Number: AZ-204
Passing Score: 800
Time Limit: 120
File Version: 14.0

Exam Code: AZ-204
Exam Name: Developing Solutions for Microsoft Azure

01 - Develop Azure compute solutions

Case study

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

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

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

To start the case study

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

Current environment

Windows Server 2016 virtual machine

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

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

Inland Transport - This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.

The VM supports the following REST API calls:

Container API - This API provides container information including weight, contents, and other attributes.

Location API - This API provides location information regarding shipping ports of call and trucking stops.

Shipping REST API - This API provides shipping information for use and display on the shipping website.

Shipping Data

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

Shipping Web Site

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

Proposed solution

The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations.

You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:

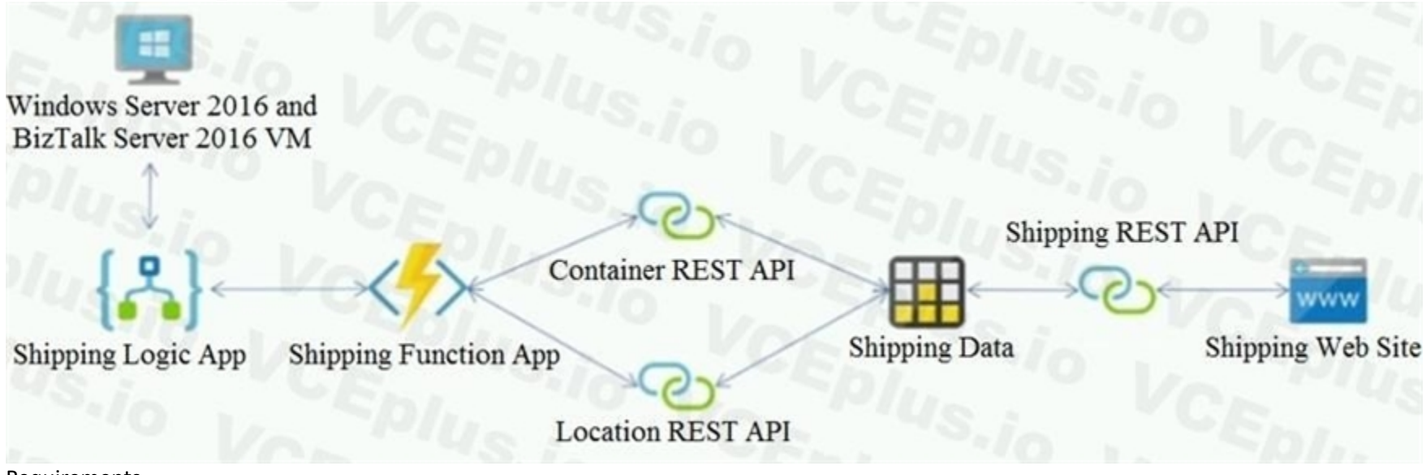

Requirements

Shipping Logic app

The Shipping Logic app must meet the following requirements:

Support the ocean transport and inland transport workflows by using a Logic App.

Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.

Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.

Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.
Shipping Function app
Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).
REST APIs
The REST API's that support the solution must meet the following requirements:
Secure resources to the corporate VNet.
Allow deployment to a testing location within Azure while not incurring additional costs.
Automatically scale to double capacity during peak shipping times while not causing application downtime.
Minimize costs when selecting an Azure payment model.
Shipping data
Data migration from on-premises to Azure must minimize costs and downtime.
Shipping website
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.
Issues
Windows Server 2016 VM
The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.
Shipping website and REST APIs
The following error message displays while you are testing the website:
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource. Origin 'http:// test.wideworldimporters.com/' is therefore not allowed access.

QUESTION 1
HOTSPOT
You need to correct the VM issues.
Which tools should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

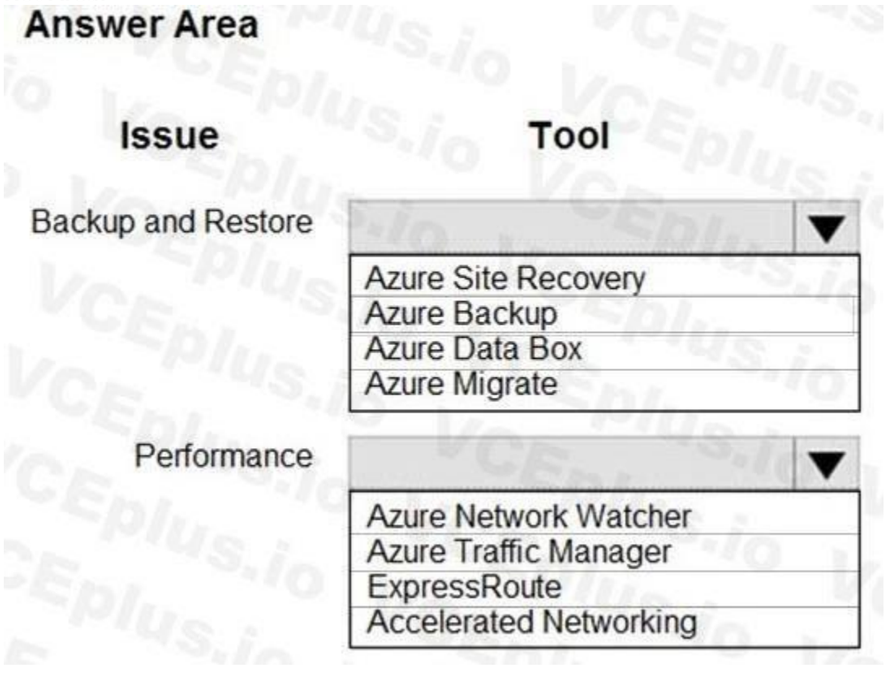

Answer Area:

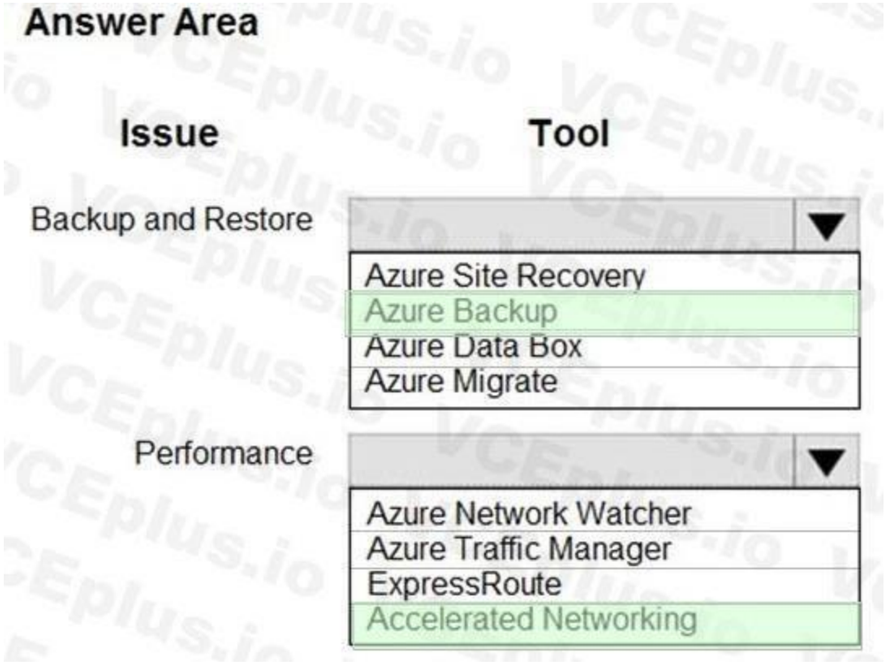

Section:

Explanation:

Box 1: Azure Backup The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.
In-Place restore of disks in IaaS VMs is a feature of Azure Backup.
Performance: Accelerated Networking
Scenario: The VM shows high network latency, jitter, and high CPU utilization.
Box 2: Accelerated networking
The VM shows high network latency, jitter, and high CPU utilization.
Accelerated networking enables single root I/O virtualization (SR-IOV) to a VM, greatly improving its networking performance. This high-performance path bypasses the host from the datapath, reducing latency, jitter, and CPU utilization, for use with the most demanding network workloads on supported VM types.
Reference:
https://azure.microsoft.com/en-us/blog/an-easy-way-to-bring-back-your-azure-vm-with-in-place-restore/

QUESTION 2

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Answer Area:

Answer Area

Option	Value
Tier	Standard
Profile	Akamai
Optimization	dynamic site acceleration

Tier: Standard
Premium

Profile: Akamai
Microsoft

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

Section:
Explanation:
Scenario: Shipping website
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.
Tier: Standard
Profile: Akamai
Optimization: Dynamic site acceleration Dynamic site acceleration (DSA) is available for Azure CDN Standard from Akamai, Azure CDN Standard from Verizon, and Azure CDN Premium from Verizon profiles.
DSA includes various techniques that benefit the latency and performance of dynamic content. Techniques include route and network optimization, TCP optimization, and more.
You can use this optimization to accelerate a web app that includes numerous responses that aren't cacheable. Examples are search results, checkout transactions, or real-time data. You can continue to use core Azure CDN caching capabilities for static data.
Reference:
https://docs.microsoft.com/en-us/azure/cdn/cdn-optimization-overview

02 - Develop Azure compute solutions

Case study

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

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

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

To start the case study

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

Background

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

Current environment

Architecture overview

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

API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration.

Logic apps are used to orchestrate the data processing while Service Bus and Event Grid handle messaging and events.

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

Architecture diagram

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

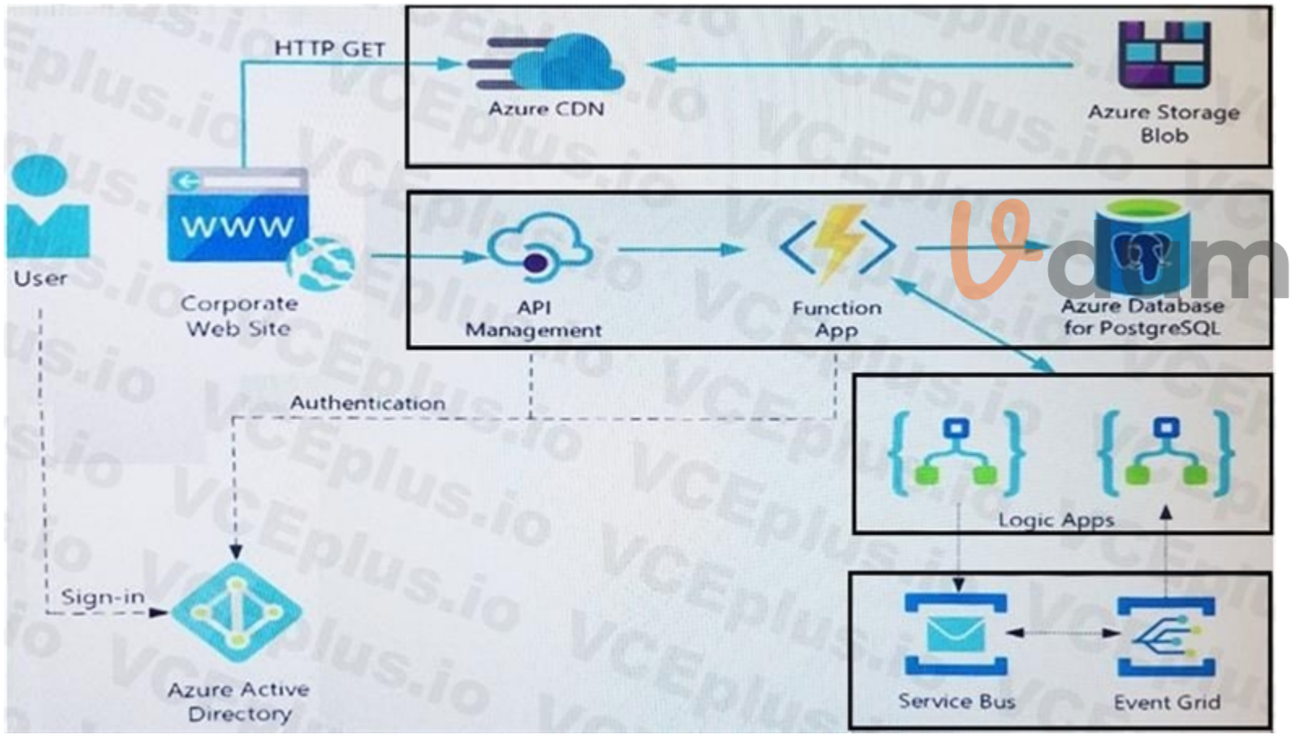

User authentication

The following steps detail the user authentication process:

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

Requirements

Corporate website

Communications and content must be secured by using SSL.

Communications must use HTTPS.

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

Data storage costs must be minimized.

Azure Database for PostgreSQL

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

Azure Key Vault name: cpandlkeyvault

Secret name: PostgreSQLConn

Id: 80df3e46ffcd4f1cb187f79905e9a1e8

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

Azure Service Bus and Azure Event Grid

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

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

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

Security

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

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

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

Compliance

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

Issues

Corporate website

While testing the site, the following error message displays:

CryptographicException: The system cannot find the file specified.

Function app

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

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

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

FunctionAppLogs

| where FunctionName = = "RequestUserApproval"

Logic app

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

'400 Bad Request'

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

Code

Corporate website

Security.cs:

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

Function app

RequestUserApproval.cs:

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

QUESTION 1
You need to correct the RequestUserApproval Function app error.
What should you do?

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

Correct Answer: C
Section:
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"
References:
https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-http-features

03 - Develop Azure compute solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.
Background
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.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22   private void EnsureLogging(string resource)
EG23   {
EG24     . . .
EG25   }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "input1",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
```

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.cs
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

QUESTION 1
You need to resolve a notification latency issue.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Set Always On to 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.

Correct Answer: A, B
Section:
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

04 - Develop Azure compute solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.
Background

Overview
You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded content from blob storage.
You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:
Messages are sent to ContentUploadService.
Content is processed by ContentAnalysisService.
After processing is complete, the content is posted to the social network or a rejection message is posted in its place.
The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.
The solution will use eight CPU cores.
Azure Active Directory
Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.
Requirements
ContentAnalysisService
The company's data science group built ContentAnalysisService which accepts user generated content as a string and returns a probable value for inappropriate content. Any values over a specific threshold must be reviewed by an employee of Contoso, Ltd.
You must create an Azure Function named CheckUserContent to perform the content checks.
Costs
You must minimize costs for all Azure services.
Manual review
To review content, the user must authenticate to the website portion of the ContentAnalysisService using their Azure AD credentials. The website is built using React and all pages and API endpoints require authentication. In order to review content a user must be part of a ContentReviewer role. All completed reviews must include the reviewer's email address for auditing purposes.
High availability
All services must run in multiple regions. The failure of any service in a region must not impact overall application availability.
Monitoring
An alert must be raised if the ContentUploadService uses more than 80 percent of available CPU cores.
Security
You have the following security requirements:
Any web service accessible over the Internet must be protected from cross site scripting attacks.
All websites and services must use SSL from a valid root certificate authority.
Azure Storage access keys must only be stored in memory and must be available only to the service.
All Internal services must only be accessible from internal Virtual Networks (VNets).
All parts of the system must support inbound and outbound traffic restrictions.
All service calls must be authenticated by using Azure AD.
User agreements
When a user submits content, they must agree to a user agreement. The agreement allows employees of Contoso, Ltd. to review content, store cookies on user devices, and track user's IP addresses.
Information regarding agreements is used by multiple divisions within Contoso, Ltd.
User responses must not be lost and must be available to all parties regardless of individual service uptime. The volume of agreements is expected to be in the millions per hour.
Validation testing
When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.
Issues
Users of the ContentUploadService report that they occasionally see HTTP 502 responses on specific pages.
Code
ContentUploadService

```
apiVersion: '2018-10-01'
type: Microsoft.ContainerInstance/containerGroups
location: westus
name: contentUploadService
properties:
    containers:
    - name: service
        properties:
        image: contoso/contentUploadService:latest
        ports:
        -  port: 80
           protocol: TCP
        resources:
          requests:
            cpu: 1.0
            memoryInGB: 1.5

ipAddress:
    ip: 10.23.121.112
    ports:
      - port: 80
        protocol: TCP

networkProfile:
id: /subscriptions/98...19/resourceGroups/container/providers/Microsoft.Network/networkProfiles/subnet
```

ApplicationManifest

```
{
    "id" : "2b079f03-9b06-2d44-98bb-e9182901fcb6",
    "appId" : "7118a7f0-b5c2-4c9d-833c-3d711396fe65",

    "createdDateTime" : "2019-12-24T06:01:44Z",
    "logoUrl" : null,
    "logoutUrl" : null,
    "name" : "ContentAnalysisService",


    "orgRestrictions" : [],
    "parentalControlSettings" : {
     "countriesBlockedForMinors" : [],
      "legalAgeGroupRule" : "Allow"
    },
    "passwordCredentials" : []
}
```

QUESTION 1
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.

Hot Area:

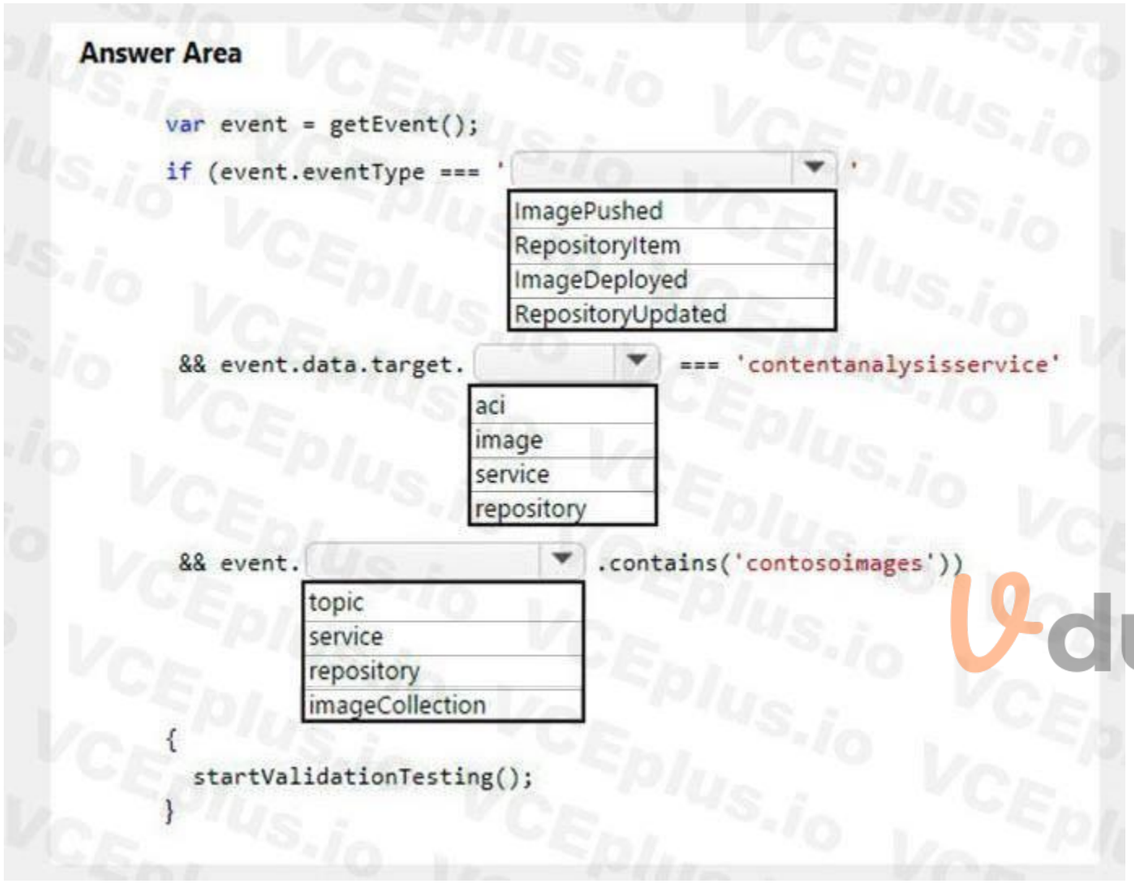

Answer Area:

Answer Area

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

Section:
Explanation:
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 2
You need to deploy the CheckUserContent Azure Function. The solution must meet the security and cost requirements.
Which hosting model should you use?

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

Correct Answer: B
Section:
Explanation:
Scenario:
You must minimize costs for all Azure services.

All Internal services must only be accessible from internal Virtual Networks (VNets).
Best for long-running scenarios where Durable Functions can't be used. Consider an App Service plan in the following situations:
You have existing, underutilized VMs that are already running other App Service instances.
You want to provide a custom image on which to run your functions.
Predictive scaling and costs are required.
Note: When you create a function app in Azure, you must choose a hosting plan for your app. There are three basic hosting plans available for Azure Functions: Consumption plan, Premium plan, and Dedicated (App Service) plan.
Incorrect Answers:
A: A Premium plan would be more costly.
C: Need the VNET functionality.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale

05 - Develop Azure compute solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.
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

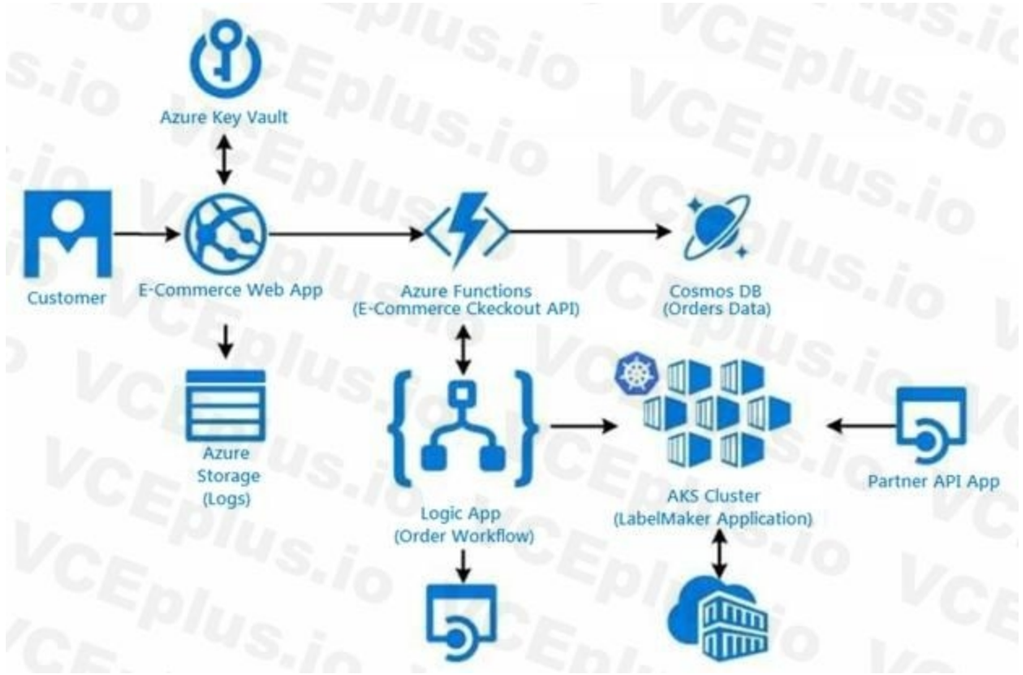

Issues
Calls to the Printer API App fail periodically due to printer communication timeouts.
Printer communication timeouts occur after 10 seconds. The label printer must only receive up to 5 attempts within one minute.
The order workflow fails to run upon initial deployment to Azure.
Order.json

