web communications must be secured by using TLS/HTTPS.
- Web content must be restricted by country/region to support corporate compliance standards
- The principle of least privilege must be applied when providing any user rights or process access rights
- Managed identities for Azure resources must be used to authenticate services that support Microsoft Entra ID authentication.

Corporate website

- Farmers report HTTP 503 errors at the same time as internal staff report that CPU and memory usage are high.
- Distributors report HTTP 502 errors at the same time as internal staff report that average response times and networking traffic are high.
- Internal staff report webpage load sizes are large and take a long time to load.
- Developers receive authentication errors to Service Bus when they debug locally. **Distributors**
- Many API telemetry values are sent in a short period of time. Telemetry traffic, data costs, and storage costs must be reduced while preserving a statistically correct analysis of the data points sent by the APIs.

Question: 66

HOTSPOT

You need to resolve the authentication errors for developers.

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

NOTE: Each correct selection is worth one point.

Azure Service Bus security configuration

Security configuration setting

Security configuration value

Azure role-based access control (RBAC) role Service Bus Data Owner

Owner Contributor

Service Bus Data Owner

Service Bus Data Sender

Service Bus scope

Namespace

Queue

Namespace

Subscription

Resource group

Answer:

Explanation:

Azure Service Bus security configuration

Security configuration setting

Security configuration value

Azure role-based access control (RBAC) role

Service Bus Data Owner

Service Bus scope

Namespace

Question: 67

You need to secure the corporate website to meet the security requirements.

What should you do?

- A. Create an App Service instance with a standard plan. Configure the custom domain with a TLS/SSL certificate.
- B. Create an Azure Application Gateway with a Web Application Firewall (WAF). Configure end-to-end TLS encryption and the WAF.
- C. Create an Azure Cache for Redis instance. Update the code to support the cache.
- D. Create an Azure Content Delivery Network profile and endpoint. Configure the endpoint.

Answer: A

Explanation:

Question: 68

HOTSPOT

You need to configure App Service to support the corporate website migration.

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

NOTE: Each correct selection is worth one point.


Azure App Service configuration

Configuration setting

App Service plan

Code change validation feature

Configuration value



Configuration setting	Configuration value
App Service plan	Basic
Code change validation feature	Deployment slot

Custom container

Domain certificate

Deployment credentials

Answer:

Explanation:

Azure App Service configuration

Configuration setting

Configuration value

App Service plan

Code change validation feature

Basic

Deployment slot

Question: 69

You need to implement farmer authentication.

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

NOTE: Each correct selection is worth one point.

- A. Add the shared access signature (SAS) token to the app
- B. Create a shared access signature (SAS) token.
- C. Register the app in Microsoft Entra ID.
- D. Create a user flow.
- E. Add the app to the user flow.

Answer: C,D,E

Explanation:

Question: 70

You need to configure all site configuration settings for the corporate website.

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

NOTE: Each correct selection is worth one point.

- A. Create a managed identity.
- B. Update the role assignments for the Azure App Configuration store
- C. Create an Azure Key Vault.
- D. Create an Azure App Configuration store.
- E. Update the role assignments for the Azure Key Vault.

Answer: A,C,E

Explanation:

Question: 71

HOTSPOT

You need to provide internal staff access to the production site after a validation.


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

Go back
to production app

Explanation:



Answer:

Answer Area

<a href="https://www.munsonspicklasandprasarvesfarm.com/? "x-ms-routin^name self
>Go back to production app

Question: 72

DRAG DROP

You need to correct the internal staff issue with webpages.

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

- 1 :: Add the Azure Content Delivery Network origin to the origin group.
- 2 :: Create an Azure Content Delivery Network origin group.
- 3 :: Configure a new Azure Content Delivery Network origin.
- 4 :: Create an Azure Content Delivery Network profile.
- 5 :: Create an Azure Content Delivery Network endpoint.
- 6 :: Configure Azure Content Delivery Network compression.

Answer area

Answer:

Explanation:

Actions

- 1 :: Add the Azure Content Delivery Network origin to the origin group.
- 2 :: Create an Azure Content Delivery Network origin group.
- 3 :: Configure a new Azure Content Delivery Network origin.

Answer area

- 1 :: Create an Azure Content Delivery Network profile.
- 2 :: Create an Azure Content Delivery Network endpoint.
- 3 :: Configure Azure Content Delivery Network compression.

Question: 73

You need to implement an aggregate of telemetry values for distributor API calls.

Which Application Insights API method should you use?

- A. TrackTrace
- B. Track Exception

C. TrackEvent

D. TrackDependency

E. TrackMetric

Answer: E

Explanation:

Question: 74

HOTSPOT

You need to correct the errors for farmers and distributors.

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

NOTE: Each correct selection is worth one point.

Farmer and distributor errors

Issue

Farmers' errors

Solution

Scale up the App Service plan to Premium A

Scale up the App Service plan to Premium

Add an App Service staging deployment slot.

Configure the App Service Local Cache feature.

Create an Azure Content Delivery Network profile and endpoint.

Restart the application from the App Service portal

Scale up the App Service plan to Premium

Configure the App Service Local Cache feature

Restart the application from the App Service portal

Create a custom autoscale rule to increase the instance count

Distributors errors

Answer:

Explanation:

Farmer and distributor errors

Issue

Farmers' errors

Distributors errors

Solution

Scale up the App Service, plan to Premium

Restart the application from the App Service portal

Question: 75

HOTSPOT

You need to display the profile photo and email for signed-in internal staff on the website.

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

NOTE: Each correct selection is worth one point.

Microsoft Graph configuration

Configuration setting	Configuration value
End point	/v1.0/me/people/?\$search=me /v1.0/users?\$select=profilePhoto,mail
Permission	User.Read User.Export.All User.ReadWrite User.ManageIdentities.All

Answer:

Explanation:

Microsoft Graph configuration

Configuration setting	Configuration value
Endpoint	/v1.0/me
Permission	User.Read

Topic 9, Fourth Coffee

Background

Current Environment

Fourth Coffee is a global coffeehouse chain and coffee company recognized as one of the world's most influential coffee brands. The company is renowned for its specialty coffee

beverages, including a wide range of espresso-based drinks, teas, and other beverages. Fourth Coffee operates thousands of stores worldwide.

The company is developing cloud-native applications hosted in Azure.

Corporate website

The company hosts a public website located at <http://www.fourthcoffee.com/>. The website is

used to place orders as well as view and update inventory items.

Inventory items

In addition to its core coffee offerings, Fourth Coffee recently expanded its menu to include inventory items such as lunch items, snacks, and merchandise. Corporate team members constantly update inventory. Users can customize items. Corporate team members configure inventory items and associated images on the website.

Orders

Associates in the store serve customized beverages and items to customers. Orders are placed on the website for pickup.

The application components process data as follows:


1. Azure Traffic Manager routes a user order request to the corporate website hosted in Azure App Service.
2. Azure Content Delivery Network serves static images and content to the user.
3. The user signs in to the application through a Microsoft Entra ID for customers tenant.
4. Users search for items and place an order on the website as item images are pulled from

Azure Blob Storage

5. Item customizations are placed in an Azure Service Bus queue message.
6. Azure Functions processes item customizations and saves the customized items to Azure Cosmos DB.

7. The website saves order details to Azure SQL Database.
8. SQL Database query results are cached in Azure Cache for Redis to improve performance.

The application consists of the following Azure services:



Requirements

The application components must meet the following requirements:

- Azure Cosmos DB development must use a native API that receives the latest updates and stores data in a document format.
- Costs must be minimized for all Azure services.
- Developers must test Azure Blob Storage integrations locally before deployment to Azure.
- Testing must support the latest versions of the Azure Storage APIs.

Corporate website

- User authentication and authorization must allow one-time passcode sign-in methods and social identity providers (Google or Facebook).
- Static web content must be stored closest to end users to reduce network latency.

Inventory items

- Customized items read from Azure Cosmos DB must maximize throughput while ensuring data is accurate for the current user on the website.
- Processing of inventory item updates must automatically scale and enable updates across an entire Azure Cosmos DB container.
- Inventory items must be processed in the order they were placed in the queue.
- Inventory item images must be stored as JPEG files in their native format to include exchangeable image file format (data) stored with the blob data upon upload of the image file.
- The Inventory Items API must securely access the Azure Cosmos DB data.

Orders

- Orders must receive inventory item changes automatically after inventory items are updated or saved.

Issues

- Developers are storing the Azure Cosmos DB credentials in an insecure clear text manner within the Inventory Items API code.
- Production Azure Cache for Redis maintenance has negatively affected application performance.

Question: 76

HOTSPOT

You need to implement a function by using Azure Functions to process customized items.

How should you implement the function? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Azure Functions implementation

Implementation option

Implementation value

Trigger

Input binding

Output binding

Event source

Cosmos DB

Blob HTTP

Service Bus

Cosmos DB

Answer:

Explanation:

Azure Functions implementation

Implementation option

Implementation value

Azure Functions feature

Trigger

Event source

Cosmos DB

Question: 77

DRAG DROP

You need to secure the corporate website for users.

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.

Corporate website security configuration

- 1. Register the function app with Microsoft Entra.
- 2. Create a user flow and associate the function app.
- 3. Grant Microsoft Graph API permissions to the function app.
- 4. Create a user flow and associate the web app.
- 5. Register the web app with Microsoft Entra.
- 6. Grant Microsoft Graph API permissions to the web app.
- 7. Create a client secret.

Answer Area

Answer:

Explanation:

Corporate website security configuration

- 1. Register the function app with Microsoft Entra.
- 2. Create a user flow and associate the function app.
- 3. Grant Microsoft Graph API permissions to the function app.

Answer Area

1. Create a user flow and associate the web app.
2. Register the web app with Microsoft Entra.
3. Grant Microsoft Graph API permissions to the web app.
4. Create a client secret.

Question: 78

You need to implement the processing of enqueueer inventory items.

Which message value should you use?

- A. Session identifier
- B. Sequence number
- C. Partition key
- D. Time stamp

Answer: B

Explanation:

Question: 79

DRAG DROP

You need to store inventory item images.

Which Azure Blob Storage feature should you use? To answer, move the appropriate Azure Blob Storage features to the correct requirements. You may use each Azure Blob Storage feature once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content

NOTE: Each correct selection is worth one point.

Azure Blob Storage features

Index tags

Change feed

System properties

User-defined metadata

Azure Blob Storage requirements

Requirement

Store images in native format

Store image Exif data

Azure Blob Storage feature

Answer:

Explanation:

Azure Blob Storage features

Index tags

Change feed

System properties

User-defined metadata

Azure Blob Storage requirements

Requirement

Store images in native format

Store image Exif data

Azure Blob Storage feature

(System properties, User-defined metadata)

Question: 80

You need to securely access inventory items when developing the Inventory Items API.

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

NOTE: Each correct selection is worth one point.

A. Create a custom Microsoft Entra roleAssign the custom role to the Azure Cosmos DB account Update the function app to use certificate-based authentication.

- B. Create a SQL role definition under the Azure Cosmos DB account. Create a user-assigned managed identity and assign the identity to the function app. Assign the user-assigned managed identity the SQL role definition. Update the function app code to implement the DefaultAzureCredential class and reference the user-assigned managed identity.
- C. Create a system-assigned managed identity for the function app with read access to secrets in Azure Key Vault. Store the Azure Cosmos DB primary key and URI in Azure Key Vault as secrets. Use function app settings to reference the secret values.
- D. Create a custom Microsoft Entra role. Assign the custom role to Azure Key Vault. Assign the custom role to the function app. Reference the custom role in the function app code when accessing Azure Key Vault values.
- E. Create a SQL role definition under the Azure Cosmos DB account. Assign the role to the function app's system-assigned managed identity. Programmatically access the Azure Cosmos DB keys from the function app.

Answer: B,C,E

Explanation:

Question: 81

You need to mitigate the Azure Cache for Redis issue.

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

NOTE: Each correct selection is worth one point.

- A. Increase the maxmemory-reserved and maxfragmentationmemory-reserved values
- B. Modify the maxmemory policy to evict the least frequently used keys out of all keys.
- C. Configure client connections to retry commands with exponential backoff.
- D. Test application code by purging the cache in the test environment.
- E. Test application code by rebooting all nodes in the test environment.

Answer: A,C

Explanation:

Question: 82

HOTSPOT

You need to save customized items to Azure Cosmos DB.

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

NOTE: Each correct selection is worth one point.

Azure Cosmos DB configuration

Configuration option

API

Consistency level

Configuration value

The image shows two dropdown menus. The first dropdown, labeled 'API', has a list of options: 'NoSQL', 'MongoDB', 'PostgreSQL', and 'Apache Gremlin'. The second dropdown, labeled 'Consistency level', has a list of options: 'Strong', 'Session', 'Eventual', and 'Bounded Staleness'.

Answer:

Explanation:

1. API:

NoSQL

Reasoning: Since the development must use a native API that stores data in a document format, the NoSQL API is the correct choice. Azure Cosmos DB's NoSQL API stores data in JSON documents and is the most commonly used API for document-based use cases, which aligns with the storage of customized items in the scenario.

2. Consistency Level:

Session

Reasoning: The scenario requires that customized items maximize throughput while ensuring that the data is accurate for the current user. Session consistency provides the best balance between performance and data

accuracy for scenarios where multiple reads and writes are performed by the same user session. It ensures that a user reads their own writes, which is important for ensuring that customized items are accurate for the current user.

Final Answer:

API: NoSQL

Consistency level: Session

Question: 83

You need to support local development testing for developers. Which tool should you use?

- A. Azurite
- B. Azure Storage Emulator
- C. SQL Server Management Studio (SSMS)
- D. Azure Storage Explorer

Answer: D

Explanation:

Question: 84

You need to serve static content from the corporate website.

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. Store all static content in Azure Blob Storage. Enable Azure Content Delivery Network for the storage account.
- B. Configure App Service networking to create a Content Delivery Network profile and endpoint.
- C. Configure the Azure App Service Local Cache feature and set the app setting

WEBSITE_LOCAL_CACHE_SIZEINMB value.

- D. Update the Azure Traffic Manager routing method to priority.
- E. Create a nested Azure Traffic Manager profile. Configure the parent profile to the performance traffic routing method and the child profile to the priority traffic routing method.

Answer: A,B

Explanation:

Topic 10, Misc. Questions

Question: 85

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: 86

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

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

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

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of

the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

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

Solution: 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: 87

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: 88

DRAG DROP

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

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

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

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

Actions

Answer Area

Use AZCopy to copy the data to the new storage account

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

Export a Resource Manager template.

Create a new template deployment

Modify the template by changing the storage account name and region



Answer:

Explanation:

Export a Resource Manager template.

Create a new template deployment.

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

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

Use AZCopy to copy the data to the new storage account

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

Question: 89

HOTSPOT

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

The application needs to store runtime diagnostic data that must be persisted across application

restarts. You have the following code:

```
public void SaveDiagData(string data)

    var path = Environment.GetEnvironmentVariable("DLAGDATA")
    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

LOCALAPPDATA
WEBSITE_LOCALCACHE_ENABLED
DOTNET_HOSTING_OPT1MIZATION_CACHE WEBSITES
ENABLE APP SERVICE STORAGE

DIAGDATA

Value

true

	▼
/home	
/local	
D:\home	
D:\local	

Answer:

Explanation:

App setting

Value

LOCALAPPDATA	true
WEBSITE_LOCALCACHE_ENABLED	
DOTNET_HOSTING_OPTIMIZATION_CACHE	
WEBSITES_ENABLE_APP_SERVICE_STORAGE	/home /local D:\home
DIAGDATA	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: 90

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:

E. 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: 91

HOTSPOT

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

WebJob that processes data updates. Four customers will use the web service.

Each instance of the WebJob processes data for a single customer and must run as a singleton instance.

Each deployment must be tested by using deployment slots prior to serving production data.

Azure costs must be minimized.

Azure resources must be located in an isolated network.

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

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

NOTE: Each correct selection is worth one point.

App service plan setting

Value

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

Answer:

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).

Reference:

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

Question: 92

DRAG DROP

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Explanation:

Answer:

Setting

CRD type

Azure Function code

Deployment

Polling interval

ScaledObject

Azure Storage connection string

Secret

Box 1: Deployment

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

Box 2: ScaledObject

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

Example of ScaledObject with polling interval

`apiVersion: keda.k8s.io/v1alpha1`

`kind: ScaledObject`

`metadata:`

`name: transformer-fn`

`namespace: tt`

`labels:`

`deploymentName: transformer-fn`

`spec:`

scaleTargetRef:

deploymentName: transformer-fn

pollingInterval: 5

minReplicaCount: 0

maxReplicaCount: 100

Box 3: Secret

Store connection strings in Kubernetes Secrets.

Example: to create the Secret in our demo Namespace:

- create the k8s demo namespace

kubectl create namespace tt

- grab connection string from Azure Service Bus

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

- g \$RG_NAME \
- -namespace-name \$SBN_NAME \
- -queue-name inbound \
- n keda-scaler \
- -query "primaryConnectionString" \
- o tsv)

to create the kubernetes secret

kubectl create secret generic tt-keda-auth \

- --from-literal KedaScaler=\$KEDA_SCALER_CONNECTION_STRING \

Reference:

Question: 93

NOTE: Each correct selection is worth one point.

az group create --location westeurope --name myResourceGroup

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

--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

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

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

source config --name \$webappname

--resource-group myResourceGroup

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

Answer:

Explanation:

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

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: 94

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: 95

HOTSPOT

You are developing a ticket reservation system for an airline.

The storage solution for the application must meet the following requirements:
Ensure at least 99.99% availability and provide low latency.

Accept reservations 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:

Explanation:

```
resourceGroupName- 4-airlineResourceGroup.p'
name- +docdb-airline-reservations'
da tabas eName- 'docdb-tickets-database'
collectionName- +docdb-tickets-collection'
consistencyLevel- w
```

Strong Eventual ConsistentPrefix BoundedStaleness

```
az cosmosdb create \
--name Sname \
enable-virtual-network true\ enable-automatic-failover
true\ kind 'GlobalDocumentDB' \ _____ kind
MongoDB\
-resource group SresourceGroupName \ -max
interval 5 \

locations 'southcentralus'
locations 'eastus'
locations'southcentralus=0 eastus=1 westus=2'
locations ^d'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: 96

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: 97

DRAG DROP

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

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

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

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

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

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

NOTE: Each correct selection is worth one point.