```
{
  “id” : 1,
  “customers” : [
  {
    “familyName” : “Doe”,
    “givenName” : “John”,
    “customerid” : 5
  }
  ],
  “line_items” : [
  {
    “fulfillable_quantity” : 1,
    “id” : 6,
    “price” : “199.99” ,
    “product_id” : 7513594,
    “quantity”: 1,
    “requires_shipping” : true ,
    “sku” : “SFC-342-N” ,
    “title”: “Surface Go” ,
    “vendor” : “Microsoft” ,
    “name” : “Surface Go – 8GB” ,
    “taxable” : true ,
    “tax_lines” : [
    {
      “title” : “State Tax” ,
      “price” : “3.98” ,
      “rate” : 0.06
     }
    ],
    “total_discount” : “5.00” ,
    “discount_allocations” : [
     {
       “amount” : “5.00” ,
       “discount_application_index” : 2
       }
    ]
   }
  ],
  “address” : {
  “state” : “NY” ,
  “state”: “Manhattan” ,
  “city” : “NY”
  }
}
```

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

QUESTION 1
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.

Correct Answer: C
Section:
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-azure-functions-app-df98c40424bb

QUESTION 2
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.

Select and Place:

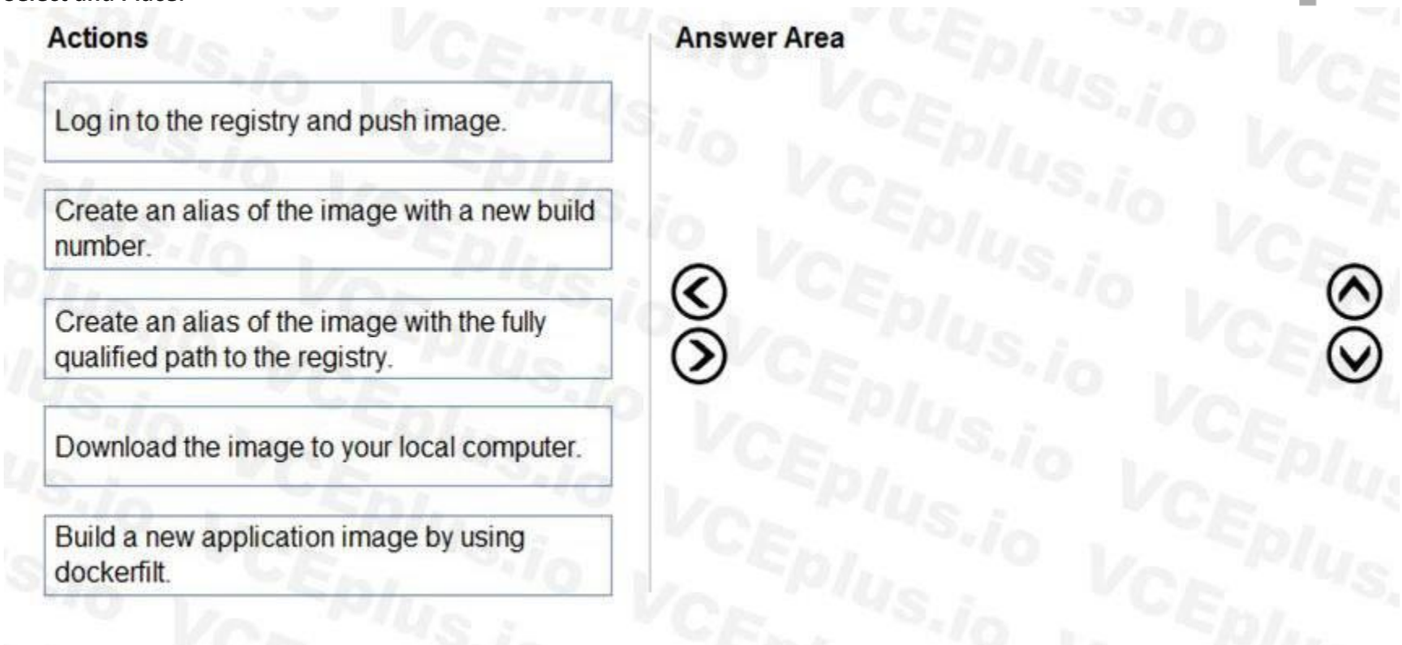

Correct Answer:

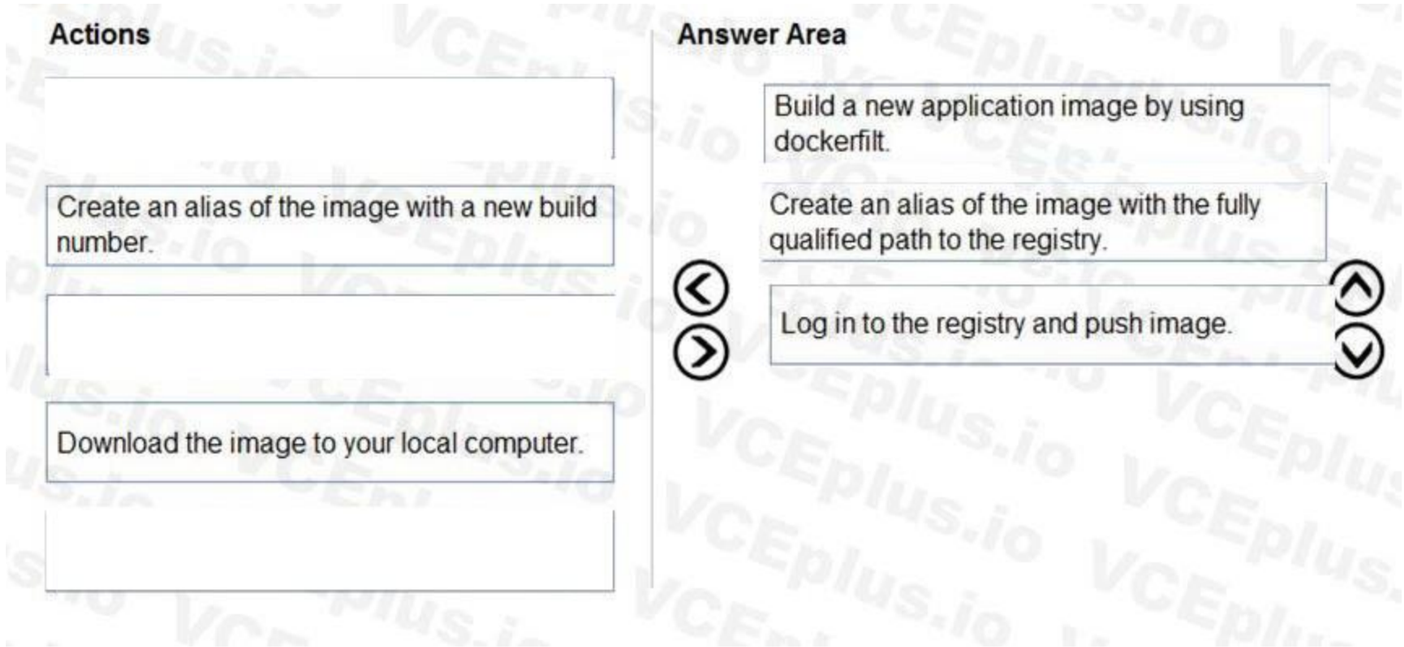

Section:
Explanation:
Step 1: Build a new application image by using dockerfile
Step 2: Create an alias if the image with the fully qualified path to the registry Before you can push the image to a private registry, you've to ensure a proper image name. This can be achieved using the docker tag command. 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 form 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

06 - Develop Azure compute solutions

QUESTION 1
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.

Hot Area:

Answer Area

Property	Value
Client certificate location	HTTP request header Client cookie HTTP message body URL query string
Encoding type	HTML URL Unicode Base64

Answer Area:

Answer Area

Property	Value
Client certificate location	HTTP request header Client cookie HTTP message body URL query string
Encoding type	HTML URL Unicode Base64

Section:

Explanation:

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 2
DRAG DROP
You are developing a Docker/Go using Azure App Service Web App for Containers. You plan to run the container in an App Service on Linux. You identify a Docker container image to use.
None of your current resource groups reside in a location that supports Linux. You must minimize the number of resource groups required.
You need to create the application and perform an initial deployment.
Which three Azure CLI commands should you use to develop the solution? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

Select and Place:

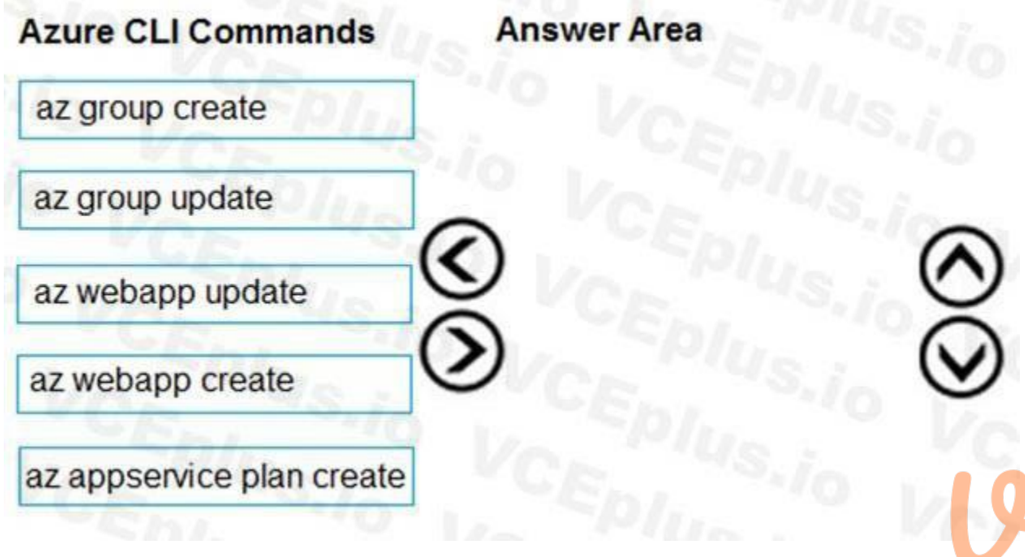

Correct Answer:

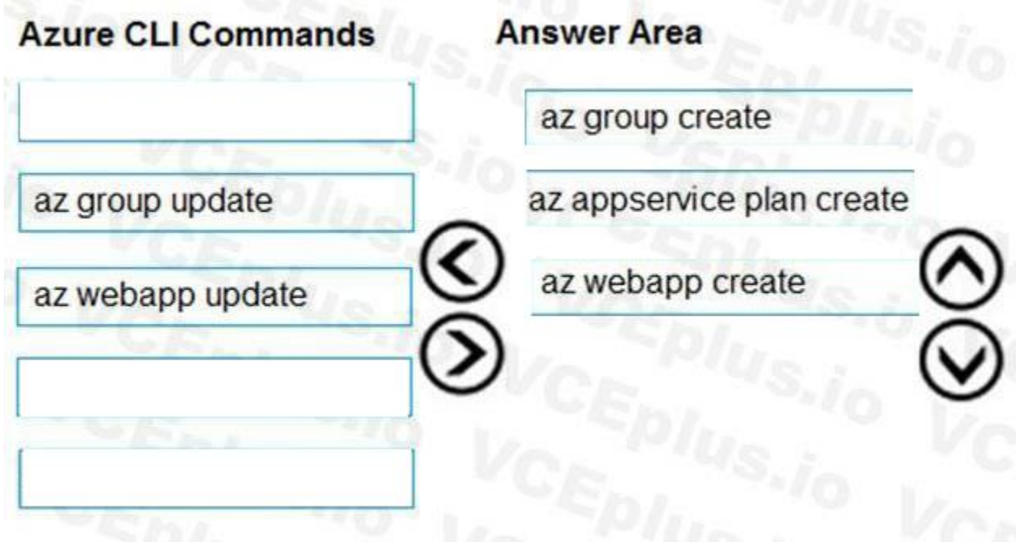

Section:
Explanation:
You can host native Linux applications in the cloud by using Azure Web Apps. To create a Web App for Containers, you must run Azure CLI commands that create a group, then a service plan, and finally the web app itself.
Step 1: az group create
In the Cloud Shell, create a resource group with the az group create command.
Step 2: az appservice plan create

In the Cloud Shell, create an App Service plan in the resource group with the az appservice plan create command.
Step 3: az webapp create
In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command. Don't forget to replace with a unique app name, and <docker-ID> with your Docker ID.
Reference:
https://docs.microsoft.com/mt-mt/azure/app-service/containers/quickstart-docker-go?view=sql-server-ver15

QUESTION 3
DRAG DROP
Fourth Coffee has an ASP.NET Core web app

Select and Place:

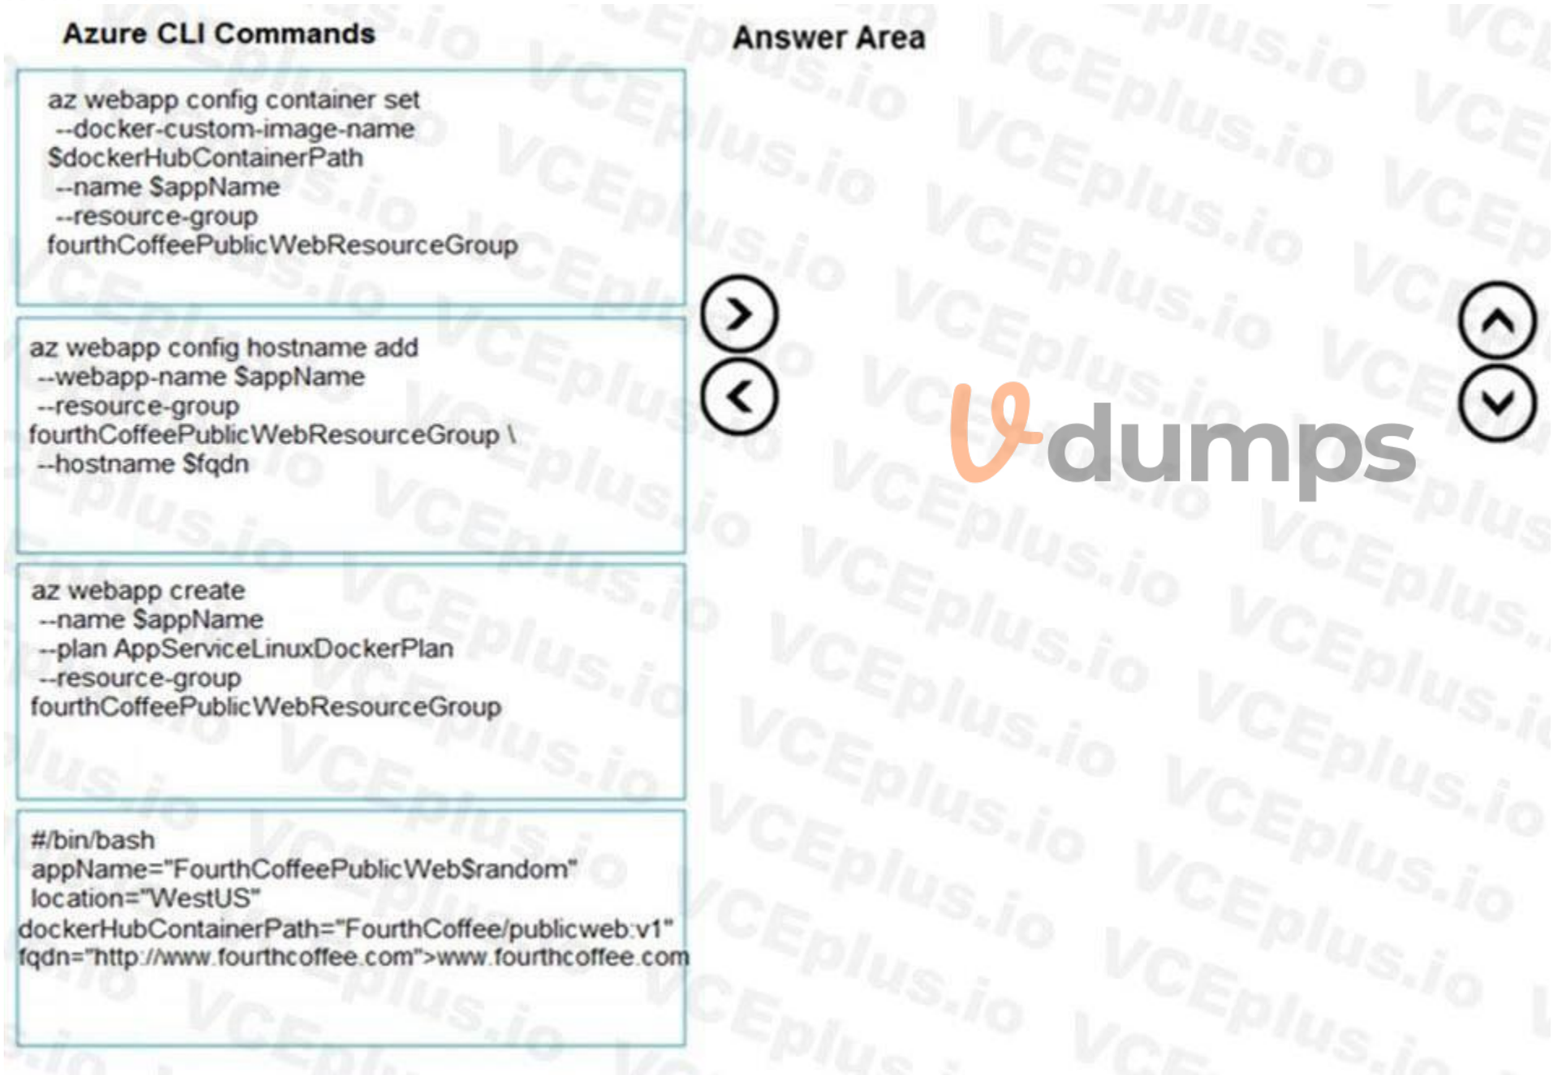

Correct Answer:

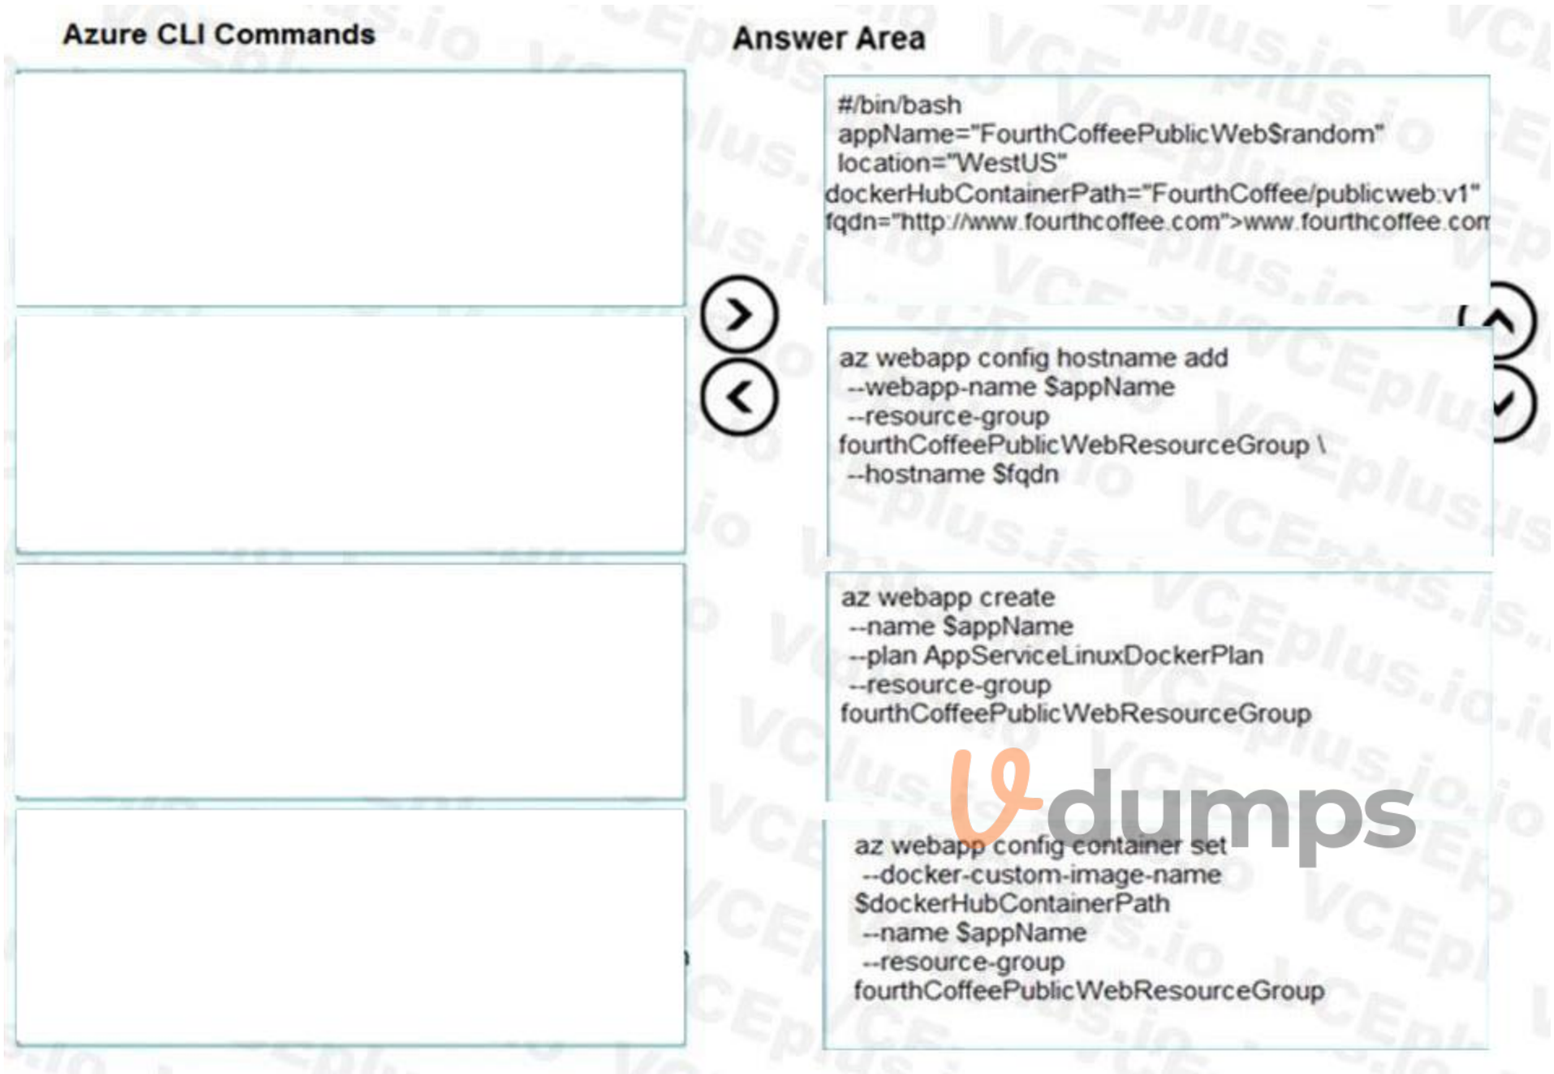

Section:
Explanation:
Step 1: #bin/bash
The appName is used when the webapp-name is created in step 2.
Step 2: az webapp config hostname add
The webapp-name is used when the webapp is created in step 3.
Step 3: az webapp create
Create a web app. In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command.
Step : az webapp confing 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 4

DRAG DROP
You are developing a serverless Java application on Azure. You create a new Azure Key Vault to work with secrets from a new Azure Functions application.
The application must meet the following requirements:
Reference the Azure Key Vault without requiring any changes to the Java code.
Dynamically add and remove instances of the Azure Functions host based on the number of incoming application events.
Ensure that instances are perpetually warm to avoid any cold starts.
Connect to a VNet.
Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.
You need to grant the Azure Functions application access to the Azure Key Vault.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

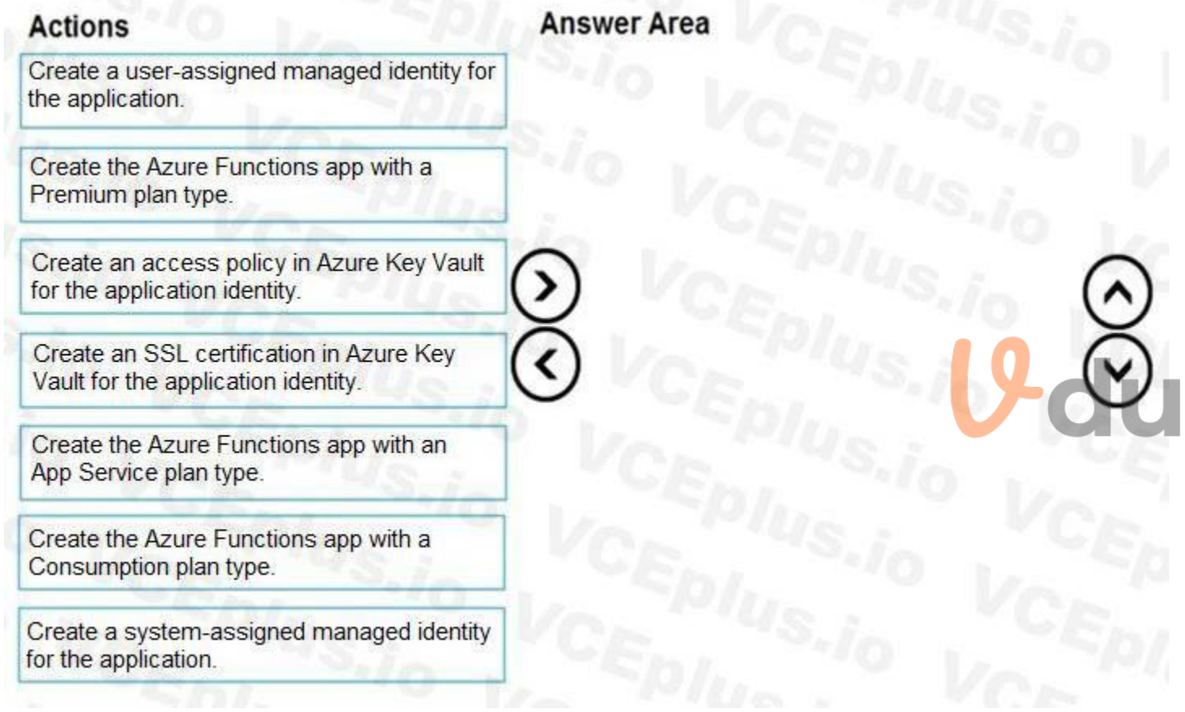

Correct Answer:

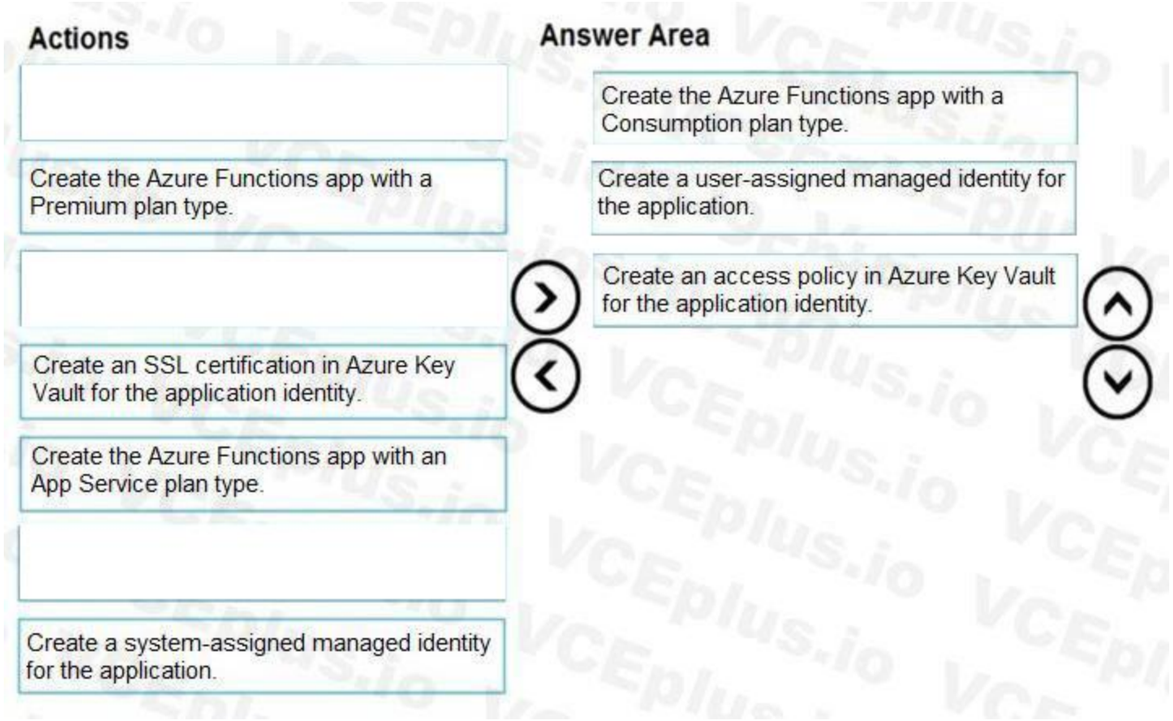

Section:

Explanation:
Step 1: Create the Azure Functions app with a Consumption plan type.
Use the Consumption plan for serverless.
Step 2: Create a system-assigned managed identity for the application.
Create a system-assigned managed identity for your application.
Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.
Step 3: Create an access policy in Key Vault for the application identity.
Create an access policy in Key Vault for the application identity you created earlier. Enable the "Get" secret permission on this policy. Do not configure the "authorized application" or applicationId settings, as this is not compatible with a managed identity.
Reference:
https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references

QUESTION 5
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Move photo processing to an Azure Function triggered from the blob upload.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: A

Section:
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 6
You are developing an application that uses Azure Blob storage.
The application must read the transaction logs of all the changes that occur to the blobs and the blob metadata in the storage account for auditing purposes. The changes must be in the order in which they occurred, include only create, update, delete, and copy operations and be retained for compliance reasons.
You need to process the transaction logs asynchronously.
What should you do?

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

Correct Answer: B
Section:
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 7
You are developing a web app that is protected by Azure Web Application Firewall (WAF). All traffic to the web app is routed through an Azure Application Gateway instance that is used by multiple web apps. The web app address is contoso.azurewebsites.net.
All traffic must be secured with SSL. The Azure Application Gateway instance is used by multiple web apps.
You need to configure the Azure Application Gateway for the web app.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

Correct Answer: A, D
Section:
Explanation:
D: The ability to specify a host override is defined in the HTTP settings and can be applied to any back-end pool during rule creation.
The ability to derive the host name from the IP or FQDN of the back-end pool members. HTTP settings also provide an option to dynamically pick the host name from a back-end pool member's FQDN if configured with the option to derive host name from an individual back-end pool member.
A (not C): SSL termination and end to end SSL with multi-tenant services.
In case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway. Therefore, there is no need to add any authentication certificates.

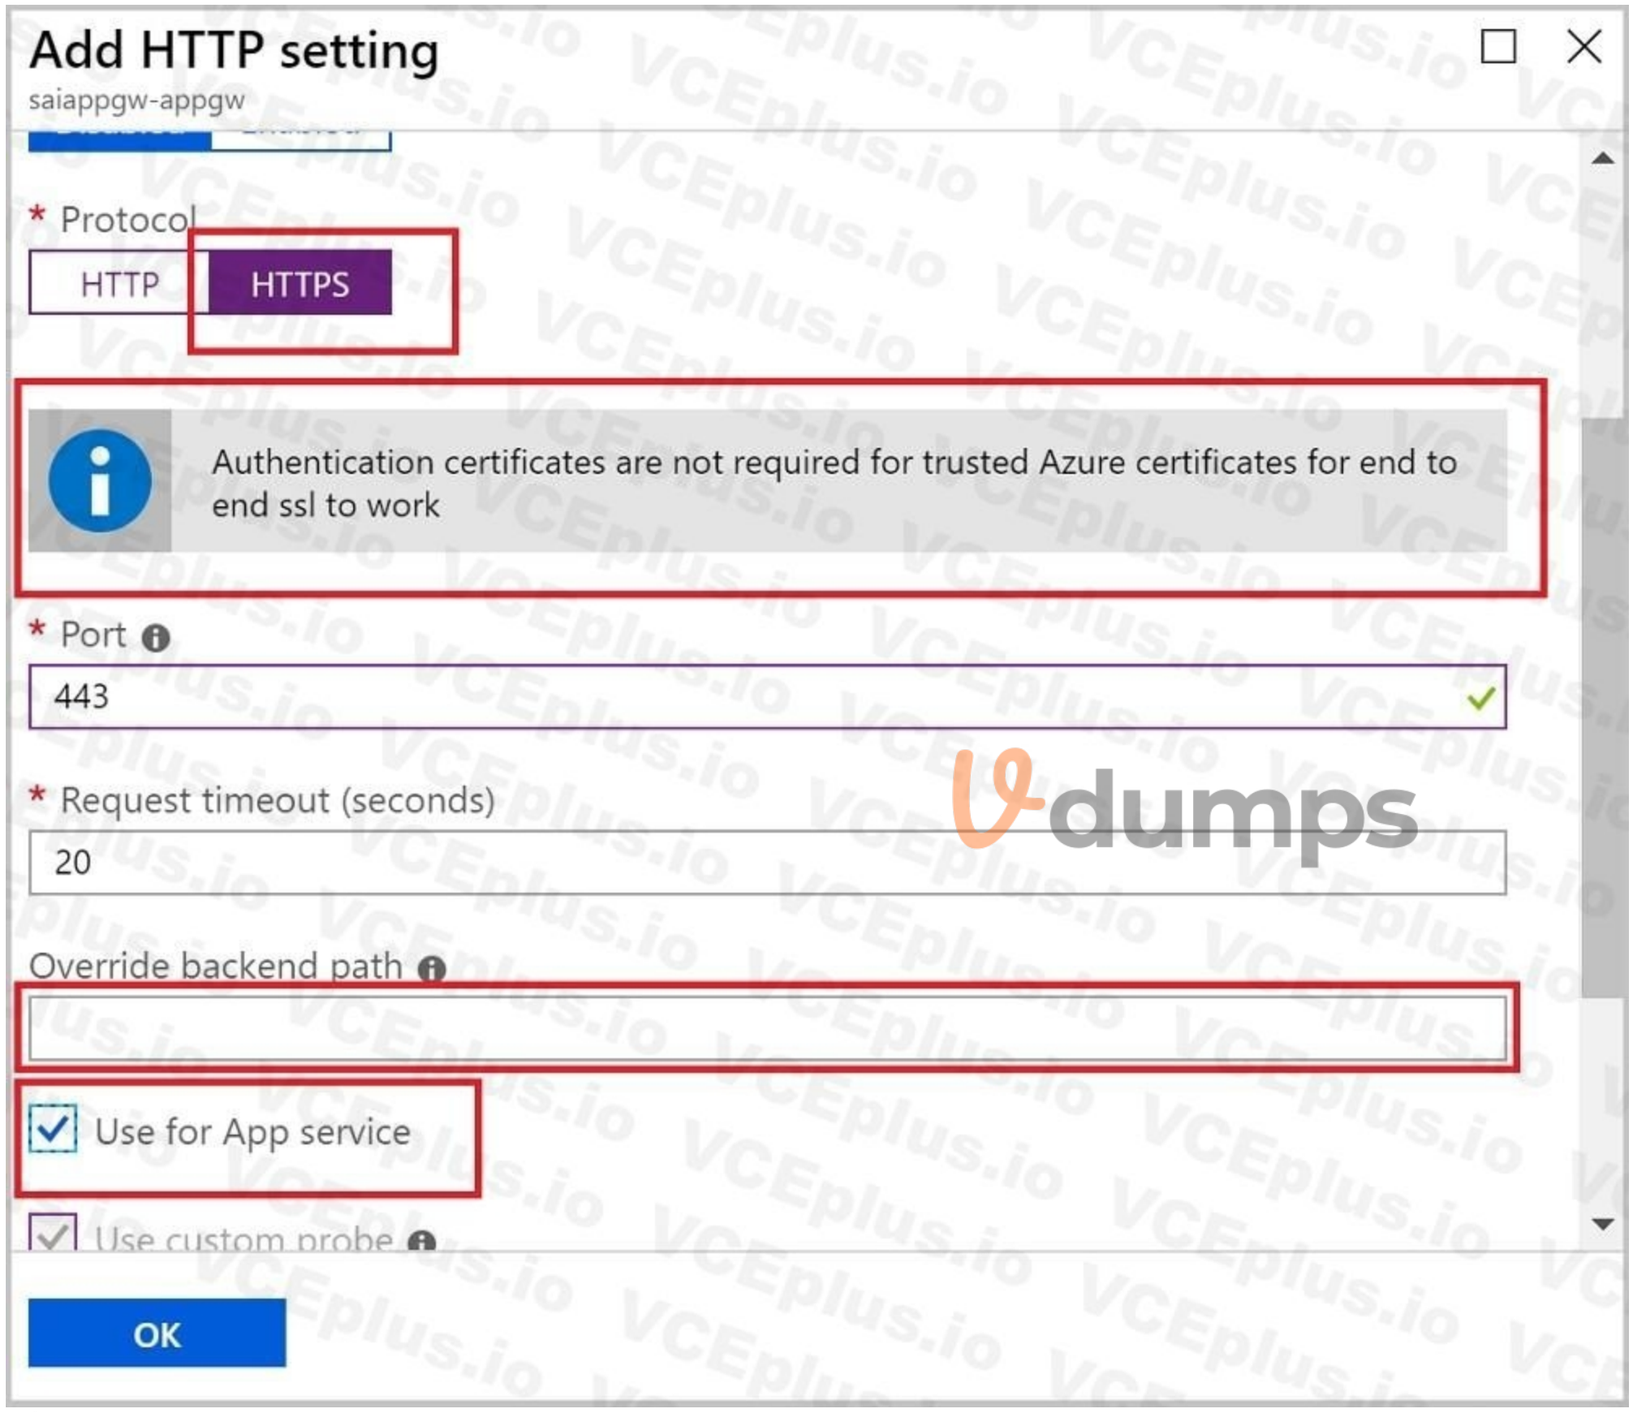

Reference:
https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-web-app-overview

QUESTION 8
Note: This question-is part of a series of questions that present the same scenario. Each question-in the series contains a unique solution that might meet the stated goals. Some question-sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question-in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Trigger the photo processing from Blob storage events.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
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 9
Note: This question-is part of a series of questions that present the same scenario. Each question-in the series contains a unique solution that might meet the stated goals. Some question-sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question-in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop 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

Correct Answer: A
Section:
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 10
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.

Select and Place:

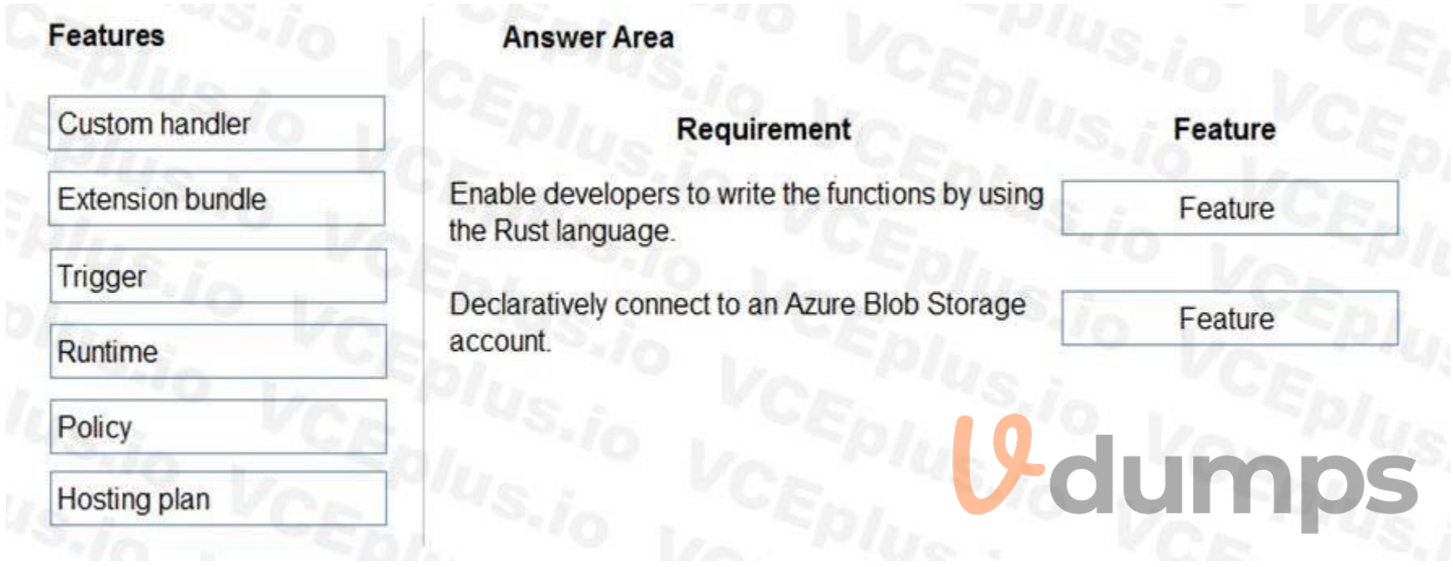

Correct Answer:

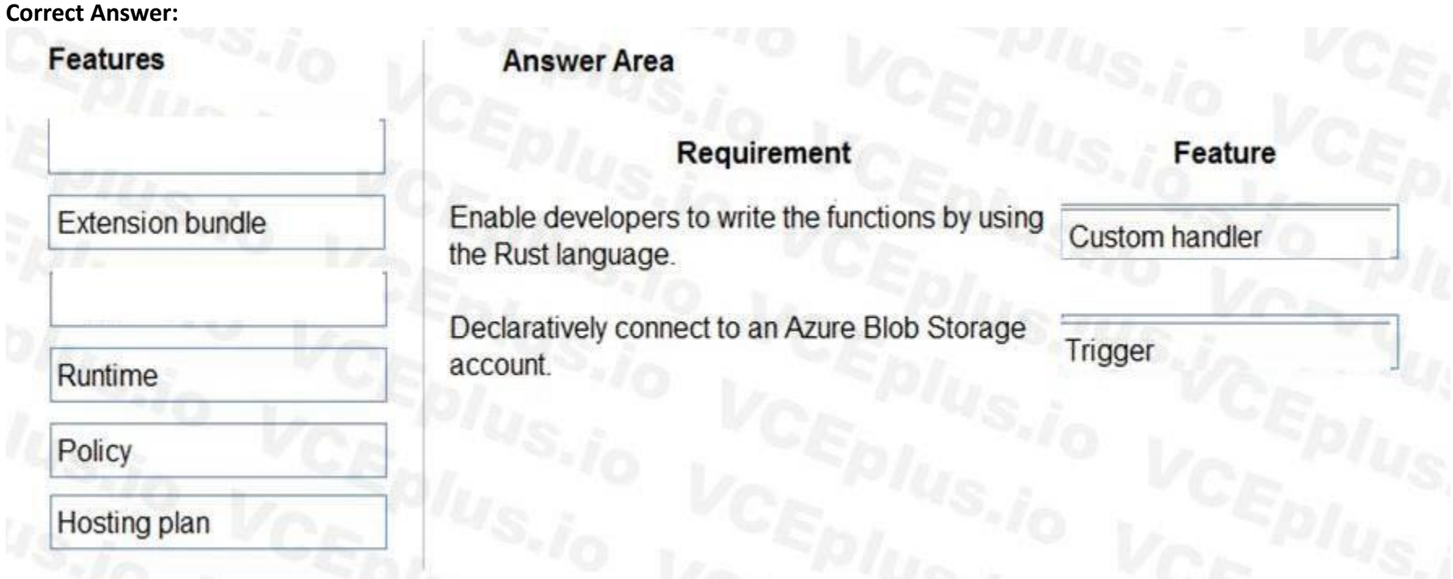

Section:
Explanation:
Box 1: Custom handler

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 11
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You 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

Correct Answer: B
Section:
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 12
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. Add the customer ID for the signed in user to the CorrelationContext in the web application
B. On the current SpanContext, set the TraceId to the customer ID for the signed in user
C. Set the header Ocp-Apim-Trace to the customer ID for the signed in user
D. Create a new SpanContext with the TraceFlags value set to the customer ID for the signed in user

Correct Answer: A
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/correlation

QUESTION 13

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

Correct Answer: B
Section:
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 14
Note: This question-is part of a series of questions that present the same scenario. Each question-in the series contains a unique solution that might meet the stated goals. Some question-sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question-in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You 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

Correct Answer: B
Section:
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 15
Note: This question-is part of a series of questions that present the same scenario. Each question-in the series contains a unique solution that might meet the stated goals. Some question-sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question-in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Convert the Azure Storage account to a BlockBlobStorage storage account.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
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 16
DRAG DROP
You are developing an application to use Azure Blob storage. You have configured Azure Blob storage to include change feeds.
A copy of your storage account must be created in another region. Data must be copied from the current storage account to the new storage account directly between the storage servers.
You need to create a copy of the storage account in another region and copy the data.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

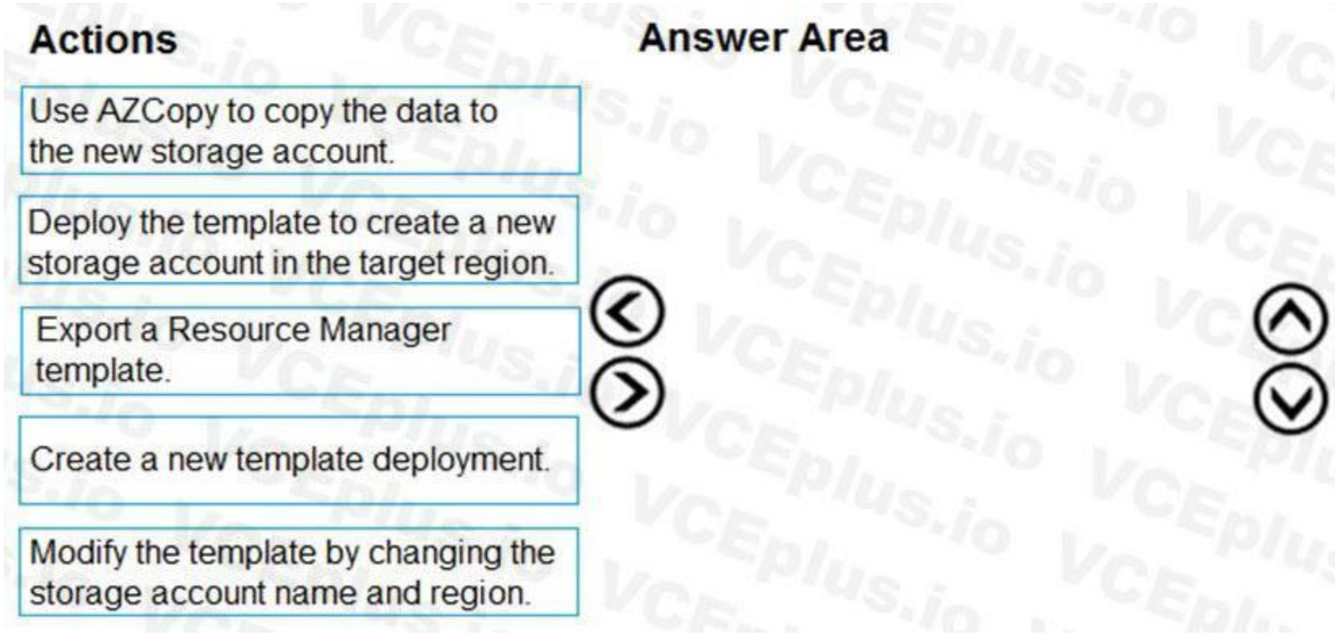

Correct Answer:

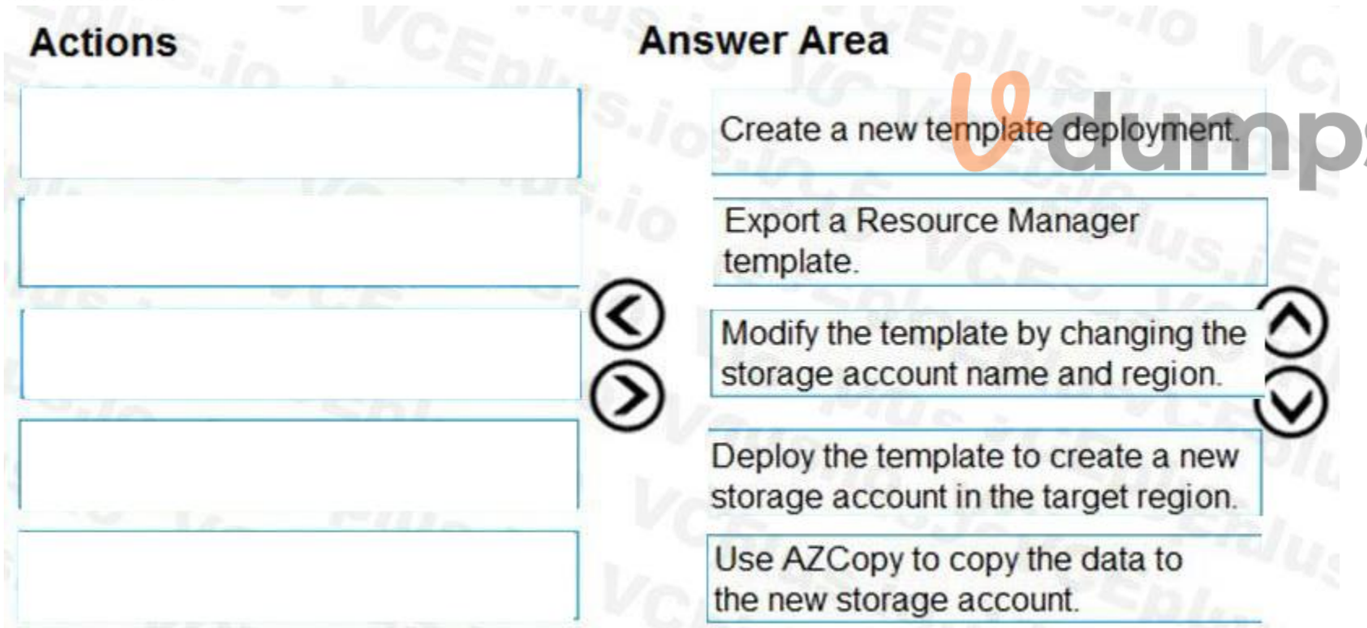

Section:

Explanation:

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

The steps are:

Export a template.

Modify the template by adding the target region and storage account name.

Deploy the template to create the new storage account.

Configure the new storage account.

Move data to the new storage account.

Delete the resources in the source region.
Note: You must enable the change feed on your storage account to begin capturing and recording changes. You can enable and disable changes by using Azure Resource Manager templates on Portal or Powershell.
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed

QUESTION 17
HOTSPOT
You are developing an ASP.NET Core web application. You plan to deploy the application to Azure Web App for Containers.
The application needs to store runtime diagnostic data that must be persisted across application restarts. You have the following code:

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

You need to configure the application settings so that diagnostic data is stored as required.
How should you configure the web app's settings? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

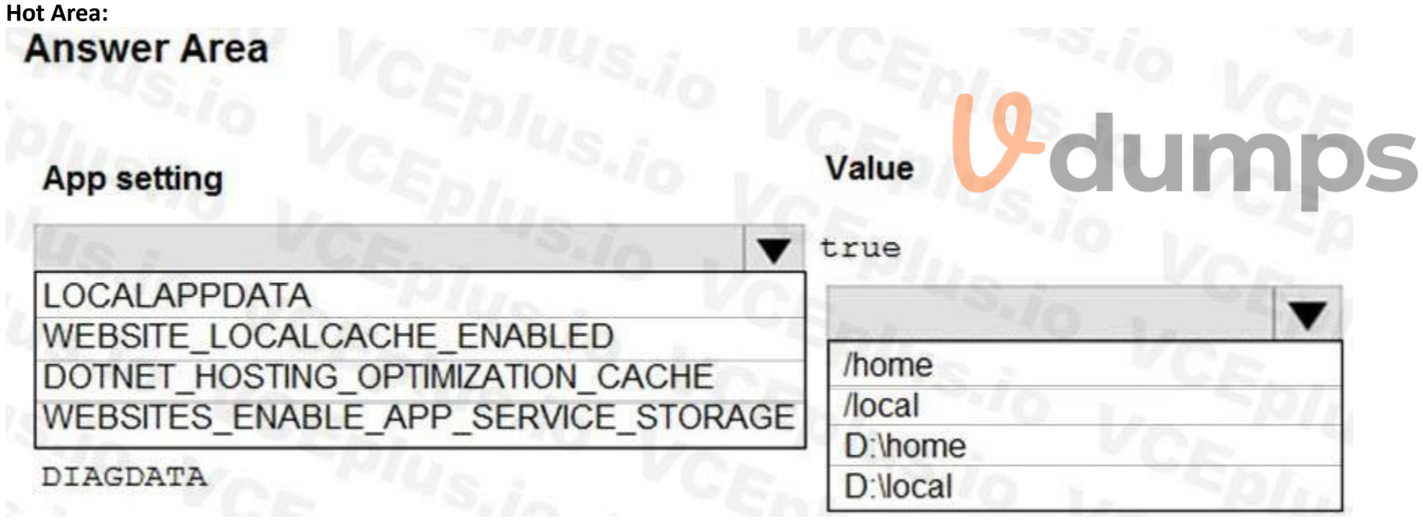

Answer Area:

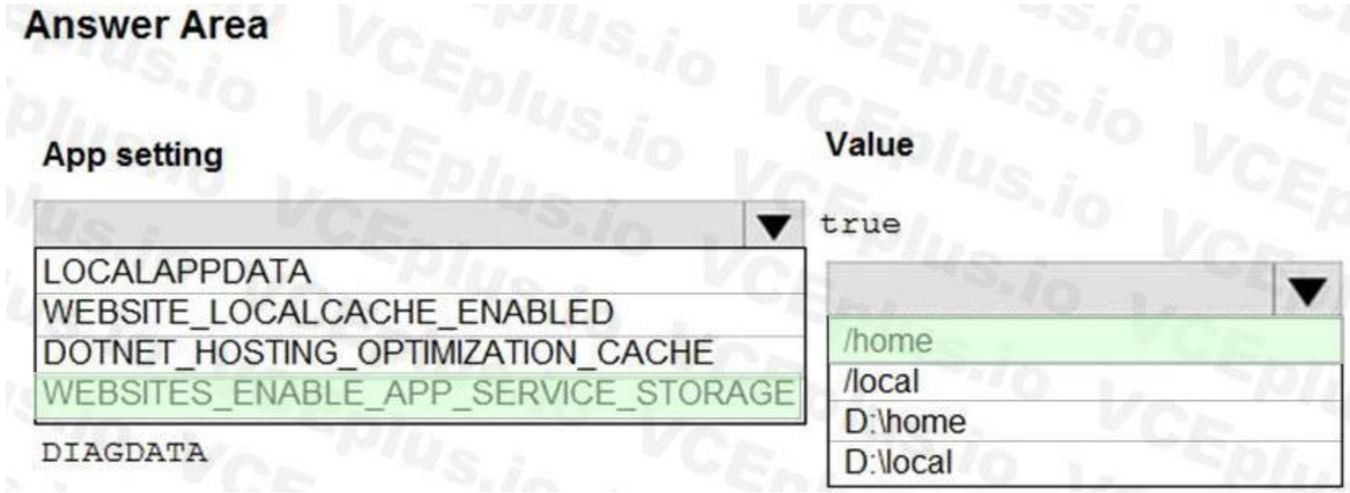

Section:
Explanation:
Box 1: If WEBSITES_ENABLE_APP_SERVICE_STORAGE
If WEBSITES_ENABLE_APP_SERVICE_STORAGE setting is unspecified or set to true, the /home/ directory will be shared across scale instances, and files written will persist across restarts
Box 2: /home
Reference:
https://docs.microsoft.com/en-us/azure/app-service/containers/app-service-linux-faq

QUESTION 18
HOTSPOT
You are implementing a software as a service (SaaS) ASP.NET Core web service that will run as an Azure Web App. The web service will use an on-premises SQL Server database for storage. The web service also includes a WebJob that processes data updates. Four customers will use the web service.
Each instance of the WebJob processes data for a single customer and must run as a singleton instance.
Each deployment must be tested by using deployment slots prior to serving production data.
Azure costs must be minimized.
Azure resources must be located in an isolated network.
You need to configure the App Service plan for the Web App.
How should you configure the App Service plan? To answer, select the appropriate settings in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

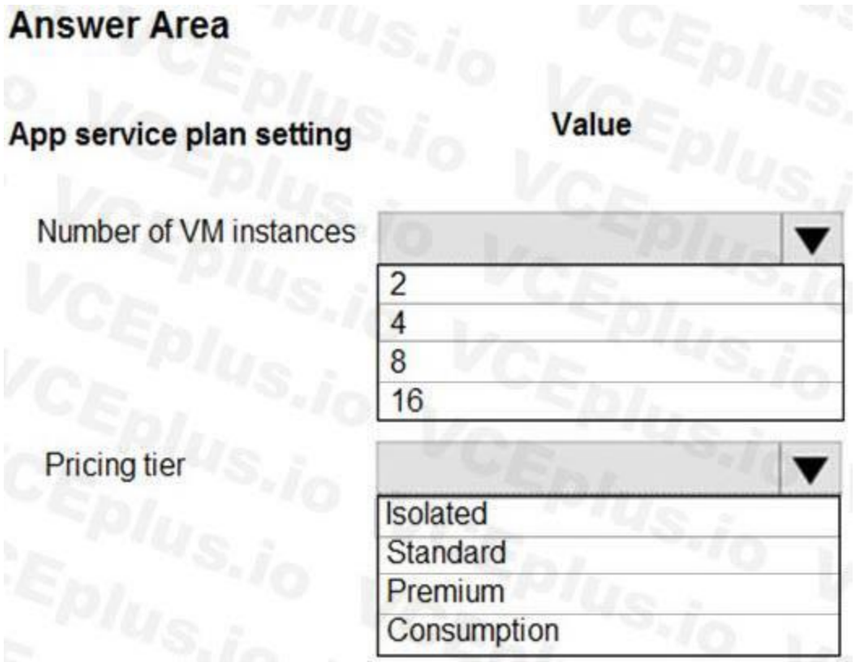

Answer Area:

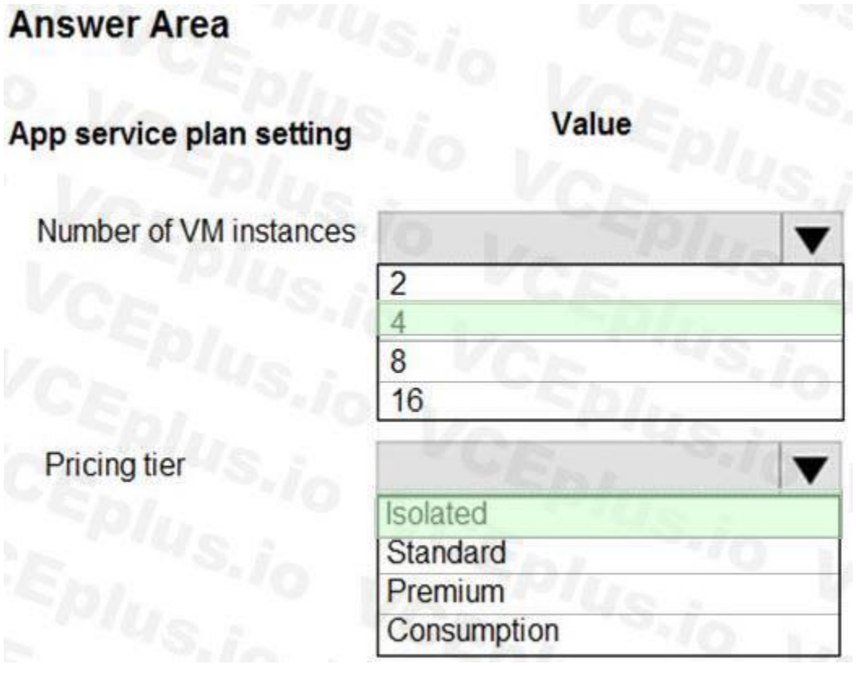

Section:

Explanation:

Number of VM instances: 4

You are not charged extra for deployment slots.

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

Azure Virtual Network (VNet).

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

QUESTION 19
DRAG DROP
You are a developer for a software as a service (SaaS) company that uses an Azure Function to process orders. The Azure Function currently runs on an Azure Function app that is triggered by an Azure Storage queue.
You are preparing to migrate the Azure Function to Kubernetes using Kubernetes-based Event Driven Autoscaling (KEDA).
You need to configure Kubernetes Custom Resource Definitions (CRD) for the Azure Function.
Which CRDs should you configure? To answer, drag the appropriate CRD types to the correct locations. Each CRD type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Select and Place:

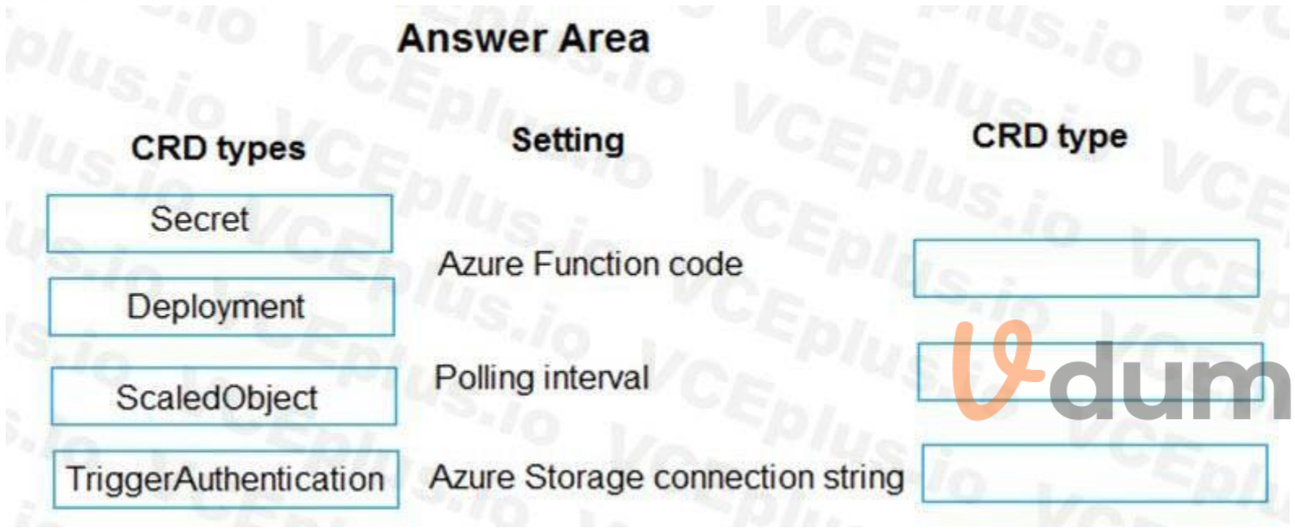

Correct Answer:

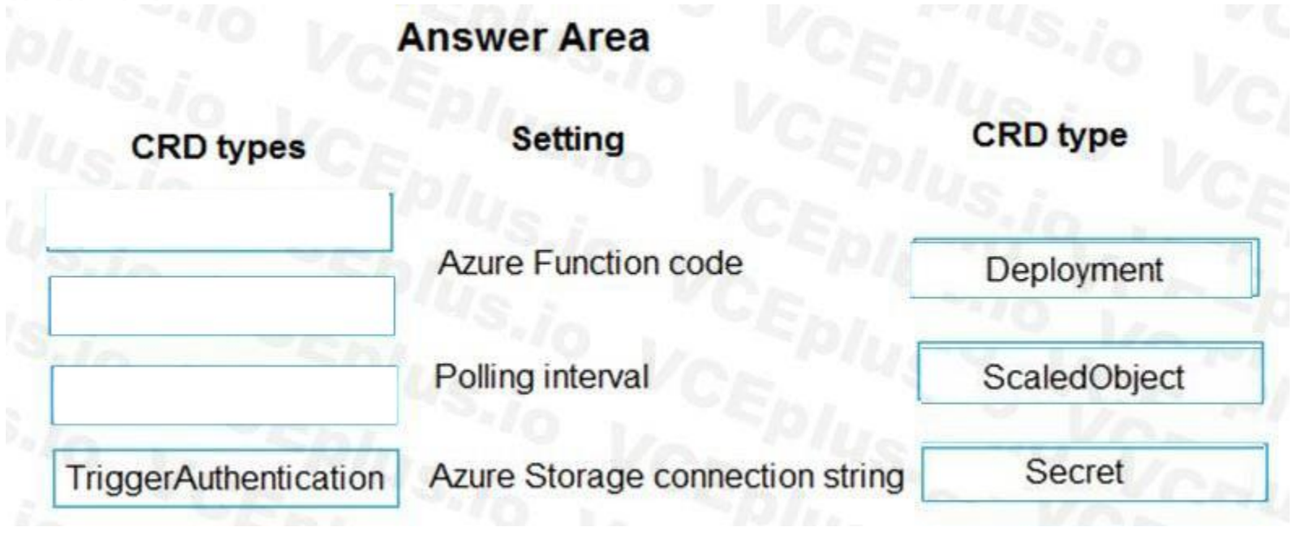

Section:
Explanation:

Box 1: Deployment
To deploy Azure Functions to Kubernetes use the func kubernetes deploy command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.
Box 2: ScaledObject
With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.
Example of ScaledObject with polling interval
apiVersion: keda.k8s.io/v1alpha1
kind: ScaledObject
metadata:
name: transformer-fn
namespace: tt
labels:
deploymentName: transformer-fn
spec:
scaleTargetRef:
deploymentName: transformer-fn
pollingInterval: 5
minReplicaCount: 0
maxReplicaCount: 100
Box 3: Secret
Store connection strings in Kubernetes Secrets.
Example: to create the Secret in our demo Namespace:
create the k8s demo namespace
kubectl create namespace tt
# grab connection string from Azure Service Bus
KEDA_SCALER_CONNECTION_STRING=$(az servicebus queue authorization-rule keys list \
-g $RG_NAME \
--namespace-name $SBN_NAME \
--queue-name inbound \
-n keda-scaler \
--query "primaryConnectionString" \
-o tsv)
create the kubernetes secret
kubectl create secret generic tt-keda-auth \
--from-literal KedaScaler=$KEDA_SCALER_CONNECTION_STRING \
--namespace tt
Reference:
https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/

QUESTION 20
HOTSPOT
You are creating a CLI script that creates an Azure web app and related services in Azure App Service. The web app uses the following variables:

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

You need to automatically deploy code from GitHub to the newly created web app.
How should you complete the script? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

[▼] `--name $webappname --resource-group myResourceGroup --sku FREE`

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

`--resource-group myResourceGroup` [▼]

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

Answer Area:

Answer Area

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

[▼] `--name $webappname --resource-group myResourceGroup --sku FREE`

- az webapp
- 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
- 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

Section:
Explanation:
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 21
You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published.
You must ensure that the website remains available and responsive while minimizing cost.
You need to deploy the website.
What should you do?

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

Correct Answer: D
Section:
Explanation:
Windows Azure Web Sites (WAWS) offers 3 modes: Standard, Free, and Shared.
Standard mode carries an enterprise-grade SLA (Service Level Agreement) of 99.9% monthly, even for sites with just one instance.
Standard mode runs on dedicated instances, making it different from the other ways to buy Windows Azure Web Sites.
Incorrect Answers:
B: Shared and Free modes do not offer the scaling flexibility of Standard, and they have some important limits.
Shared mode, just as the name states, also uses shared Compute resources, and also has a CPU limit. So, while neither Free nor Shared is likely to be the best choice for your production environment due to these limits.

QUESTION 22
HOTSPOT
A company is developing a Java web app. The web app code is hosted in a GitHub repository located at https://github.com/Contoso/webapp.
The web app must be evaluated before it is moved to production. You must deploy the initial code release to a deployment slot named staging.
You need to create the web app and deploy the code.
How should you complete the commands? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

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

az [▼] create --location centralus --name $resourcegroupname

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

az [▼] create --name $webappname --resource-group $resourcegroupname --sku S3

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

az [▼] create --name $webappname --resource-group $resourcegroupname --plan $webappname

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

az [▼] create --name $webappname --resource-group $resourcegroupname --slot staging

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

az [▼] config --name $webappname --resource-group $resourcegroupname \
--slot staging --repo-url $gitrepo --branch master --manual-integration

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

Answer Area:

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

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] create --location centralus --name $resourcegroupname

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] create --name $webappname --resource-group $resourcegroupname --sku S3

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] create --name $webappname --resource-group $resourcegroupname --plan $webappname

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] create --name $webappname --resource-group $resourcegroupname --slot staging

az [group / webapp / appservice plan / webapp deployment slot / webapp deployment source] config --name $webappname --resource-group $resourcegroupname \
--slot staging --repo-url $gitrepo --branch master --manual-integration

Section:
Explanation:
Box 1: group
Create a resource group.
az group create --location westeurope --name myResourceGroup
Box 2: appservice plan
Create an App Service plan in STANDARD tier (minimum required by deployment slots).
az appservice plan create --name $webappname --resource-group myResourceGroup --sku S1
Box 3: webapp
Create a web app.
az webapp create --name $webappname --resource-group myResourceGroup \
--plan $webappname
Box 4: webapp deployment slot
#Create a deployment slot with the name "staging".

az webapp deployment slot create --name $webappname --resource-group myResourceGroup \
--slot staging
Box 5: webapp deployment source
Deploy sample code to "staging" slot from GitHub.
az webapp deployment source config --name $webappname --resource-group myResourceGroup \
--slot staging --repo-url $gitrepo --branch master --manual-integration
Reference:
https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment

QUESTION 23
HOTSPOT
You have a web service that is used to pay for food deliveries. The web service uses Azure Cosmos DB as the data store.
You plan to add a new feature that allows users to set a tip amount. The new feature requires that a property named tip on the document in Cosmos DB must be present and contain a numeric value.
There are many existing websites and mobile apps that use the web service that will not be updated to set the tip property for some time.
How should you complete the trigger?
NOTE: Each correct selection is worth one point.

Hot Area:

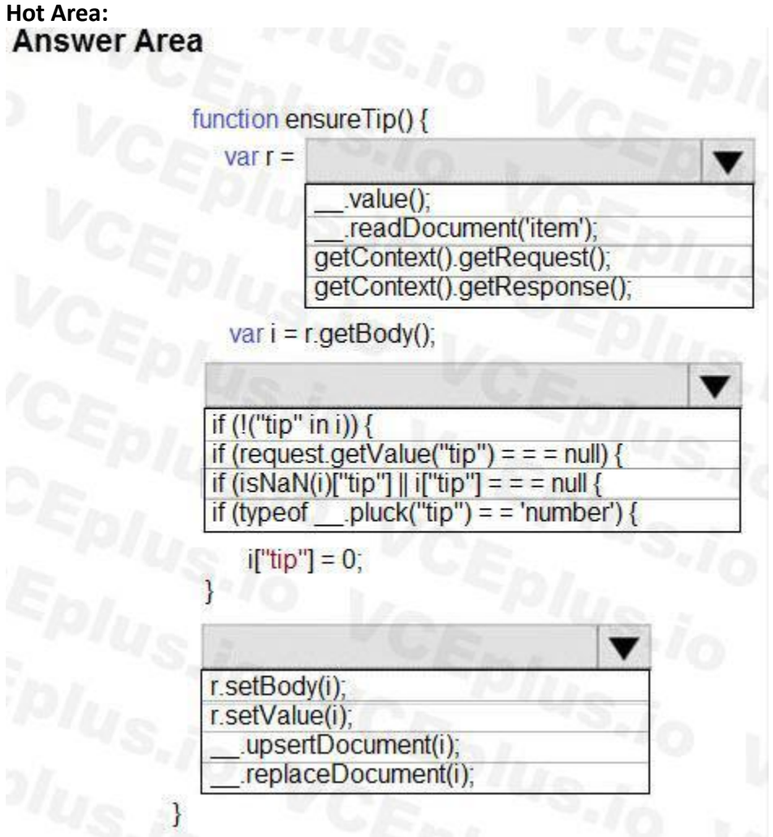

Answer Area:

Answer Area

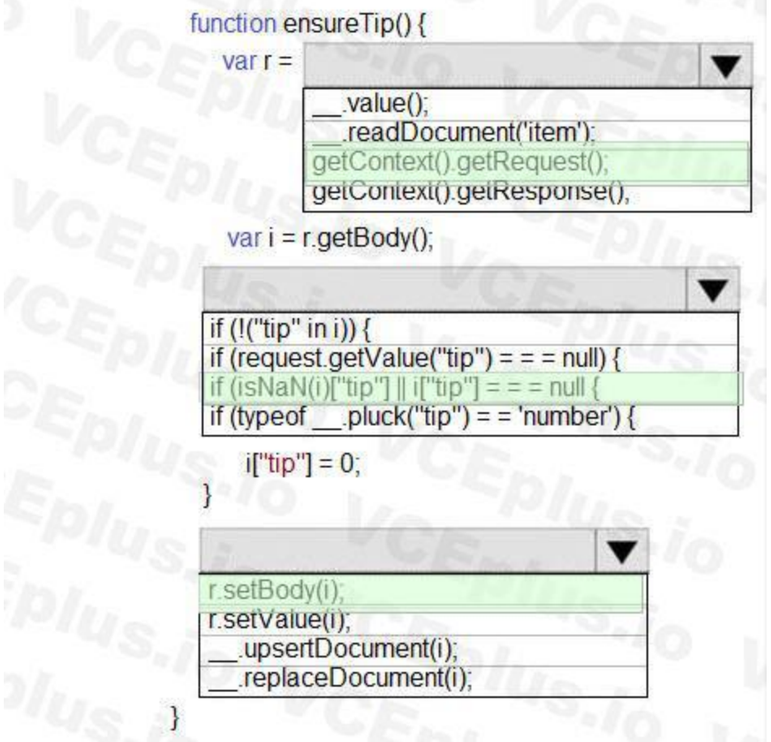

Section:
Explanation:

QUESTION 24
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data.
You need to ensure the app does not time out and processes the blob data.
Solution: Use the Durable Function async pattern to process the blob data.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
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 25
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data.
You need to ensure the app does not time out and processes the blob data.
Solution: Pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: A
Section:
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 26
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob.
The app continues to time out after four minutes. The app must process the blob data.
You need to ensure the app does not time out and processes the blob data.
Solution: Configure the app to use an App Service hosting plan and enable the Always On setting.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
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 27
DRAG DROP
You plan to create a Docker image that runs an ASP.NET Core application named ContosoApp. You have a setup script named setupScript.ps1 and a series of application files including ContosoApp.dll.
You need to create a Dockerfile document that meets the following requirements:
Call setupScripts.ps1 when the container is built.
Run ContosoApp.dll when the container starts.
The Dockerfile document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.
Which five commands should you use to develop the solution? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

Select and Place:

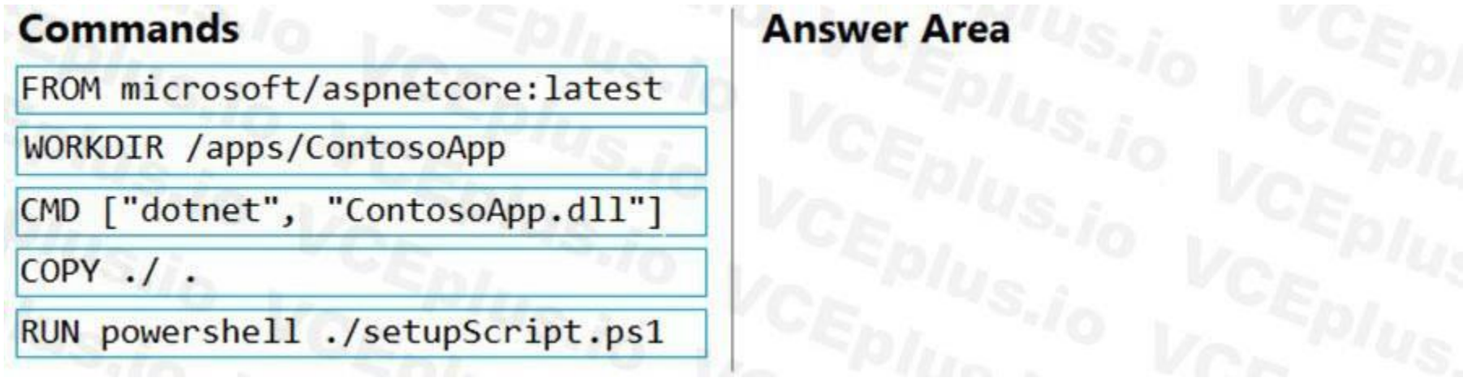

Correct Answer:

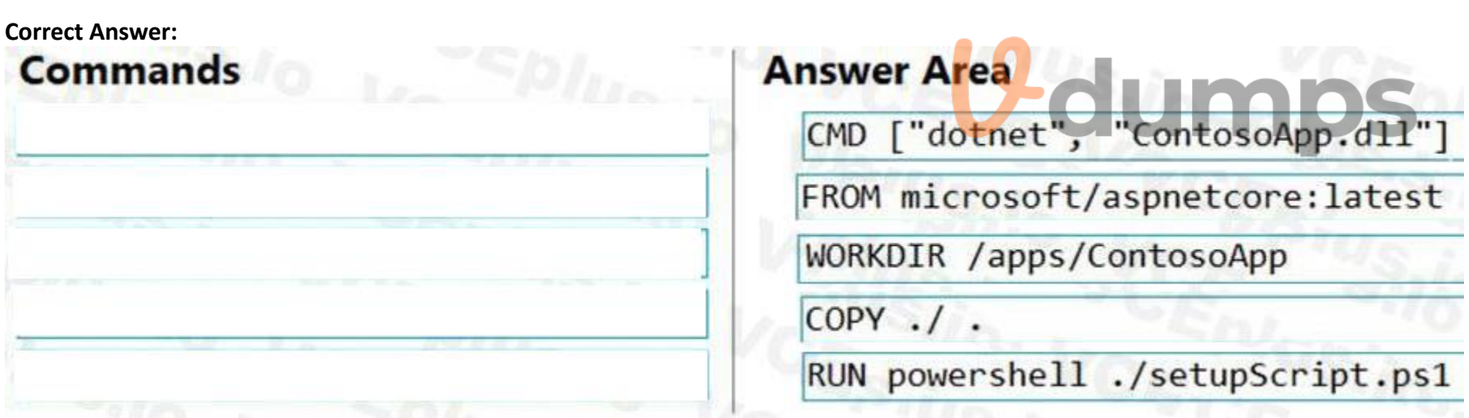

Section:
Explanation:
Box 1: CMD [..]
Cmd starts a new instance of the command interpreter, Cmd.exe.
Syntax: CMD <string>
Specifies the command you want to carry out.
Box 2: FROM microsoft/aspnetcore-build:latest
Box 3: WORKDIR /apps/ContosoApp
Bxo 4: COPY ./ .
Box 5: RUN powershell ./setupScript.ps1

QUESTION 28
You are developing an Azure Function App that processes images that are uploaded to an Azure Blob container.
Images must be processed as quickly as possible after they are uploaded, and the solution must minimize latency. You create code to process images when the Function App is triggered.

You need to configure the Function App.
What should you do?

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

Correct Answer: B
Section:
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 29
HOTSPOT
You are configuring a new development environment for a Java application.
The environment requires a Virtual Machine Scale Set (VMSS), several storage accounts, and networking components.
The VMSS must not be created until the storage accounts have been successfully created and an associated load balancer and virtual network is configured.
How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

copy
copyIndex
priority
dependsOn

copy
copyIndex
priority
dependsOn

copy
copyIndex
priority
dependsOn

Answer Area:

Answer Area

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

copy
copyIndex
priority
dependsOn

copy
copyIndex
priority
dependsOn

copy
copyIndex
priority
dependsOn

Section:
Explanation:
Box 1: copyIndex
Notice that the name of each resource includes the copyIndex() function, which returns the current iteration in the loop. copyIndex() is zero-based.
Box 2: copy
By adding the copy element to the resources section of your template, you can dynamically set the number of resources to deploy.
Box 3: dependsOn
Example:
"type": "Microsoft.Compute/virtualMachineScaleSets",
"apiVersion": "2020-06-01",
"name": "[variables('namingInfix')]",
"location": "[parameters('location')]",
"sku": {
"name": "[parameters('vmSku')]",
"tier": "Standard",
"capacity": "[parameters('instanceCount')]"
},
"dependsOn": [
"[resourceId('Microsoft.Network/loadBalancers', variables('loadBalancerName'))]",
"[resourceId('Microsoft.Network/virtualNetworks', variables('virtualNetworkName'))]"
],
Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/copy-resources
https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/quick-create-template-windows

QUESTION 30
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
{
  [FunctionName("ProcessOrders")]
  public static void ProcessOrders([QueueTrigger("incoming-orders")]CloudQueueMessage myQueueItem, [Table("Orders")]ICollector<Order> tableBindings, TraceWriter log)
  {
    log.Info($"Processing Order: {myQueueItem.Id}");
    log.Info($"Queue Insertion Time: {myQueueItem.InsertionTime}");
    log.Info($"Queue Expiration Time: {myQueueItem.ExpirationTime}");
    tableBindings.Add(JsonConvert.DeserializeObject<Order>(myQueueItem.AsString));
  }
  [FunctionName("ProcessOrders-Poison")]
  public static void ProcessFailedOrders([QueueTrigger("incoming-orders-poison")]CloudQueueMessage myQueueItem, TraceWriter log)
  {
    log.Error($"Failed to process order: {myQueueItem.AsString}");
    . . .
  }
}
```