JSON segments

orderBy
sortOrder
ascending
descending
compositeIndexes

Answer Area

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

Answer:

Explanation:

```

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

```

```

"path": " ", "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: 98

HOTSPOT

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

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

Data throughput 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:

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

Reference:

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

Question: 99

DRAG DROP

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

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

reverse proxy capabilities

configurable traffic routing

TLS termination with a custom certificate

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

NOTE: Each correct selection is worth one point.

Components

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

Answer area

Action

- Deploy solution.
- View cluster And external IP addressing
- implement a single public IP endpoint that is routed to multiple microservices.

Component

- Component
- Component
- Component

b

Answer:

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: 100

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: 101

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.

Incorrect Answers:

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

Reference:

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

Question: 102

DRAG DROP

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

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

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

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

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

NOTE: Each correct selection is worth one point

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

Answer:

Explanation:

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

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: 103

DRAG DROP

You develop a web application.

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

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

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:

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: 104

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:

Reference:

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

Question: 105

DRAG DROP

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 SsecretName -SecretValue
$ secretvalue

Get-AzStorageAccountKey -
ResourceGroupName SresGroup -Name $ storAcct

Set-AzContext -SubscriptionId
SsubscriptionId

Get-AzKeyVaultSecret -VaultName
$ vaultName

Get-AzSubscription
```



Answer:

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: 106

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: A,C

Explanation:

Reference:

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

Example:

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

```
{
```

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

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

```
var operation = telemetryClient.StartOperation<DependencyTelemetry>("enqueue " + queueName);
```

```
operation.Telemetry.Type = "Azure Service Bus";
```

```
operation.Telemetry.Data = "Enqueue " + queueName;
```

```
var message = new BrokeredMessage(payload);
```

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

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

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

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

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

Reference:

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

Question: 107

HOTSPOT

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:

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: 108

HOTSPOT

You are developing an Azure App Service hosted ASP.NET Core web app to deliver video on-demand streaming media. You enable an Azure Content Delivery Network (CDN) Standard for the web endpoint. Customer videos are downloaded from the web app by using the following example URL.: `http://www.contoso.com/content.mp4?quality=1`

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

You need to configure Azure CDN caching rules.

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

NOTE: Each correct selection is worth one point.

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:

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: 109

DRAG DROP

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

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:

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: 110

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: 111

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 for order processing and financial transactions.

Reference:

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

Question: 112

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: 113

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

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

You are developing an Azure 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: 114

DRAG DROP

You manage several existing Logic Apps.

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

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

NOTE: Each correct selection is worth one point.

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

Answer:

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.

Reference:

<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: 115

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: 116

HOTSPOT

You are developing an application that uses Azure Storage Queues.

You have the following code:

```
1 creds = StorageCredentials.FromConnectionString(connectionString);  
2 CloudStorageAccount storageAccount = CloudStorageAccount.Parse(  
3     creds.ToString());  
4 CloudConfigurationManager.GetSetting("StorageConnectionString");  
5 queueClient = storageAccount.CreateCloudQueueClient();  
6  
7 CloudQueue queue = queueClient.GetQueueReference("appqueue");  
8 await queue.CreateIfNotExistsAsync();  
9  
10 peekedMessage = await queue.PeekMessageAsync();  
11 if (peekedMessage != null)  
12 {  
13     Console.WriteLine("The peeked message is: {0}", peekedMessage.AsString());  
14 }  
15  
16 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:

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: 117

HOTSPOT

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.

0

0

If the body size is >256k, an error will occur.

0

0

If the request is http://contoso.com/api/9 2/, the policy will retain the higher version

0

Answer:

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.

0

0

If the request is http://contoso.com/api/9 2/, the policy will retain the higher version. 0

0

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: 118

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: 119

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);
```

Reference:

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

Question: 120

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.

Reference:

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

Question: 121

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);
}
```


Reference:

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

Question: 122

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

Explanation:

Question: 123

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"));
```

Reference:

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

Question: 124

DRAG DROP

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

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

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

You set the value for Indexing Mode to Consistent.

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

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

NOTE: Each correct selection is worth one point.

Consistency levels		Answer Area
Strong	Bounded Staleness	Return the most recent patient status.
Consistent Prefix	Eventual	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:

Explanation:

Return the most recent patient status | Strong

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

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

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.

Reference:

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

Question: 125

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.


```

        email phone region }

    public string Phone { get; set; } }

public class Player

    protected PlayerEntity player;
    async void GetPlayer(string cs,
        ▼ table, string pk, string rk)

        CloudTable CloudTableClient TableEntity
        TableEntityAdapter

        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 =..

Reference:

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

Question: 126

HOTSPOT

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

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

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


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

NOTE: Each correct selection is worth one point.

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

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

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




Answer:

Explanation:

Answer Area

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



ExecuteBatch
Execute
Insert InsertOrMerge

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.

Reference:

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

Question: 127

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.

Incorrect Answers:

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

Reference:

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

Question: 128

DRAG DROP

You are deploying an Azure Kubernetes Services (AKS) cluster that will use multiple containers.

You need to create the cluster and verify that the services for the containers are configured correctly and available.

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

Command segments

az aks get-credentials

az appservice plan create

az aks create

az group create

kubectl apply

Answer Area



Answer:

Explanation:

az group create

az aks create

kubectl apply

I 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

Reference:

<https://docs.bitnami.com/azure/get-started-aks/>

Question: 129

HOTSPOT

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

You add the following parameters to the function.

Parameter name	Description
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<CloudTask> StartTask(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.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(batchAccountName, batchAccountKey, (outputContainerSasUrl));
                OutputFileBlobContainerDestination failedContainer =
                    new OutputFileBlobContainerDestination(batchAccountName, batchAccountKey, 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
                    TaskCompletion)));
                task = new CloudTask(taskId, fileTask.Command,
                    new OutputFileDestination(outputContainer),
                    new OutputFileDestination(failedContainer),
                    new OutputFileUploadOptions (OutputFileUploadCondition.
                    TaskSuccess TaskFailure TaskCompletion
                    TaskCompletion));
                tasks.Add(task);
            });
    return tasks;
}
```

Answer

Explanation:

* = catch client job operations.

▼ 0;

GetJob
GetTask
EnableJob
CreateJob

```
job.Id = jobId,  
job.PoolInformation = new PoolInfo { PoolId = poolId }; job.Commit();  
fileTasks.ForEach((fileTask) =>
```

```
    string taskId = $"Task{DateTime.Now.ToFileTimeUtc().ToString("O")};  
    CloudTask task = new CloudTask(taskId, fileTask.Command);  
    List<OutputFile> outputFileList = new List<OutputFile>();  
    OutputFileBlobContainerDestination outputContainer =  
        new OutputFileBlobContainerDestination(fileTask.Output, (outputContainerSasUrl) ;  
    if (outputContainerSasUrl == null) failedContainer =  
        new OutputFileBlobContainerDestination(failedContainerSasUrl);  
    outputFileList.Add(new OutputFile(taskId, fileTask.Output, failedContainer,  
        new OutputFileDestination(outputContainer),  
        new OutputFileUploadOptions(OutputFileUploadCondition.
```

TaskSuccess
TaskFailure
TaskCompletion

```
    outputFileList.Add(new OutputFile(taskId, fileTask.Output, failedContainer,  
        new OutputFileDestination(outputContainer),  
        new OutputFileUploadOptions(OutputFileUploadCondition.
```

TaskSuccess
TaskFailure
TaskCompletion

task



=outputFileList;

OutputFiles
FilesToStage
ResourceFiles
StageFiles

```
task.Add(task);
```

```
});
```

```
}
```

```
return tasks,
```

```
}
```


Box 1: CreateJob

Box 2: TaskSuccess

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

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

Box 3: TaskFailure

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

Box 4: OutputFiles

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

Reference:

<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: 130

DRAG DROP

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:

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)
    {
        ILoggerFactory.AddApplicationInsights (app. AppiCatlOnSeMCeS app.UseMvc(); LogLevel.Trace);
    }
}
```

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);
```

Reference:

<https://blog.computedcloud.com/enabling-application-insights-trace-logging-in-asp-net-core/>

Question: 131

DRAG DROP

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

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

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

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:

Explanation:

A user requests the innage 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.

i 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 innage to the client.

Subsequent requests for the file maybe 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.

Reference:

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

Question: 132

HOTSPOT

You are creating an app that uses Event Grid to connect with other services. Your app's event data will be sent to a serverless function that checks compliance. This function is maintained by your company.

You write a new event subscription at the scope of your resource. The event must be invalidated after a specific period of time. You need to configure Event Grid to ensure security.

What should you implement? To answer, select the appropriate options in the 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:

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.

Reference:

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

Question: 133

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

Incorrect Answers:

A: The Import-AzureRmApiManagementApi cmdlet imports an Azure API Management API from a file or a URL in Web Application Description Language (WADL), Web Services Description Language (WSDL), or Swagger format.

B: New-AzureRmApiManagementBackend creates a new backend entity in Api Management.

C: The New-AzureRmApiManagement cmdlet creates an API Management deployment in Azure API Management.

Reference:

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

Question: 134

DRAG DROP

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

Property	Description
ShipLocation	the country/region where the order will be shipped
Correlation Id	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

SQLFilter

CorrelationFilter

No Filter

Answer Area

Subscription Filter type

FutureOrders

HighPriorityOrders

InternationalOrders

HighQuantityOrders

AllOrders

Explanation:

Answer:

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.

Reference:

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

Question: 135

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

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

You need to implement a reply trail auditing solution.

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

NOTE: Each correct selection is worth one point.

- A. Assign the value of the hazard message SessionId property to the ReplyToSessionId property.
- B. Assign the value of the hazard message MessageId property to the DeliveryCount property.

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

Answer: A,C

Explanation:

Reference:

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

Question: 136

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,B

Explanation:

Question: 137

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");
```

Reference:

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

Question: 138

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

You need to capture the data required to implement the Usage Analytics feature of Application Insights. Which three data values should you capture? Each correct answer presents part of the solution

NOTE: Each correct selection is worth one point.

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

Answer: A,D,E

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.

Reference:

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

Question: 139

DRAG DROP

You are developing Azure WebJobs.

You need to recommend a WebJob type for each scenario.

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

NOTE: Each correct selection is worth one point.

WebJob types

Scenario

WebJob type

Triggered

Continuous

Run on all instances that the web app runs on.
Optionally restrict the WebJob to a single instance.

Run on a single instance that Azure selects for load balancing.

Supports remote debugging

Answer:

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 selects 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.

Reference:

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

Question: 140

DRAG DROP

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

You develop the following code:

```
services.AddAuthorization (options =>
```

```

options.AddPolicy("Providerpartner", policy => {
    .policy.AddAuthenticationSchemes("Cookie,
    Bearer") policy.RequireAuthenticatedUser();
    policy.RequireRole("ProviderAdmin",
    "SysAdmin"); policy.Required aim ("edi tor",
    "partner") ;
D;
}

```

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.

Code Seaments

```
(
    (Policy - "ProviderEditor")
    (Role - "SysAdmin"))

(Role - "ProviderAdmin")
(Role - "SysAdmin")

(Role - "SysAdmin" ^ "ProviderAdmin")
```

Answer Area

public class

Public class

Manage (>

Answer/oz*

Answer

Explanation:

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

public class
```

{Authorize(Policy * "ProviderEditor" ^ (Role ^ SysAdmin))}

Public class 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: 141

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:

Reference:

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

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

Question: 142

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Property

Value

Client certificate location

HTTP request header
Client cookie
HTTP message body
URL query string

Encoding type

HTML
URL
Unicode
Base64

Answer:

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.

Reference:

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

Question: 143

DRAG DROP

Fourth Coffee has an ASP.NET Core web app that runs in Docker. The app is mapped to the

www.fourthcoffee.com domain.

Fourth Coffee is migrating this application to Azure.

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

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

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

Azure CLI commands

```
az webapp config hostname add
--webapp-name SappName
--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 SappName
-plan AppServiceLinuxDockerPlan
-resource-group fourthCoffeePublicWebResourceGroup

az webapp config container set
-docker-custom-image-name SdockerHibContainerPath
-name SappName
-resource-group fourthCoffeePublicWebResourceGroup
```

Answer area



Explanation:

```
#fbn/bash
appName="FourthCoffeePublicWebSrandom'.
```

Answer:

location "WestUS"

**dockerHubContainerPath="FourthCoffee/publicweb:vr
fqdn=http://www.fourthcoffee.com>www.foijrthcoffee.com**

az webapp config hostname add

-webapp-name SappName

-resource-grcup fourthCoffeePublic WebResourceGroup

-■hostname Sfqn

azwebapp create

-name SappName

-plan AppServiceLinuxDockerPlan

-resource-group fourthCoffeePublicWebResourceGroup

azwebapp config container set

-docker-custom-image-name SdockerHibCcntanerPath

-name SappName

-resource-group fcurthCcffeePublicWebResourceGroup

Step 1: #bin/bash

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

Step 2: az webapp config hostname add

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

Step 3: az webapp create

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

Step : az webapp config container set

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

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

Reference:

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

Question: 144

HOTSPOT

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

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

You need to implement push notifications.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

hub=
NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails



GetInstallation
CreateClientFromConnectionString
CreateOrUpdateInstallation
PatchInstallation

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


Answer:

Explanation:

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



```
(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.
```



Box 1: NotificationHubClient

Box 2: NotificationHubClient

Box 3: CreateClientFromConnectionString

```
// Initialize the Notification Hub
```

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

Box 4: SendWindowsNativeNotificationAsync

Send the push notification.

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

Reference:

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

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

Question: 145

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.

Reference:

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

Question: 146

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.

Reference:

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

Question: 147

HOTSPOT

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

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

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

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

```

01. namespace ContosoCradt
02. {
03.     public class PlayerDbContext : DbContext
04.     {
05.         public PlayerDbContext() : base ("name=dbConnstring") { }
06.         public DbSet<Player> Players { get ; set ; }
07.         public DbSet<Game> Games { get ; set ; }
08.         protected override void OnModelCreating(DBModelBuilder modelBuilder)
09.         {
10.             modelBuilder.Entity<Player>().HasMany(x => x.Games). WithMany (x => x.Players);
11.         }
12.     }
13.     internal static class dbCn {
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.             public class Player
42.             {
43.                 public int PlayerId { get ; set; }
44.                 public string PlayerName { get ; set; }
45.                 public virtual List<Game> Games { get ; set; }
46.             }
47.             public class Game
48.             {
49.                 public int Gamers { get ; set }
50.                 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 gameId value.

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

Answer:

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 gameId 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: 148

HOTSPOT

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

You have the following requirements:

- Code should be scalable and allow for growth.
- Each record must consist of a 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.)

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

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

PartitionKey	RowKey	Email
Harp	Walter	wharp^contoso com
Smith	Steve	ssmith@con(oso com
Smith	Jeff	]smith@contoso com

```
01 CloudTableC. lent tableClient = account.CreateCloudTableClient();
02 CloudTable table = tableclient.GetTableReference("people");
03 TableQuery<Cust>I r.erEr.ti?y> query = new TableQuery<Cust:merEr.'iity>()
04     .Where(TableQuery.CombineFilters(
05         , .Generate..And, TableQuery.GenerateFilterCondition(Email, Qu-;iyC ?e. ... .Equal,
06 Tab' aOperstcrs.And, Tar '.eQuery .GeneratePilterCondition (Email, QueryCompar: sons .Equal,
07 "ssmith@contoso.com")
08 await table.ExecuteQuerySegmentedA8ync<Custc.-i"iEntity>(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 exists with the same playerId and gameId.

already

\$

☐

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:

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

Reference:

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

Question: 149

HOTSPOT

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

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

```
1  public async Task<List<PlayerScore>> GetScore (string gameId, int score, string playerName)
2  {
3      TableQuery<PlayerScore> query = new TableQuery<PlayerScore>().Where("Score > {score}").OrderBy("Score").Take(20);
4      EntityResolver<PlayerScore> resolver = new EntityResolver<PlayerScore>((rowKey, props) => new PlayerScore {
5          PartitionKey = gameId,
6          RowKey = rowKey,
7          Score = props["Score"] as int?
8      });
9      public class PlayerScore : TableEntity
10     {
11         public PlayerScore(string gameId, string playerId, int score, long timePlayed)
12         {
13             PartitionKey = gameId;
14             RowKey = playerId;
```

```

15 Score = score;
16 TimePlayed = timePlayed;
17 }
18 public int Score { get; set; }
19 public long TimePlayed { get; set; }
20 }

```

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

```

01 public void SaveScore(string gameId, string playerId, int score, long timePlayed) 02 (
03     {
04         // Create table client
05         TableClient tableClient = storageAccount.CreateCloudTableClient();
06         Table table = tableClient.GetTableReference("players");
07         Table.CreateIfNotExists(table);
08         var scoreRecord = new PlayerScore(gameId, playerId, score, timePlayed);
09         Table.Insert(scoreRecord);
10     }
11 public class PlayerScore : TableEntity
12 {
13     public Playerscore(string gameId, string playerId, int score, long timePlayed)
14     {
15         this.PartitionKey = gameId;
16         this.RowKey = playerId;
17         Score = score;
18         TimePlayed = timePlayed;
19     }
20 public int Score { get; set; }
    public long TimePlayed { get; set; }
}

```

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>().Where(
04     TableQuery.CombineFilters(
05         TableQuery.GenerateFilterCondition("Email", QueryCondition.Equal, "smith"),
06         TableQuery.CombineFilters(
07             TableQuery.GenerateFilterCondition("Email", QueryCondition.Equal, "ssmithgcontoso.com")
08         )
09     );
10 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 KeyValueCollection that ExecuteQuery returns will contain a value for PlayerID. ☐ Yes ☐ No

Answer:

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. ☐ Yes ☐ No

All records will be sent to the client. The client will display records for scores greater than or equal to 15,000. ☐ Yes ☐ No

The scoreItem.Key property of the KeyValueCollection that ExecuteQuery returns will contain a value for PlayerID. ☐ Yes ☐ No

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

Reference:

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

Question: 150

DRAG DROP

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

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

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

Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Azure Resource
HTTP(s) endpoint
Basic
Client cert

Configuration parameter	Value
Target	value
Gateway credentials	value

Answer:

Explanation:

Configuration parameter Value

Target

Azure Resource

Gateway credentials Client cert

Box 1: Azure Resource

Box 2: Client cert

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

Reference:

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

Question: 151

HOTSPOT

You are developing a .NET Core MVC application for customers to research hotels. The application will use Azure Search. The application will search the index by using various criteria to locate

documents related to hotels. The index will include search fields for rate, a list of amenities, and distance to the nearest airport.

The application must support the following scenarios for specifying search criteria and organizing results:

- Search the index by using regular expressions.
- Organize results by counts for name-value pairs.
- List hotels within a specified distance to an airport and that fall within a specific price range.

You need to configure the SearchParameters class.

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

NOTE Each correct selection is worth one point.

\$CtfW4«

Sewcft lbe index by wing rogUv wprriyan'.

Org»«» IKIXN by axuw ft* name-vita cum

Lil! MWi within j ipwr'ird dHldnce la an jnpal and l'm!

MWEfoA d WKJfe pKi rl^r

Property



Answer:

Explanation:

Scenario

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.

Property

	▼
QueryType	
OrderBy	
SearchMode	

	▼
Facets	
Filter	
SearchMode	

	▼
Order by	
Top	
Filter	

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.

Reference:

<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: 152

DRAG DROP

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

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

You need to implement an Azure Service Bus solution.

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

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:

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.

Box 1: Create a single Service Bus Namespace

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

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

Create topics.

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

Topics can have multiple, independent subscriptions.

Reference:

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

Question: 153

HOTSPOT

A company runs an international travel and bookings management service. The company plans to begin offering restaurant bookings. You must develop a solution that uses Azure Search and meets the following requirements:

- Users must be able to search for restaurants by name, description, location, and cuisine.
- Users must be able to narrow the results further by location, cuisine, rating, and family friendliness.
- All words in 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. Is Filterable.IsSortable, IsFacetable]

] s F i l t e r a b l e. I s F a c e t a b l e, R e q u i r e d]

[Is Searchable]

[IsSearchable, Required]

public string location { get; set; }

public string Phone { get; set; }

	Required]	
[IsSearchable] IsFilterable, IsFacetable, Required] [sFilterable, IsFacetable, IsSortable]		

public

[IsFiltrable, IsSortable, IsSearchable]	
[IsFilterable, IsSortable, IsFacetable]	
[IsFiltrable, IsSortable, Key]	
[IsFilteranle, IsSortable, IsSearchable, Required]	

public double Rating { get;

[IsSearchable, IsFilterable, IsFacetable	
[IsFilterbale, IsSortable, Key]	
[IsFilterable, IsSortable, IsSearchable]	
[IsFilterable, IsSortable, Key, Required]	

Answer:

Explanation:

Answer Area

[SerializePropertyNameAsCamelCase]

```
public class Restaurant
```

```
[Key, IsFilterable]
```

```
public int Restaurant Id { 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; }
```

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.

Reference:

<https://www.henkboelman.com/azure-search-the-basics/>

Question: 154

DRAG DROP

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

Management to publish APIs.

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

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

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

NOTE: Each correct selection is worth one point.

Policy types

Inbound

Outbound

Backend

Answer Area

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

Remove formatting text from responses.

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

Requirement Policy type

policy type

policy type

Answer:

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: 155

HOTSPOT

A company is developing a gaming platform. Users can join teams to play online and see leaderboards that include player statistics. The solution includes an entity named Team.

You plan to implement an Azure Redis Cache instance to improve the efficiency of data operations for entities that rarely change.

You need to invalidate the cache when team data is changed.

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

NOTE: Each correct selection is worth one point.

```
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:

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));
```

Reference:

<https://azure.microsoft.com/sv-se/blog/lap-around-azure-redis-cache-preview/>

Question: 156

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: 157

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.

Reference:

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

Question: 158

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();  
}
```

Reference:

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

Question: 159

HOTSPOT

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

The service agents of the IT department must only have view resources and create support ticket permissions to all subscriptions. 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:

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*, 'MicrosoftSupport/*" "/read* "/Microsoft. Support/"

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

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/custom-roles-powershell>

Question: 160

DRAG DROP

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

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

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

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-AzureRmVM0isk`.

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:

Explanation:

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

Run the Azure PowerShell command `set-AzmeRmVMDSDisk`.

Run the Azure PowerShell command set-

Azur eRmVHD i s kE no ryp t i o nExt e ns i an

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.

Reference:

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

Question: 161

DRAG DROP

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

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

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

The 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"]

Answer Area

EXPOSE ./ContosoApp/ /apps/ContosoApp

COPY ./

FROM microsoft/aspnetcore:2.0

WORKDIR /apps/ContosoApp

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

Explanation:

WORKDIR /apps/ContosoApp

COPY J

EXPOSE ./ContosoApp/ /apps/ContosoApp

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

Step 1: WORKDIR /apps/ContosoApp

Step 2: COPY ./

Answer:

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

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

Step 4: CMD powershell ./setupScript.ps1

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

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

Call setupScript.ps1 when the container is built.

Run ContosoApp.dll when the container starts.

Reference:

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

Question: 162

DRAG DROP

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

You have the following parameters:

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

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

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

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

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 mytaskSi

—job-id myjob

—command-Line ^script

az batch, job

create

—id myjob

—pool-id mypool

for i in {1..SnumberOfJobs} 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

Reference:

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

Question: 163

HOTSPOT

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

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

```

public static class OrderProcessor

[FundijnNa™("ProcessOrders")]
public static void ProcessOrders([QueueTrigger("incoming-orders")](loudQueueMessage myQueueItem, (Tabb("Orders")) [tollod r <Order> tableBindings, Trate-Writer log)
{
    log.Info("Processing Order: {myQueueItem.Id}");
    log.Info("Queue Insertion Time: {myQueueItem.InsertionTime}");
    log.Info("Queue Expiration Time: {myQueueItem.ExpirationTime}");
    tableBindings.Add(orkk onvert.Deserializerject<order>(myQueueItemfl.AsString));
    (Furn 11onNa ("ProcessOrders-Poison"))
    public static void Proce5sFailedOrders((OJ.?u.: iriEger("incoming-orders-poison"))CloidQue s:-:ssagr myQueueItem, TraceWrite log)
    {
        log.Error($"Failed to process order: {myQueueItem.AsString}");
    }
}

```

NOTE: Each correct selection is worth one point.

Yes No

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

O

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

Q

When there are multiple orders in the queue, a batch of orders will be Q

Q

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 O
in Azure Table Storage.

O

Answer:

Explanation:

Yes No

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

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

Reference:

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

Question: 164

DRAG DROP

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

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

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

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

Azure CLI Commands

az group create

az group update

az webapp update

az webapp create

az appservice plan create

Answer Area



Answer:

Explanation:

az group create

az appservice plan create

1 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.

Reference:

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

Question: 165

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.

Incorrect Answers:

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

Reference:

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

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

Question: 166

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 (

Reference:

<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: 167

HOTSPOT

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

A rule already exists to scale up the App Service when the average queue length of unprocessed and valid queue messages is greater than 1000.

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

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

NOTE: Each correct selection is worth one point.

Answer Area

Scale rule



Metric

Storage queue

Service Bus queue Current

resource Storage queue

(classic) Resource type

Service Bus Namespaces

Resource

MessageQueue103

* Queue item

Criteria

* Metric

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

Action

* Operation

Explanation:

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

* 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: 168

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

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

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

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

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

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

Answer:

Explanation:

gitrepo=https://github.com/Contoso/webapp
web appname=bus i ne s webapp

resourcegroupname=BusinessAppResourceGroup

az

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

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

az

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

Box 1: group

Create a resource group.

az

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

az group create --location westeurope --name myResourceGroup

Box 2: appservice plan

az

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

az

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

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

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

Box 3: webapp

Create a web app.

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

-plan \$webappname

Box 4: webapp deployment slot

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

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

```
# -slot staging
```

Box 5: webapp deployment source

```
# Deploy sample code to "staging" slot from GitHub.
```

```
az webapp deployment source config --name $webappname --resource-group myResourceGroup \
```

```
# -slot staging --repo-url $gitrepo --branch master --manual-integration
```

Reference:

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

Question: 169

DRAG DROP

You have a web app named MainApp. You are developing a triggered App Service background task by using the WebJobs SDK. This task automatically invokes a function 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:

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

Reference:

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

Question: 170

HOTSPOT

A company is developing a mobile app for field service employees using Azure App Service Mobile Apps as the backend.

The company's network connectivity varies throughout the day. The solution must support offline use and synchronize changes in the background when the app is online app.

You need to implement the solution.

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

NOTE: Each correct selection is worth one point.

```
var client = new MobileServiceClient("MOBILE APPURL"); 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.UpdateAsyncQ;
```

Answer:

Explanation:

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

```
var todoTable = client.GetSyncTable<ToDoItem>();
var todoTable = client.GetTable<ToDoItem>();
var todoTable = client.SyncTable;
var todoTable = client.Table;
```

```
await client.SyncContext.PushAsync();
```

```
await todoTable.PullAsync("allToDoItems",todoTable.CreateQuery());
await todoTable.UpdateAsync();
todoTable.PullAsync("allToDoItems", todoTable.CreateQuery());
todoTable.UpdateAsync();
```

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();
```

Reference:

<https://azure.microsoft.com/es-es/blog/offline-sync-for-mobile-services/>

Question: 171

DRAG DROP

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

You must automatically move blocks to Archive tier after they have not been accessed for 180 days. The path for any item that is not archived must be placed in an existing queue. This operation must be performed automatically once a month.

You set the value of TierAgeInDays to 180.

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

NOTE: Each correct selection is worth one point.

Triggers and Action

Insert Entity		0
Table	processing	X
Entity	PathX	
Show advanced options		

Tier blob		0
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 Connection Change connection	

1 Recurrence	
Interval	Frequency
	Month
Show advanced options	

Answer Area

{x}	Set tier age variable	

	Set tier age variable	

For each	
Scan all blobs in this folder	
* Select an output from previous steps	
0 value x	

When there are messages in a queue	
Queue Name	processing
Show advanced options	
Connected to tableStorageAccountConnection Change connection	

if true



[Add an action](#)

[Add an action](#)

Answer

Explanation:

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

I When there are messages in a queue

Queue Name processing

Show advanced options

Connected to tableStorageAccountConnection. [Change connection](#)

If true

Recurrence

Interval

Frequency

Month

Show advanced options v

X

Box 1: Recurrence

Box 2: Insert Entity

Box 3 (if true): Tier Blob

Box 4: (if false):

Leave blank.

Reference:

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

Question: 172

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

A

Save

iscard 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
Sales LT	Customer	First Name
Sales LT	Customer	LastName
Sales LT	Customer	Email Address
Sales LT	Customer	Phone
Sales LT	CustomerAddress	AddressID

Reference:

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

Question: 173

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=2.0> 14-04-01
- 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.

Reference:

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

Question: 174

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.

Reference:

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

Question: 175

HOTSPOT

You plan to deploy a new application to a Linux virtual machine (VM) that is hosted in Azure.

The entire VM must be secured at rest by using industry-standard encryption technology to address organizational security and compliance requirements.

You need to configure Azure Disk Encryption for the VM.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
HT prn?jd'r rjgiitcr -n l'-ic'nsoft. FTyVan 1 T
re i DU r ^ -v " - up- "my 1!&*QU '■ ttGrOup "
ai gr^ut Create --nam= ir±.CL teAPaup —ilotalian .TiLuL
keye 5ul: _'ife-nryv aul tnaneJ RAJJDO"
```

```
WB keyvauit
keyvauit key
■.m erioiyplion
'--y.silw-1
ur-
```

```
vm
kewau l
-rTysij τ
```

BZ

```
vu
neyvau't
keyvault key
vm
enayptic
```

```
keyvault
keyvauit
key
```

--volume-type

```
create- X
it_nint X
jp ^rrspurregroup \ Lus \
iisk-encryption "rue
```

* create X

(Byvault_n6M k

Create X

```
jp Jreiou' cearcup
.
```

:al:Ubufltu5er7er:16.a4-U5:late5t

```
* rnnh> \
ip S^tfiour tegr'oup \
```

```
.on-keyvault Skeyvaultjname \ iH-^,
Xairrl \
```



Answer:

Explanation:

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



vm

keyvault

keyvault key

vm encryption

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



vm

keyvault

keyvault key

vm encryption

- -vault-name \$keyvault name\ - -name Name1 \
- -protection software



vm

keyvault

keyvault key

vm encryption

- -resource -group Sresourcegroup \
- -name Name 2
- -image Canonical:UbuntuServer:16.04=LTS:latest \
- -admin-username azureuser \
- -generate-ssh-keys \
- -data-disk-sizes-gb 5

▼ 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 an error in the question. Should have enable instead of create.

Box 5: all

Encrypt both data and operating system.

Reference:

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

Question: 176

DRAG DROP

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:

Explanation:

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.

Yes

No

Yes

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.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

Question: 177

DRAG DROP

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

You need to identify trends in application usage.

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

NOTE: Each correct selection is worth one point.

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

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 users decision to purchase a product?	Impact
Which events most influence a users 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?

Reference:

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

Question: 178

HOTSPOT

You are validating the configuration of an Azure Search indexer.

The service has been configured with an indexer that uses the Import Data option. The index is configured using options as shown in the Index Configuration exhibit. (Click the Index Configuration tab.)

Import data

OMS

tablesource

Costive Swth

Add cognitive skillH (Optional)

Customize target index

Import your data

X

War

Mp^MIM.

cy (* MM* the h^MI yew Mat M^4 I.e^V^") + sMK. W W^* *MI + u"

azuretable-index


RowKey

a

PartitcnKey	UriMaf v . ' ✓ v	.M
RowKey	llnVm)	...
InStockCount	Eorru:	...
SemDescription	UrID=TFm	...
temName	Idm-rt^ 3/ v	...
LocatcnRow	Earns	...
LocatonShelf	eauff	...
SKU	IMvUs^xj	...
	(MM*) v	

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.

RowKey	Timestamp	InStockCount	ItemDescription	ItemName	LocationRow	LocationShelf	SKU
Food	2018-08-25T15:47:26.123Z	2	A box of chocolate candy bars	Chocolate	5	3	573421
Hardware	2018-08-25T15:48:38.402Z	2	A bag of bolts	Bolts	1	4	878964
Hardware	2018-08-25T15:48:41.402Z	20	A box of nails	Nails	2	1	654388



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 bag of to return the correct results

Answer:

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, Q and selecting the SEARCHABLE option for the field.

You can resolve the issue by running the indexer.

O

O

You can resolve the issue by changing the query string in Search explorer to _
bag of to return the correct results

H

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

Reference:

<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: 179

HOTSPOT

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:

Explanation:

Code segment

Value

Attribute

Attribute
Authorize
AllowAnonymous
AutoValidateAntiforgeryToken

Request Header

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

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",  
}
```

Reference:

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

Question: 180

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.

Reference:

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

Question: 181

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.

Reference:

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

Question: 182

HOTSPOT

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

```

fgitrepo="http://github.com/TailSpinToy/webapp" fwebappname-"TailSpinToycekb" fiocation-"westUS2"

New AzWebAppSlot New* AIW -Name myResourceGroup -Location flotation

tb App NewAzAppServicePlsn

New-AzResourceGroup

NewAzWebAppSlot New-Az
Web App New -Name fwebappname "Location flotation -ReiOurceGroupName myReiourceGroup -Tier Standard
AzAppServicePlan
NewAzResourceGroup

NewAiWebAppSlot New-
AzWeb App -Name fwebappnaiae "Location flotation -AppServiceRlan fwebappname "RetOurctGroupName myResourteGroup

New AzAp p serv t cePl an
NewAzReiourfteGroup

New AzWebAppSloi New - Az
Web App -Name fwebappname -ResourceGroupName myResourceGroup -Slot review

NewAzAppServicePlan New-Az
Resourced roup

^PropertiesObject - ^{repoUrl - "fgitrepo";branch - "master";}
Set-AzResource -PropertyObject fPropertiesObject - RoovrccGrowpName myResourceGroup -ResourceType Microsoft.Web/sites/slots/sourcecontrols -ReaourceName
fwebappname/review/web -ApiVersion 2015-0801 -Force Switch-AzWebAppSlot Name fwebappname - ResourceGroupName BryRetOurceGroup ' "SourceSlotHaae review -
DestinationSlotName production

```

Answer:

Explanation:

```

Hitrepo="https://github.com/TailSpinToy9/webapp" fwebappname-"TailSpinToyeWeb" fiocation-"westUS2"

New- AzWebAppSlot New* -Name myRetourceGroup -Location flotation

AzWeb App
NewAzAppServicePlan
NewAjResourceGroup

New AzWebAppSlot New Az
Web App -Name fwebappname "Location flotation -ResourteGroupName myRrsourcergroup -Tier Standard
New-AzAppServicePbn New-
AzResourceG roup

New-Az Web App Slot New Az -Name fwebappname -Location flotation -AppServicePlan fwebappname -Re»ourteGroupNa*e myRetourceOroup
Web App Ne w * A Z A p p
ServiceR I an
NewAzRewurceGfOug

New - AzWebAppSlot New -Name fwebappname -ResourceOroupName myResourceGroup -Slot review
AzWebApp
NewAzAppServicePlan New-
AzReiOurceGfgup

fPropertiesObject -E{repoUrl - "fgltrepo";branch - "master";}
Set-AzResource -PropertyObject SPropertiesObject "ResourceGroupName myResourceGroup -RcsourceType Microsoft.Web/sites/slots/sourcecontrols -RelourceName
fwebappname/review/web -ApiVersion 2015 08 01 -Force Switch-AzWebAppSlot -Name fwebappname -ReiourceGroupName myRetourceGroup ' -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.

Reference:

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

<https://docs.microsoft.com/en-us/powershell/module/az.websites/new-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: 183

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: 184

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: 185

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:

Reference:

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

Question: 186

HOTSPOT

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

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

You need to save the customized VM for future provisioning.

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

NOTE: Each correct selection is worth one point.

Answer Area

Action
Generalize the VM.

Tool or service

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:

Explanation:

Action


Generalize the VM.

Tool or service



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

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource#create-an-image-of-a-vm-using-powershell>

Question: 187

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

App service plan setting

Value

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

Answer:

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).

Reference:

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

Question: 188

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).

Reference:

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

Question: 189

DRAG DROP

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

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

You need to deploy the application.

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

NOTE: Each correct selection is worth one point.

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:

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
```

Reference:

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

Question: 190

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

Reference:

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

Question: 191

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.

Reference:

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

Question: 192

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.

Reference:

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

Question: 193

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.

Reference:

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

Question: 194

HOTSPOT

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:

Explanation:

```
{
  ...
  "appId": "d61126e3-089b-4adb-b721-
d5023213df7d",
```

```
    : "All",
    "optionalClaims": {
      "groupMembershipClaims":
```

```
      : true
      "allowPublicClient"
      "oauth2Permissions"
      "requiredResourceAccess"
      "oauth2AllowImplicitFlow"
    }
}
```

Box 1: groupMembershipClaims

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

Group claims can also be configured in the Optional Claims section of the Application Manifest. Enable group membership claims by changing the groupMembershipClaim

The valid values are:

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

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

Box 2: oauth2AllowImplicitFlow

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

auth2Permissions can only accept collections value like an array, not a boolean.

oauth2AllowImplicitFlow accepts boolean value.

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

Question: 195

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.

Reference:

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

Question: 196

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.

Reference:

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

Question: 197

DRAG DROP

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

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

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

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

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

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

NOTE: Each correct selection is worth one point.

Settings

client_id

delegated

profile

application

user_impersonation

Answer Area

API

Azure Storage

Microsoft Graph

Permission

Setting

User.Read

Type

Setting

Setting

Answer:

Explanation:

API

Azure Storage

Microsoft Graph

Permission

user_impersonation

User.Read

Type

delegated

delegated

Box 1: user_impersonation

Box 2: delegated

Example:

1. Select the API permissions section
2. Click the Add a permission button and then:

Ensure that the My APIs tab is selected

3. In the list of APIs, select the API TodoListService-aspnetcore.
4. In the Delegated permissions section, ensure that the right permissions are checked: user_impersonation.
5. Select the Add permissions button.

Box 3: delegated

Example

1. Select the API permissions section
2. Click the Add a permission button and then,

Ensure that the Microsoft APIs tab is selected

3. In the Commonly used Microsoft APIs section, click on Microsoft Graph
4. In the Delegated permissions section, ensure that the right permissions are checked: User.Read. Use the search box if necessary.
5. Select the Add permissions button

Reference:

<https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect-aspnetcore/calling-a-web-api-in-an-aspnet-core-web-application-using-azure-ad/>

Question: 198

HOTSPOT

You plan to deploy a web app to App Service on Linux. You create an App Service plan. You create and push a custom Docker image that 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:

Explanation:

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

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

Reference:

<https://docs.microsoft.com/en-us/cli/azure/webapp/log>

Question: 199

HOTSPOT

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

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

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

How should you complete the trigger?