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Answer Area:

Answer Area

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

Section:

Explanation:

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 31
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.

Select and Place:

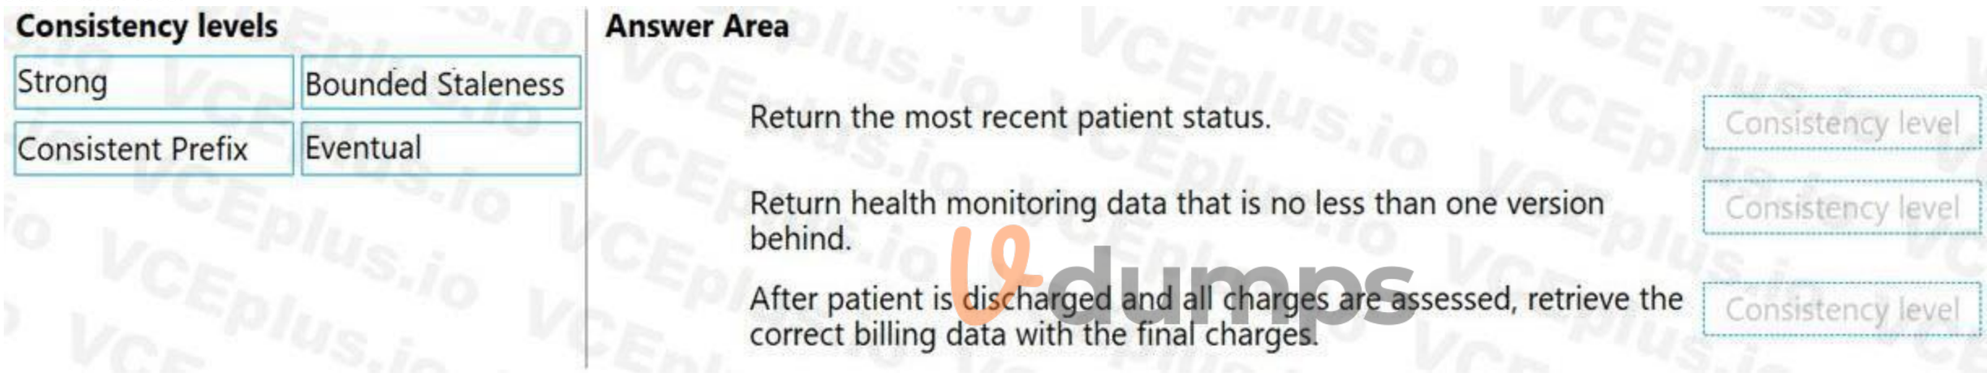

Correct Answer:

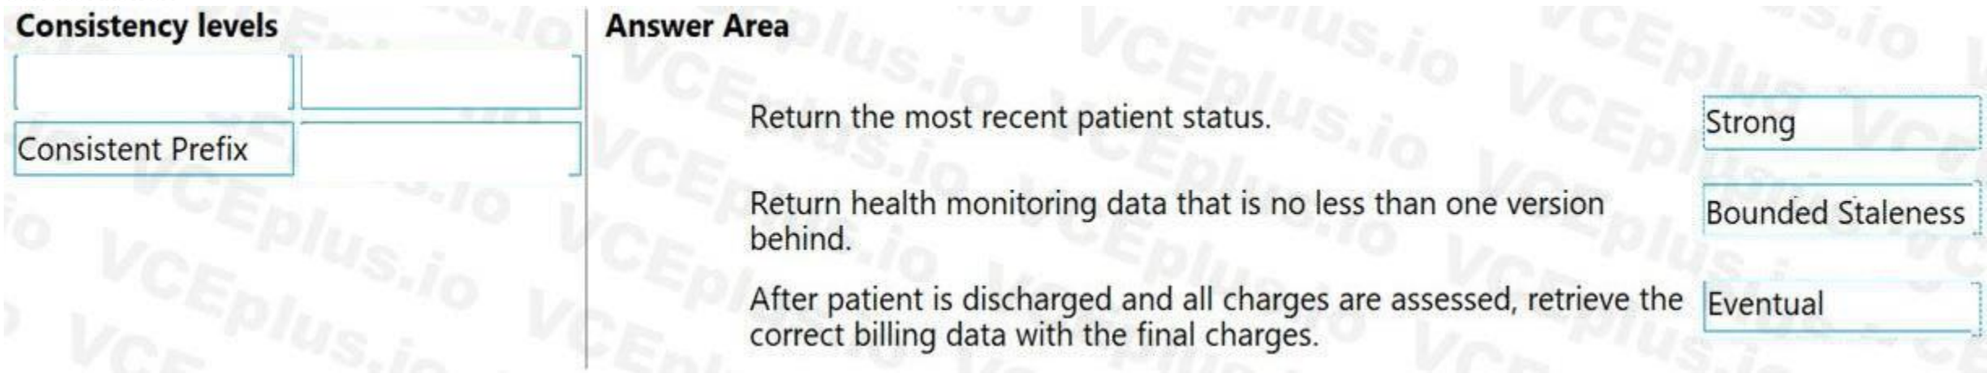

Section:
Explanation:
Box 1: Strong
Strong: Strong consistency offers a linearizability guarantee. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.
Box 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.
Incorrect Answers:
Consistent prefix: Updates that are returned contain some prefix of all the updates, with no gaps. Consistent prefix guarantees that reads never see out-of-order writes.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels

QUESTION 32
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.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.
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.
Hot Area:

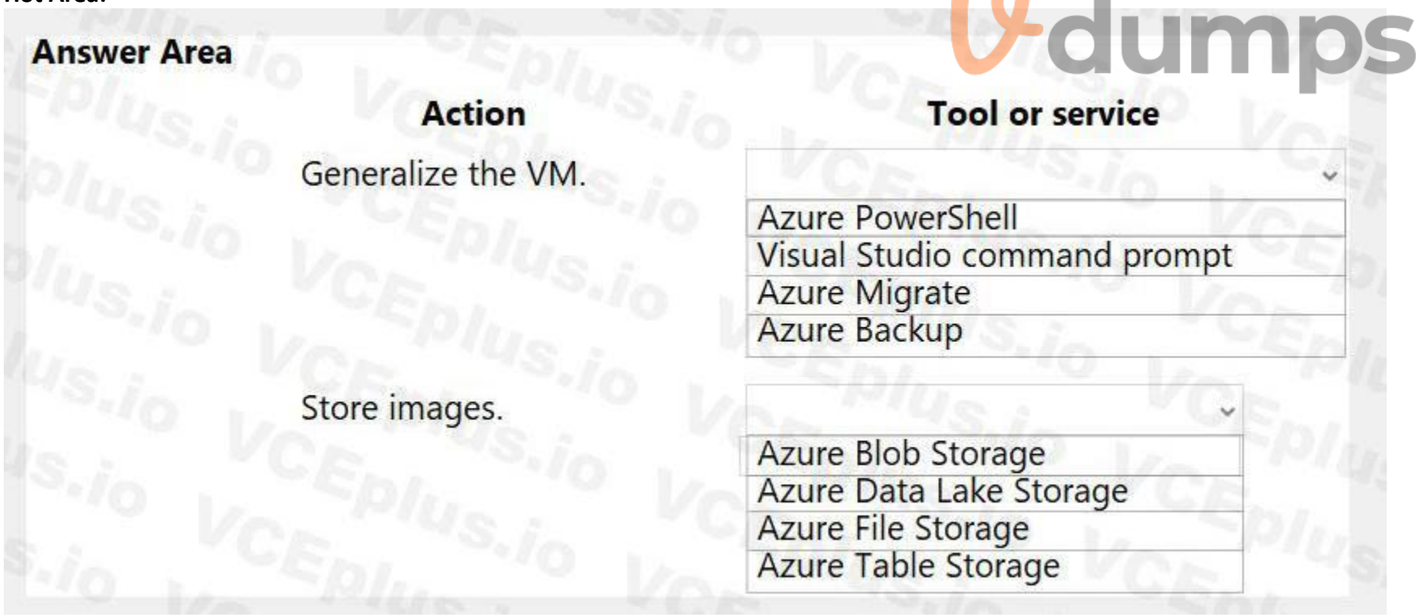

Answer Area:

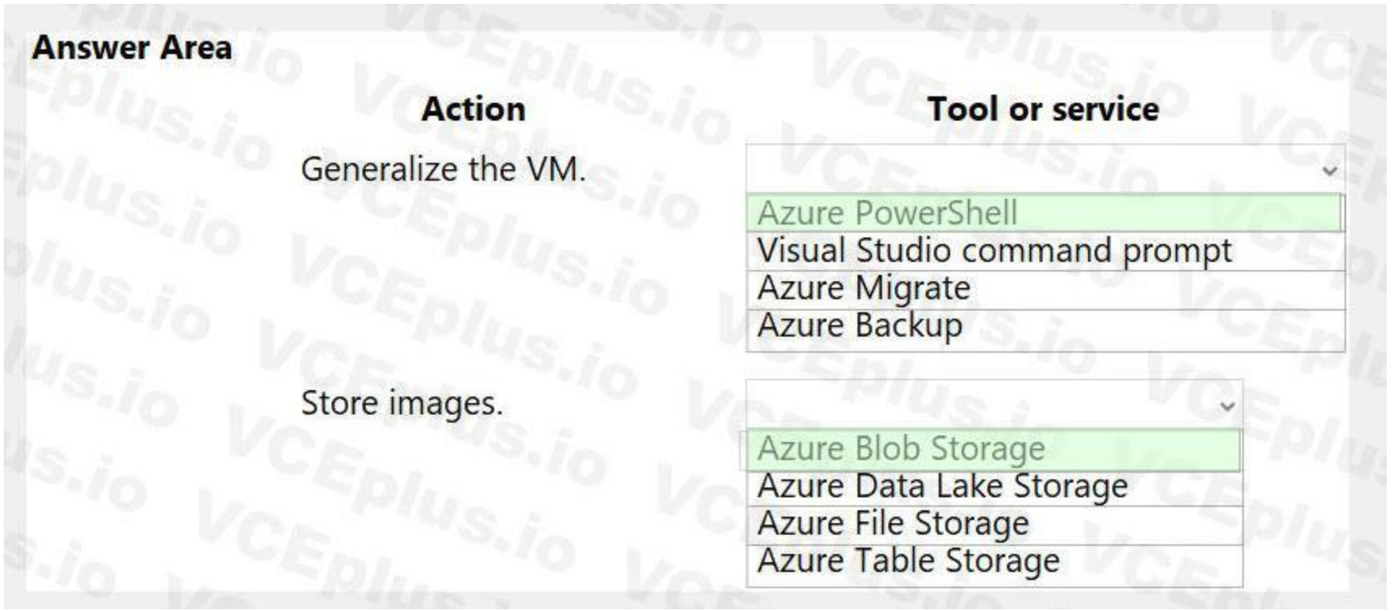

Section:
Explanation:
Box 1: Azure Powershell
Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks.
Before you begin, make sure that you have the latest version of the Azure PowerShell module.
You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image.
Box 2: Azure Blob Storage
You can store images in 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 33
You are preparing to deploy a website to an Azure Web App from a GitHub repository. The website includes static content generated by a script.
You plan to use the Azure Web App continuous deployment feature.
You need to run the static generation script before the website starts serving traffic.
What are two possible ways to achieve this goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

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

Correct Answer: A, D
Section:
Explanation:
A: In Azure, you can run your functions directly from a deployment package file in your function app. The other option is to deploy your files in the d:\home\site\wwwroot directory of your function app (see A above).
To enable your function app to run from a package, you just add a WEBSITE_RUN_FROM_PACKAGE setting to your function app settings.
Note: The host.json metadata file contains global configuration options that affect all functions for a function app.
D: To customize your deployment, include a .deployment file in the repository root.

You just need to add a file to the root of your repository with the name .deployment and the content:
[config]
command = YOUR COMMAND TO RUN FOR DEPLOYMENT
this command can be just running a script (batch file) that has all that is required for your deployment, like copying files from the repository to the web root directory for example.
Reference:
https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script
https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package

QUESTION 34
DRAG DROP
You are preparing to deploy an Azure virtual machine (VM)-based application.
The VMs that run the application have the following requirements:
When a VM is provisioned the firewall must be automatically configured before it can access Azure resources.
Supporting services must be installed by using an Azure PowerShell script that is stored in Azure Storage.
You need to ensure that the requirements are met.
Which features should you use? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Select and Place:

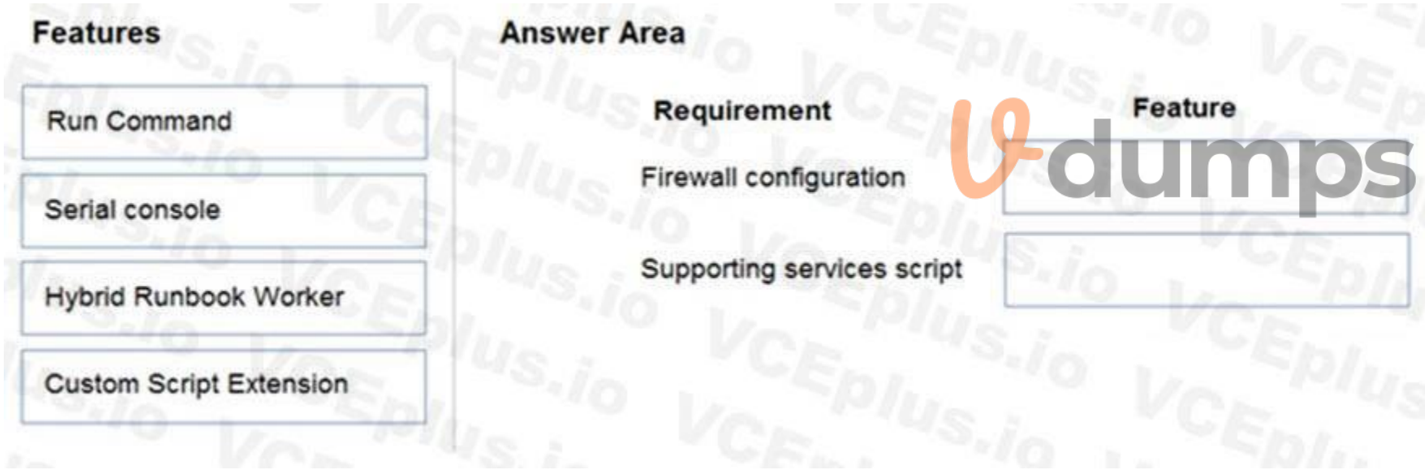

Correct Answer:

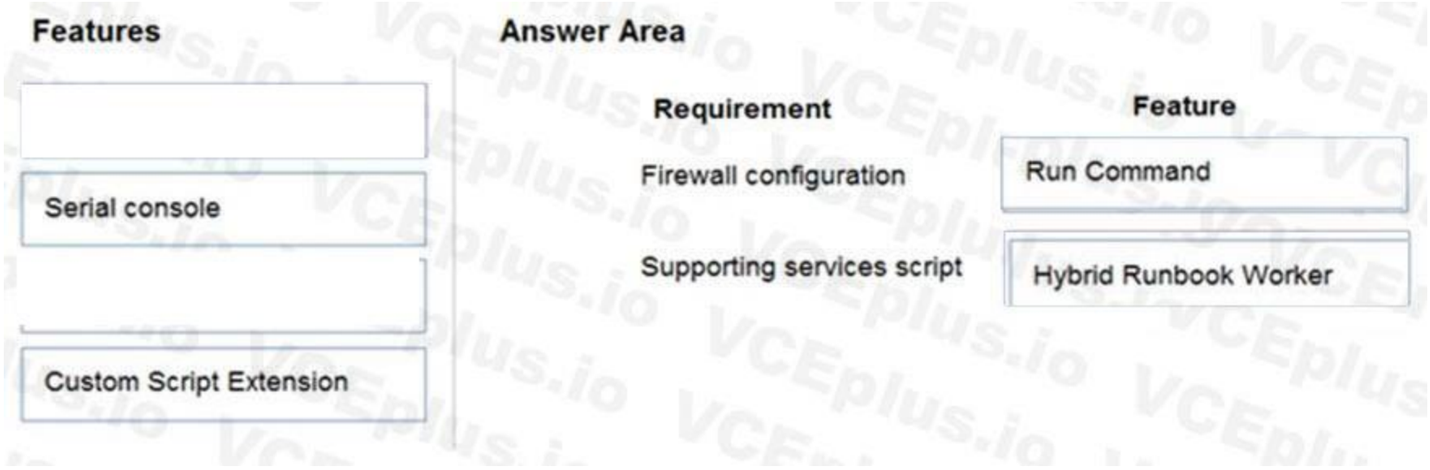

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-command

QUESTION 35
HOTSPOT
A company is developing a Node.js web app. The web app code is hosted in a GitHub repository located at https://github.com/TailSpinToys/webapp.
The web app must be reviewed before it is moved to production. You must deploy the initial code release to a deployment slot named review.
You need to create the web app and deploy the code.
How should you complete the commands? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"

[ ▼ ] -Name myResourceGroup -Location $location
[ ▼ ] -Name $webappname -Location $location -ResourceGroupName myResourceGroup -Tier Standard
[ ▼ ] -Name $webappname -Location $location -AppServicePlan $webappname -ResourceGroupName myResourceGroup
[ ▼ ] -Name $webappname -ResourceGroupName myResourceGroup -Slot review

$PropertiesObject = @{repoUrl = "$gitrepo";branch = "master";}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup
-SourceSlotName review -DestinationSlotName production
```

New-AzWebAppSlot
New-AzWebApp
New-AzAppServicePlan
New-AzResourceGroup

Answer Area:

Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
```

[▼] -Name myResourceGroup -Location $location

- New-AzWebAppSlot
- New-AzWebApp
- New-AzAppServicePlan
- New-AzResourceGroup

[▼] -Name $webappname -Location $location -ResourceGroupName myResourceGroup -Tier Standard

- New-AzWebAppSlot
- New-AzWebApp
- New-AzAppServicePlan
- New-AzResourceGroup

[▼] -Name $webappname -Location $location -AppServicePlan $webappname -ResourceGroupName myResourceGroup

- New-AzWebAppSlot
- New-AzWebApp
- New-AzAppServicePlan
- New-AzResourceGroup

[▼] -Name $webappname -ResourceGroupName myResourceGroup -Slot review

- New-AzWebAppSlot
- New-AzWebApp
- New-AzAppServicePlan
- New-AzResourceGroup

```
$PropertiesObject = @{repoUrl = "$gitrepo";branch = "master";}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup
-SourceSlotName review -DestinationSlotName production
```

Section:

Explanation:

Box 1: New-AzResourceGroup

The New-AzResourceGroup cmdlet creates an Azure resource group.