NOTE: Each correct selection is worth one point.

```
function ensureTip() { var r = 1
    | _value(); _.readDocument('item');
    getContext().getRequest();
    getContext().getResponse():
    var i = r.getBody();
```

```
    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:

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

Reference:

<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: 200

HOTSPOT

You are developing an application that use an Azure blob named data to store application data. The application creates blob snapshots to allow application state to be reverted to an earlier state. The Azure storage account has soft deleted enabled.

The system performs the following operations in order:

- The blob is updated
- Snapshot 1 is created.
- Snapshot 2 is created.
- Snapshot 1 is deleted.

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:

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.

Reference:

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

Question: 201

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.

E. 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.

Reference:

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

[https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment- package](https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package)

Question: 202

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)
```

Reference:

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

Question: 203

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: A

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: 204

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: 205

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: 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: 206

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-vm-access-arm>

Question: 207

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-vm-access-arm>

Question: 208

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-vm-access-arm>

Question: 209

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: 210

HOTSPOT

You are developing an application that needs access to an Azure virtual machine (VM). The access lifecycle for the application must be associated with the VM service instance. You need to enable managed identity for the VM.

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


NOTE Each correct selection is worth one point.



Answer:

Explanation:

Answer:



```
$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM
```

```
Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -AssignIdentity:$SystemAssigned
```

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

Question: 211

You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms. If the stored intake forms are downloaded from storage by a third party, the content of the forms must not be compromised.

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

Question: 212

You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms. If the stored intake forms are downloaded from storage by a third party, the content of the forms must not be compromised.

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

Solution:

Create a Azure Key Vault key named skey.

Encrypt the intake forms using the public key portion of skey.

Store the encrypted data in Azure Blob storage

Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:

Question: 213

You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms. If the stored intake forms are downloaded from storage by a third party, the content of the forms must not be compromised.

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

Solution:

Create an Azure Cosmos DB database with Storage Service Encryption enabled.

Store the intake forms in the Azure Cosmos DB database.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Question: 214

DRAG DROP

You are developing an application to securely transfer data between on-premises file systems and Azure Blob storage.

The application stores keys, secrets, and certificates in Azure Key Vault. The application uses the Azure Key Vault

APIs.

The application must allow recovery of an accidental deletion of the key vault or key vault objects.
Key vault objects must be retained for 90 days after deletion.

You need to protect the key vault and key vault objects.

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

NOTE: Each correct selection is worth one point.

Features

Access policy

Purge protection

Soft delete

Shared access signature

Answer Area

Enable retention period and accidental deletion.

Enforce retention period and accidental deletion.

Action

Feature

Answer:

Explanation:

Action

Feature

Enable retention period and accidental deletion.

Soft delete

Enforce retention period and accidental deletion.

Purge protection

Box 1: Soft delete

When soft-delete is enabled, resources marked as deleted resources are retained for a specified period (90 days by default).

The service further provides a mechanism for recovering the deleted object, essentially undoing the deletion.

Box 2: Purge protection

Purge protection is an optional Key Vault behavior and is not enabled by default. Purge protection can only be enabled once soft-delete is enabled.

When purge protection is on, a vault or an object in the deleted state cannot be purged until the retention period has passed. Soft-deleted vaults and objects can still be recovered, ensuring that the retention policy will be followed.

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/general/soft-delete-overview>

Question: 215

DRAG DROP

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

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

You need to ensure that the requirements are met.

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

Features		Requirement	Feature
Run Command	Serial console	Firewall configuration	
Hypno Runtime* Worker Custom Script Intension		Supporting services script	

Answer:

Explanation:

Requirement	Feature
Firewall configuration	Run Command
Supporting services script	hybrid Run book Worker

Reference:

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

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

Question: 216

DRAG DROP

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

The application must meet the following requirements:

Reference the Azure Key Vault without requiring any changes to the Java code.

Dynamically add and remove instances of the Azure Functions host based on the number of incoming application events.

Ensure that instances are perpetually warm to avoid any cold starts.

Connect to a VNet.

Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.

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

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

Actions

Answer Area

Create a user-assigned managed identity for the application.

Create the Azure Functions app with a Premium plan type.

Create an access policy in Azure Key Vault for the application identity.

Create an SSL certification in Azure Key Vault for the application identity.

Create the Azure Functions app with an App Service plan type.

Create the Azure Functions app with a Consumption plan type.

Create a system-assigned managed identity for the application.



Answer:

Explanation:

Create the Azure Functions app with a Consumption plan type.

Create a user-assigned managed identity for the application.

Create an access policy in Azure Key Vault for the application identity.

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

Use the Consumption plan for serverless.

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

Create a system-assigned managed identity for your application.

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

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

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

Reference:

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

Question: 217

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

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

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

output binding on the blob.

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

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

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

Reference:

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

Question: 218

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

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

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

Question: 219

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

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

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

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

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

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

Reference:

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

Question: 220

HOTSPOT

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

```
01 CloudBlockBlob src = null;
02 try
03 {
04     src = container.ListBlobs().OfType<CloudBlockBlob>().FirstOrDefault();
05     var id = await src.AcquireLeaseAsync(null);
06     var dst = container.GetBlockBlobReference(src.Name);
07     string cpid = await dst.StartCopyAsync(src);
08     await dst.FetchAttributeAsync();
09     return id;
10 }
11 catch (Exception e)
12 {
13     throw;
```

```
14 I
15 finally
16 I
17 if (src != null)
18     await src.FetchAttributesAsync();
19 if (src.Properties.LeaseState != LeaseState.Available)
20     await src.BreakLeaseAsync(new TimeSpan(0));
21 }
```

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

NOTE: Each correct selection is worth one point.

Statement

Yes

No

The code creates an infinite lease

☐☐

The code at line 06 always creates a new blob

☐

The finally block releases the lease

☐☐

Answer:

Explanation:

Statement

Yes

No

The code creates an infinite lease

The code at line 06 always creates a new blob

The finally block releases the lease

Box 1: Yes

AcquireLeaseAsync does not specify leaseTime.

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

Box 2: No

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

Box 3: Yes

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

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.getblockblobreference>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync>

Question: 221

D18912E1457D5D1DDCBD40AB3BF70D5D

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

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

You need to document the minimum SLA for data recovery.

Which SLA should you use?

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

Answer: B

Explanation:

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

Reference:

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

Question: 222

You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage containers named Container1 and Container2.

Uploading of videos occurs on an irregular basis.

You need to copy specific blobs from Container1 to Container2 when a new video is uploaded.

What should you do?

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

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.

Reference:

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

Question: 223

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

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

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

Which type of cache purge should you use?

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

Answer: A

Explanation:

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

Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example, /pictures/strasbourg.png;

Wildcard purge: Asterisk (*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /*, for example, /pictures/*.

Root domain purge: Purge the root of the endpoint with "/" in the path.

Reference:

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

Question: 224

HOTSPOT

You are debugging an application that is running on Azure Kubernetes cluster named cluster1. The cluster uses Azure Monitor for containers to monitor the cluster.

The application has sticky sessions enabled on the ingress controller.

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

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

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

NOTE: Each correct selection is worth one point.

let startTimestamp =

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

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

top ContainerID
union ContainerID
sample ContainerID
distinct ContainerID

ContainerLog

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

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

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

Answer:

Explanation:

let startTimestamp =

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

let ContainerIDs = Ku bePod Inventory | where ClusterName == 'ulusterr

top ContainerID union ContainerID sample ContainerID distinct ContainerID

ContainerLog

fork containerIDs

where ContainerID in (ContainerIDs) restrict ContainerID in (ContainerIDs) join

ContainerID == ContainerIDs.ContainerID

| where TimeGenerated > startTimestamp | where LogEntrySource == 'stderr'¹

project by Computer summarize by Computer

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

Box 1: ago(1d)

Box 2: distinct containerID

Box 3: where ContainerID in (ContainerIDs)

Box 4: summarize Count by Computer

Summarize: aggregate groups of rows

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

Reference:

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

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

Question: 225

You develop and deploy an ASP.NET web app to Azure App Service. You use Application Insights telemetry to monitor the app.

You must test the app to ensure that the app is available and responsive from various points around the world

and at regular intervals. If the app is not responding, you must send an alert to support staff.

You need to configure a test for the web app.

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

NOTE: Each correct selection is worth one point.

- A. integration
- B. multi-step web
- C. URL ping
- D. unit
- E. load

Answer: B,C

Explanation:

There are three types of availability tests:

URL ping test: a simple test that you can create in the Azure portal.

Multi-step web test: A recording of a sequence of web requests, which can be played back to test more complex scenarios.

Multi-step web tests are created in Visual Studio Enterprise and uploaded to the portal for execution.

Custom Track Availability Tests: If you decide to create a custom application to run availability tests, the `TrackAvailability()` method can be used to send the results to Application Insights.

Reference:

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

Question: 226

DRAG DROP

A web service provides customer summary information for e-commerce partners. The web service is implemented as an Azure Function app with an HTTP trigger. Access to the API is provided by an Azure API Management instance. The API Management instance is configured in consumption plan mode. All API calls are authenticated by using OAuth.

API calls must be cached. Customers must not be able to view cached data for other customers.

You need to configure API Management policies for caching.

How should you complete the policy statement?

Targets

Expect
Public
Private
Internal
External
Authorization

Answer Area

```
<policies>  
<inbound>  
<base />  
<cache-lookup caching-type="Target" downstream-caching-type = "Target" >  
  <vary-by-header>  
    Target  
  </vary-by-header>  
</cache-lookup>  
</inbound>  
</policies>
```

Answer:

Explanation:

```
<policies>  
<inbound>  
<base />  
<cache-lookup caching-type="Internal" downstream-caching-type = "Private" >  
  <vary-by-header>  
    Authorization  
  </vary-by-header>  
</cache-lookup>  
</inbound>  
</policies>
```

Box 1: internal

caching-type

Choose between the following values of the attribute:

internal to use the built-in API Management cache,

external to use the external cache as Azure Cache for Redis

prefer-external to use external cache if configured or internal cache otherwise.

Box 2: private

downstream-caching-type

This attribute must be set to one of the following values.

none - downstream caching is not allowed.

private - downstream private caching is allowed.

public - private and shared downstream caching is allowed.

Box 3: Authorization

<vary-by-header>Authorization</vary-by-header>

<!-- should be present when allow-private-response-caching is "true"-->

Note: Start caching responses per value of specified header, such as Accept, Accept-Charset, Accept-Encoding, Accept-Language, Authorization, Expect, From, Host, If-Match

Reference:

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

Question: 227

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

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

You need to ensure that the subscription client processes all messages.

Which code segment should you use?

A. `await subscriptionClient.AddRuleAsync(new RuleDescription(RuleDescription.DefaultRuleName, new TrueFilter()));`

- B. `subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName); D18912E1457D5D1DDCBD40AB3BF70D5D`
- C. `await subscriptionClient.CloseAsync();`
- D. `subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);`

Answer: D

Explanation:

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

`subscriptionClient.RegisterMessageHandler(ReceiveMessagesAsync, messageHandlerOptions);`

Reference:

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

Question: 228

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

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

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

You have the following requirements:

Queue size must not grow larger than 80 gigabytes (GB).

Use first-in-first-out (FIFO) ordering of messages.

Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application.

Create an Azure VM that is triggered from Azure Storage Queue events.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Don't use a VM, instead create an Azure Function App that uses an Azure Service Bus Queue trigger.

Reference:

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

Question: 229

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

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

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

You have the following requirements:

Queue size must not grow larger than 80 gigabytes (GB).

Use first-in-first-out (FIFO) ordering of messages.

Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Service Bus Queue from the mobile application. Create an Azure Windows VM that is triggered from Azure Service Bus Queue.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Don't use a VM, instead create an Azure Function App that uses an Azure Service Bus Queue trigger.

Reference:

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

Question: 230

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

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

What should you use?

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

Answer: A

Explanation:

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

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

Reference:

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

Question: 231

HOTSPOT

You are preparing to deploy a Python website to an Azure Web App using a container. The solution will use multiple containers in the same container group. The Dockerfile that builds the container is as follows:

```
FROM python:3
ADD website.py
CMD [ "python", "./website.py"]
```

You build a container by using the following command. The Azure Container Registry instance named images is a private registry.

```
docker build -t images.azurecr.io/website:v1.0.0
```

The user name and password for the registry is admin.

The Web App must always run the same version of the website regardless of future builds.

You need to create an Azure Web App to run the website.

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

NOTE: Each correct selection is worth one point.

```
az configure --defaults web=website
az configure --defaults group=website
az appservice plan create --name websitePlan
```

V

-sku SHARED '
-tags container
-sku B1 -hyper-v
-sku B1 -is-hnux

```
az webapp create --plan websitePlan
```

y

-deployment-source-uri images.azurecr.io/websitev1.0.0
-deployment-source-url images.azurecr.io/website latest -
deployment-container-image-name images.azurecr
io/website v1.0.0 -deployment-container-image-name images
azurecr.io website latest

```
az webapp config
```

y

set --python-version 2.7 --generic-configurations user=admin password=admin
set --python-version 3.6 --generic-configurations user=admin password=admin

container set -docker-registry-server-url https //images azurecr io -u admin -p admin
admin container set -docker-registry-seiver-url https //images azurecr io/wsebsite -u admin -p admin

Answer:

Explanation:

```
az configure --defaults web=website  
az configure --defaults group=website  
az appservice plan create --name websitePlan
```

-Sku SHARED¹

-tags container

-sku B1 -hyper-v

-sku B1 -is-hnux

```
az webapp create --plan websitePlan
```

-deployment-source-url images.azurecr.io/website:v1.0.0

-deployment-source-url images azurecr io website latest

-deployment-container-name images azurecr i o
'website v1 00

-deployment-container-image-name images azurecr
io/website latest

```
az webapp config
```

set -python-version 2.7 -generic-configurations user=admin password=admin

set -python-version 3.6 -generic-configurations user=admin password=admin

container set -docker-registry-server-url https //images azurecr io -u admin -p

admin container set -docker-registry-server-url https ' images azurecr io wsebsite -u

admin -p admin

Box 1: --SKU B1 --hyper-v

--hyper-v

Host web app on Windows container.

Box 2: --deployment-source-url images.azurecr.io/website:v1.0.0

--deployment-source-url -u

Git repository URL to link with manual integration.

The Web App must always run the same version of the website regardless of future builds.

Incorrect:

```
--deployment-container-image-name -i
```

Linux only. Container image name from Docker Hub, e.g. publisher/image-name:tag.

Box 3: az webapp config container set -url https://images.azurecr.io -u admin -p admin

az webapp config container set

Set a web app container's settings.

Parameter: --docker-registry-server-url -r

The container registry server url.

The Azure Container Registry instance named images is a private registry.

Example:

```
az webapp config container set --docker-registry-server-url https://{azure-container-registry-name}.azurecr.io
```

Reference:

<https://docs.microsoft.com/en-us/cli/azure/appservice/plan>

Question: 232

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

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

You need to restrict access to the Azure AD group.

Which role-based access control should you use?

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

Answer: C

Explanation:

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

Reference:

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

Question: 233

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

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

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

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

You need to configure authorization.

Solution:

Create a new Azure AD application. In the application's manifest, define application roles that match the required permission levels for the application.

Assign the appropriate Azure AD group to each role. In the website, use the value of the roles claim from the JWT for the user to determine permissions.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

To configure Manifest to include Group Claims in Auth Token

Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:

Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.

Locate the "groupMembershipClaims" setting. Set its value to either "SecurityGroup" or "All". To help you decide which:

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

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

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

Reference:

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

Question: 234

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

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

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

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

You have the following requirements:

Queue size must not grow larger than 80 gigabytes (GB).

Use first-in-first-out (FIFO) ordering of messages.

Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application.

Create an Azure Function App that uses an Azure Storage Queue trigger.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

Reference:

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

Question: 235

DRAG DROP

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

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

Call setupScripts.ps1 when the container is built.

Run ContosoApp.dll when the container starts.

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

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

Commands

Answer Area

FROM microsoft/aspnetcore:latest WORKDIR /apps/ContosoApp

CMD ["dotnet", "ContosoApp.dll"] COPY ./.

RUN powershell ./setupscript.psi

Answer:

Explanation:

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

FROM microsoft/aspnetcore:latest

WORKDIR /apps/ContosoApp COPY

RUN powershell ./setupScript.psi

Box 1: CMD [..]

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

Syntax: CMD <string>

Specifies the command you want to carry out.

Box 2: FROM microsoft/aspnetcore-build:latest

Box 3: WORKDIR /apps/ContosoApp

Box 4: COPY ./.

Box 5: RUN powershell ./setupScript.ps1

Question: 236

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

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

You need to configure the Function App.

What should you do?

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

Answer: B

Explanation:

The Blob storage trigger starts a function when a new or updated blob is detected. The blob contents are provided as input to the function.

The Consumption plan limits a function app on one virtual machine (VM) to 1.5 GB of memory.

Reference:

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

Question: 237

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

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

Answer:

Explanation:

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



```
"sku": { "name": "Standard LRS"
```

```
    "kind": "Storage",  
    "properties": {},
```

```
    copyindex  
    priority  
    dependsOn  
  
    "name": "storagesetup", "count": 3  
  ]
```

```
    "apiVersion": "2015-06-15",  
    "type": "Microsoft.Compute/virtualMachines",  
    "name": "[concat('VM', uniqueString(resourceGroup().id))]",  
    copy  
    copyindex  
    priority  
    dependsOn  
    "[variables('loadBalancerName')]",  
    "[variables('virtualNetworkName')]",  
    "storagesetup",
```

Box 1: copyIndex

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

Box 2: copy

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

Box 3: dependsOn

Example:

```
"type": "Microsoft.Compute/virtualMachineScaleSets",
```

```
"apiVersion": "2020-06-01",

"name": "[variables('namingInfix')]",

"location": "[parameters('location')]",

"sku": {
  "name": "[parameters('vmSku')]",

  "tier": "Standard",

  "capacity": "[parameters('instanceCount')]"
},

"dependsOn": [

  "[resourceId('Microsoft.Network/loadBalancers', variables('loadBalancerName'))]",

  "[resourceId('Microsoft.Network/virtualNetworks', variables('virtualNetworkName'))]"
],
```

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/copy-resources>

<https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/quick-create-template-windows>

Question: 238

DRAG DROP

You have an application that uses Azure Blob storage.

You need to update the metadata of the blobs.

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

Methods

Metadata.Add

SetMetadataAsync

FetchAttributesAsync

UploadFileStream

SetPropertiesAsync

Answer Area

Answer:

Explanation:

Metadata.Add

SetMetadataAsync

SetPropertiesAsync

Metadata.Add example:

```
// Add metadata to the dictionary by calling the Add method metadata.Add("docType", "textDocuments");
```

SetMetadataAsync example:

```
// Set the blob's metadata.
```

```
await blob.SetMetadataAsync(metadata);
```

```
// Set the blob's properties.
```

```
await blob.SetPropertiesAsync();
```

Reference:

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

Question: 239

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

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

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

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

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

Solution: Provision an Azure Event Grid. Configure the machine identifier as the partition key and enable capture.

Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:

Reference:

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

Question: 240

DRAG DROP

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

Storage containers.

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

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

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

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

NOTE: Each correct selection is worth one point.

Settings

client_id
profile
delegated
application
user_impersonation

Answer Area

API	Permission	Type
Azure Storage		
Microsoft Graph User.Read		

Answer:

Explanation:

API

Azure Storage

Microsoft Graph

Permission

user_impersonation

User.Read

Type

delegated

delegated

Box 1: user_impersonation

Box 2: delegated

Example:

1. Select the API permissions section
2. Click the Add a permission button and then:

Ensure that the My APIs tab is selected

3. In the list of APIs, select the API TodoListService-aspnetcore.
4. In the Delegated permissions section, ensure that the right permissions are checked: user_impersonation.
5. Select the Add permissions button.

Box 3: delegated

Example

1. Select the API permissions section
2. Click the Add a permission button and then,

Ensure that the Microsoft APIs tab is selected

3. In the Commonly used Microsoft APIs section, click on Microsoft Graph
4. In the Delegated permissions section, ensure that the right permissions are checked: User.Read. Use the search box if necessary.
5. Select the Add permissions button

Reference:

<https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect-aspnetcore/calling-a-web-api-in-an-aspnet-core-web-application-using-azure-ad/>

Question: 241

HOTSPOT

You are developing an ASP.NET Core app that includes feature flags which are managed by Azure App Configuration. You create an Azure App Configuration store named AppFeatureFlagStore that contains a feature flag named Export.

You need to update the app to meet the following requirements:

Use the Export feature in the app without requiring a restart of the app.

Validate users before users are allowed access to secure resources.

Permit users to access secure resources.

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 void Configure(IApplicationBuilder app, IWebHostEnvironment env) {  
    if (env.IsDevelopment())  
    {  
        app.UseDeveloperExceptionPage();  
    }  
    else  
    {  
        app.UseExceptionHandler("/Error");  
    }  
}
```

app.UseExceptionHandler("/Error");

UseAuthentication

UseStaticFiles

UseSession

UseCookiePolicy


```

app- * 0;
    UseAuthorization
    UseHttpsRedirection
    UseSession
    UseCookiePolicy

app- _____ * ();
    UseAzureAppConfiguration
    UseRequestLocalization _____
    UseCors
    UseStaticFiles
app.UseEndpoint(endpoints =>
    endpoints.MapRazorPages());
}

```

Answer:

Explanation:

```
}
```

```
app.
```

```
();
```

UseAuthentication

UseStaticFiles

UseSession

UseCookiePolicy

```
app.
```

```
();
```

UseAuthorization

UseHttpsRedirection

UseSession

UseCookiePolicy

```
app.
```

```
();
```

UseAzureAppConfiguration

UseRequestLocalization

UseCors

UseStaticFiles

```
app.UseEndpoint(endpoints =>
```

```
{
```

```
    endpoints.MapRazorPages();
```

```
});
```

```
}
```

Box 1: UseAuthentication

Need to validate users before users are allowed access to secure resources.

UseAuthentication adds the AuthenticationMiddleware to the specified IApplicationBuilder, which enables authentication capabilities.

Box 2: UseAuthorization

Need to permit users to access secure resources.

UseAuthorization adds the AuthorizationMiddleware to the specified IApplicationBuilder, which enables authorization capabilities.

Box 3: UseStaticFiles

Need to use the Export feature in the app without requiring a restart of the app.

UseStaticFiles enables static file serving for the current request path

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.aspnetcore.builder.iapplicationbuilder?view=aspnetcore-5.0>

Question: 242

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

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

You need to design the approach to loading application secrets.

What should you do?

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

Answer: C

Explanation:

Use Key Vault references for App Service and Azure Functions.

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

Reference:

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

Question: 243

DRAG DROP

A company has multiple warehouse. Each warehouse contains IoT temperature devices which deliver temperature data to an Azure Service Bus queue.

You need to send email alerts to facility supervisors immediately if the temperature at a warehouse goes above or below specified threshold temperatures.

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

Actions

Add a logic app trigger that fires when one or more messages arrive in the queue.

Add a Recurrence trigger that schedules the app to run every 15 minutes.

Add an action that sends an email to specified personnel if the temperature is outside of those thresholds.

Add a trigger that reads IoT temperature data from a Service Bus queue.

Add a logic app action that fires when one or more messages arrive in the queue.

Add a condition that compares the temperature against the upper and lower thresholds.

Create a blank Logic app.

Add an action that reads IoT temperature data from the Service Bus queue.

Answer Area

Answer:

Explanation:

Create a blank Logic app.

Add a logic app action that fires when one or more messages arrive in the queue.

Add an action that reads IoT temperature data from the Service Bus queue.

Add a condition that compares the temperature against the upper and lower thresholds.

Add an action that sends an email to specified personnel if the temperature is outside of those thresholds.

Step 1: Create a blank Logic app.

Create and configure a Logic App.

Step 2: Add a logical app trigger that fires when one or more messages arrive in the queue.

Configure the logic app trigger.

Under Triggers, select When one or more messages arrive in a queue (auto-complete).

Step 3: Add an action that reads IoT temperature data from the Service Bus queue

Step 4: Add a condition that compares the temperature against the upper and lower thresholds.

Step 5: Add an action that sends an email to specified personnel if the temperature is outside of

those thresholds

Reference: [https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-monitoring-notifications-with- azure-logic-apps](https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-monitoring-notifications-with-azure-logic-apps)

Question: 244

DRAG DROP

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

A driver selects the restaurants from which they will deliver orders.

Orders are sent to all available drivers in an area.

Only orders for the selected restaurants will appear for the driver.

The first driver to accept an order removes it from the list of available orders.

You need to implement an Azure Service Bus solution.

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

Actions

Answer Area

Create a single Service Bus topic.

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

Create a single Service Bus subscription.

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

Create a single Service Bus Namespace.

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

Answer:

Explanation:

Create a single Service Bus Namespace.

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

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

Box 1: Create a single Service Bus Namespace

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

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

Create topics.

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

Topics can have multiple, independent subscriptions.

Reference:

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

Question: 245

HOTSPOT

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

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

You need to implement push notifications.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
string notificationHubName = "contosoHub";
string notificationHubConnection = "connectionstring";

NotificationHubClient client =
    NotificationHubClient.GetInstallation(
        notificationHubName,
        notificationHubConnection,
        NotificationHubClientSettings.Default,
        NotificationHubJob.SubmitNotification,
        NotificationDetails.Default);

await client.CreateClientFromConnectionString(
    notificationHubName,
    notificationHubConnection,
    NotificationHubClientSettings.Default,
    NotificationHubJob.SubmitNotification,
    NotificationDetails.Default);

await client.PatchInstallation(
    notificationHubName,
    notificationHubConnection,
    NotificationHubClientSettings.Default,
    NotificationHubJob.SubmitNotification,
    NotificationDetails.Default);

string toastPayload =
    @"<toastxvisualxbinding template=""ToastText01""xtext id=""1"">" +
    @"New item to view" + @"</textx/bindingx/visualx/toast>";

try
{
    var result =
        await hub.
            SendWindowsNativeNotificationAsync(
                SubmitNotificationHubJobAsync,
                ScheduleNotificationAsync,
                SendAppleNativeNotificationAsync);
}
catch (System.Exception ex) { }
```

Answer:

Explanation:

NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails
NotificationHubClient
NotificationHubClientSettings
NotificationHubJob
NotificationDetails

GetInstallation
CreateClientFromConnectionString
CreateOrUpdateInstallation
PatchInstallation

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

Box 1: NotificationHubClient

Box 2: NotificationHubClient

Box 3: CreateClientFromConnectionString

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

Box 4: SendWindowsNativeNotificationAsync

Send the push notification.

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

Reference:

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

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

Question: 246

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

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

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

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

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

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

Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:

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

Reference:

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

Question: 247

You develop Azure solutions.

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

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

Which object should you use?

A. QueueClient

B. SubscriptionClient

C. TopicClient

D. CloudQueueClient

Answer: A

Explanation:

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

Reference:

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

Question: 248

DRAG DROP

You are maintaining an existing application that uses an Azure Blob GPv1 Premium storage account. Data older than three months is rarely used.

Data newer than three months must be available immediately. Data older than a year must be saved but does not need to be available immediately.

You need to configure the account to support a lifecycle management rule that moves blob data to archive storage for data not modified in the last year.

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

Actions

Answer Area

Upgrade the storage account to GPv2

Create a new GPv2 Standard account and set its default access tier level to cool

Change the storage account access tier from hot to cool

Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account



Answer:

Explanation:

Upgrade the storage account to GPv2

Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account

Change the storage account access tier from hot to cool

Step 1: Upgrade the storage account to GPv2

Object storage data tiering between hot, cool, and archive is supported in Blob Storage and General Purpose v2 (GPv2) accounts. General Purpose v1 (GPv1) accounts don't support tiering.

You can easily convert your existing GPv1 or Blob Storage accounts to GPv2 accounts through the Azure portal.

Step 2: Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account

Step 3: Change the storage account access tier from hot to cool

Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Only the hot and cool access tiers can be set at the account level. The archive access tier can only be set at the blob level.

Reference:

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

Question: 249

You are developing an Azure function that connects to an Azure SQL Database instance. The function is triggered by an Azure Storage queue.

You receive reports of numerous System.InvalidOperationExceptions with the following message: "Timeout expired. The timeout period elapsed prior to obtaining a connection from the pool. This may have occurred because all pooled connections were in use and max pool size was reached."

You need to prevent the exception.

What should you do?

- A. In the host.json file, decrease the value of the batchSize option
- B. Convert the trigger to Azure Event Hub
- C. Convert the Azure Function to the Premium plan
- D. In the function.json file, change the value of the type option to queueScaling

Answer: A

Explanation:

With the Premium plan the max outbound connections per instance is unbounded compared to the 600 active (1200 total) in a Consumption plan.

Note: The number of available connections is limited partly because a function app runs in a sandbox environment. One of the restrictions that the sandbox imposes on your code is a limit on the number of outbound connections, which is currently 600 active (1,200 total) connections per instance. When you reach this limit, the functions runtime writes the following message to the logs: Host thresholds exceeded: Connections.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/manage-connections>

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#service-limits>

Question: 250

DRAG DROP

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

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

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

Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

Answer Area

Azure Resource

Configuration parameter

Value

HTTP(s) endpoint

Target

Gateway credentials Basic

Client cert

Explanation:

Answer:

Configuration parameter

Value

Target

Azure Resource

Gateway credentials

Client cert

Box 1: Azure Resource

Box 2: Client cert

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

Reference:

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

Question: 251

HOTSPOT

You have a single page application (SPA) web application that manages information based on data returned by Microsoft Graph from another company's Azure Active Directory (Azure AD) instance.

Users must be able to authenticate and access Microsoft Graph by using their own company's Azure AD instance.

You need to configure the application manifest for the app registration.

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

NOTE: Each correct selection is worth one point.

```

{
  "oauth2AllowImplicitFlow": true,
  "addIns": {
    "orgRestrictions": null,
    "availableToOtherTenants": false,
    "requiredResourceAccess": [
      {
        "resourceAppId": "00000003-0000-0000-c000-000000000000",
        "resourceAccess": [
          {
            "id": "24a6cdd6-fab1-4aaf-91b8-3cc8225e90d0",
            "type": "Scope"
          }
        ]
      }
    ]
  },
  "signInAudience": "All"
}

```

Answer:

Explanation:

`"oauth2AllowImplicitFlow":`

	▼
add	
false	
spa	
true	

`"`

	▼
addIns	
orgRestrictions	
availableToOtherTenants	
requiredResourceAccess	

`"resourceAppId": "00000003-0000-0000-c000-000000000000",`

`"resourceAccess": [{`

`"id": "24a6cdd6-fab1-4aaf-91b8-3cc8225e90d0",`

`"type": "Scope"`

`}]},`

`"signInAudience": "`

	▼
All	
AzureADMyOrg	
AzureADMultipleOrgs	
AzureADandPersonalMicrosoftAccount	

Box 1: true

The `oauth2AllowImplicitFlow` attribute Specifies whether this web app can request OAuth2.0 implicit flow access tokens. The default is false. This flag is used for browser-based apps, like JavaScript single-page apps.

In implicit flow, the app receives tokens directly from the Azure Active Directory (Azure AD) authorize endpoint, without any server-to-server exchange. All authentication logic and session handling is done entirely in the JavaScript client with either a page redirect or a pop-up box.

Box 2: `requiredResourceAccess`

With dynamic consent, `requiredResourceAccess` drives the admin consent experience and the user consent experience for users who are using static consent. However, this parameter doesn't drive the user consent experience for the

general case.

resourceAppId is the unique identifier for the resource that the app requires access to. This value should be equal to the appId declared on the target resource app.

resourceAccess is an array that lists the OAuth2.0 permission scopes and app roles that the app requires from the specified resource. Contains the id and type values of the specified resources.

Example:

```
"requiredResourceAccess": [

  {
    "resourceAppId": "00000002-0000-0000-c000-000000000000",
    "resourceAccess": [

      {
        "id": "311a71cc-e848-46a1-bdf8-97ff7156d8e6",
        "type": "Scope"
      }
    ]
  },
],
```

Box 3: AzureADMyOrg

The signInAudience attribute specifies what Microsoft accounts are supported for the current application. Supported values are:

AzureADMyOrg - Users with a Microsoft work or school account in my organization's Azure AD tenant (for example, single tenant)

AzureADMultipleOrgs - Users with a Microsoft work or school account in any organization's Azure AD tenant (for example, multi-tenant)

AzureADandPersonalMicrosoftAccount - Users with a personal Microsoft account, or a work or school account in any

organization's Azure AD tenant

Reference:

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

<https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-implicit-grant-flow>

Question: 252

You are developing a web application that runs as an Azure Web App. The web application stores data in Azure SQL Database and stores files in an Azure Storage account. The web application makes HTTP requests to external services as part of normal operations.

The web application is instrumented with Application Insights. The external services are OpenTelemetry compliant.

You need to ensure that the customer ID of the signed in user is associated with all operations throughout the overall system.

What should you do?

- A. Create a new SpanContext with the TraceRags value set to the customer ID for the signed in user.
- B. On the current SpanContext, set the Traceld to the customer ID for the signed in user.
- C. Add the customer ID for the signed in user to the CorrelationContext in the web application.
- D. Set the header Ocp-Apim-Trace to the customer ID for the signed in user.

Answer: C

Explanation:

Reference:

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

Question: 253

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

The Azure Logic app must securely access the Azure Blob storage account. Azure AD resources must remain if the Azure Logic app is deleted.

You need to secure the Azure Logic app.

What should you do?

- A. Create an Azure AD custom role and assign role-based access controls.
- B. Create an Azure AD custom role and assign the role to the Azure Blob storage account.
- C. Create an Azure Key Vault and issue a client certificate.
- D. Create a user-assigned managed identity and assign role-based access controls.
- E. Create a system-assigned managed identity and issue a client certificate.

Answer: D

Explanation:

To give a managed identity access to an Azure resource, you need to add a role to the target resource for that identity.

Note: To easily authenticate access to other resources that are protected by Azure Active Directory (Azure AD) without having to sign in and provide credentials or secrets, your logic app can use a managed identity (formerly known as Managed Service Identity or MSI). Azure manages this identity for you and helps secure your credentials because you don't have to provide or rotate secrets.

If you set up your logic app to use the system-assigned identity or a manually created, user-assigned identity, the function in your logic app can also use that same identity for authentication.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/create-managed-service-identity>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates-for-clients>

Question: 254

You develop a solution that uses Azure Virtual Machines (VMs).

The VMs contain code that must access resources in an Azure resource group. You grant the VM access to the resource group in Resource Manager.

You need to obtain an access token that uses the VMs system-assigned managed identity.

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

- A. Use PowerShell on a remote machine to make a request to the local managed identity for Azure resources endpoint.
- B. Use PowerShell on the VM to make a request to the local managed identity for Azure resources endpoint.
- C. From the code on the VM. call Azure Resource Manager using an access token.
- D. From the code on the VM. call Azure Resource Manager using a SAS token.
- E. From the code on the VM. generate a user delegation SAS token.

Answer: B,C

Explanation:

Question: 255

HOTSPOT

You are developing an application that uses a premium block blob storage account. You are optimizing costs by automating Azure Blob Storage access tiers.

You apply the following policy rules to the storage account. You must determine the implications of applying the rules to the data. (Line numbers are included for reference only.)

```
01 {  
02 "rules":  
33 [  
04   "name": "usingDoTaRule"  
85   "Enabled*": true  
06   "type": "ifect'zle\
```

Answer Area

Yes No

Block blobs prefixed with container1/salesorders or container2/inventory  which have not been modified

in over 60 days are moved to cool storage.
Blobs that have not been modified in 120 days are moved to the archive tier.
Blobs are moved to cool storage if they have not been accessed for 30 days.
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.
All block blobs older than 730 days will be deleted.

O

O O

Answer:

Explanation:

1. Yes
2. Yes
3. Yes
4. No

<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#move-aging-data-to-a-cooler-tier>

Question: 256

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

```
conditions": [ { "name": "IsDevice", "parameters": {  
  "@odata.type": "#Microsoft.Azure.Cdn.Models." "operator": "Equal" "matchvalues": [ "▼" ]
```

iOS Mobile iPhone Desktop

```
"matchValues": [
```

iOS Mobile iPhone Desktop

DeliveryRulesDeviceConditionParameter
S
DeliveryRuleCookiesConditionParameter
S
DeliveryRulePostArgsConditionParameter

Answer:

```
{  
  name": "RequestHeader", parameters": {  
    "@odata.type": "#Microsoft.Azure.Cdn.Models,  
    "operator": "Contains", "selector": " ▼
```

DeliveryRulesDeviceConditionParameter
S
DeliveryRuleCookiesConditionParameter
S
DeliveryRulePostArgsConditionParameter

FROM
PRAGMA
X-POWERED-BY
HTTP_USER_AGENT

Explanation:

```
conditions": [ { "name": "IsDevice" "parameters": {  
  "@odata.type": "#Microsoft.Azure.Cdn.Models.  
  "operator": "Equal"  
  "matchvalues": [ " ▼
```

iOS
Mobile
iPhone
Desktop

DeliveryRulesDeviceConditionParameters
DeliveryRuleCookiesConditionParameters
DeliveryRulePostArgsConditionParameters
DeliveryRuleRequestHeaderConditionPara

```
name": "RequestHeader", parameters": {  
  "@odata.type": "#Microsoft.Azure.Cdn.Models.  
  "operator": "Contains",  
  "selector": " * in J
```

FROM
PRAGMA

DeliveryRulesDeviceConditionParameters
DeliveryRuleCookiesConditionParameters
DeliveryRulePostArgsConditionParameters
DeliveryRuleRequestHeaderConditionPara

X-POWERED-BY
HTTP_USER_AGE

"matchvalues": [" _____ ▼ "]

}}
]

[iOS

Mobile

iPhone Desktop

Box 1: iOS

Azure AD Conditional Access supports the following device platforms:

Android

iOS

Windows Phone

Windows

macOS

Box 2: DeliveryRulesDeviceConditionParameters

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

Box 3: HTTP_USER_AGENT

Box 4: DeliveryRuleRequestHeaderConditionParameters

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

Box 5: iOS

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

Reference:

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

Question: 257

HOTSPOT

You are developing a .NET application that communicates with Azure Storage.

A message must be stored when the application initializes.

You need to implement the message.