Box 2: New-AzAppServicePlan

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

Box 3: New-AzWebApp

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

Box 4: New-AzWebAppSlot

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

Reference:

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

https://docs.microsoft.com/en-us/powershell/module/az.websites/new-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 36
HOTSPOT
You are developing an application that needs access to an Azure virtual machine (VM).
The access lifecycle for the application must be associated with the VM service instance.
You need to enable managed identity for the VM.
How should you complete the PowerShell segment? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

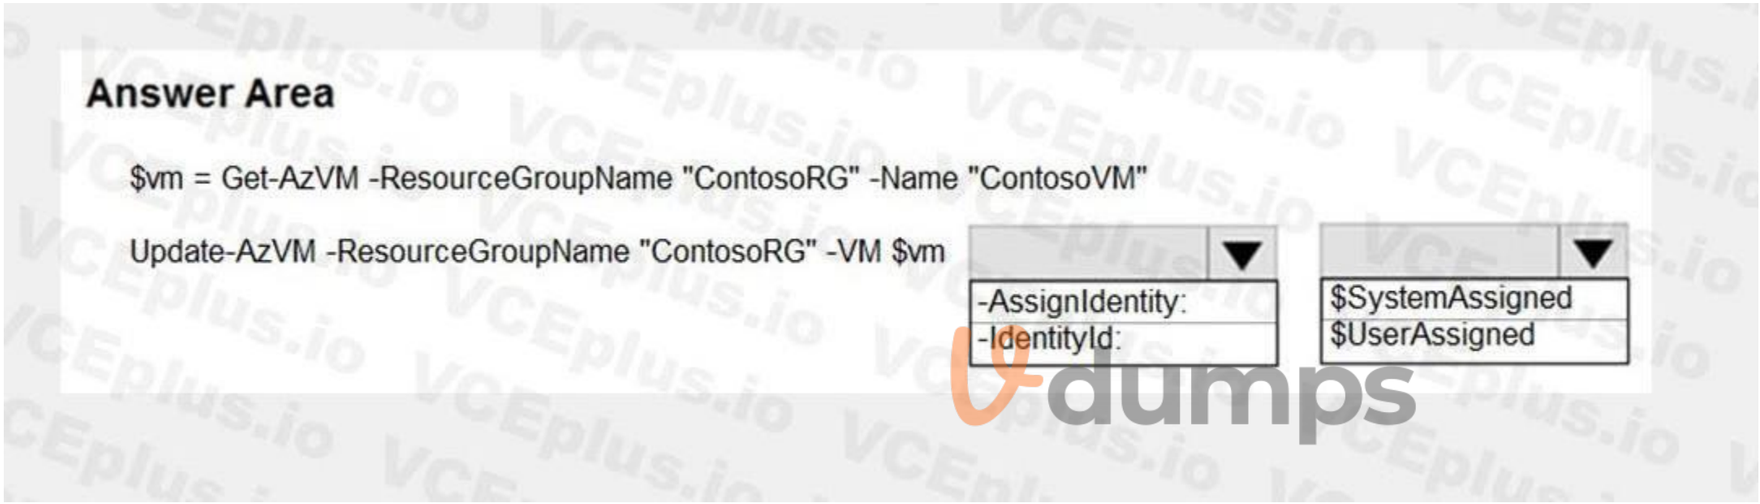

Answer Area:

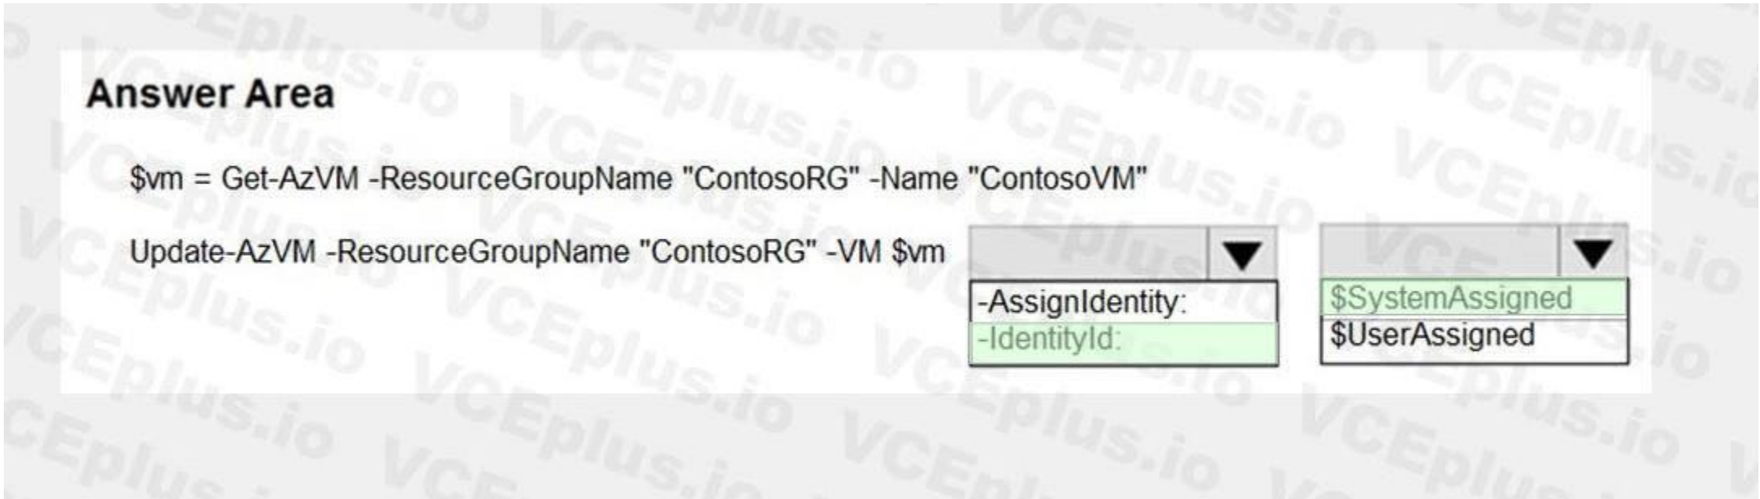

Section:
Explanation:
Box 1: -IdentityType
Enable system-assigned managed identity on an existing Azure VM:

To enable a system-assigned managed identity, use the -IdentityType switch on the Update-AzVM cmdlet (see below).
Box 2: $SystemAssigned
$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM
Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -IdentityType SystemAssigned
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm

QUESTION 37
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Create an Azure Function app that uses the Consumption hosting model and that is triggered from the blob upload.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: A
Section:
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 38
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.
You need to ensure that scripts run and resources are available before a swap operation occurs.
Solution: Update the app with a method named statuscheck to run the scripts. Update the app settings for the app. Set the WEBSITE_SWAP_WARMUP_PING_PATH and WEBSITE_SWAP_WARMUP_PING_STATUSES with a path to the new method and appropriate response codes.
Does the solution meet the goal?

A. No
B. Yes

Correct Answer: B
Section:
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 39
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.

Hot Area:

Answer Area

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 Area:

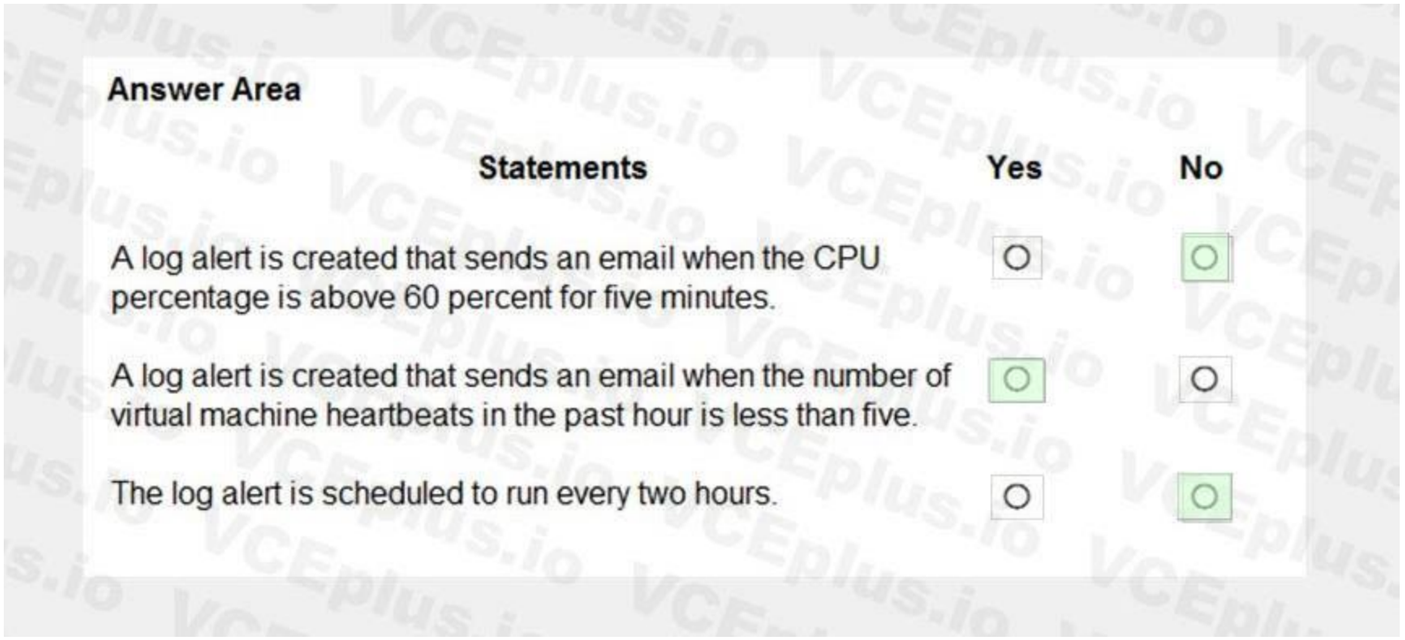

Section:
Explanation:
Box 1: No
The AzScheduledQueryRuleSource is Heartbeat, not CPU.
Box 2: Yes
The AzScheduledQueryRuleSource is Heartbeat!
Note: New-AzScheduledQueryRuleTriggerCondition creates an object of type Trigger Condition. This object is to be passed to the command that creates Alerting Action object.
Box 3: No
The schedule is 60 minutes, not two hours.
-FrequencyInMinutes: The alert frequency.
-TimeWindowInMinutes: The alert time window
The New-AzAscheduledQueryRuleSchedule command creates an object of type Schedule. This object is to be passed to the command that creates Log Alert Rule.
Reference:
https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryrule
https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryruletriggercondition

01 - Monitor troubleshoot and optimize Azure solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.
Background
Overview
You are a developer for Contoso, Ltd. The company has a social networking website that is developed as a Single Page Application (SPA). The main web application for the social networking website loads user uploaded

content from blob storage.
You are developing a solution to monitor uploaded data for inappropriate content. The following process occurs when users upload content by using the SPA:
• Messages are sent to ContentUploadService.
• Content is processed by ContentAnalysisService.
• After processing is complete, the content is posted to the social network or a rejection message is posted in its place.
The ContentAnalysisService is deployed with Azure Container Instances from a private Azure Container Registry named contosoimages.
The solution will use eight CPU cores.
Azure Active Directory
Contoso, Ltd. uses Azure Active Directory (Azure AD) for both internal and guest accounts.
Requirements
ContentAnalysisService

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

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

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

QUESTION 1

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

Correct Answer: C
Section:
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:
1. Observe and monitor application behavior
2. Collect data
3. Mitigate the issue
Reference:
https://docs.microsoft.com/en-us/cli/azure/monitor/activity-log

QUESTION 2
You need to monitor ContentUploadService according to the requirements.
Which command should you use?

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

Correct Answer: B
Section:
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

02 - Monitor troubleshoot and optimize Azure solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

Background
City Power & Light company provides electrical infrastructure monitoring solutions for homes and businesses. The company is migrating solutions to Azure.
Current environment
Architecture overview
The company has a public website located at http://www.cpandl.com/. The site is a single-page web application that runs in Azure App Service on Linux. The website uses files stored in Azure Storage and cached in Azure Content Delivery Network (CDN) to serve static content.
API Management and Azure Function App functions are used to process and store data in Azure Database for PostgreSQL. API Management is used to broker communications to the Azure Function app functions for Logic app integration.
Logic apps are used to orchestrate the data processing while Service Bus and Event Grid handle messaging and events.
The solution uses Application Insights, Azure Monitor, and Azure Key Vault.
Architecture diagram
The company has several applications and services that support their business. The company plans to implement serverless computing where possible. The overall architecture is shown below.

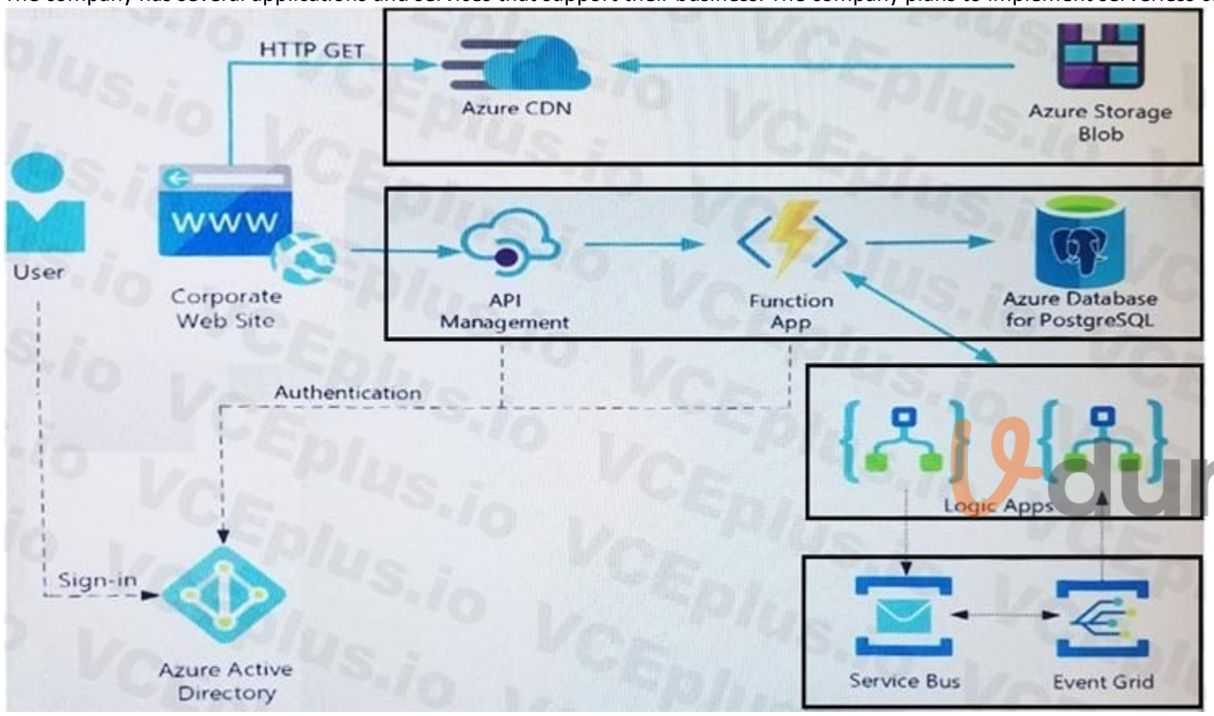

User authentication
The following steps detail the user authentication process:
The user selects Sign in in the website.
The browser redirects the user to the Azure Active Directory (Azure AD) sign in page.
The user signs in.
Azure AD redirects the user's session back to the web application. The URL includes an access token.
The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token.
The back-end API validates the access token.
Requirements
Corporate website
Communications and content must be secured by using SSL.
Communications must use HTTPS.
Data must be replicated to a secondary region and three availability zones.
Data storage costs must be minimized.
Azure Database for PostgreSQL
The database connection string is stored in Azure Key Vault with the following attributes:
Azure Key Vault name: cpandlkeyvault
Secret name: PostgreSQLConn

Id: 80df3e46ffcd4f1cb187f79905e9a1e8
The connection information is updated frequently. The application must always use the latest information to connect to the database.
Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.
Security
All SSL certificates and credentials must be stored in Azure Key Vault.
File access must restrict access by IP, protocol, and Azure AD rights.
All user accounts and processes must receive only those privileges which are essential to perform their intended function.
Compliance
Auditing of the file updates and transfers must be enabled to comply with General Data Protection Regulation (GDPR). The file updates must be read-only, stored in the order in which they occurred, include only create, update, delete, and copy operations, and be retained for compliance reasons.
Issues
Corporate website
While testing the site, the following error message displays:
CryptographicException: The system cannot find the file specified.
Function app
You perform local testing for the RequestUserApproval function. The following error message displays:
'Timeout value of 00:10:00 exceeded by function: RequestUserApproval'
The same error message displays when you test the function in an Azure development environment when you run the following Kusto query:
FunctionAppLogs
| where FunctionName = = "RequestUserApproval"
Logic app
You test the Logic app in a development environment. The following error message displays:
'400 Bad Request'
Troubleshooting of the error shows an HttpTrigger action to call the RequestUserApproval function.
Code
Corporate website
Security.cs:

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

Function app
RequestUserApproval.cs:

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

QUESTION 1
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.

Hot Area:

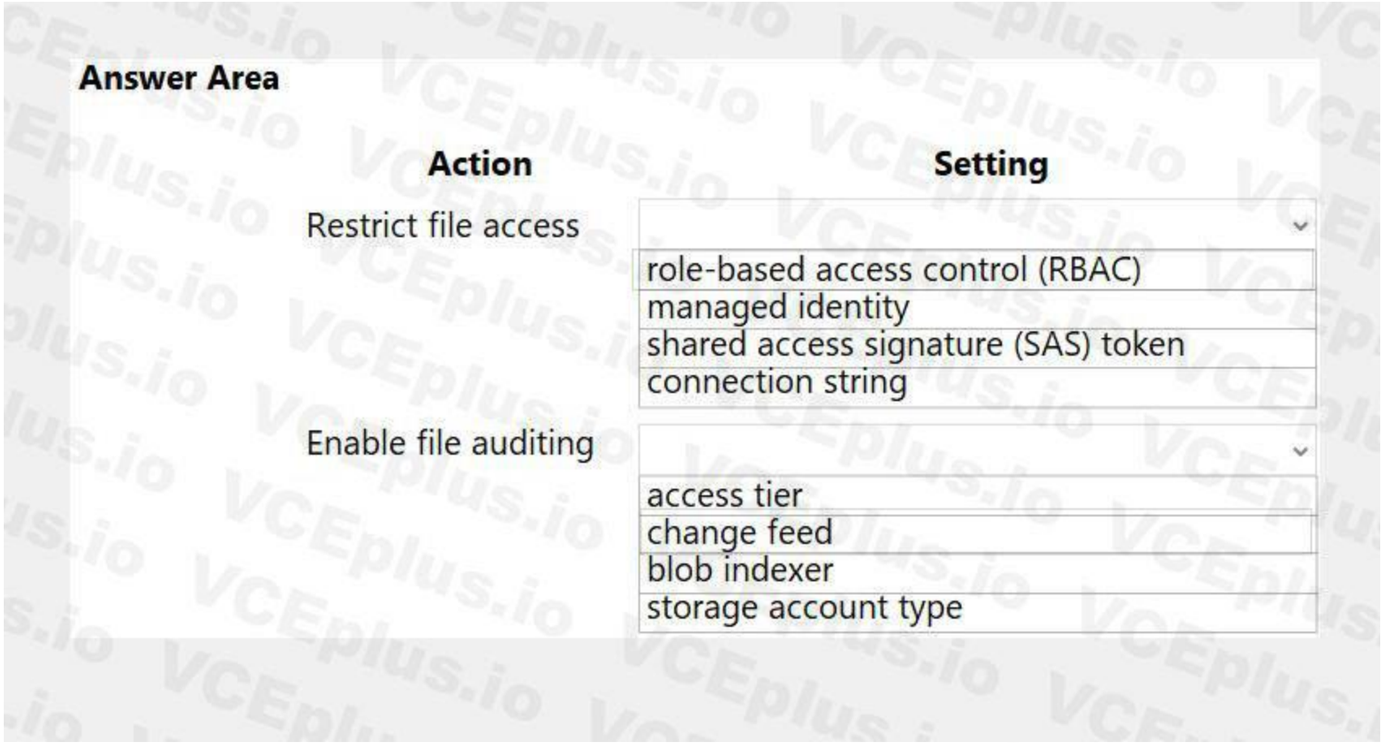

Answer Area:

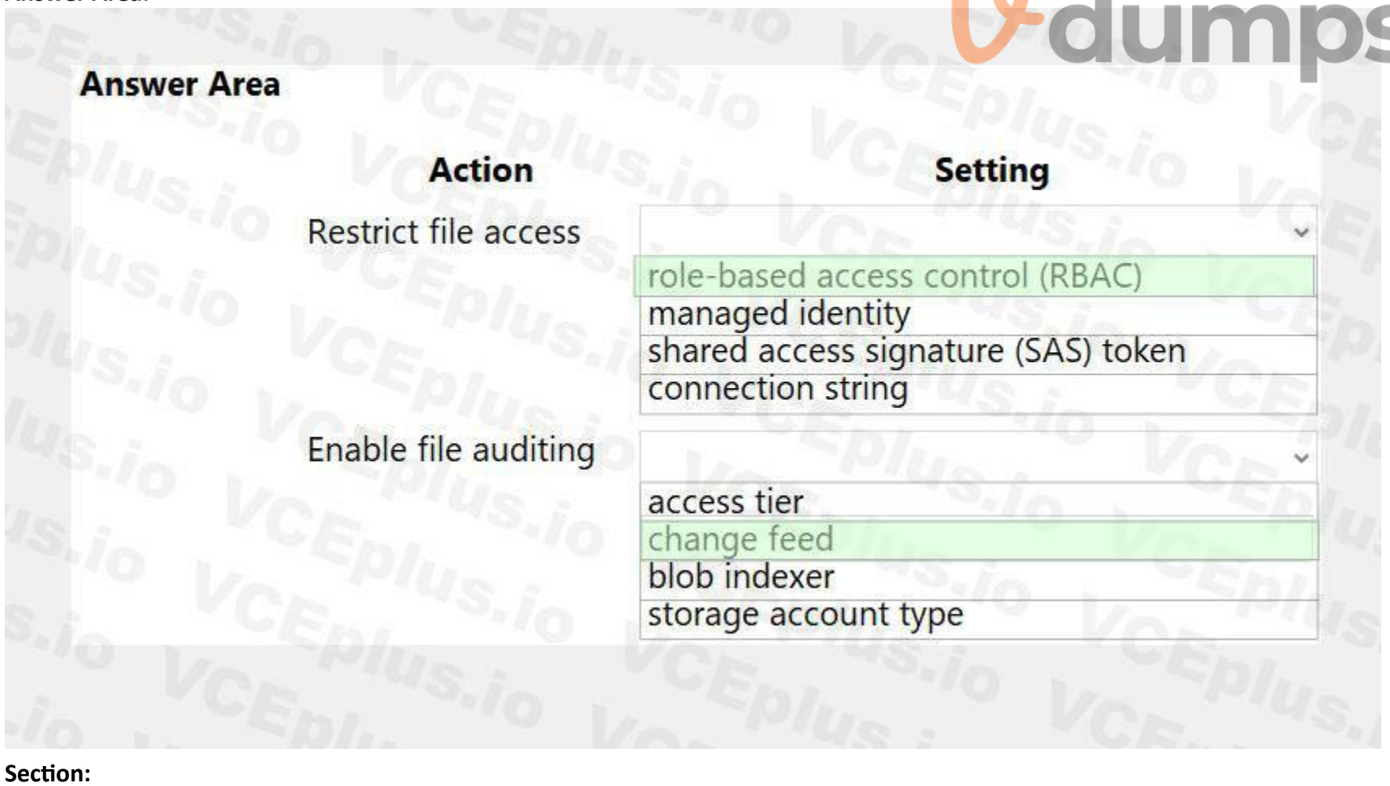

Section:
Explanation:

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 2
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.

Correct Answer: A
Section:
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

03 - Monitor troubleshoot and optimize Azure solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.
Background
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.cs

```
EG01 public class EventGridController : Controller
EG02 {
EG03   public static AsyncLocal<string> EventId = new AsyncLocal<string>();
EG04   public IActionResult Process([fromBody] string eventsJson)
EG05   {
EG06     var events = JArray.Parse(eventsJson);
EG07
EG08     foreach (var @event in events)
EG09     {
EG10       EventId.Value = @event["id"].ToString();
EG11       if (@event["topic"].ToString().Contains("providers/Microsoft.Storage"))
EG12       {
EG13         SendToAnomalyDetectionService(@event["data"]["url"].ToString());
EG14       }
EG15
EG16       {
EG17         EnsureLogging(@event["subject"].ToString());
EG18       }
EG19     }
EG20     return null;
EG21   }
EG22    private void EnsureLogging(string resource)
EG23    {
EG24      . . .
EG25    }
EG26   private async Task SendToAnomalyDetectionService(string uri)
EG27   {
EG28     var content = GetLogData(uri);
EG29     var scoreRequest = new
EG30     {
EG31       Inputs = new Dictionary<string, List<Dictionary<string, string>>>()
EG32       {
EG33         {
EG34           "input1",
EG35           new List<Dictionary<string, string>>()
EG36           {
EG37             new Dictionary<string, string>()
EG38             {
EG39               {
EG40                 "logcontent", content
EG41               }
EG42             }
EG43           }
EG44         },
EG45       },
EG46       GlobalParameters = new Dictionary<string, string>() { }
EG47     };
EG48     var result = await (new HttpClient()).PostAsJsonAsync("...", scoreRequest);
EG49     var rawModelResult = await result.Content.ReadAsStringAsync();
EG50     var modelResult = JObject.Parse(rawModelResult);
EG51     if (modelResult["notify"].HasValues)
EG52     {
EG53       . . .
EG54     }
EG55   }
```

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.cs
LE01 public class LoginEvent
LE02 {
LE03
LE04   public string subject { get; set; }
LE05   public DateTime eventTime { get; set; }
LE06   public Dictionary<string, string> data { get; set; }
LE07   public string Serialize()
LE08   {
LE09     return JsonConvert.SerializeObject(this);
LE10   }
LE11 }
```

QUESTION 1

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.

Select and Place:

Code segments

- /health
- /status
- RequestTelemetry
- PageViewTelemetry
- ITelementryProcessor
- ITelementryInitializer

Answer Area

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

Correct Answer:

Code segments

- /health
- /status
- RequestTelemetry
- PageViewTelemetry
- ITelementryProcessor
- ITelementryInitializer

Answer Area

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

Section:

Explanation:

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 2

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.

Select and Place:

Code segment

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

Answer Area

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

Correct Answer:

Code segment

- All
- (blank)
- EventHub
- subjectEndsWith
- Mictosoft.Storage
- (blank)
- (blank)

Answer Area

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

Section:

Explanation:

Box 1:WebHook

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

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

Box 2: SubjectBeginsWith

Box 3: Microsoft.Storage.BlobCreated

Scenario: Log Policy

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

Example subscription schema

```
{
"properties": {
"destination": {
"endpointType": "webhook",
"properties": {
"endpointUrl": "https://example.azurewebsites.net/api/HttpTriggerCSharp1?code=VXbGWce53l48Mt8wuotr0GPmyJ/nDT4hgdFj9DpBiRt38qqnnm5OFg=="
}
},
```

"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 3
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

Correct Answer: D
Section:
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 4
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.

Select and Place:

Code segments

- Process
- Initialize
- telemetry.Sequence
- ITelemetryProcessor
- ITelemetryInitializer
- telemetry.Context
- EventGridController.EventId.Value
- ((EventTelemetry)telemetry).Properties["EventId"]

Answer Area

```
public class IncludeEventId : [code segment]
{
    public void [code segment] (ITelemetry telemetry)
    {
        [code segment].Properties["EventId"] =
        [code segment];
    }
}
```

Correct Answer:

Code segments

- Process
- telemetry.Sequence
- ITelemetryProcessor
- EventGridController.EventId.Value

Answer Area

```
public class IncludeEventId : ITelemetryInitializer
{
    public void Initialize (ITelemetry telemetry)
    {
        telemetry.Context .Properties["EventId"] =
        ((EventTelemetry)telemetry).Properties["EventId"] ;
    }
}
```

Section:

Explanation:

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

04 - Monitor troubleshoot and optimize Azure solutions

Case study

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

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Background
You are a developer for Litware Inc., a SaaS company that provides a solution for managing employee expenses. The solution consists of an ASP.NET Core Web API project that is deployed as an Azure Web App.
Overall architecture
Employees upload receipts for the system to process. When processing is complete, the employee receives a summary report email that details the processing results. Employees then use a web application to manage their receipts and perform any additional tasks needed for reimbursement.
Receipt processing
Employees may upload receipts in two ways:
Uploading using an Azure Files mounted folder
Uploading using the web application
Data Storage
Receipt and employee information is stored in an Azure SQL database.
Documentation
Employees are provided with a getting started document when they first use the solution. The documentation includes details on supported operating systems for Azure File upload, and instructions on how to configure the mounted folder.
Solution details
Users table

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

Web Application
You enable MSI for the Web App and configure the Web App to use the security principal name WebAppIdentity.
Processing
Processing is performed by an Azure Function that uses version 2 of the Azure Function runtime. Once processing is completed, results are stored in Azure Blob Storage and an Azure SQL database. Then, an email summary is sent to the user with a link to the processing report. The link to the report must remain valid if the email is forwarded to another user.
Logging
Azure Application Insights is used for telemetry and logging in both the processor and the web application. The processor also has TraceWriter logging enabled. Application Insights must always contain all log messages.
Requirements
Receipt processing
Concurrent processing of a receipt must be prevented.
Disaster recovery
Regional outage must not impact application availability. All DR operations must not be dependent on application running and must ensure that data in the DR region is up to date.
Security
User's SecurityPin must be stored in such a way that access to the database does not allow the viewing of SecurityPins. The web application is the only system that should have access to SecurityPins.
All certificates and secrets used to secure data must be stored in Azure Key Vault.
You must adhere to the principle of least privilege and provide privileges which are essential to perform the intended function.
All access to Azure Storage and Azure SQL database must use the application's Managed Service Identity (MSI).
Receipt data must always be encrypted at rest.
All data must be protected in transit.
User's expense account number must be visible only to logged in users. All other views of the expense account number should include only the last segment, with the remaining parts obscured.
In the case of a security breach, access to all summary reports must be revoked without impacting other parts of the system.
Issues

Upload format issue
Employees occasionally report an issue with uploading a receipt using the web application. They report that when they upload a receipt using the Azure File Share, the receipt does not appear in their profile. When this occurs, they delete the file in the file share and use the web application, which returns a 500 Internal Server error page.
Capacity issue
During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.
Log capacity issue
Developers report that the number of log messages in the trace output for the processor is too high, resulting in lost log messages.
Application code
Processing.cs

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

PC39 }
```

Database.cs

```
DB01 public class Database
DB02 {
DB03   private string ConnectionString =
DB04
DB05   public async Task<object> LoadUserDetails(string userId)
DB06   {
DB07
DB08     return await policy.ExecuteAsync(async () =>
DB09     {
DB10       using (var connection = new SqlConnection(ConnectionString))
DB11       {
DB12         await connection.OpenAsync();
DB13         using (var command = new SqlCommand("…", connection))
DB14         using (var reader = command.ExecuteReader())
DB15         {
DB16           …
DB17         }
DB18       }
DB19     });
DB20   }

DB21 }
```

ReceiptUploader.cs

```
RU01 public class ReceiptUploader
RU02 {
RU03   public async Task UploadFile(string file, byte[] binary)
RU04   {
RU05     var httpClient = new HttpClient();
RU06     var response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU07     while (ShouldRetry(response))
RU08     {
RU09       response = await httpClient.PutAsync("…", new ByteArrayContent(binary));
RU10     }
RU11   }
RU12   private bool ShouldRetry(HttpResponseMessage response)
RU13   {
RU14
RU15   }
RU16 }
```

ConfigureSSE.ps1

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

QUESTION 1
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

Correct Answer: D
Section:
Explanation:
If you want to read the files in parallel, you cannot use forEach. Each of the async callback function calls does return a promise. You can await the array of promises that you'll get with Promise.all.
Scenario: Capacity issue: During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

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

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

QUESTION 2
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

Correct Answer: B
Section:
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 3
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

Correct Answer: C
Section:
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:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling

05 - Monitor troubleshoot and optimize Azure solutions

QUESTION 1
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

Correct Answer: B, C
Section:
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 2
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?

Select and Place:

Correct Answer:

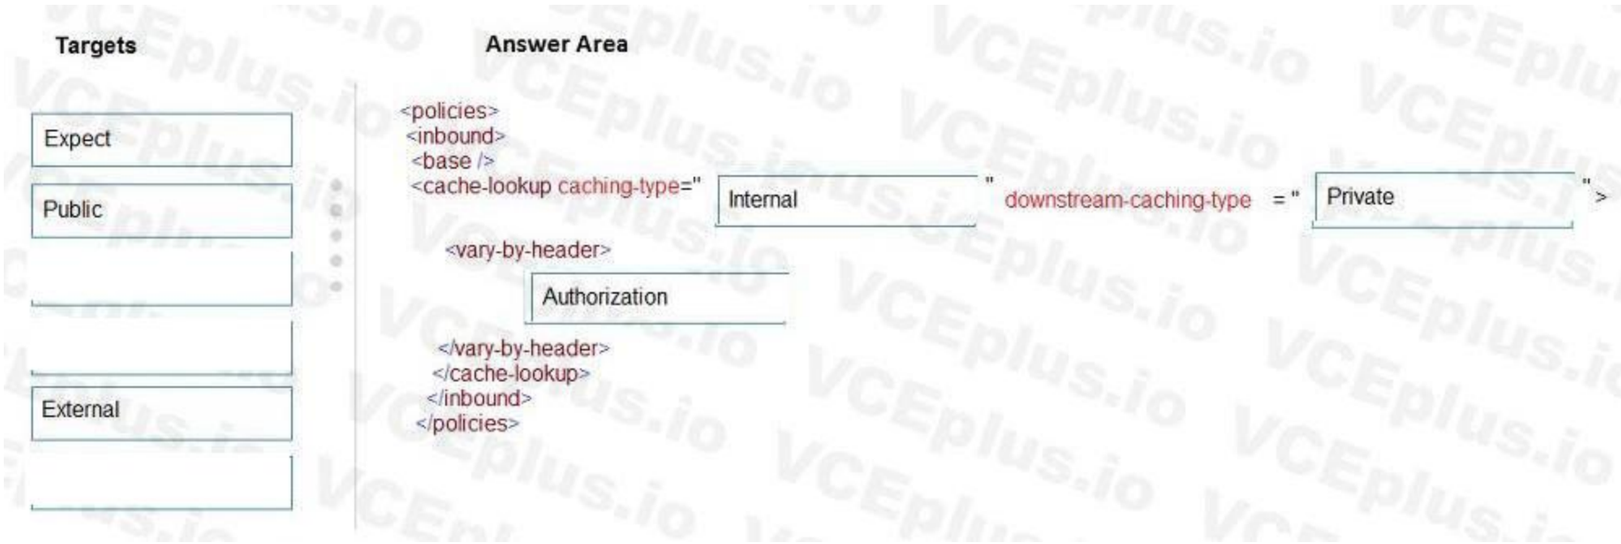

Section:
Explanation:
Box 1: internal
caching-type
Choose between the following values of the attribute:
internal to use the built-in API Management cache,
external to use the external cache as Azure Cache for Redis
prefer-external to use external cache if configured or internal cache otherwise.
Box 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 3
DRAG 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.