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

```
CloudQueueClient storageAccount = CloudStorageAccount.Parse(CloudConfigurationManager.GetSetting  
("StorageConnectionString"));  
  
CloudQueueClient pVar1 = storageAccount.  
CloudTableClient pVar2 = pVar1.  
CreateCloudQueueClient  
CreateCloudTableClient GetQueue  
Reference GetTableReference  
0;  
  
try  
{  
    pVar2.CreateIfNotExistsAsync();  
    pVar1.QueueMessage("cantoso-storage");  
}
```

Answer:

Explanation:

Answer Area

```
CloudQueueClient pVar1 = storageAccount.CreateCloudQueueClient();  
CloudTableClient pVar2 = pVar1.CreateCloudTableClient("cantoso-storage");  
try  
{  
    await pVar2.CreateIfNotExistsAsync();  
}
```

Question: 258

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

Instead update the web.config file to include the applicationInitialization configuration element.

Specify custom initialization actions to run the scripts.

Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.

```
<system.webServer>  
  
<applicationInitialization>  
  
  <add initializationPage="/" hostname="[app hostname]" />  
  <add initializationPage="/Home/About" hostname="[app hostname]" />  
</applicationInitialization>  
  
</system.webServer>
```

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps>

Question: 259

You have an existing Azure storage account that stores large volumes of data across multiple containers.

You need to copy all data from the existing storage account to a new storage account. The copy process must meet the following requirements:

Automate data movement.

Minimize user input required to perform the operation.

Ensure that the data movement process is recoverable.

What should you use?

- A. AzCopy
- B. Azure Storage Explorer
- C. Azure portal
- D. .NET Storage Client Library

Answer: A

Explanation:

You can copy blobs, directories, and containers between storage accounts by using the AzCopy v10 command-line utility.

The copy operation is synchronous so when the command returns, that indicates that all files have been copied.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy>

Question: 260

DRAG DROP

You are developing a web service that will run on Azure virtual machines that use Azure Storage. You configure all virtual machines to use managed identities.

You have the following requirements:

Secret-based authentication mechanisms are not permitted for accessing an Azure Storage account.

Must use only Azure Instance Metadata Service endpoints.

You need to write code to retrieve an access token to access Azure Storage. To answer, drag the appropriate code segments to the correct locations. Each code segment may be used 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 segment 1

```
http://localhost:50342/oauth2/token  
http://169.254.169.254:50432/oauth2/token  
http://localhost/metadata/identity/oauth2/token  
http://169.254.169.254/metadata/identity/oauth2/token
```

Code segment 2

```
XDocument.Parse(payload);  
new MultipartContent(payload);  
new NetworkCredential("Azure", payload);  
JsonConvert.DeserializeObject<Dictionary<string, string>>(payload);
```

Answer Area

```
var url = "Code segment 1";  
var queryString = "...";  
var client = new HttpClient();  
var response = await client.GetAsync(url + queryString);  
var payload = await response.Content.ReadAsStringAsync();  
return Code segment 2
```

Answer:

Explanation:

```
var url = "http://169.254.169.254/metadata/identity/oauth2/token";  
var queryString = "...";  
var client = new HttpClient();  
var response = await client.GetAsync(url + queryString);  
var payload = await response.Content.ReadAsStringAsync();  
return JsonConvert.DeserializeObject<Dictionary<string, string>>(payload);
```

Azure Instance Metadata Service endpoints "/oauth2/token"

Box 1: http://169.254.169.254/metadata/identity/oauth2/token

Sample request using the Azure Instance Metadata Service (IMDS) endpoint (recommended):

GET 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' HTTP/1.1 Metadata: true

Box 2: JsonConvert.DeserializeObject<Dictionary<string,string>>(payload);

Deserialized token response; returning access code.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token>

<https://docs.microsoft.com/en-us/azure/service-fabric/how-to-managed-identity-service-fabric-app-code>

Question: 261

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.

You need to configure authorization.

Solution:

Configure and use Integrated Windows Authentication in the website.
In the website, query Microsoft Graph API to load the group to which the user is a member.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Microsoft Graph is a RESTful web API that enables you to access Microsoft Cloud service resources.

Instead in the Azure AD application's manifest, set value of the groupMembershipClaims option to All. In the website, use the value of the groups claim from the JWT for the user to determine permissions.

Reference:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

Question: 262

DRAG DROP

You are developing a REST web service. Customers will access the service by using an Azure API Management instance.

The web service does not correctly handle conflicts. Instead of returning an HTTP status code of 409, the service returns a status code of 500. The body of the status message contains only the word **conflict**.

You need to ensure that conflicts produce the correct response.

How should you complete the policy? To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Policy segments

server
context
on-error
set-status
when-error
override-status

Answer Area

```
< Policy segment >
<base />
<choose>
  <when condition = " @ context .Response.StatusCode == 500
    && Policy segment .LastError.Message.Contains
      " conflict = " ) ) " >
    <return-response> (
      < Policy segment >
    )
  </return-response>
</when>
<otherwise />
</choose>
< Policy segment >
```

Answer:

Explanation:

```
< on-error >
  <base />
  <choose>
    <when condition = " @ context .Response.StatusCode == 500
      • Icontext .LastError.Message.Contains
        " conflict = " ) ) ">
      <return-response>
        < | set-status >
      </return-response>
    </when>
    <otherwise />
  </choose>
</on-error>
```

Box 1: on-error

Policies in Azure API Management are divided into inbound, backend, outbound, and on-error.

If there is no on-error section, callers will receive 400 or 500 HTTP response messages if an error condition occurs.

Box 2: context

Box 3: context

Box 4: set-status

The return-response policy aborts pipeline execution and returns either a default or custom response to the caller.

Default response is 200 OK with no body.

Custom response can be specified via a context variable or policy statements.

Syntax:

```
<return-response response-variable-name="existing context variable">
```

```
<set-header/>
```

```
<set-body/>
```

```
<set-status/>
```

```
</return-response>
```

Box 5: on-error

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-error-handling-policies>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies>

Question: 263

DRAG DROP

You are developing an Azure Function app.

The app must meet the following requirements:

Enable developers to write the functions by using the Rust language.

Declaratively connect to an Azure Blob Storage account.

You need to implement the app.

Which Azure Function app features should you use? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Features

Answer Area

Custom handler

Extension bundle

Trigger

Runtime

Policy

Hosting plan

Requirement

Enable developers to write the functions by using the Rust language.

Declaratively connect to an Azure Blob Storage account.

Feature

Feature

Feature

Answer:

Explanation:

Requirement

Feature

Enable developers to write the functions by using the Rust language.

Custom handler

Declaratively connect to an Azure Blob Storage account.

Trigger

Box 1: Custom handler

Custom handlers can be used to create functions in any language or runtime by running an HTTP server process, for example Go or Rust.

Box 2: Trigger

Functions are invoked by a trigger and can have exactly one. In addition to invoking the function, certain triggers also serve as bindings. You may also define multiple bindings in addition to the trigger. Bindings provide a declarative way to connect data to your code.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/create-first-function-vs-code-other>

<https://docs.microsoft.com/en-us/dotnet/architecture/serverless/azure-functions>

Question: 264

HOTSPOT

You create the following PowerShell script:

```
{source - New-AzScheduledQueryRuleSource -Query 'Heartbeat | where TimeGenerated > ago(1h)' -DataSourceId "contoso"
{schedule = New-AzScheduledQueryRuleSchedule -FrequencyInMinutes 60 -TimeWindowInMinutes 60
{triggercondition = New-AzScheduledQueryRuleTriggerCondition -Thresholdoperator "LessThan" -Threshold 5
{aznsActionGroup = New-AzScheduledQueryRuleAznsActionGroup -ActionGroup "contoso" -EmailSubject "Custom email subject"
-CustomwebhookPayload "{ \"alert\": \"alertrulename\", \"IncludeSearchResults\": true }"
{alertingAction - New-AzScheduledQueryRuleAlertingAction -AznsAction {aznsActionGroup -Severity "3" -Trigger {triggerCondition
New-AzScheduledQueryRule -ResourceGroupName "contoso" -Location "eastus" -Action {alertingAction -Enabled {true
-Description "Alert description" -Schedule {schedule -Source {source -Name "Alert Name"
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No,

NOTE: Each correct selection is worth one point.

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	☐	☐
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	☐	☐
The log alert is scheduled to run every two hours.	☐	☐

Answer:

Explanation:

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes	☐	☐
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five	☐	☐
The log alert is scheduled to run every two hours	☐	☐

Box 1: No

The AzScheduledQueryRuleSource is Heartbeat, not CPU.

Box 2: Yes

The AzScheduledQueryRuleSource is Heartbeat!

Note: New-AzScheduledQueryRuleTriggerCondition creates an object of type Trigger Condition. This object is to be passed to the command that creates Alerting Action object.

Box 3: No

The schedule is 60 minutes, not two hours.

-FrequencyInMinutes: The alert frequency.

-TimeWindowInMinutes: The alert time window

The New-AzScheduledQueryRuleSchedule command creates an object of type Schedule. This object is to be passed to the command that creates Log Alert Rule.

Reference:

<https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryrule>

<https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryruletriggercondition>

Question: 265

HOTSPOT

You are developing an Azure-hosted e-commerce web application. The application will use Azure Cosmos DB to store sales orders. You are using the latest SDK to manage the sales orders in the database.

You create a new Azure Cosmos DB instance. You include a valid endpoint and valid authorization key to an appSettings.json file in the code project.

You are evaluating the following application code: (Line number are included for reference only.)

```

01 using System;
02 using System.Threading.Tasks;
03 using Microsoft.Azure.Cosmos;
04 using Microsoft.Extensions.Configuration;
05 using Newtonsoft.Json;
06 namespace SalesOrders
07 {
08     public class SalesOrder
09     {
10         ...
11     }
12     internal class ManageSalesOrders
13     {
14         private static async Task GenerateSalesOrders() (
15             configuration = new ConfigurationBuilder().AddJsonFile("appSettings.json").Build();
16             string endpoint = configuration["EndPointUrl"];
17             string authKey = configuration["AuthorizationKey"];
18             using CosmosClient client = new CosmosClient(endpoint, authKey);
19             Database database = null;
20             using (await client.GetDataCase("SalesOrders").DeleteStreamAsync() {})
21             database = await client.CreateDatabaseIfNotExistsAsync("SalesOrders");
22             container1 = await database.CreateContainerAsync(id: "Container1", partitionKeyPath: "/AccountNumber"); container2 = await database.CreateContainerAsync(id: "Container2",
23             partitionKeyPath: "/AccountNumber");
24             SalesOrder salesOrder1 = new SalesOrder { AccountNumber = "123456" };
25             await container1.CreateItemAsync(salesOrder1, new PartitionKey(salesOrder1.AccountNumber));
26             SalesOrder salesOrder2 = new SalesOrder { AccountNumber = "654321" };
27             await container1.CreateItemAsync(salesOrder2, new PartitionKey(salesOrder2.AccountNumber));
28             SalesOrder salesOrder3 = new SalesOrder { AccountNumber = "109876" };
29             await container2.CreateItemAsync(salesOrder3, new PartitionKey(salesOrder3.AccountNumber));
30             User user1 = await database.CreateUserAsync("User1");
31             User user2 = await database.GetUserAsync("User1");
32             User user3 = await user1.ReadAsync();
33         }
34     }
35 }
36

```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Statements

Yes No

A database named SalesOrders is created. The database will include two containers.

Container1 will contain two items.

☐ Yes ☒ No

Container2 will contain one item.

☒ Yes ☐ No

Answer:

Explanation:

Statements

Yes

No

A database named SalesOrders is created. The database will include two containers.

☒☐

Container1 will contain two items.

☐☐

Container2 will contain one item.

☒☐

Box 1: Yes

The `createDatabaseIfNotExistsAsync` method checks if a database exists, and if it doesn't, create it.

The `Database.CreateContainerAsync` method creates a container as an asynchronous operation in the Azure Cosmos service.

Box 2: Yes

The `CosmosContainer.CreateItemAsync` method creates an item as an asynchronous operation in the Azure Cosmos service.

Box 3: Yes

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient.createdatabaseifnotexistsasync>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.database.createcontainerasync>

<https://docs.microsoft.com/en-us/dotnet/api/azure.cosmos.cosmoscontainer.createitemasync>

Question: 266

DRAG DROP

You are developing an Azure solution.

You need to develop code to access a secret stored in Azure Key Vault.

How should you complete the code segment? To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Code segments

DefaultAzureCredential

ClientSecretCredential

CloudClients

SecretClient

Answer Area

```
string var1 = Environment.GetEnvironmentVariable("KEY_VAULT_URI");
```

```
var var2 = new Code segment ( new Uri(var1), new Code segment ());
```

Answer:

Explanation:

```
string var1 = Environment.GetEnvironmentVariable("KEY_VAULT_URI");  
var var2 = new SecretClient ( new Uri(var1), new DefaultAzureCredential ());
```

Box 1: SecretClient

Box 2: DefaultAzureCredential

In below example, the name of your key vault is expanded to the key vault URI, in the format "https://<your-key-vault-name>.vault.azure.net". This example is using 'DefaultAzureCredential()' class from Azure Identity Library, which allows to use the same code across different environments with different options to provide identity.

```
string keyVaultName = Environment.GetEnvironmentVariable("KEY_VAULT_NAME");
```

```
var kvUri = "https://" + keyVaultName + ".vault.azure.net";
```

```
var client = new SecretClient(new Uri(kvUri), new DefaultAzureCredential());
```

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/secrets/quick-create-net>

Question: 267

You are developing an Azure App Service REST API.

The API must be called by an Azure App Service web app. The API must retrieve and update user profile information stored in Azure Active Directory (Azure AD).

You need to configure the API to make the updates.

Which two tools should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Microsoft Graph API
- B. Microsoft Authentication Library (MSAL)
- C. Azure API Management
- D. Microsoft Azure Security Center
- E. Microsoft Azure Key Vault SDK

Answer: A,C

Explanation:

A: You can use the Azure AD REST APIs in Microsoft Graph to create unique workflows between Azure AD resources and third-party services.

Enterprise developers use Microsoft Graph to integrate Azure AD identity management and other services to automate administrative workflows, such as employee onboarding (and termination), profile maintenance, license deployment, and more.

C: API Management (APIM) is a way to create consistent and modern API gateways for existing backend services.

API Management helps organizations publish APIs to external, partner, and internal developers to unlock the potential of their data and services.

Reference:

<https://docs.microsoft.com/en-us/graph/azuread-identity-access-management-concept-overview>

Question: 268

You develop a REST API. You implement a user delegation SAS token to communicate with Azure Blob storage.

The token is compromised.

You need to revoke the token.

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. Revoke the delegation keys
- B. Delete the stored access policy.
- C. Regenerate the account key.
- D. Remove the role assignment for the security principle.

Answer: A,B

Explanation:

A: Revoke a user delegation SAS

To revoke a user delegation SAS from the Azure CLI, call the `az storage account revoke-delegation- keys` command. This command revokes all of the user delegation keys associated with the specified storage account. Any shared access signatures associated with those keys are invalidated.

B: To revoke a stored access policy, you can either delete it, or rename it by changing the signed identifier.

Changing the signed identifier breaks the associations between any existing signatures and the stored access policy. Deleting or renaming the stored access policy immediately effects all of the shared access signatures associated with it.

D18912E1457D5D1DDCBD40AB3BF70D5D

Reference:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/storage/blobs/storage-blob-user-delegationsas-create-cli.md>

<https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy#modifying- or-revoking-astored-access-policy>

Question: 269

You are developing applications for a company. You plan to host the applications on Azure App Services.

The company has the following requirements:

Every five minutes verify that the websites are responsive.

Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly.

Generate alerts if a website is experiencing issues.

If a website fails to load, the system must attempt to reload the site three more times.

You need to implement this process with the least amount of effort.

What should you do?

D18912E1457D5D1DDCBD40AB3BF70D5D

A. Create a Selenium web test and configure it to run from your workstation as a scheduled task.

B. Set up a URL ping test to query the home page.

- C. Create an Azure function to query the home page.
- D. Create a multi-step web test to query the home page.
- E. Create a Custom Track Availability Test to query the home page.

Answer: D

Explanation:

You can monitor a recorded sequence of URLs and interactions with a website via multi-step web tests.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-multistep>

Question: 270

You develop and add several functions to an Azure Function app that uses the latest runtime host. The functions contain several REST API endpoints secured by using SSL. The Azure Function app runs in a Consumption plan.

You must send an alert when any of the function endpoints are unavailable or responding too slowly.

You need to monitor the availability and responsiveness of the functions.

What should you do?

- A. Create a URL ping test.
- B. Create a timer triggered function that calls TrackAvailability() and send the results to Application Insights.
- C. Create a timer triggered function that calls GetMetric("Request Size") and send the results to Application Insights.
- D. Add a new diagnostic setting to the Azure Function app. Enable the FunctionAppLogs and Send to Log Analytics options.

Answer: B

Explanation:

You can create an Azure Function with `TrackAvailability()` that will run periodically according to the configuration given in `TimerTrigger` function with your own business logic. The results of this test will be sent to your Application Insights resource, where you will be able to query for and alert on the availability results data. This allows you to create customized tests similar to what you can do via **Availability Monitoring** in the portal.

Customized tests will allow you to write more complex availability tests than is possible using the portal UI, monitor an app inside of your Azure VNET, change the endpoint address, or create an availability test even if this feature is not available in your region.

D18912E1457D5D1DDCBD40AB3BF70D5D

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-azure-functions>

Question: 271

DRAG DROP

You are developing an application to retrieve user profile information. The application will use the Microsoft Graph SDK.

The app must retrieve user profile information by using a Microsoft Graph API call.

You need to call the Microsoft Graph API from the application.

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

Create an authentication provider.

Create a new instance of the GraphServiceClient.

Invoke the request to the Microsoft Graph API.

Register the application with the Microsoft identity platform.

Build a client by using the client app ID.

Answer Area



Answer:

Explanation:

Register the application with the Microsoft identity platform

Build a client by using the client app ID

Create an authentication provider

) Create a new instance of the GraphServiceClient

Invoke the request to the Microsoft Graph API

Step 1: Register the application with the Microsoft identity platform.

To authenticate with the Microsoft identity platform endpoint, you must first register your app at the Azure app registration portal

Step 2: Build a client by using the client app ID

Step 3: Create an authentication provider

Create an authentication provider by passing in a client application and graph scopes.

Code example:

```
DeviceCodeProvider authProvider = new DeviceCodeProvider(publicClientApplication, graphScopes);
```

```
// Create a new instance of GraphServiceClient with the authentication provider.
```

```
GraphServiceClient graphClient = new GraphServiceClient(authProvider);
```

Step 4: Create a new instance of the GraphServiceClient

Step 5: Invoke the request to the Microsoft Graph API

Reference:

<https://docs.microsoft.com/en-us/graph/auth-v2-service>

<https://docs.microsoft.com/en-us/graph/sdks/create-client>

Question: 272

DRAW DROP

You develop and deploy an Azure Logic App that calls an Azure Function app. The Azure Function App includes an OpenAPI (Swagger) definition and uses an Azure Blob storage account. All resources are secured by using Azure Active Directory (Azure AD).

The Logic App must use Azure Monitor logs to record and store information about runtime data and events. The logs must be stored in the Azure Blob storage account.

You need to set up Azure Monitor logs and collect diagnostics data for the Azure Logic App.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

Create action groups and alert rules.

Create a Log Analytics workspace.

Install the Logic Apps Management

Add a diagnostic setting to the Azure Function App.

Create an Azure storage account.

Add a diagnostic setting to the Azure Logic App.



Answer:

Explanation:

Create a Log Analytics workspace

install the Logic Apps Management solution

Add a diagnostic setting to the Azure Logic App

Step 1: Create a Log Analytics workspace

Before you start, you need a Log Analytics workspace.

Step 2: Install the Logic Apps Management solution

To set up logging for your logic app, you can enable Log Analytics when you create your logic app, or you can install the Logic Apps Management solution in your Log Analytics workspace for existing logic apps.

Step 3: Add a diagnostic setting to the Azure Logic App

Set up Azure Monitor logs

In the Azure portal, find and select your logic app.

On your logic app menu, under Monitoring, select Diagnostic settings > Add diagnostic setting.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/monitor-logic-apps-log-analytics>

Question: 273

DRAG DROP

You are a developer for a Software as a Service (SaaS) company. You develop solutions that provide the ability to send notifications by using Azure Notification Hubs.

You need to create sample code that customers can use as a reference for how to send raw notifications to Windows Push Notification Services (WNS) devices. The sample code must not use external packages.

How should you complete the code segment? To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Code segments

raw

windows

windowsphone

application/xml

application/json

application/octet-stream

Answer Area

```
var endpoint = "___";
var payload = *___*;
request = HttpRequestMessage(HttpMethod.Post, endpoint);
request.Headers.Add("X-WNS-Type", "wns/raw");
request.Headers.Add("ServiceBusNotification-Format",
request.Content = StringContent(payload, Encoding.UTF8, var
client = : - HttpClient();
client.SendAsync(request);
```

Code segment

Code segment

Answer:

Explanation:

```
var endpoint =
var payload = "
var request = new
request.Headers.Add("X-WNS-Type", "wns/raw");
request.Headers.Add("ServiceBusNotification-Format",
request.Content = StringContent(payload, Encoding.UTF8, var
client = new HttpClient();
client.SendAsync(request);
```

windows

application/octet-stream

Box 1: windows

Example code:

```
var request = new HttpRequestMessage(method, $"{resourceUri}?api-version=2017-04");
request.Headers.Add("Authorization", createToken(resourceUri, KEY_NAME, KEY_VALUE));
request.Headers.Add("X-WNS-Type", "wns/raw");
request.Headers.Add("ServiceBusNotification-Format", "windows");
return request;
```

Box 2: application/octet-stream

Example code capable of sending a raw notification:

```
string resourceUri =
$"https://{NH_NAMESPACE}.servicebus.windows.net/{HUB_NAME}/messages/";
```

```
using (var request = CreateHttpRequest(HttpMethod.Post, resourceUri))
```

```
{  
  
    request.Content = new StringContent(content, Encoding.UTF8,  
    "application/octet-stream");  
    request.Content.Headers.ContentType.CharSet = string.Empty;
```

```
    var httpClient = new HttpClient();  
  
    var response = await httpClient.SendAsync(request);  
  
    Console.WriteLine(response.StatusCode);  
}
```

Reference:

<https://stackoverflow.com/questions/31346714/how-to-send-raw-notification-to-azure-notification-hub/31347901>

Question: 274

You develop and deploy an Azure App Service web app named App1. You create a new Azure Key Vault named Vault 1. You import several API keys, passwords, certificates, and cryptographic keys into Vault1.

You need to grant App1 access to Vault1 and automatically rotate credentials. Credentials must not be stored in code.

What should you do?

- A. Enable App Service authentication for App1. Assign a custom RBAC role to Vault1.
- B. Add a TLS/SSL binding to App1.
- C. Assign a managed identity to App1.

D. Upload a self-signed client certificate to Vault1. Update App1 to use the client certificate.

Answer: D

Explanation:

Question: 275

You deploy an Azure App Service web app. You create an app registration for the app in Azure Active Directory (Azure AD) and Twitter. the app must authenticate users and must use SSL for all communications. The app must use Twitter as the identity provider. You need to validate the Azure AD request in the app code. What should you validate?

- A. HTTP response code
- B. ID token header
- C. ID token signature
- D. Tenant ID

Answer: B

Explanation:

Question: 276

HOTSPOT

You are developing a web application that makes calls to the Microsoft Graph API. You register the application in the Azure portal and upload a valid X509 certificate.

You create an appsettings.json file containing the certificate name, client identifier for the application, and the tenant identifier of the Azure active Directory (Azure AD). You create a method named ReadCertificate to return the X509 certificate by name.

You need to implement code that acquires a token by using the certificate.

How should you complete the code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

        config = .ReadFromIsonFile("appsettings.json");
        certificate = ReadCertificate(config.CertificateName);
        .Create(config.ClientId)

        ConfidentialClientApplicationBuilder
        GetAccountAsyncO
        GetAccountsAsyncO
        ConfidentialClientAppiication

        .WithCertificate(certificate)
        .WithAuthority(new (config.Authority))
        .Build();
        string[] scopes = new string[] { $"{config.ApiUrl}.default" };
        nt result = await app.AcquireTokenFordient(
        * ).ExecuteAsync();
        scopes app config

```

Answer:

Explanation:

```

        config = .ReadFromIsonFile("appsettings.json");
        certificate = ReadCertificate(config.CertificateName);
        .Create(config.ClientId)

        ConfidentialClientApplicationBuilder
        GetAccountAsyncO
        GetAccountsAsyncO
        ConfidentialClientAppiication

        .WithCertificate(certificate)
        .WithAuthority(new (config.Authority))
        .Build();
        string[] scopes = new string[] { $"{config.ApiUrl}.default" };
        tnr ti aticr ->. st result = await app.AcquireTokenForClient(
        * ).ExecuteAsync();
        scopes app config

```

<https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-app-configuration?tabs=dotnet#instantiate-the-confidential-client-application-with-a-client-certificate>

<https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-acquire-token?tabs=dotnet#acquiretokenforclient-api>

Question: 277

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: Use the Azure Blob Storage change feed to trigger photo processing.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

The change feed is a log of changes that are organized into hourly segments but appended to and updated every few minutes.

These segments are created only when there are blob change events that occur in that hour.

Instead catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed>

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

Question: 278

DRAG DROP

You develop an Azure solution that uses Cosmos DB.

The current Cosmos DB container must be replicated and must use a partition key that is optimized for queries.

You need to implement a change feed processor solution.

Which change feed processor 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 the content.

NOTE: Each correct selection is worth one point.

Components	Requirement	Component
Host	Store the data from which the change feed is generated.	J
Delegate	Coordinate processing of the change feed across multiple workers.	
Lease container	Use the change feed processor to listen for changes.	
Monitored container	Handle each batch of changes.	

Answer:

Explanation:

Requirement

Store the data from which the change feed is generated.

Coordinate processing of the change feed across multiple workers.

Use the change feed processor to listen for changes.

Handle each batch of changes.

Component

Monitored container

Lease container

Host

Delegate

Box 1: The monitored container

The monitored container has the data from which the change feed is generated. Any inserts and updates to the monitored container are reflected in the change feed of the container.

Box 2: The lease container

The lease container acts as a state storage and coordinates processing the change feed across multiple workers. The lease container can be stored in the same account as the monitored container or in a separate account.

Box 3: The host: A host is an application instance that uses the change feed processor to listen for changes. Multiple instances with the same lease configuration can run in parallel, but each instance should have a different instance

name.

Box 4: The delegate

The delegate is the code that defines what you, the developer, want to do with each batch of changes that the change feed processor reads.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-processor>

Question: 279

DRAG DROP

You are developing an Azure-hosted application that must use an on-premises hardware security module (HSM) key.

The key must be transferred to your existing Azure Key Vault by using the Bring Your Own Key (BYOK) process.

You need to securely transfer the key to Azure Key Vault.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

Generate a key transfer blob file by using the HSM vendor-provided tool.

Generate a Key Exchange Key (KEK).

Create a custom policy definition in Azure Policy.

Run the az keyvault key import Command.

Run the az keyvault key restore Command.

Retrieve the Key Exchange Key (KEK) public key.

Answer:

Explanation:

Generate a Key Exchange Key (KEK).

Retrieve the Key Exchange Key (KEK) public key.

Generate a key transfer blob file by using the HSM vendor-provided tool.

Run the az keyvault key import command.

To perform a key transfer, a user performs following steps:

Generate KEK.

Retrieve the public key of the KEK.

Using HSM vendor provided BYOK tool - Import the KEK into the target HSM and exports the Target Key protected by the KEK.

Import the protected Target Key to Azure Key Vault.

Step 1: Generate a Key Exchange Key (KEK).

Step 2: Retrieve the Key Exchange Key (KEK) public key.

Step 3: Generate a key transfer blob file by using the HSM vendor-provided tool.

Generate key transfer blob using HSM vendor provided BYOK tool

Step 4: Run the az keyvault key import command

Upload key transfer blob to import HSM-key.

Customer will transfer the Key Transfer Blob (".byok" file) to an online workstation and then run a az keyvault key import command to import this blob as a new HSM-backed key into Key Vault.

To import an RSA key use this command:

```
az keyvault key import
```

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/keys/byok-specification>

Question: 280

You are developing a solution that will use a multi-partitioned Azure Cosmos DB database. You plan to use the latest Azure Cosmos DB SDK for development.

The solution must meet the following requirements:

- Send insert and update operations to an Azure Blob storage account.

- Process changes to all partitions immediately.

- Allow parallelization of change processing.

You need to process the Azure Cosmos DB operations.

What are two possible ways to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Create an Azure App Service API and implement the change feed estimator of the SDK. Scale the API by using multiple Azure App Service instances.
- B. Create a background job in an Azure Kubernetes Service and implement the change feed feature of the SDK.
- C. Create an Azure Function to use a trigger for Azure Cosmos DB. Configure the trigger to connect to the container.
- D. Create an Azure Function that uses a FeedIterator object that processes the change feed by using the pull model on the container. Use a FeedRange object to parallelize the processing of the change feed across multiple functions.

Answer: C,D

Explanation:

Azure Functions is the simplest option if you are just getting started using the change feed. Due to its simplicity, it is also the recommended option for most change feed use cases. When you create an Azure Functions trigger for Azure Cosmos DB, you select the container to connect, and the Azure Function gets triggered whenever there is a change in the container. Because Azure Functions uses the change feed processor behind the scenes, it automatically parallelizes change processing across your container's partitions.

Note: You can work with change feed using the following options:

Using change feed with Azure Functions

Using change feed with change feed processor

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed>

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model>

<https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed#azure-functions>

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model#using-feedrange-for-parallelization>

Question: 281

HOTSPOT

You have an Azure Web app that uses Cosmos DB as a data store. You create a CosmosDB container by running the following PowerShell script:

```
$resourceGroupName = "testResourceGroup"
```

```
$accountName = "testCosmosAccount"
```

```
$databaseName = "testDatabase"
```

```
$containerName = "testContainer"
```

```
$partitionKeyPath = "/EmployeeId"
```

```
$autoscaleMaxThroughput = 5000
```

```
New-AzCosmosDBSqlContainer
```

```
-ResourceGroupName $resourceGroupName
```

```
-AccountName $accountName
```

```
-DatabaseName $databaseName
```

```
-Name $containerName
```

```
-PartitionKeyKind Hash
```

```
-PartitionKeyPath $partitionKeyPath
```

```
-AutoscaleMaxThroughput $autoscaleMaxThroughput
```

You create the following queries that target the container:

SELECT * FROM c WHERE c.EmployeeId > '12345'

SELECT * FROM c WHERE c.UserID = '12345'

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 minimum throughput for the container is 400 R/Us.

☐☐

The first query statement is an in-partition query.

☐☐

The second query statement is a cross-partition query.

☐☐

Answer:

Explanation:

Yes **No**

The minimum throughput for the container is 400 R/Us.

The first query statement is an in-partition query.

The second query statement is a cross-partition query.

Box 1: No

You set the highest, or maximum RU/s Tmax you don't want the system to exceed. The system automatically scales the throughput T such that $0.1 * Tmax \leq T \leq Tmax$.

In this example we have autoscaleMaxThroughput = 5000, so the minimum throughput for the container is 500 R/Us.

Box 2: No

First query: `SELECT * FROM c WHERE c.EmployeeId > '12345'`

Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition.

In order to be an in-partition query, the query must have an equality filter that includes the partition key:

`SELECT * FROM c WHERE c.DeviceId > 'XMS-0001'`

Box 3: Yes

Example of In-partition query:

Consider the below query with an equality filter on DeviceId. If we run this query on a container partitioned on DeviceId, this query will filter to a single physical partition.

SELECT * FROM c WHERE c.Deviceld = 'XMS-0001'

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer>

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-query-container>

Question: 282

HOTSPOT

You develop a containerized application. You plan to deploy the application to a new Azure Container instance by using a third-party continuous integration and continuous delivery (CI/CD) utility.

The deployment must be unattended and include all application assets. The third-party utility must only be able to push and pull images from the registry. The authentication must be managed by Azure Active Directory (Azure AD). The solution must use the principle of least privilege.

You need to ensure that the third-party utility can access the registry.

Which authentication options should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Authentication

Registry authentication method

Option

Service principal
Individual identity
Repository-scoped access token
Managed identity for Azure resources

RBAC role

AcrPull
Owner
AcrPush
Contributor

Answer:

Explanation:

Authentication

Registry authentication method

Option

Service principal
Individual identity
Repository-scoped access token
Managed identity for Azure resources

RBAC role

AcrPull
Owner
AcrPush
Contributor

Box 1: Service principal

Applications and container orchestrators can perform unattended, or "headless," authentication by using an Azure Active Directory (Azure AD) service principal.

Box 2: AcrPush

AcrPush provides pull/push permissions only and meets the principle of least privilege.

Reference:

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-authentication?tabs=azure-cli>

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-roles?tabs=azure-cli>

Question: 283

DRAG DROP

You develop an application. You plan to host the application on a set of virtual machines (VMs) in Azure.

You need to configure Azure Monitor to collect logs from the application.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Create a Log Analytics workspace.	
Install agents on the VM and VM scale set to be monitored.	
Send console logs.	
Add a VMInsights solution.	
Create an Application Insights resource.	

Answer:

Explanation:

Answer Area

Create a Log Analytics workspace.

Add a VMInsights solution.

Install agents on the VM and VM scale set to be monitored

Create an Application Insights resource.

Step 1: Create a Log Analytics workspace.

First create the workspace.

Step 2: Add a VMInsights solution.

Before a Log Analytics workspace can be used with VM insights, it must have the VMInsights solution installed.

Step 3: Install agents on the VM and VM scale set to be monitored.

Prior to onboarding agents, you must create and configure a workspace. Install or update the Application Insights Agent as an extension for Azure virtual machines and VM scale sets.

Step 4: Create an Application Insights resource

Sign in to the Azure portal, and create an Application Insights resource.

Home > New > Application Insights >

Application Insights

Monitor web app performance and usage

Basics Tags Review + create

Create an Application Insights resource to monitor your live web application. With Application Insights, you have full observability into your application across all components and dependencies of your complex distributed architecture. It includes powerful analytics tools to help you diagnose issues and to understand what users actually do with your app. It's designed to help you continuously improve performance and usability. It works for apps on a wide variety of platforms including .NET, Node.js and Java EE, hosted on-premises, hybrid, or any public cloud. [Learn More](#)

PROJECT DETAILS

Select a subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription *

Visual Studio Enterprise

Resource Group *

My_Resource_Group

[Create new](#)

INSTANCE DETAILS

Name *

My_AppInsights_Resource

Region *

(US) West US 2

Resource Mode *

Classic

Workspace-based

WORKSPACE DETAILS

Subscription *

Visual Studio Enterprise

Log Analytics Workspace *

my-workspace-name [westus2]

Review + create

« Previous

Next : Tags >

Once a workspace-based Application Insights resource has been created, configuring monitoring is relatively straightforward.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/vm/vminsights-configure-workspace>

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/create-workspace-resource>

Question: 284

DRAG DROP

You are developing an Azure solution to collect inventory data from thousands of stores located around the world. Each store location will send the inventory data hourly to an Azure Blob storage account for processing.

The solution must meet the following requirements:

Begin processing when data is saved to Azure Blob storage.

Filter data based on store location information.

Trigger an Azure Logic App to process the data for output to Azure Cosmos DB.

Enable high availability and geographic distribution.

Allow 24-hours for retries.

Implement an exponential back off data processing.

You need to configure the solution.

What should you implement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Technologies

Azure Event Hub

Azure Event Grid

Azure Service Bus

Azure Blob Storage

Azure App Service

Azure Logic App

Answer Area

Object

Technology

Event Source

Event Receiver

Event Handler

Technology

Technology

Technology

Answer:

Explanation:

Object

Technology

Event Source

Azure Event Grid

Event Receiver

Azure Logic App

Event Handler

Azure Service Bus


Box 1: Azure Event Grid

Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure

Logic Apps, or even to your own http listener. Event Grid provides reliable event delivery to your applications through rich retry policies and dead-lettering.

Box 2: Azure Logic App

Event Grid uses event subscriptions to route event messages to subscribers. This image illustrates the relationship between event publishers, event subscriptions, and event handlers.



Box 3: Azure Service Bus

The Event Grid service doesn't store events. Instead, events are stored in the Event Handlers, including ServiceBus, EventHubs, Storage Queue, WebHook endpoint, or many other supported Azure Services.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

<https://docs.microsoft.com/en-us/java/api/overview/azure/messaging-eventgrid-readme>

Question: 285

You are creating an app that will use CosmosDB for data storage. The app will process batches of relational data.

You need to select an API for the app.

Which API should you use?

- A. MongoDBAPI
- B. Table API
- C. SQL API
- D. Cassandra API

Answer: C

Explanation:

For relational data you will need the SQL API

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>

Question: 286

You develop and deploy an Azure App Service web app. The app is deployed to multiple regions and uses Azure Traffic Manager. Application Insights is enabled for the app.

You need to analyze app uptime for each month.

Which two solutions will achieve the goal? Each correct answer presents a complete solution

NOTE: Each correct selection is worth one point

- A. Application Insights alerts
- B. Application Insights web tests
- C. Azure Monitor logs
- D. Azure Monitor metrics

Answer: A,C

Explanation:

Reference:

<https://azure.microsoft.com/en-us/blog/creating-a-web-test-alert-programmatically-with-application-insights/>

Question: 287

HOTSPOT

You are developing an application to collect the following telemetry data for delivery drivers: first name, last name, package count, item id, and current location coordinates. The app will store the data in Azure Cosmos DB.

You need to configure Azure Cosmos DB to query the data.

Which values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Configuration Parameter

Value

Azure Cosmos DB API

	▼
Gremlin	
Table API	
Core (SQL)	

Azure Cosmos DB partition key

	▼
first name	
last name	
package count	
item id	

Answer:

Explanation:

Configuration Parameter

Value

Azure Cosmos DB API

	▼
Gremlin	
Table API	
Core (SQL)	

Azure Cosmos DB partition key

	▼
first name	
last name	
package count	
item id	

Box 1: Core (SQL)

Core(SQL) API stores data in document format. It offers the best end-to-end experience as we have full control over the interface, service, and the SDK client libraries. SQL API supports analytics and offers performance isolation between operational and analytical workloads.