Select and Place:

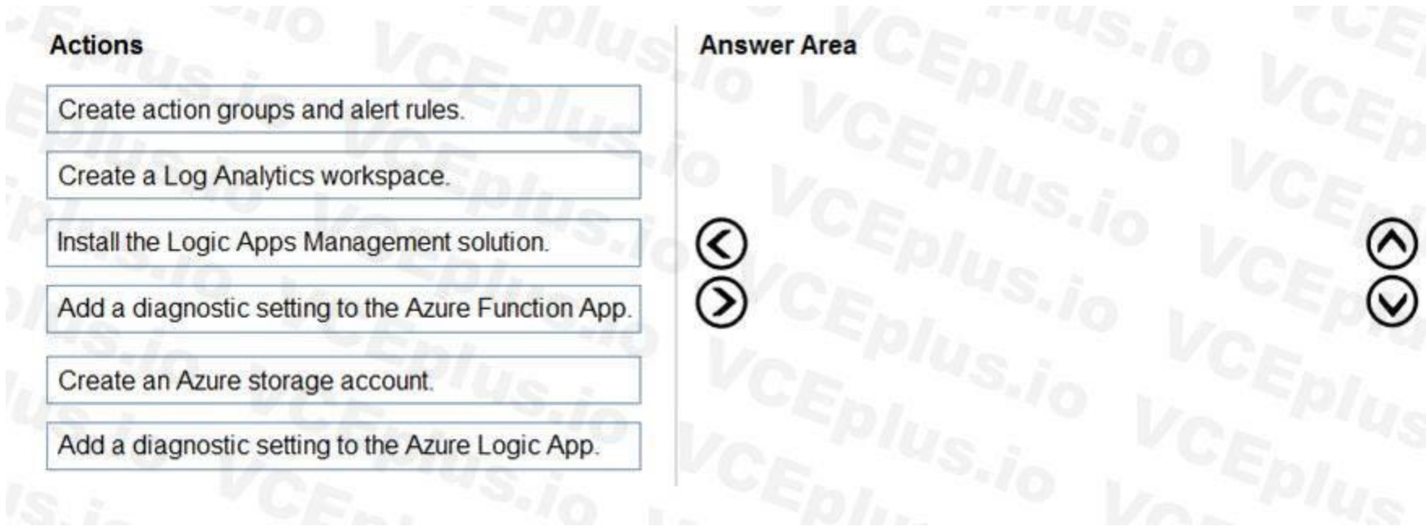

Correct Answer:

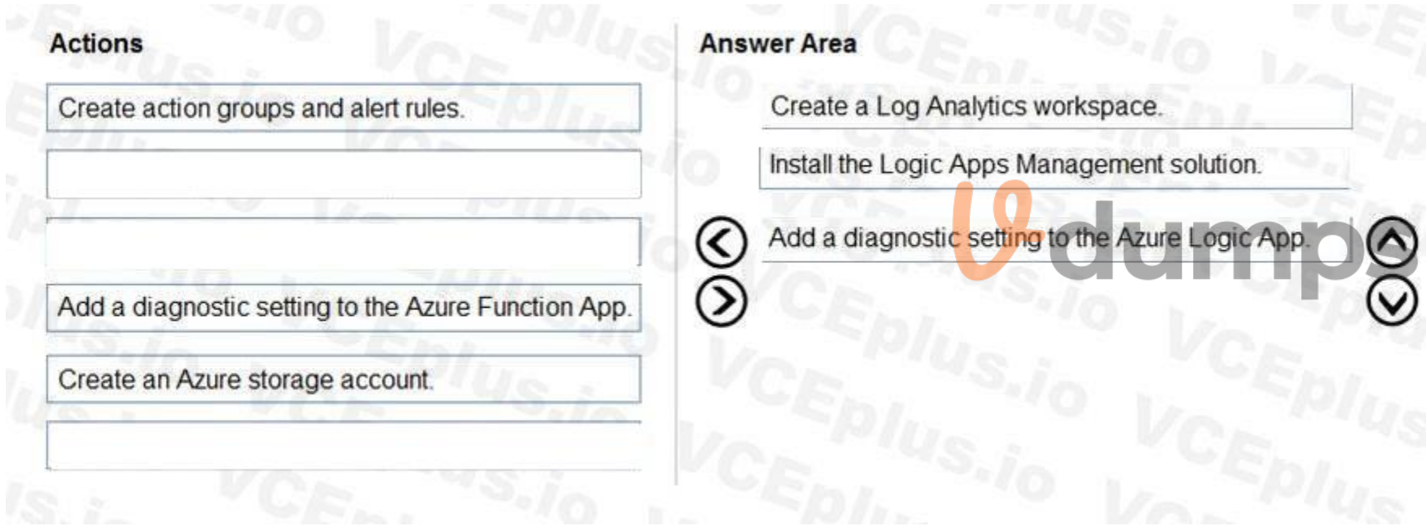

Section:
Explanation:
Step 1: Create a Log Analytics workspace
Before you start, you need a Log Analytics workspace.
Step 2: Install the Logic Apps Management solution
To set up logging for your logic app, you can enable Log Analytics when you create your logic app, or you can install the Logic Apps 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
1. In the Azure portal, find and select your logic app.
2. 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 4
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.

Select and Place:

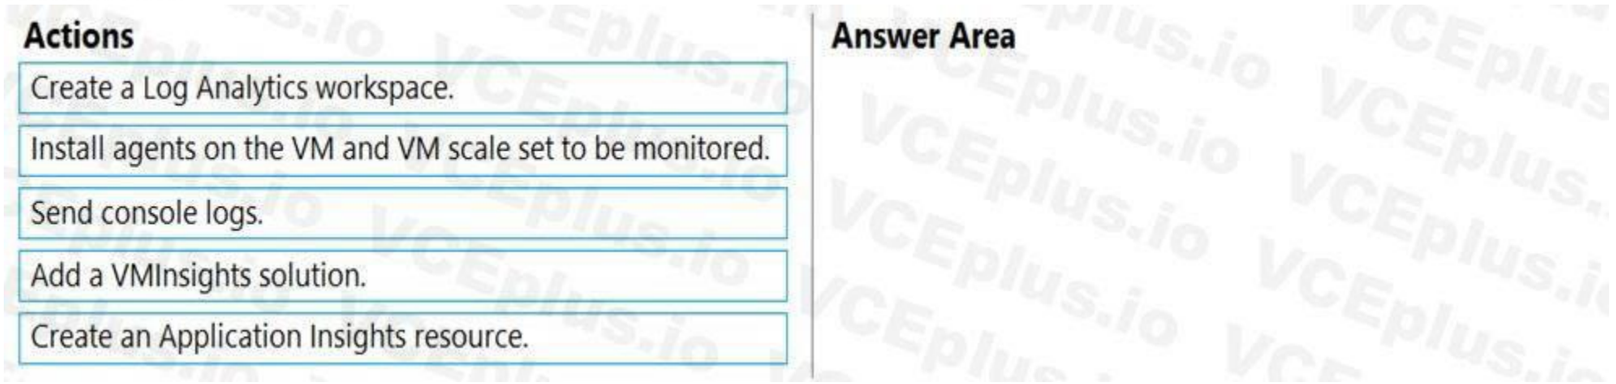

Correct Answer:

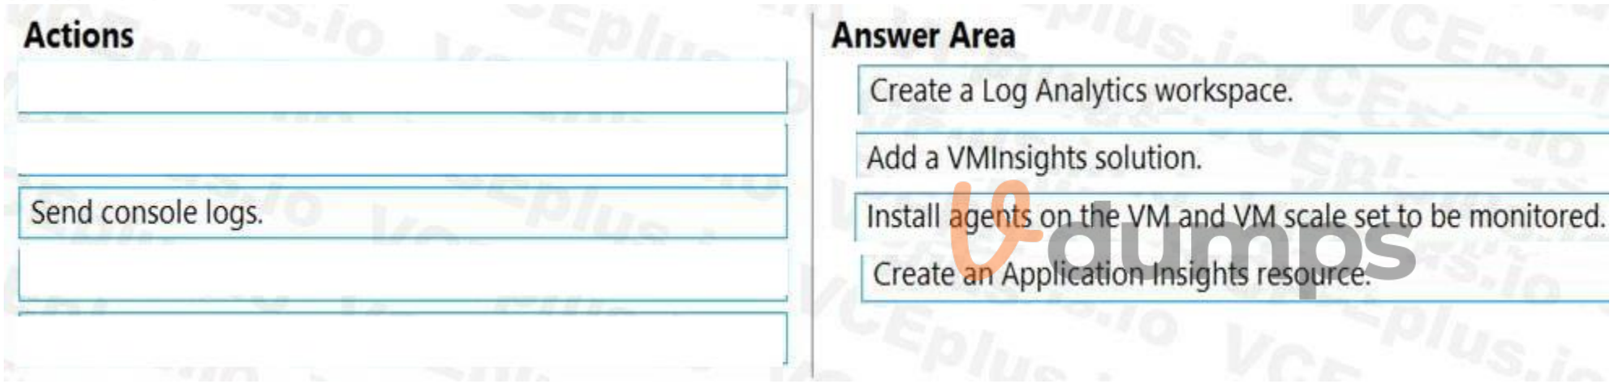

Section:
Explanation:
Step 1: Create a Log Analytics workspace.
First create the workspace.
Step 2: Add a VMInsights solution.
Before a Log Analytics workspace can be used with VM insights, it must have the VMInsights solution installed.
Step 3: Install agents on the VM and VM scale set to be monitored.
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 scalet sets.
Step 4: Create an Application Insights resource
Sign in to the Azure portal, and create an Application Insights resource.

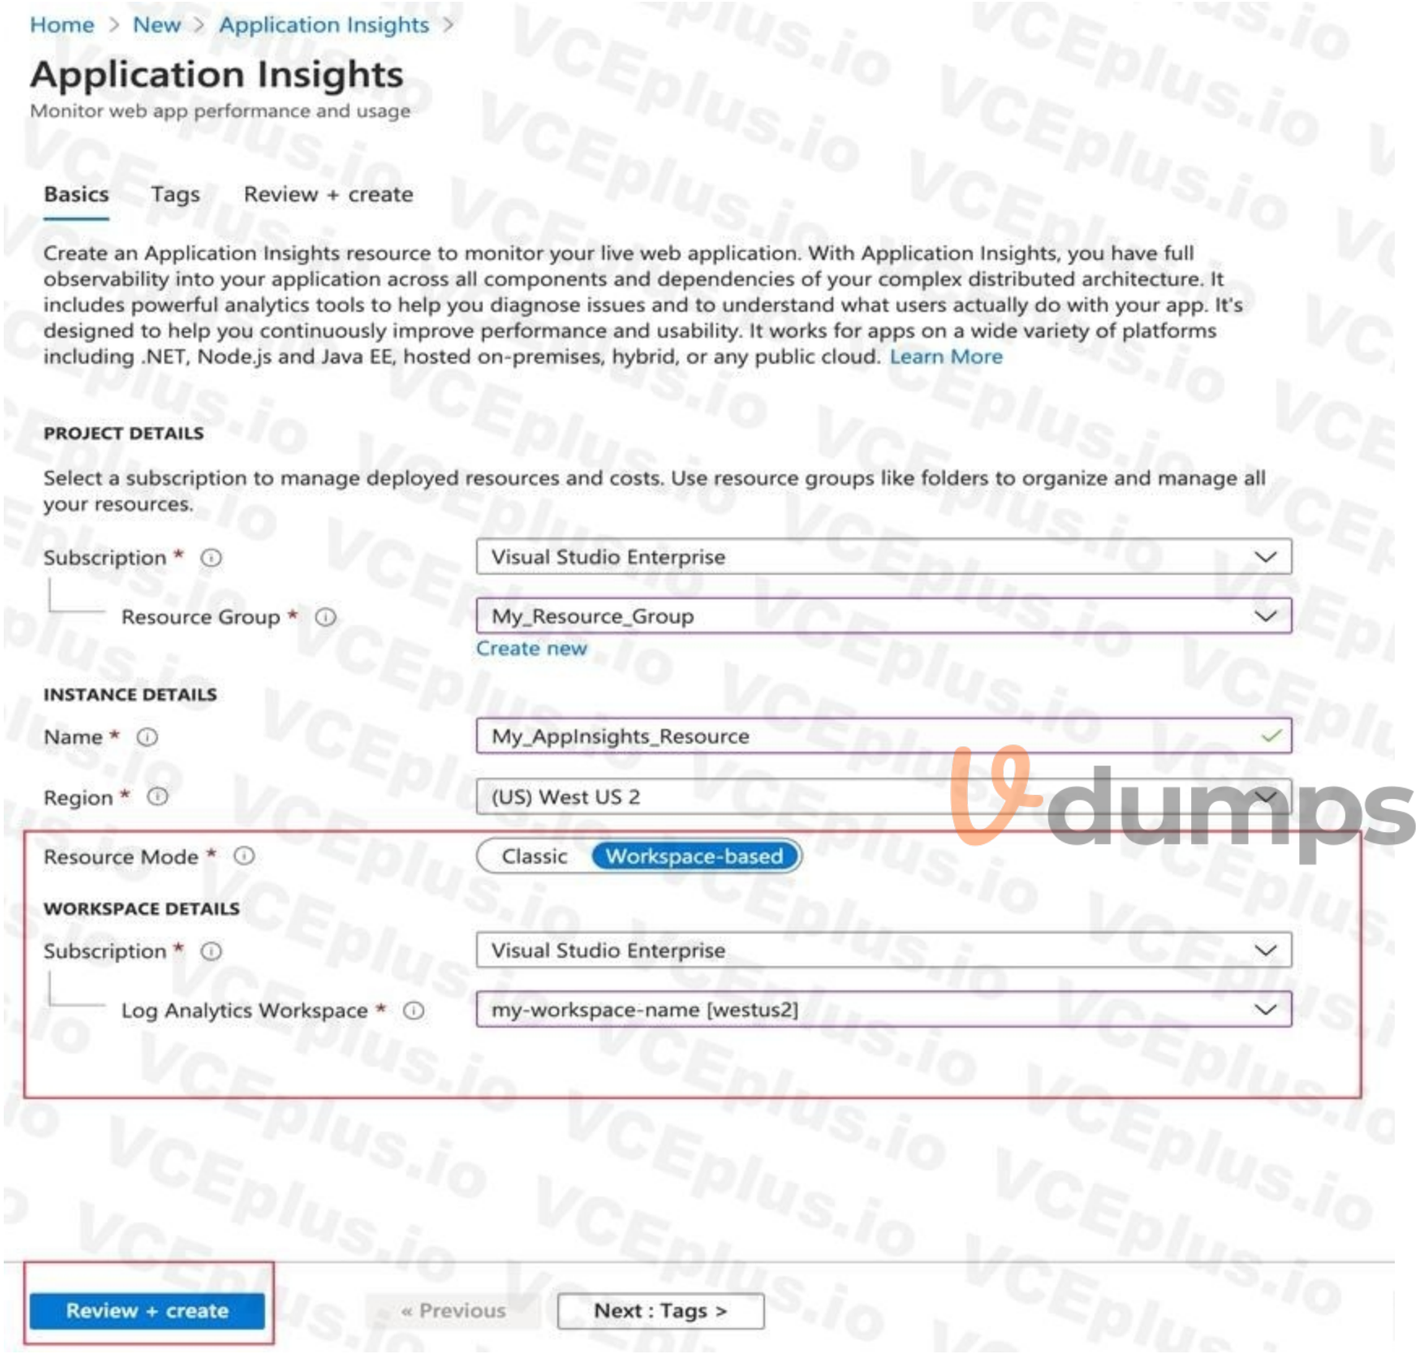

Once a workspace-based Application Insights resource has been created, configuring monitoring is relatively straightforward.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/vm/vminsights-configure-workspace
https://docs.microsoft.com/en-us/azure/azure-monitor/app/create-workspace-resource

QUESTION 5
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 analyse app uptime for each month.

Which two solutions will achieve the goal? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Azure Monitor logs
B. Application Insights alerts
C. Azure Monitor metrics
D. Application Insights web tests

Correct Answer: B, D
Section:
Explanation:
Reference:
https://azure.microsoft.com/en-us/blog/creating-a-web-test-alert-programmatically-with-application-insights/

QUESTION 6
DRAG DROP
You develop and deploy an Azure App Service web 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 drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Features	Requirement	Feature
horizontal partitioning	Create a data structure for storing collections of related items	Feature
channel	Create a data structure for the most recently accessed cache items	Feature
list	Send messages through a high-performance publisher/subscriber mechanism	Feature
set		

Correct Answer:

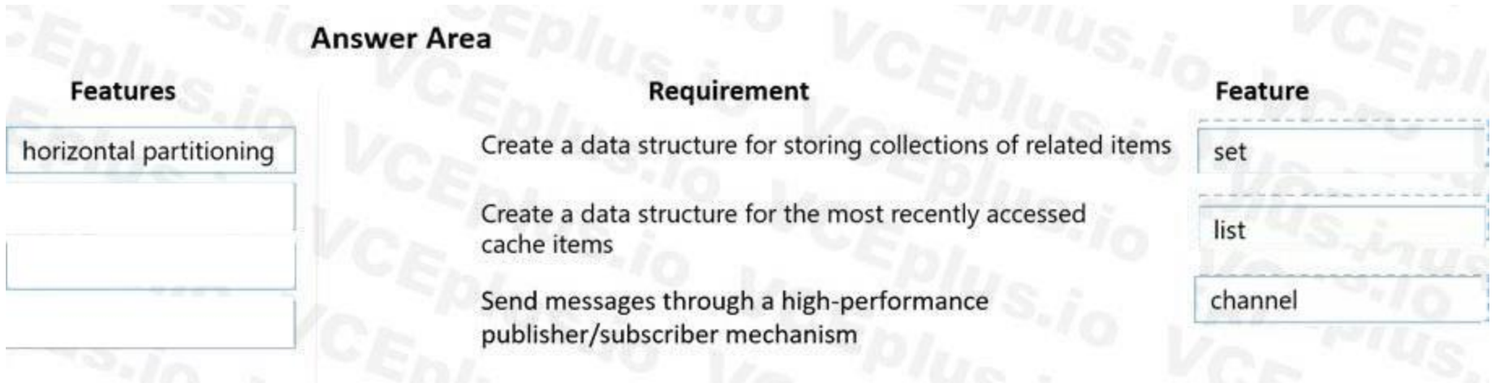

Section:
Explanation:
Reference:
https://www.red-gate.com/simple-talk/development/dotnet-development/overview-of-azure-cache-for-redis/
https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching

QUESTION 7
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.

Hot Area:

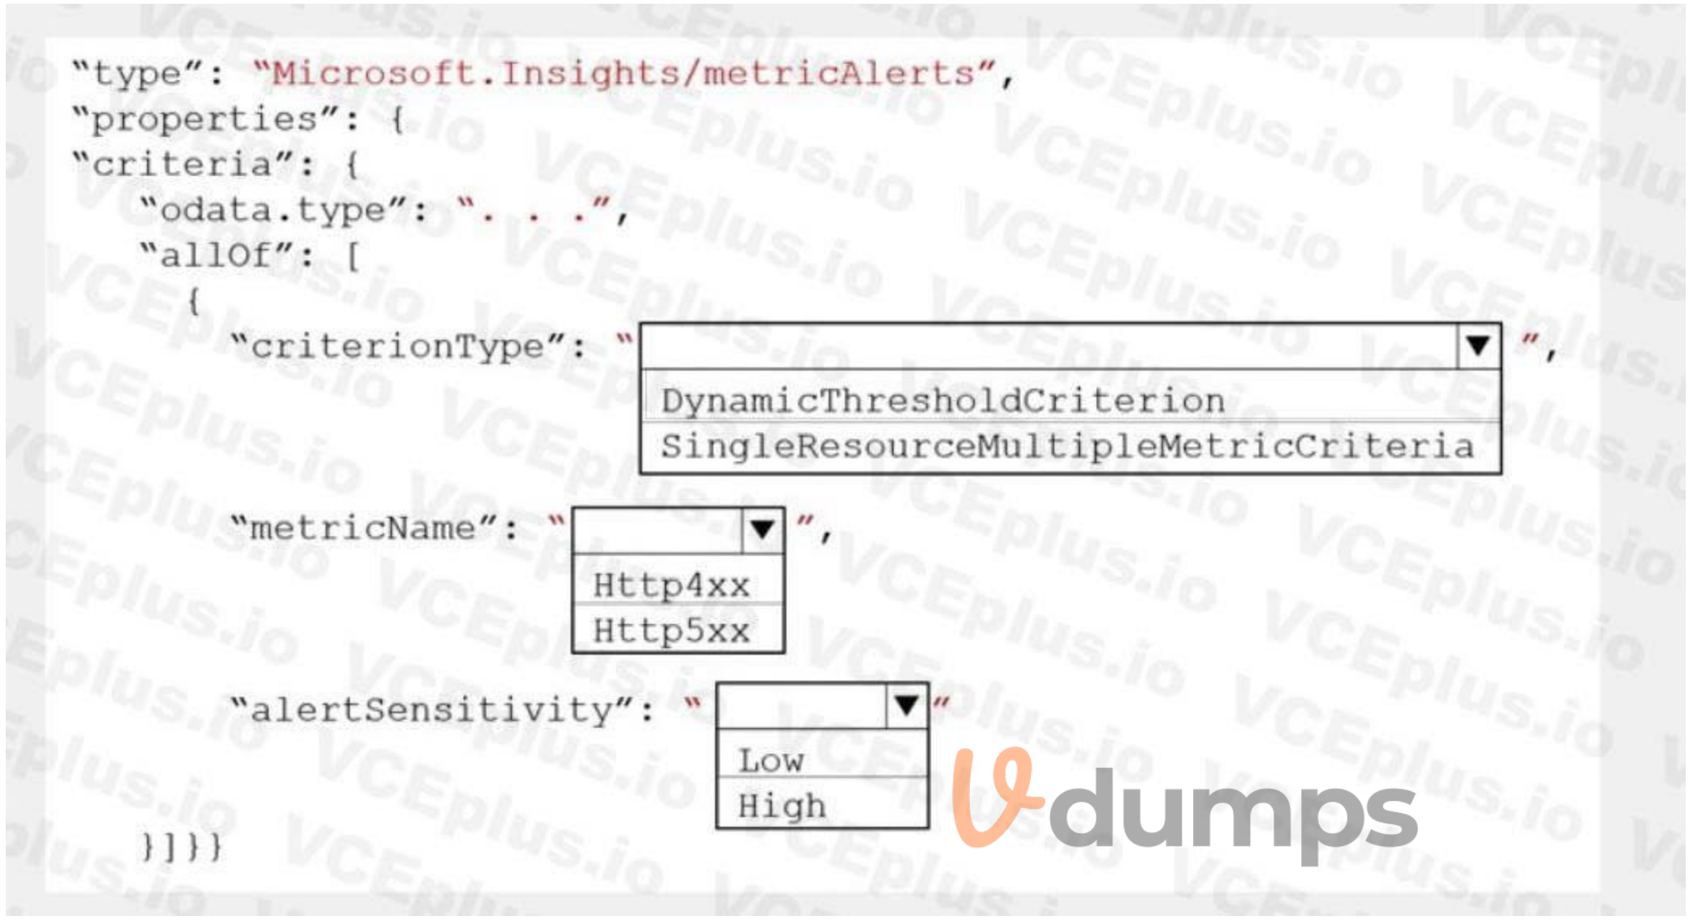
"type": "Microsoft.Insights/metricAlerts",
"properties": {
"criteria": {
"odata.type": ". . .",
"allOf": [
{
"criterionType": "
DynamicThresholdCriterion
SingleResourceMultipleMetricCriteria
",
"metricName": "
Http4xx
Http5xx
",
"alertSensitivity": "
Low
High
"
}]}}

Answer Area:

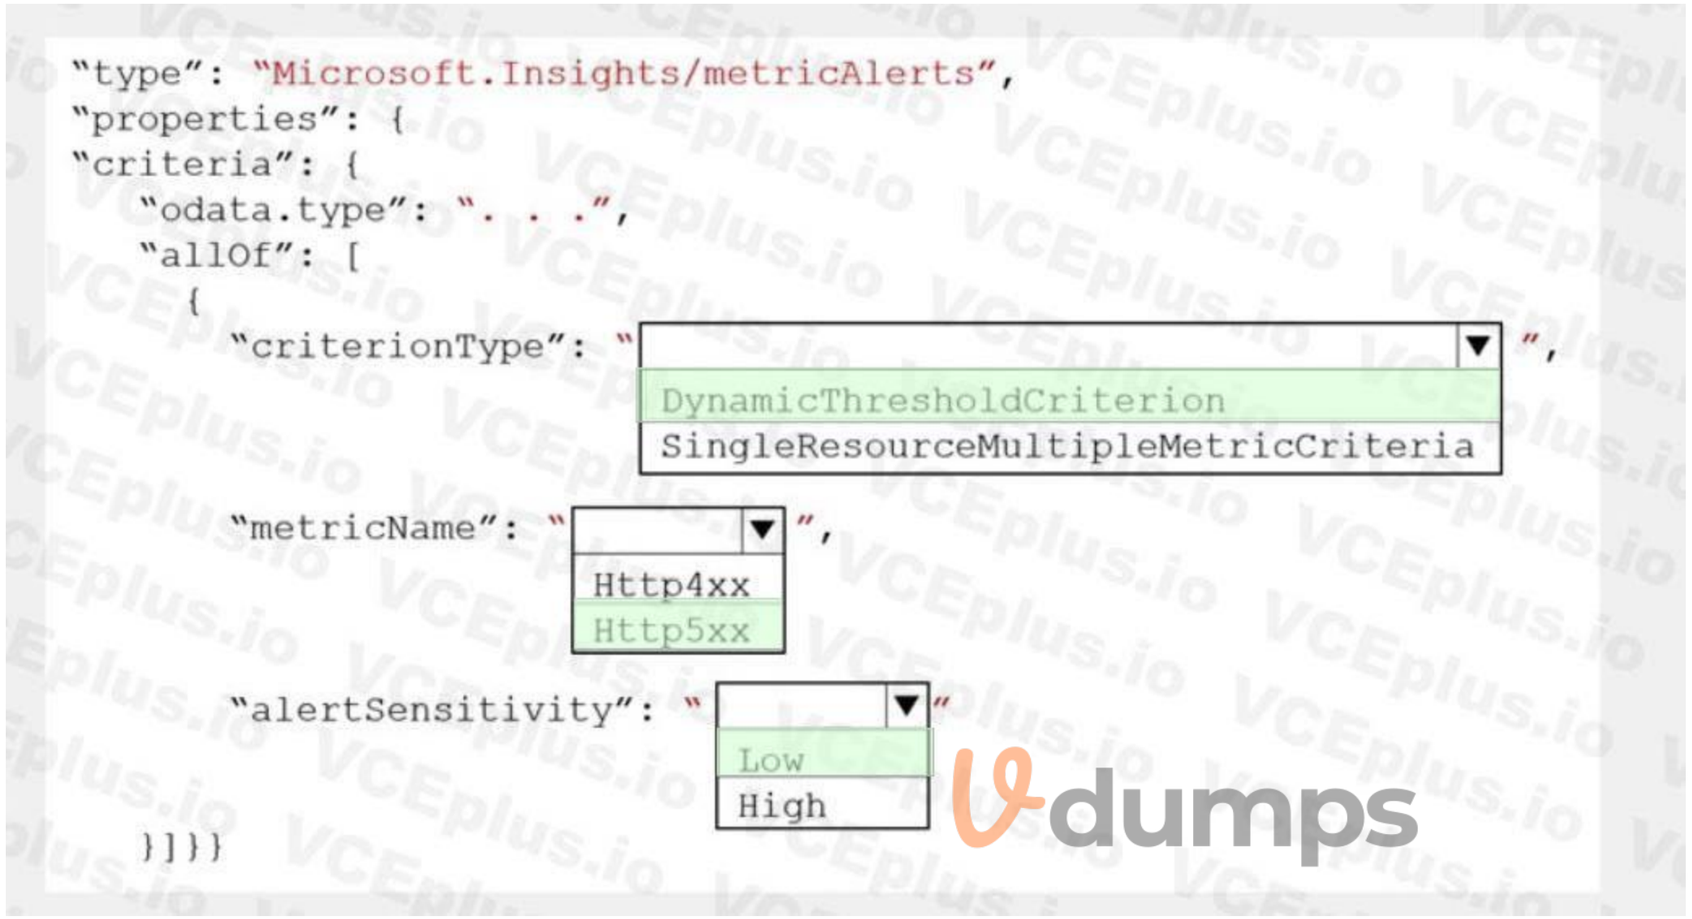

Section:
Explanation:
Box 1: DynamicThresholdCriterion
Box 2: Http5xx
Server errors are in the 5xx range.
Client errors are in the 4xx range
Box 3: Low
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-dynamic-thresholds

QUESTION 8
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.

Hot Area:

Answer Area

```
void ClearCachedTeams()
{
    [ IDatabase cache = Connection.GetDatabase();
      ICache cache = Connection.GetDatabase(); ▼ ]
    [ cache.KeyDelete("Team");
      cache.StringSet("Team", "");
      cache.ValueDelete("Team");
      cache.StringGet("Team",""); ▼ ]
    ViewBag.msg += "Team data removed from cache.";
}
```

Answer Area:

Answer Area

```
void ClearCachedTeams()
{
    [ IDatabase cache = Connection.GetDatabase();
      ICache cache = Connection.GetDatabase(); ▼ ]
    [ cache.KeyDelete("Team");
      cache.StringSet("Team", "");
      cache.ValueDelete("Team");
      cache.StringGet("Team",""); ▼ ]
    ViewBag.msg += "Team data removed from cache.";
}
```


Section:

Explanation:

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/
https://docs.microsoft.com/en-us/cli/azure/webapp/config/container

QUESTION 9
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.