Box 2: item id

item id is a unique identifier and is suitable for the partition key.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>

<https://docs.microsoft.com/en-us/azure/cosmos-db/partitioning-overview>

Question: 288

You are developing a web application by using the Azure SDK. The web application accesses data in a zone-redundant BlockBlobStorage storage account.

The application must determine whether the data has changed since the application last read the data. Update operations must use the latest data changes when writing data to the storages.

You need to implement the update operations.

Which values should you use? To answer, select the appropriate option in the answer area.

NOTE Each correct selection is worth one point.

**Answer: See the
Explanation below:**

Explanation:

See the answer in below image.

Answer Area

Code evaluation Value

HTTP Header value VeranoId

Conditional header if Match

Question: 289

DRAG DROP

You develop and deploy an Azure App Service — app. The web app accesses data in an Azure SQL database.

You must update the web app to store frequently used data in a new Azure Cache for Redis Premium instance.

You need to implement the Azure Cache for Redis features.

Which feature should you implement? To answer, drag the appropriate feature to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to - between panes or scroll to view content.

NOTE Each correct selection is worth one point

Features	Answer Area
☐ horizontal partitioning	
☐ channel	
☐ list	
☐ set	

Requirement	Feature
Create a data structure for storing collections of related items.	
Create a data structure for storing the most recently accessed cache items.	
Send messages through a high-performance publisher/subscriber mechanism.	

Answer:

Explanation:

Requirement

Feature

Create a data structure for storing collections of related items

set

Create a data structure for the most recently accessed cache items

list

Send messages through a high-performance publisher/subscriber mechanism

channel

Question: 290

An organization hosts web apps in Azure. The organization uses Azure Monitor You discover that configuration changes were made to some of the web apps. You need to identify the configuration changes. Which Azure Monitor log should you review?

- A. AppServiceEnvironmentPlatformLogs
- B. AppServiceApplogs

- C. AppServiceAuditLogs
- D. AppServiceConsoteLogs

Answer: C

Explanation:

Question: 291

DRAG DROP

You provision virtual machines (VMs) as development environments.

One VM does not have host.

The VM is stuck in a Windows update process. You attach the OS disk for the affected VM to a recovery VM.

You need to correct the issue.

In which order should you perform the actions? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer area
Open C:\temp\FixHost.txt file and locate the update that is in a pending state.	1
Run the following command at an elevated command prompt: <code>diskpart /list disk</code>	2
Run the following command at an elevated command prompt: <code>diskpart /select disk 1 /add package x: \temp\FixHost.txt</code>	3
Detach the OS disk and recreate the VM.	4

Answer:

Explanation:

Actions

Answer area

Run the Knowing command at an elevated command prompt

di MI .J^feM /get-packages > c:\temp\Pdtdl.txt

21 Open C:\temp\Wch.txt and locate the update that is in a pending state

Run the following command at an elevated command prompt

dism /image:.\Attacted OS #U>tl /Utam-t-Packet /pict*(eMe «= deletes

4 Detach the OS disk and recreate the VM.

Remove the update that causes the problem

Take a snapshot of the OS disk of the affected VM as a backup.

Attach the OS disk to a recovery VM.

Once the OS disk is attached on the recovery VM, run diskmgmt.msc to open Disk Management, and ensure the attached disk is ONLINE.

(Step 1) Open an elevated command prompt instance (Run as administrator). Run the following command to get the list of the update packages that are on the attached OS disk:

```
dism /image:<Attached OS disk>:\ /get-packages > c:\temp\Patch_level
```

(Step 2) Open the C:\temp\Patch_level.txt file, and then read it from the bottom up. Locate the update that's in Install Pending or Uninstall Pending state.

Remove the update that caused the problem:

```
dism /Image:<Attached OS disk>:\ /Remove-Package /PackageName:<PACK
```

(Step 4) Detach the OS disk and recreate the VM. Then check whether the issue is resolved.

Reference:

<https://docs.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshoot-stuck-updating-boot-error>

Question: 292

You are developing an Azure messaging solution.

You need to ensure that the solution that meets the following requirements:

- Provide duplicate detection.
- Store the messages for an unlimited period of time

* Provide transactional support

Which two technologies will meet the requirements? Each correct answer presents a complete solution NOTE Each correct selection is worth one point.

- A. Azure Service Bus Queue
- B. Azure Storage Queue
- C. Azure Service Bus TopicD Azure Event Hub

Answer: A,C

Explanation:

The Azure Service Bus Queue and Topic has duplicate detection.

Enabling duplicate detection helps keep track of the application-controlled MessageId of all messages sent into a queue or topic during a specified time window.

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/duplicate-detection>

Question: 293

HOTSPOT

You are developing a web application that will use Azure Storage. Older data will be less frequently used than more recent data.

You need to configure data storage for the application. You have the following requirements:

Retain copies of data for five years.

Minimize costs associated with storing data that is over one year old.

Implement Zone Redundant Storage for application data.

What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Requirement
Configure an Azure Storage account

Solution
☐
☒ Implement Blob Storage
☐ Implement Azure Cosmos DB
☐ Implement Storage (general purpose v1)
☒ Implement StorageV2 (general purpose)

Configure data retention

☐
☒ Snapshot blobs and move them to the archive tier
☐ Set a lifecycle management policy to move blobs to the cool tier
☐ Use AzCopy to copy the data to an on-premises device for backup
☐ Set a lifecycle management policy to move blobs to the archive tier

Answer:

Explanation:

Requirement
Configure an Azure Storage account

Solution
☐
☐ Implement Blob Storage
☐ Implement Azure Cosmos DB
☐ Implement Storage (general purpose v1)
☒ Implement StorageV2 (general purpose v2)

Configure data retention

☐
☐ Snapshot blobs and move them to the archive tier
☒ Set a lifecycle management policy to move blobs to the cool tier
☐ Use AzCopy to copy the data to an on-premises device for backup
☐ Set a lifecycle management policy to move blobs to the archive tier

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy?toc=/azure/storage/blobs/toc.json>

Question: 294

You develop and deploy a web application to Azure App Service. The application accesses data stored in an Azure Storage account. The account contains several containers with several blobs with large amounts of data. You deploy all Azure resources to a single region.

You need to move the Azure Storage account to the new region. You must copy all data to the new region.

What should you do first?

- A. Export the Azure Storage account Azure Resource Manager template
- B. Initiate a storage account failover
- C. Configure object replication for all blobs
- D. Use the AzCopy command line tool
- E. Create a new Azure Storage account in the current region
- F. Create a new subscription in the current region

Answer: A

Explanation:

To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice and finally, delete the resources in the source region.

To get started, export, and then modify a Resource Manager template.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

Question: 295

A development team is creating a new REST API. The API will store data in Azure Blob storage. You plan to deploy the API to Azure App Service.

Developers must access the Azure Blob storage account to develop the API for the next two months. The Azure Blob storage account must not be accessible by the developers after the two-month time period.

You need to grant developers access to the Azure Blob storage account.

What should you do?

- A. Generate a shared access signature (SAS) for the Azure Blob storage account and provide the SAS to all developers.
- B. Create and apply a new lifecycle management policy to include a last accessed date value. Apply the policy to the Azure Blob storage account.
- C. Provide all developers with the access key for the Azure Blob storage account. Update the API to include the Coordinated Universal Time (UTC) timestamp for the request header.
- D. Grant all developers access to the Azure Blob storage account by assigning role-based access control (RBAC) roles.

Explanation:

Answer:

A

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

Question: 296

DRAG DROP

An organization plans to deploy Azure storage services.

You need to configure shared access signature (SAS) for granting access to Azure Storage.

Which SAS types should you use? To answer, drag the appropriate SAS types to the correct requirements.

Each SAS type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

SAS types

Account-level

Service-level

User delegation

Answer Area

Requirement

Delegate access to resources in one or more of the storage services

Delegate access to a resource in a single storage service

Secure a resource by using Azure AD credentials

SAS type

Answer:

Explanation:

Requirement

SAS type

Delegate access to resources in one or more of the storage services

Account-level

Delegate access to a resource in a single storage service

Service-level

Secure a resource by using Azure AD credentials

User delegation

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

Question: 297

HOTSPOT

You are developing an ASP.NET Core time sheet application that runs as an Azure Web App. Users of the application enter their time sheet information on the first day of every month.

The application uses a third-party web service to validate data.

The application encounters periodic server errors due to errors that result from calling a third-party web server. Each request to the third-party server has the same chance of failure.

You need to configure an Azure Monitor alert to detect server errors unrelated to the third-party service.

You must minimize false-positive alerts.

How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
"type": "Microsoft.Insights/metricAlerts",
"properties": {
  "criteria": {
    "odata.type": "...",
    "allOf": [
      {
        "criterionType": "
        DynamicThresholdCriterion
        SingleResourceMultipleMetricCriteria
      ",
        "metricName": "
        Http4xx
        Http5xx
      ",
        "alertSensitivity": "
        Low
        High
      "
    ]
  }
}
```

Answer:

Explanation:

```

"type": "Microsoft.Insights/metricAlerts",
"properties": {
  "criteria": {
    "odata.type": "...",
    "allof": [
      {
        "criterionType": "
        DynamicThresholdCriterion
        SingleResourceMultipleMetricCriteria
      }
    ],
    "metricName": "
    Http4xx
    Http5xx
  ",
    "alertSensitivity": "
    Low
    High
  "
  }
}

```

Box 1: DynamicThresholdCriterion

Box 2: Http5xx

Server errors are in the 5xx range.

Client errors are in the 4xx range

Box 3: Low

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-dynamic-thresholds>

Question: 298

You are developing a web application that uses Azure Cache for Redis. You anticipate that the cache will frequently fill and that you will need to evict keys.

You must configure Azure Cache for Redis based on the following predicted usage pattern: A small subset of elements will be accessed much more often than the rest.

You need to configure the Azure Cache for Redis to optimize performance for the predicted usage pattern.

Which two eviction policies will achieve the goal?

NOTE: Each correct selection is worth one point.

A. noeviction

B. allkeys-lru

C. volatile-lru

D. allkeys-random

E. volatile-ttl

F. volatile-random

Answer: B,C

Explanation:

B: The allkeys-lru policy evict keys by trying to remove the less recently used (LRU) keys first, in order to make space for the new data added. Use the allkeys-lru policy when you expect a power-law distribution in the popularity of your requests, that is, you expect that a subset of elements will be accessed far more often than the rest.

C: volatile-lru: evict keys by trying to remove the less recently used (LRU) keys first, but only among keys that have an expire set, in order to make space for the new data added.

Note: The allkeys-lru policy is more memory efficient since there is no need to set an expire for the key to be evicted under memory pressure.

Reference:

<https://redis.io/topics/lru-cache>

Question: 299

HOTSPOT

A software as a service (SaaS) company provides document management services. The company has a service that consists of several Azure web apps. All Azure web apps run in an Azure App Service Plan named PrimaryASP.

You are developing a new web service by using a web app named ExcelParser. The web app contains a third-party library for processing Microsoft Excel files. The license for the third-party library stipulates that you can only run a single instance of the library.

You need to configure the service.

How should you complete the script? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
Set-AzAppServicePlan `
  -ResourceGroupName $rg `
  -Name "PrimaryASP" `
```

```
NumberOfSites 1
PerSiteScaling $true
TargetWorkerCount = 1
MaxNumberOfWorkers = 1
SiteConfig.NumberOfWorkers = 1
```

```
$app = Get-AzWebApp `
  -ResourceGroupName $rg `
  -Name "ExcelParser"
```

```
$app.  
NumberOfSites 1  
PerSiteScaling $true  
TargetWorkerCount = 1  
MaxNumberOfWorkers = 1  
SiteConfig.NumberOfWorkers = 1
```

```
Set-AzWebApp $app
```

Answer:

Explanation:

```
Set-AzAppServicePlan `
  -ResourceGroupName $rg `
  -Name "PrimaryASP" `
```

	▼
NumberOfSites 1	
PerSiteScaling \$true	
TargetWorkerCount = 1	
MaxNumberOfWorkers = 1	
SiteConfig.NumberOfWorkers = 1	

```
$app = Get-AzWebApp `
  -ResourceGroupName $rg `
  -Name "ExcelParser"
```

\$app.

	▼
NumberOfSites 1	
PerSiteScaling \$true	
TargetWorkerCount = 1	
MaxNumberOfWorkers = 1	
SiteConfig.NumberOfWorkers = 1	

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/manage-scale-per-app>

Question: 300

HOTSPOT

You are developing an ASP.NET Core app that includes feature flags which are managed by Azure App Configuration. You create an Azure App Configuration store named AppreaiureflagStore as shown in the exhibit:

Key	Label	State	Description	Last modified
cnetworks.com	www.atmicnetworks.com	www.atmicnetworks.com	www.atmicnetworks.com	
Export	Export	\$ ⁿ 1	Ability to export daw	6/11/2020, 9.13:26

You must be able to use the feature in the app by using the following markup:

cirntu n*ne«*Exjxjrf*>

< clns5-'n«v itpn'>

■ . class* nau link text-dark" asp-zrep. asp-rnntnrlrr*"innr" asp-artian-"Fsporr">txF-nrt Data</^> </U>

RateX

You went to update the app to use the feature flag.

Which values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area	Code section	Value
	Controller attribute	Feat reGate Route
		ServiceFilter
	Startup method	AddAzuteAppConfiguration
		AddControllersWithViews
		AddUserSecrets
	AppConfig endpoint setting	https://appfeatureflagstore.azure.io
		https://appfeatureflagstore.vault.azure.net
		https://erport.azure.io
		https://expo.rtvault.azure.net

Answer:

Explanation:

Box 1: FeatureGate

You can use the `FeatureGate` attribute to control whether a whole controller class or a specific action is enabled.

Box 2: AddAzureAppConfiguration

The extension method `AddAzureAppConfiguration` is used to add the Azure App Configuration Provider.

Box 3: <https://appfeatureflagstore.azureconfig.io>

You need to request the access token with resource=<https://<yourstorename>.azureconfig.io>

Reference:

<https://docs.microsoft.com/en-us/azure/azure-app-configuration/use-feature-flags-dotnet-core>

<https://csharp.christiannagel.com/2020/05/19/azureappconfiguration/>

<https://stackoverflow.com/questions/61899063/how-to-use-azure-app-configuration-rest-api>

Question: 301

You are developing a web application that uses the Microsoft identity platform to authenticate users and resources.

The web application calls several REST APIs.

The APIs require an access token from the Microsoft identity platform.

You need to request a token.

Which three properties should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Application secret
- B. Redirect URI/URL
- C. Application name
- D. Supported account type
- E. Application ID

Answer: A,B,E

Explanation:

Question: 302

HOTSPOT

You are developing an Azure Function App. You develop code by using a language that is not supported by the Azure Function App host. The code language supports HTTP primitives.

You must deploy the code to a production Azure Function App environment.

You need to configure the app for deployment.

Which configuration values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Configuration parameter

Publish

Runtime stack

Version

Configuration value

	▼
Code	
Docker Container	

	▼
Node.js	
Python	
PowerShell Core	
Custom Handler	

	▼
14 LTS	
7.0	
custom	

Answer:

Explanation:

Box 1: Docker container

A custom handler can be deployed to every Azure Functions hosting option. If your handler requires operating system or platform dependencies (such as a language runtime), you may need to use a custom container. You can create and deploy your code to Azure Functions as a custom Docker container.

Box 2: PowerShell core

When creating a function app in Azure for custom handlers, we recommend you select .NET Core as the stack. A "Custom" stack for custom handlers will be added in the future.

PowerShell Core (PSC) is based on the new .NET Core runtime.

Box 3: 7.0

On Windows: The Azure Az PowerShell module is also supported for use with PowerShell 5.1 on Windows.

On Linux: PowerShell 7.0.6 LTS, PowerShell 7.1.3, or higher is the recommended version of PowerShell for use with the Azure Az PowerShell module on all platforms.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-function-linux-custom-image>

<https://docs.microsoft.com/en-us/powershell/azure/install-az-ps?view=azps-7.1.0>

Question: 303

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.

The app continues to time out after four minutes. The app must process the blob data.

You need to ensure the app does not time out and processes the blob data.

Solution: Update the functionTimeout property of the host.json project file to 10 minutes.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.

Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

Question: 304

You manage a data processing application that receives requests from an Azure Storage queue.

You need to manage access to the queue. You have the following requirements:

Provide other applications access to the Azure queue.

Ensure that you can revoke access to the queue without having to regenerate the storage account keys.

Specify access at the queue level and not at the storage account level.

Which type of shared access signature (SAS) should you use?

- A. Service SAS with a stored access policy
- B. Account SAS
- C. User Delegation SAS
- D. Service SAS with ad hoc SAS

Answer: A

Explanation:

A service SAS is secured with the storage account key. A service SAS delegates access to a resource in only one of the Azure Storage services: Blob storage, Queue storage, Table storage, or Azure Files.

Stored access policies give you the option to revoke permissions for a service SAS without having to regenerate the storage account keys.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

Question: 305

HOTSPOT

You are developing an application to store and retrieve data in Azure Blob storage. The application will be hosted in an on-premises virtual machine (VM). The VM is connected to Azure by using a Site- to-Site VPN gateway connection. The application is secured by using Azure Active Directory (Azure AD) credentials.

The application must be granted access to the Azure Blob storage account with a start time, expiry time, and read permissions. The Azure Blob storage account access must use the Azure AD credentials of the application to secure data access. Data access must be able to be revoked if the client application security is breached.

You need to secure the application access to Azure Blob storage.

Which security features should you use? To answer select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Component

Security Feature

Application (Client)

	▼
Storage Account Access Key	
System-assigned Managed Identity	
Shared access signature (SAS) token	

Azure Storage (Server)

	▼
Stored Access Policy	
User-assigned Managed Identity	
Cross-Origin Resource Sharing (CORS)	

Answer:

Explanation:

Component

Security Feature

Application (Client)

	▼
Storage Account Access Key	
System-assigned Managed Identity	
Shared access signature (SAS) token	

Azure Storage (Server)

	▼
Stored Access Policy	
User-assigned Managed Identity	
Cross-Origin Resource Sharing (CORS)	

Box 1: Shared access signature (SAS) token

When your application design requires shared access signatures for access to Blob storage, use Azure AD credentials to create a user delegation SAS when possible for superior security.

Box 2: Stored access policy

Stored access policies give you the option to revoke permissions for a service SAS without having to regenerate the storage account keys.

A shared access signature can take one of the following two forms:

Service SAS with stored access policy. A stored access policy is defined on a resource container, which can be a blob container, table, queue, or file share. The stored access policy can be used to manage constraints for one or more service shared access signatures. When you associate a service SAS with a stored access policy, the SAS inherits the constraints – the start time, expiry time, and permissions – defined for the stored access policy.

Ad hoc SAS.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

Question: 306

You are developing an Azure Function App that runs in an App Service Plan. The Azure Function is triggered by a Timer object. You observe that the Azure Function does not reliably trigger when scheduled. Which two actions should you perform?

- A. Verify that Always On is enabled.
- B. Modify the trigger to use a SignalR trigger.
- C. Ensure that the function has a retry configured.
- D. Modify the trigger to use Consumption mode instead of the App Service plan.

Answer: A,C