Select and Place:

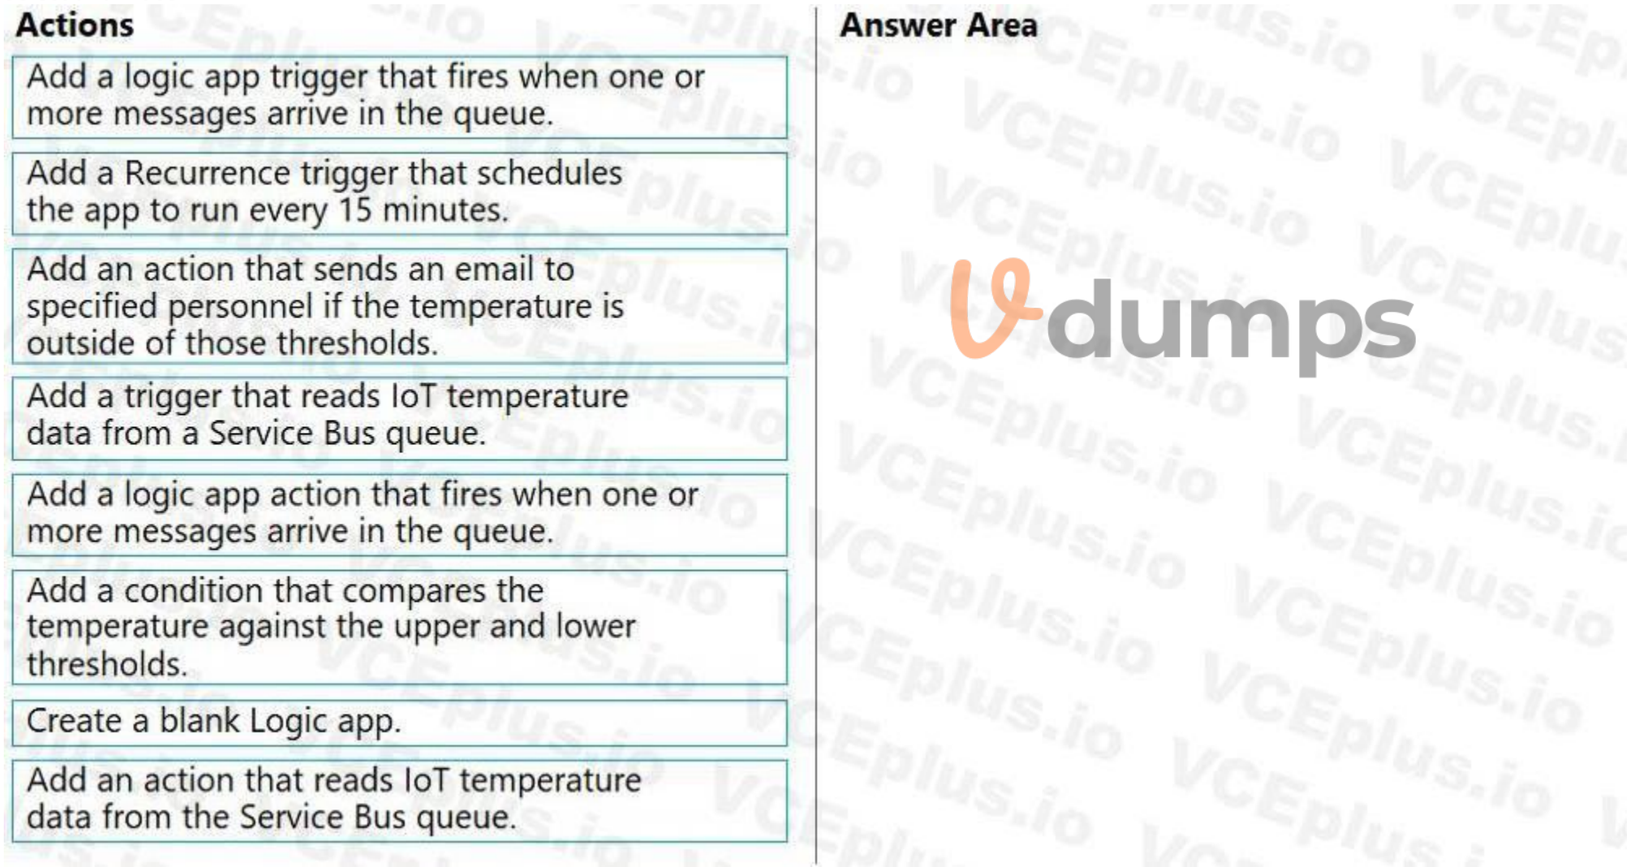

Correct Answer:

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 a trigger that reads IoT temperature data from a Service Bus queue.

Answer Area

- 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.

Section:
Explanation:
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

QUESTION 10
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.

Select and Place:

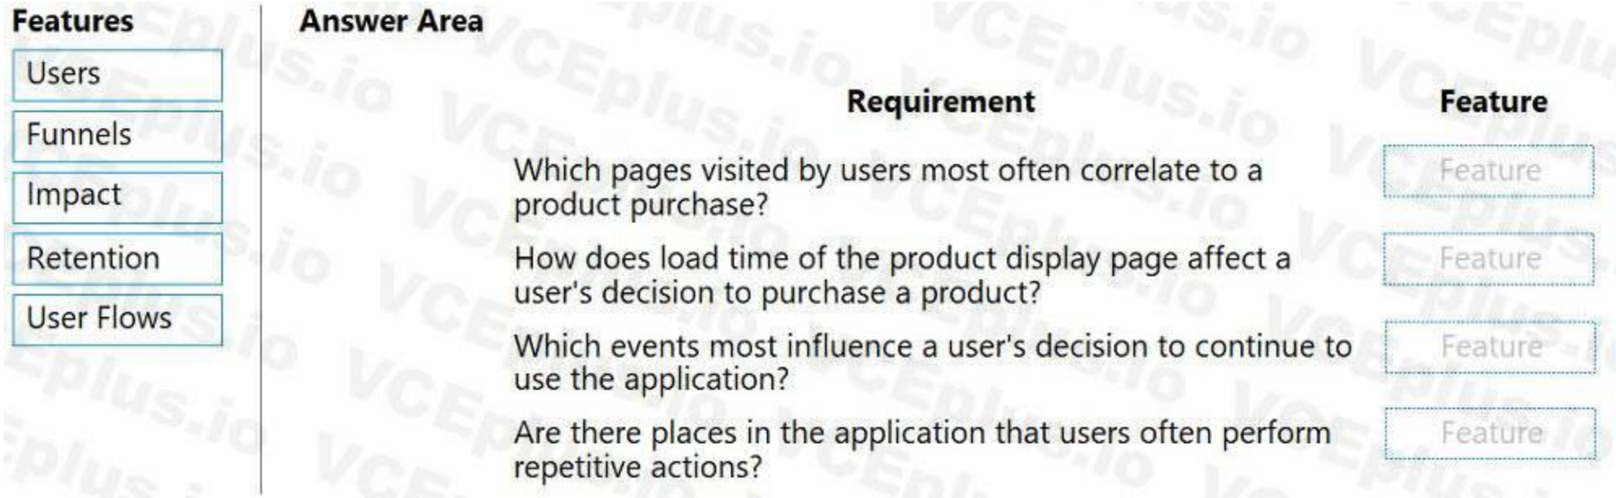

Correct Answer:

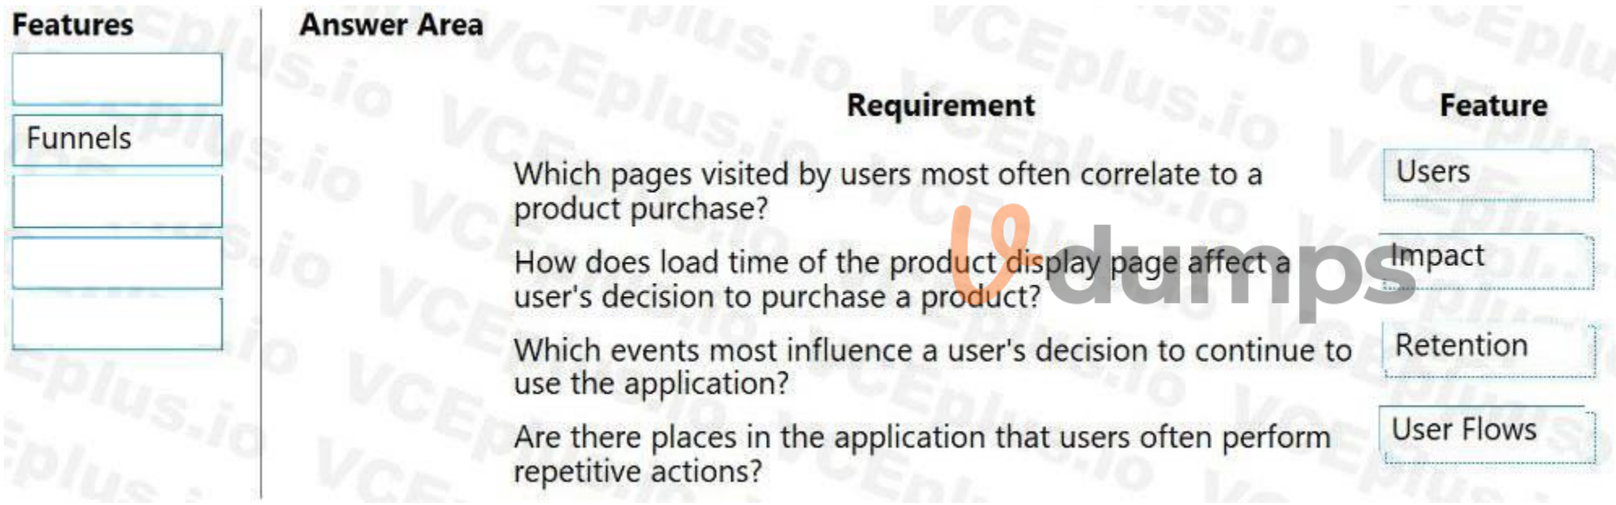

Section:

Explanation:

Box 1: 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?

Incorrect Answers:

Funnel: If your application involves multiple stages, you need to know if most customers are progressing through the entire process, or if they are ending the process at some point. The progression through a series of steps in

a web application is known as a funnel. You can use Azure Application Insights Funnels to gain insights into your users, and monitor step-by-step conversion rates.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-impact

QUESTION 11
You are developing an ASP.NET Core Web API web service. The web service uses Azure Application Insights for all telemetry and dependency tracking. The web service reads and writes data to a database other than Microsoft SQL
Server.
You need to ensure that dependency tracking works for calls to the third-party database.
Which two dependency telemetry properties should you use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

Correct Answer: B, D
Section:
Explanation:
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 12
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You are developing and deploying several ASP.NET web applications to Azure App Service. You plan to save session state information and HTML output.
You must use a storage mechanism with the following requirements:
Share session state across all ASP.NET web applications.
Support controlled, concurrent access to the same session state data for multiple readers and a single writer.
Save full HTTP responses for concurrent requests.
You need to store the information.
Proposed Solution: Enable Application Request Routing (ARR).
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
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 13
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You are developing and deploying several ASP.NET web applications to Azure App Service. You plan to save session state information and HTML output.
You must use a storage mechanism with the following requirements:
Share session state across all ASP.NET web applications.
Support controlled, concurrent access to the same session state data for multiple readers and a single writer.
Save full HTTP responses for concurrent requests.
You need to store the information.
Proposed Solution: Deploy and configure an Azure Database for PostgreSQL. Update the web applications.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
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 14
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets the stated goals.
You are developing and deploying several ASP.NET web applications to Azure App Service. You plan to save session state information and HTML output.
You must use a storage mechanism with the following requirements:
Share session state across all ASP.NET web applications.
Support controlled, concurrent access to the same session state data for multiple readers and a single writer.
Save full HTTP responses for concurrent requests.
You need to store the information.
Proposed Solution: Deploy and configure Azure Cache for Redis. Update the web applications.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: A
Section:
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 15
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.

Hot Area:

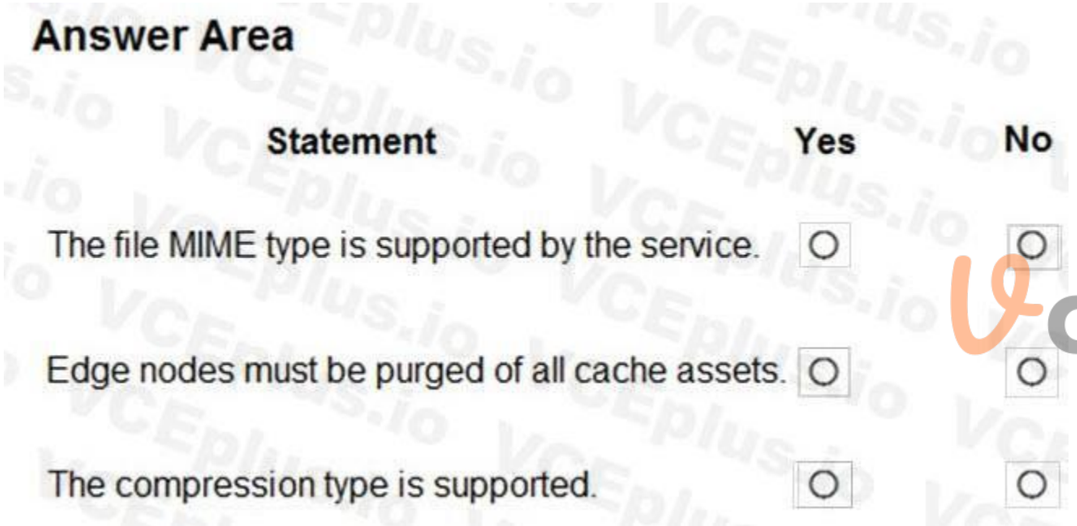

Answer Area:

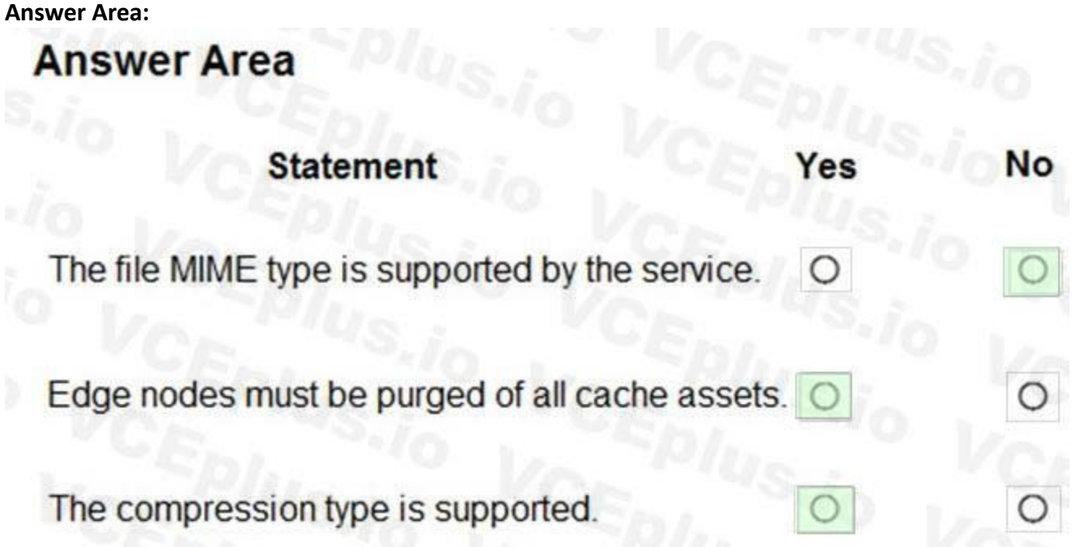

Section:
Explanation:
Box 1: No
Front Door can dynamically compress content on the edge, resulting in a smaller and faster response to your clients. All files are eligible for compression. However, a file must be of a MIME type that is eligible for compression list.
Box 2: No Sometimes you may wish to 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 16
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.

Hot Area:

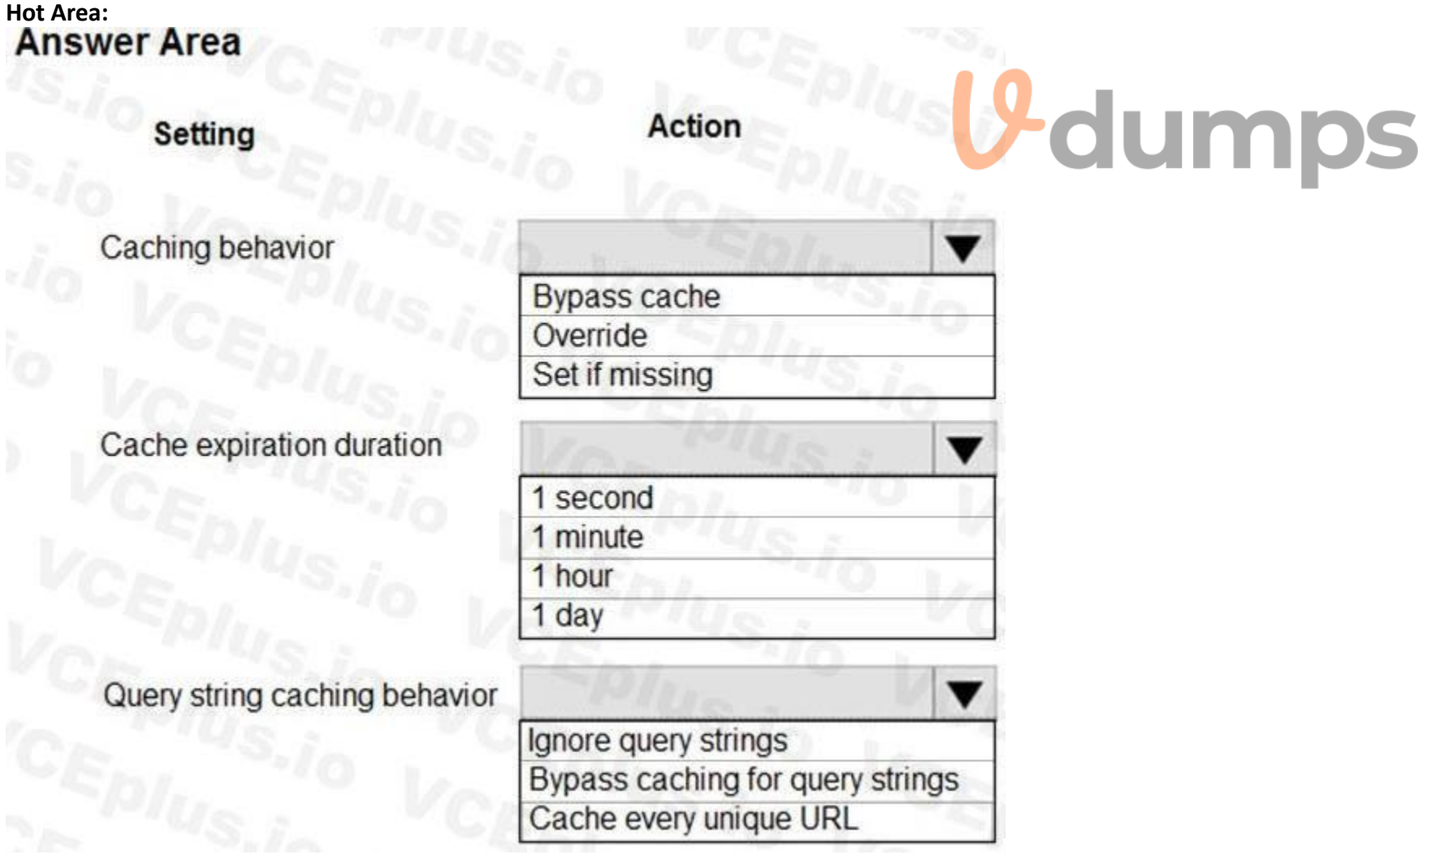

Answer Area:

Answer Area

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

Section:

Explanation:

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.

Incorrect Answers:

Bypass caching for query strings: In this mode, requests with query strings are not cached at the CDN POP node. The POP node retrieves the asset directly from the origin server and passes it to the requestor with each request.

Ignore query strings: Default mode. In this mode, the CDN point-of-presence (POP) node passes the query strings from the requestor to the origin server on the first request and caches the asset. All subsequent requests for the asset that are served from the POP ignore the query strings until the cached asset expires.

Reference:

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

QUESTION 17

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 you select.

Select and Place:

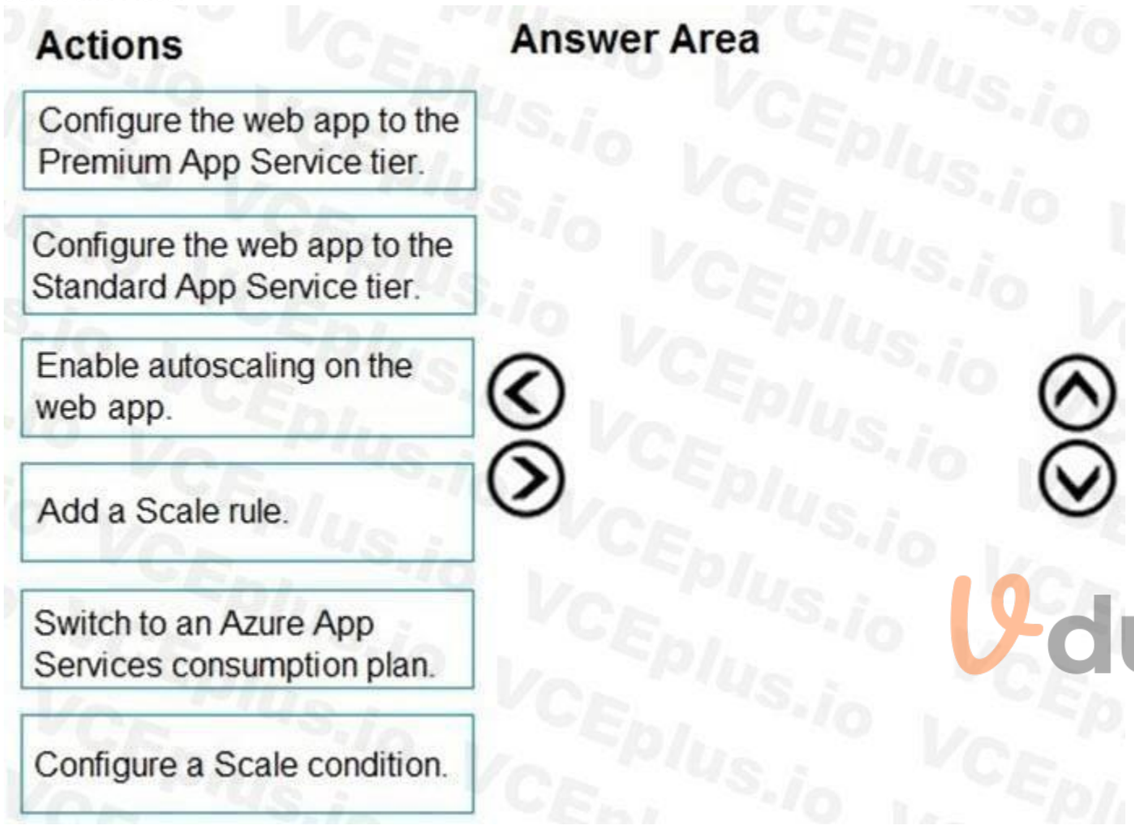

Correct Answer:

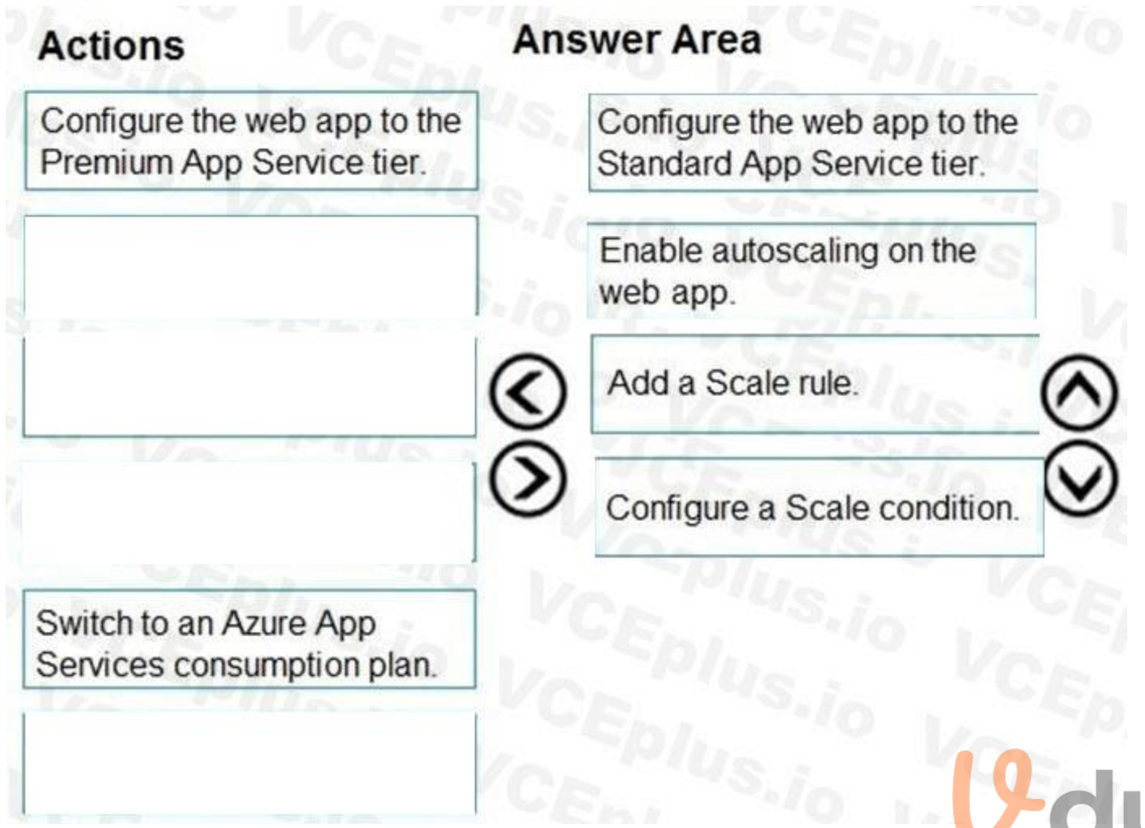

Section:

Explanation:
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 18
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.

Hot Area:

Answer Area

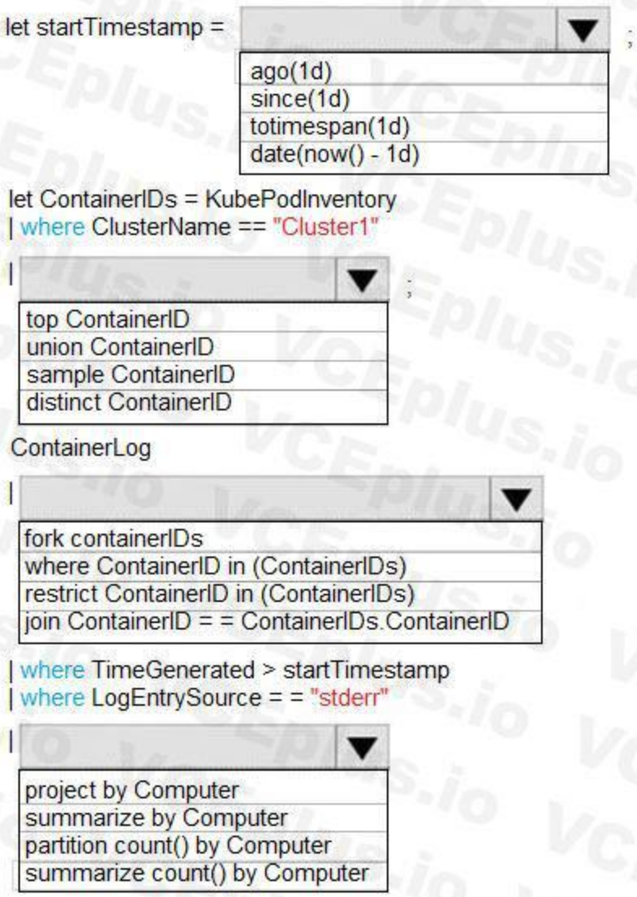

Answer Area:

Answer Area

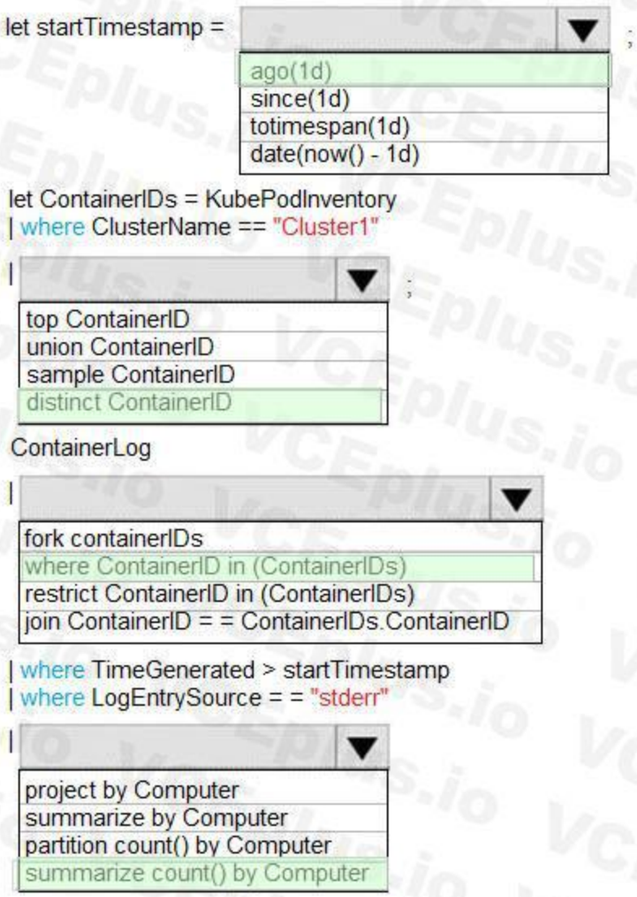

Section:
Explanation:
Box 1: ago(1d)
Box 2: distinct containerID
Box 3: where ContainerID in (ContainerIDs)
Box 4: summarize Count by Computer
Summarize: aggregate groups of rows
Use summarize to identify groups of records, according to one or more columns, and apply aggregations to them. The most common 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 19
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 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.

Hot Area:

Answer Area

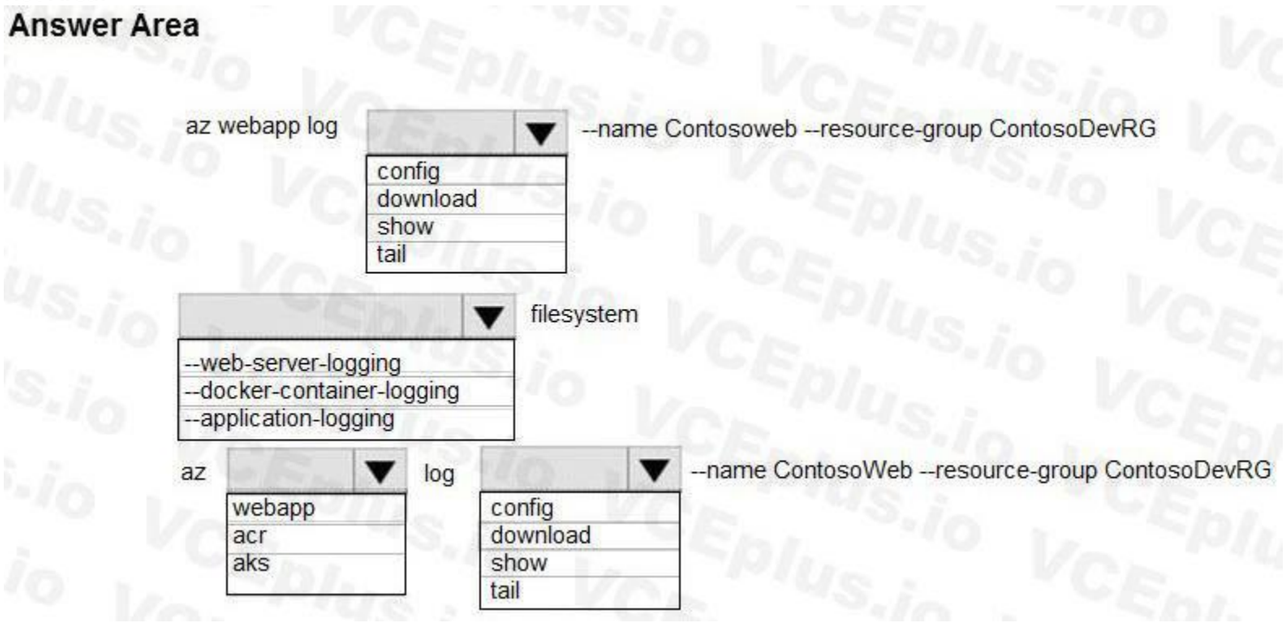

Answer Area:

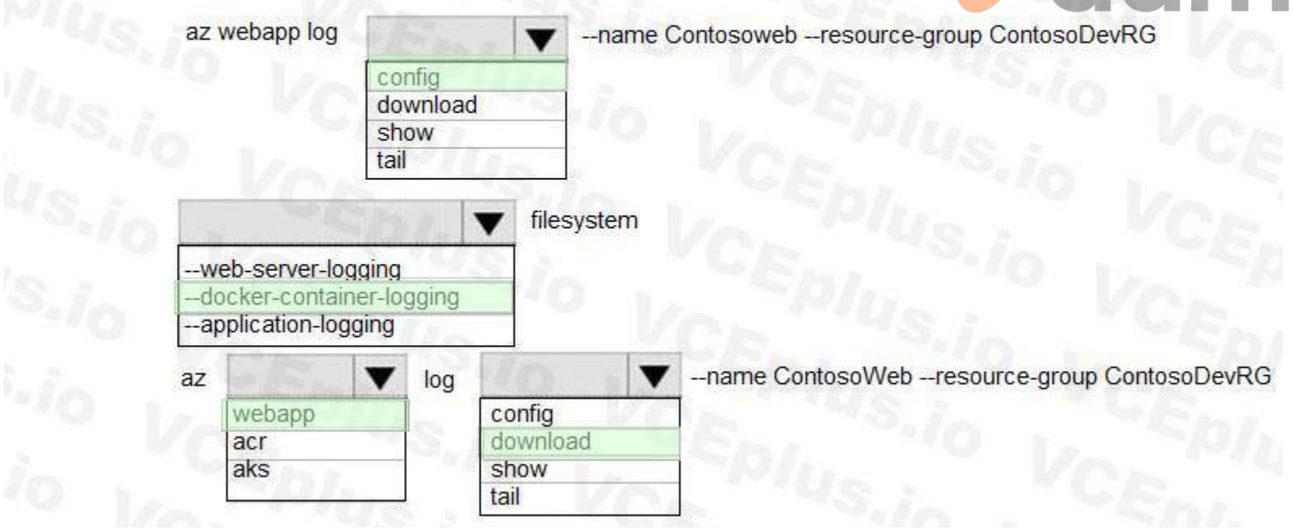

Section:

Explanation:

Box 1: config

To Configure logging for a web app use the command:

az webapp log config

Box 2: --docker-container-logging

Syntax include:

az webapp log config [--docker-container-logging {filesystem, off}]
Box 3: webapp
To download a web app's log history as a zip file use the command:
az webapp log download
Box 4: download
Reference:
https://docs.microsoft.com/en-us/cli/azure/webapp/log

QUESTION 20
You develop a gateway solution for a public facing news API. The news API back end is implemented as a RESTful service and uses an OpenAPI specification.
You need to ensure that you can access the news API by using an Azure API Management service instance.
Which Azure PowerShell command should you run?

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

Correct Answer: D
Section:
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=azurermps-6.13.0

QUESTION 21
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 resents 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 DevileryCount 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.

Correct Answer: A, C
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads

QUESTION 22
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

Correct Answer: A
Section:
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 23
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?

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.

Correct Answer: D
Section:
Explanation:
You can monitor a recorded sequence of URLs and interactions with a website via multi-step web tests.
Incorrect Answers:
A: Selenium is an umbrella project for a range of tools and libraries that enable and support the automation of web browsers.
It provides extensions to emulate user interaction with browsers, a distribution server for scaling browser allocation, and the infrastructure for implementations of the W3C WebDriver specification that lets you write interchangeable code for all major web browsers.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-multistep

QUESTION 24
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.

Correct Answer: B
Section:
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.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/availability-azure-functions

QUESTION 25
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.

Select and Place:

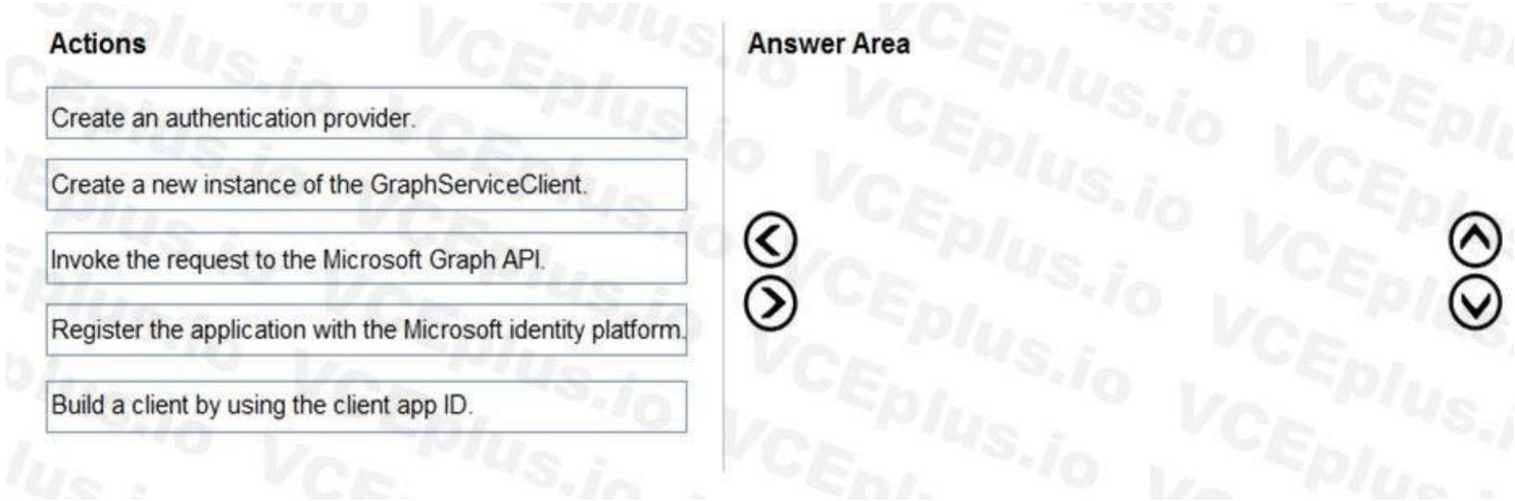

Correct Answer:

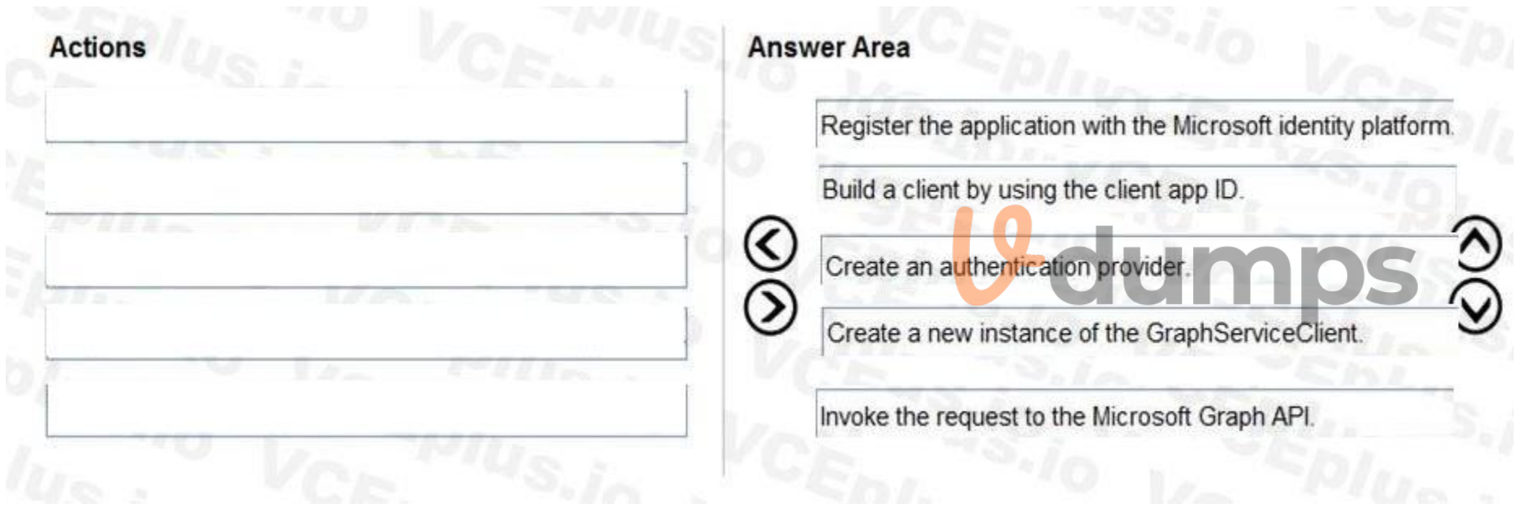

Section:

Explanation:

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 26
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

Correct Answer: B, D
Section:
Explanation:

01 - Connect to and consume Azure services and third-party services
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.
Current environment
Windows Server 2016 virtual machine
This virtual machine (VM) runs BizTalk Server 2016. The VM runs the following workflows:
Ocean Transport - This workflow gathers and validates container information including container contents and arrival notices at various shipping ports.
Inland Transport - This workflow gathers and validates trucking information including fuel usage, number of stops, and routes.
The VM supports the following REST API calls:
Container API - This API provides container information including weight, contents, and other attributes.
Location API - This API provides location information regarding shipping ports of call and trucking stops.
Shipping REST API - This API provides shipping information for use and display on the shipping website.
Shipping Data
The application uses MongoDB JSON document storage database for all container and transport information.
Shipping Web Site
The site displays shipping container tracking information and container contents. The site is located at http://shipping.wideworldimporters.com/
Proposed solution
The on-premises shipping application must be moved to Azure. The VM has been migrated to a new Standard_D16s_v3 Azure VM by using Azure Site Recovery and must remain running in Azure to complete the BizTalk component migrations.
You create a Standard_D16s_v3 Azure VM to host BizTalk Server. The Azure architecture diagram for the proposed solution is shown below:

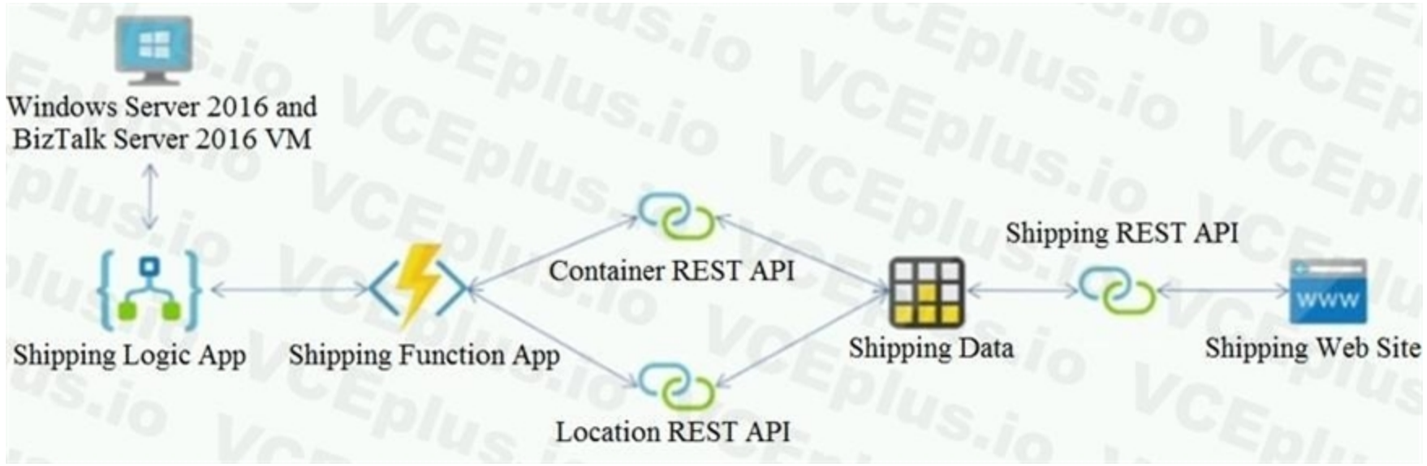

Requirements
Shipping Logic app
The Shipping Logic app must meet the following requirements:
Support the ocean transport and inland transport workflows by using a Logic App.
Support industry-standard protocol X12 message format for various messages including vessel content details and arrival notices.
Secure resources to the corporate VNet and use dedicated storage resources with a fixed costing model.
Maintain on-premises connectivity to support legacy applications and final BizTalk migrations.
Shipping Function app
Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).
REST APIs
The REST API's that support the solution must meet the following requirements:
Secure resources to the corporate VNet.
Allow deployment to a testing location within Azure while not incurring additional costs.
Automatically scale to double capacity during peak shipping times while not causing application downtime.
Minimize costs when selecting an Azure payment model.
Shipping data
Data migration from on-premises to Azure must minimize costs and downtime.
Shipping website
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.
Issues
Windows Server 2016 VM
The VM shows high network latency, jitter, and high CPU utilization. The VM is critical and has not been backed up in the past. The VM must enable a quick restore from a 7-day snapshot to include in-place restore of disks in case of failure.
Shipping website and REST APIs
The following error message displays while you are testing the website:
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin' header is present on the requested resource. Origin 'http:// test.wideworldimporters.com/' is therefore not allowed access.

QUESTION 1

DRAG DROP

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

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

Select and Place:

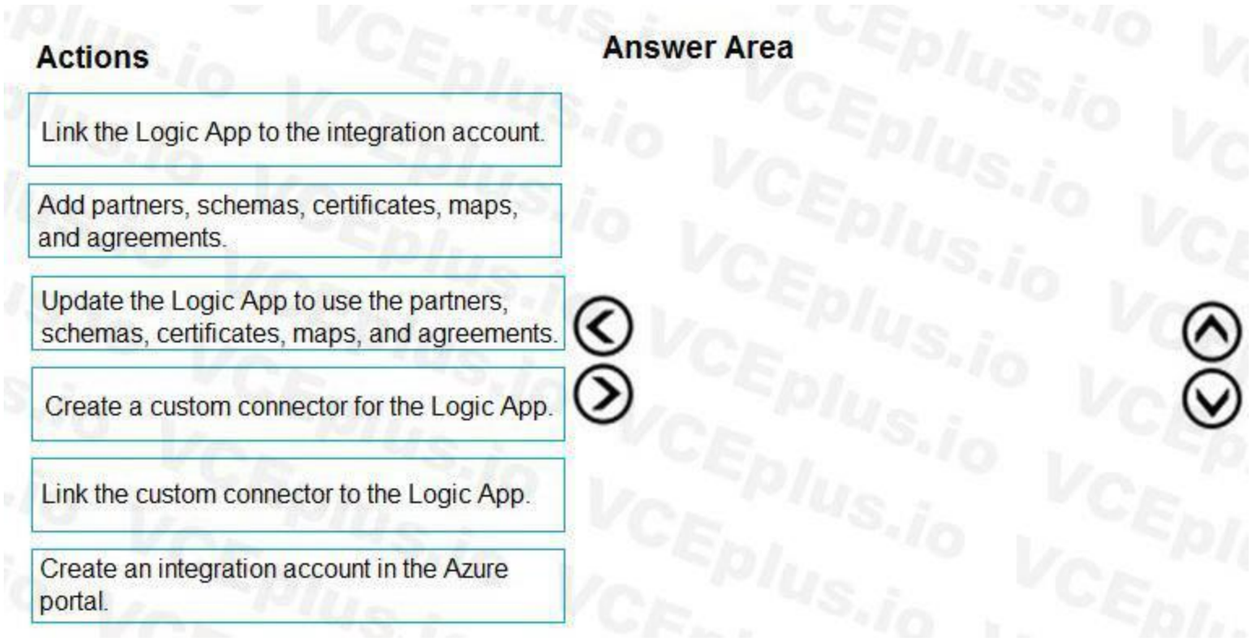

Correct Answer:

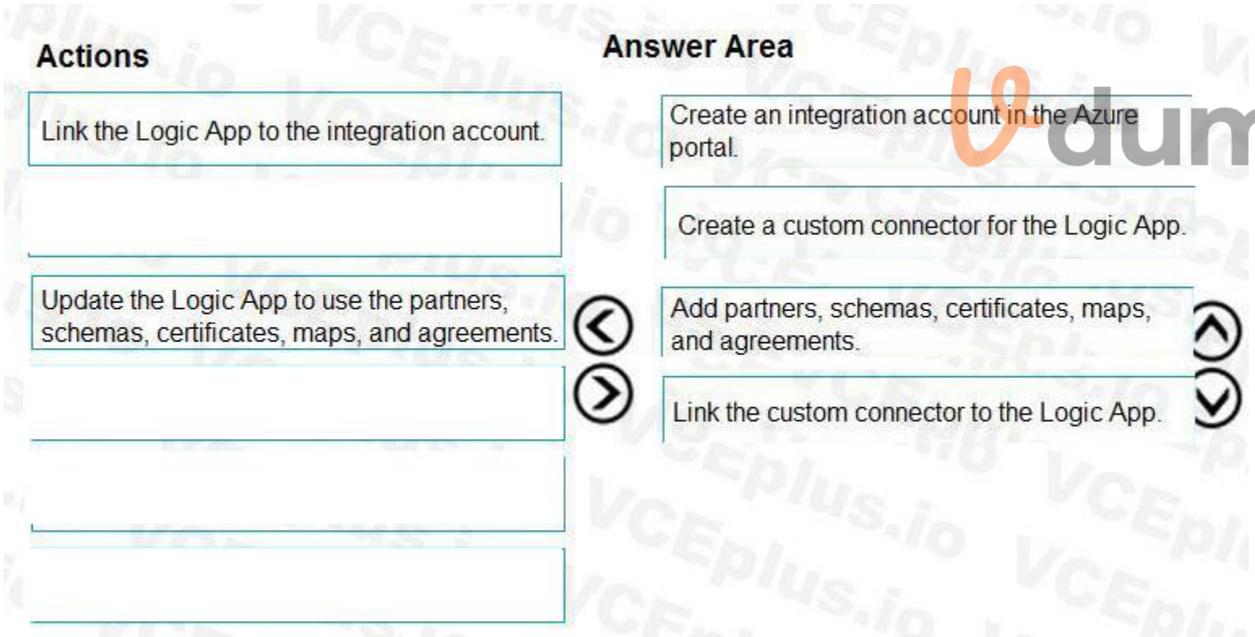

Section:
Explanation:
Step 1: Create an integration account in the Azure portal You can define custom metadata for artifacts in integration accounts and get that metadata during runtime for your logic app to use. For example, you can provide metadata for artifacts, such as partners, agreements, schemas, and maps -all store metadata using key-value pairs.
Step 2: Link the Logic App to the integration account
A logic app that's linked to the integration account and artifact metadata you want to use.
Step 3: Add partners, schemas, certificates, maps, and agreements
Step 4: Create a custom connector for the Logic App.

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

QUESTION 2
You need to troubleshoot the order workflow.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Review the API connections.
B. Review the activity log.
C. Review the run history.
D. Review the trigger history.

Correct Answer: C, D
Section:

QUESTION 3
HOTSPOT
You need to update the order workflow to address the issue when calling the Printer API App.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

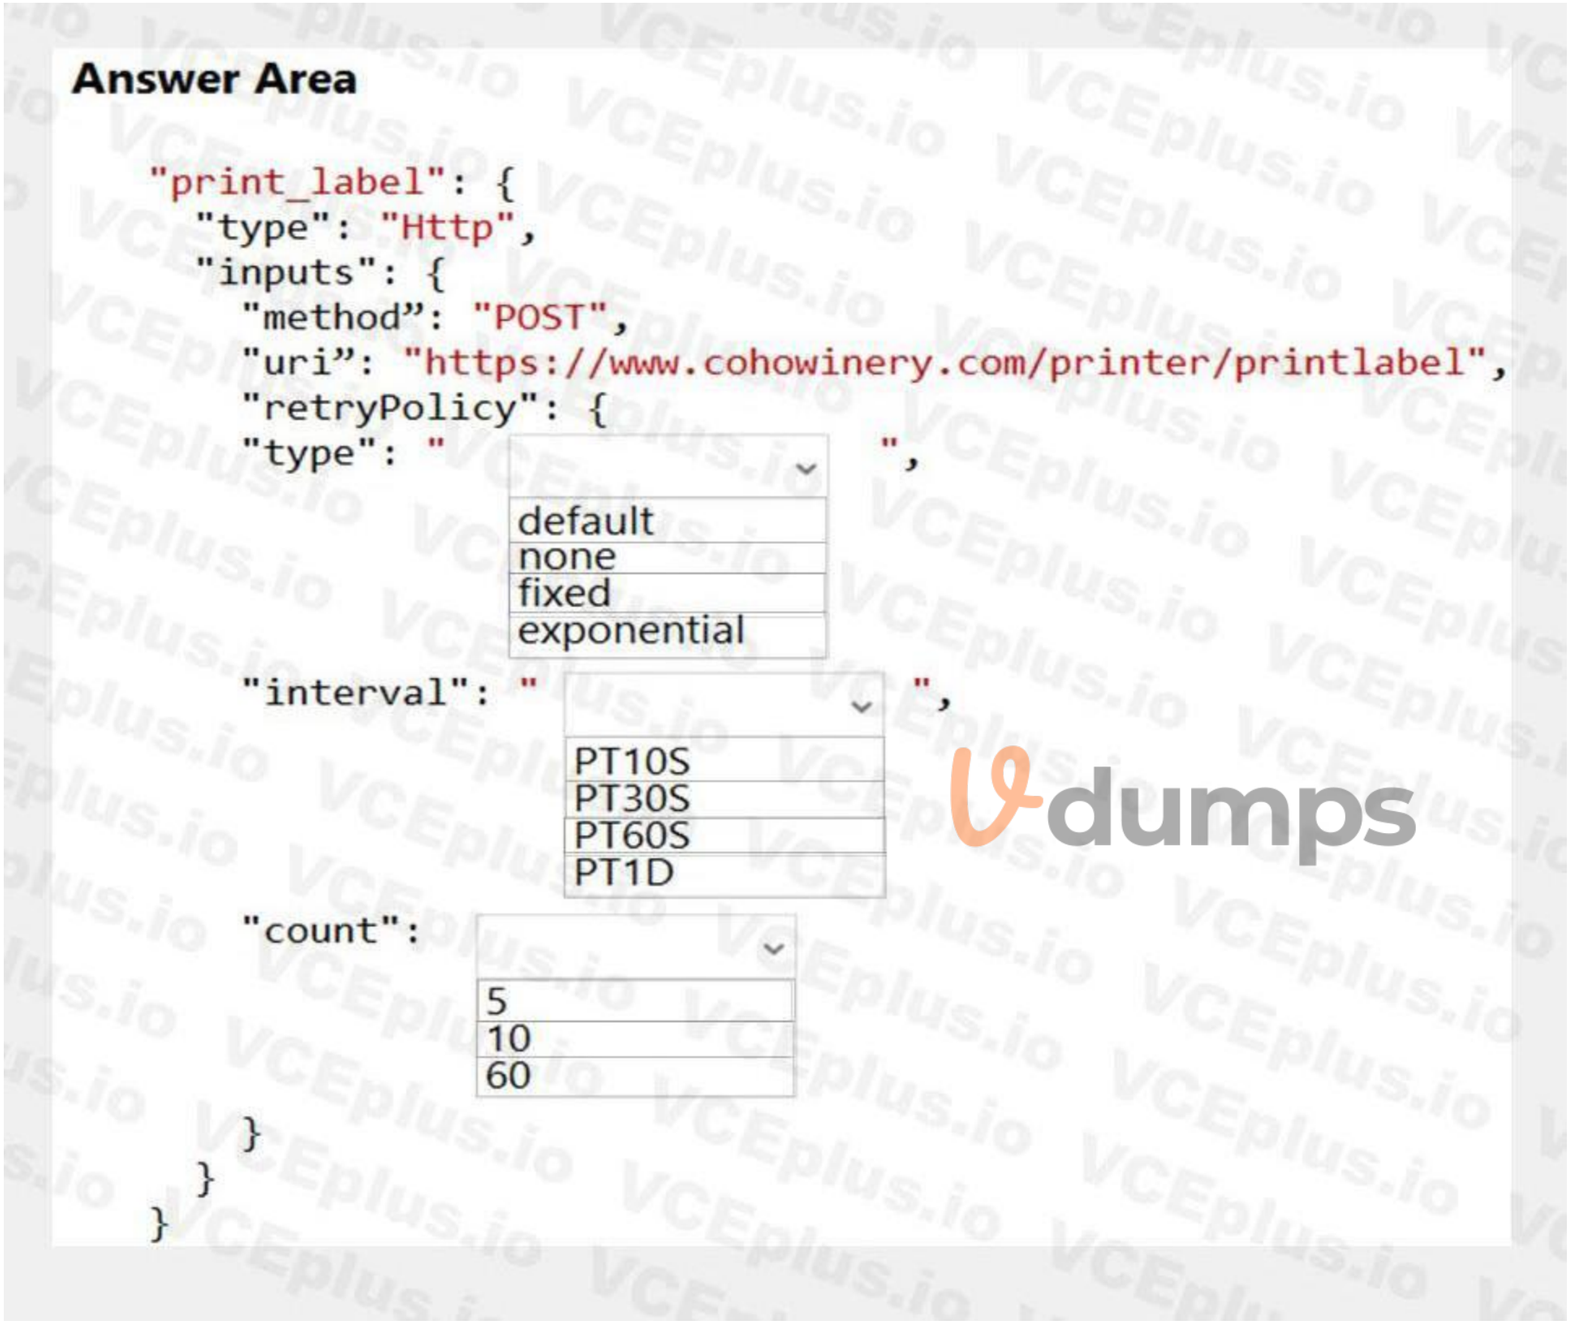

Answer Area:

Answer Area

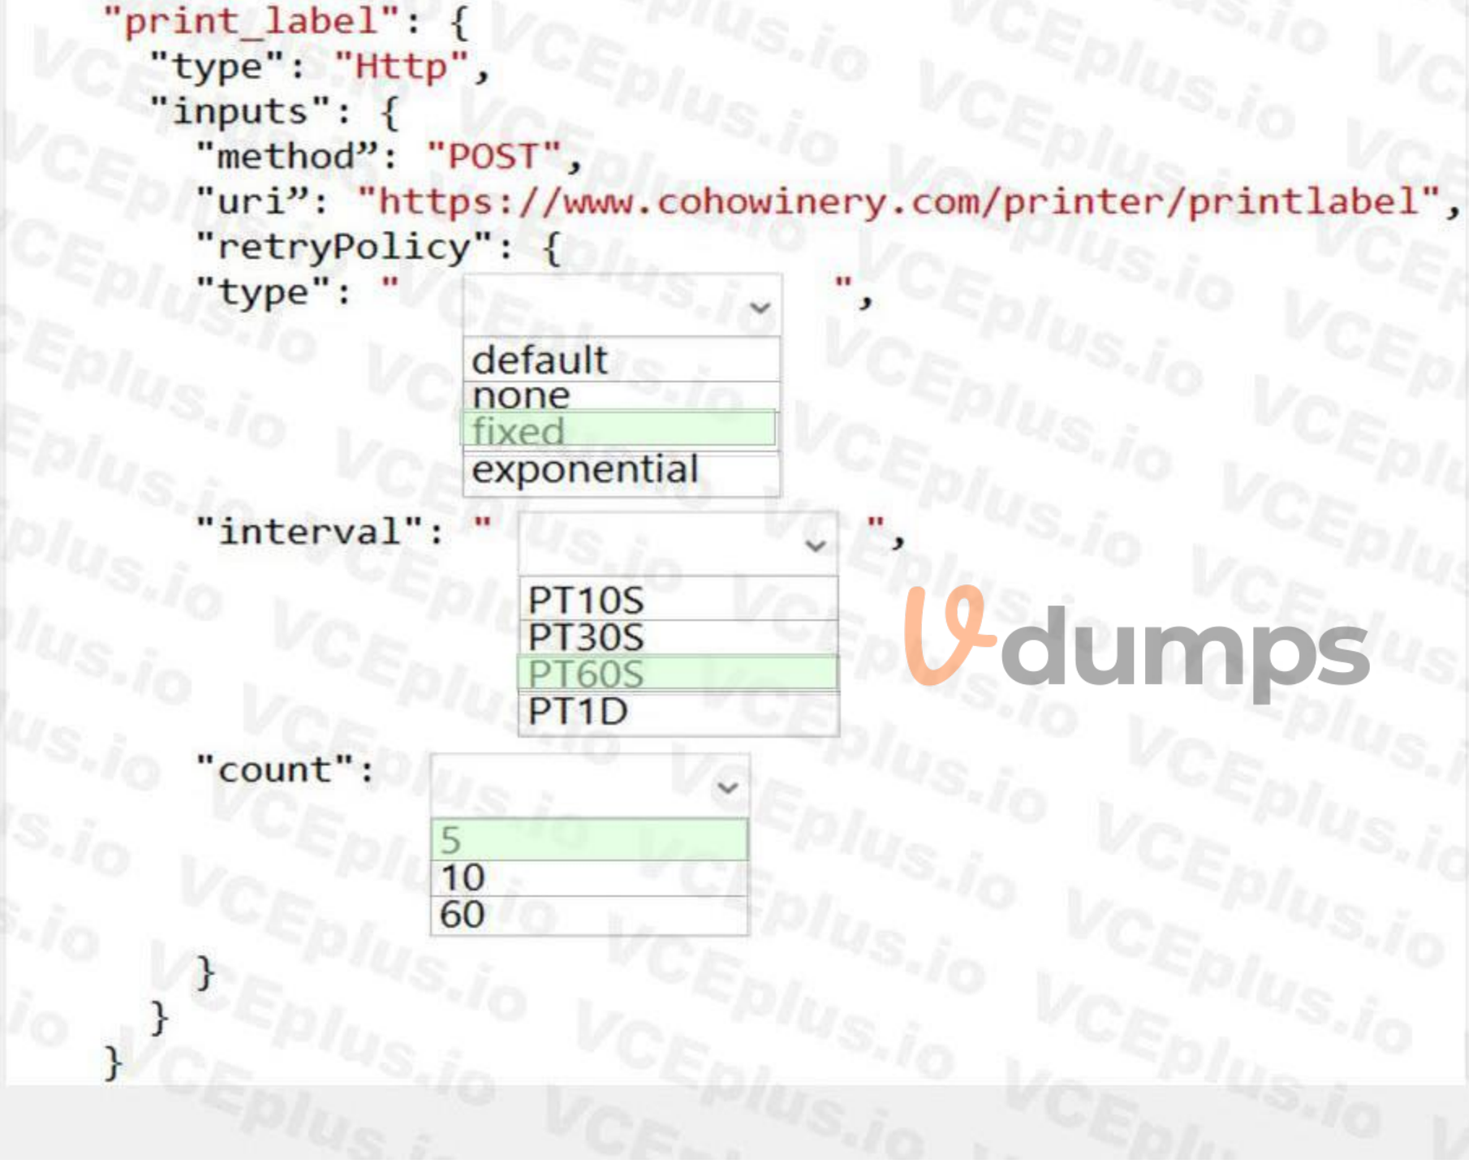

Section:
Explanation:
Box 1: fixed
The 'Default' policy does 4 exponential retries and from my experience the interval times are often too short in situations.
Box 2: PT60S
We could set a fixed interval, e.g. 5 retries every 60 seconds (PT60S).
PT60S is 60 seconds.
Scenario: Calls to the Printer API 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.
Box 3: 5
Reference:
https://michalsacewicz.com/error-handling-in-power-automate/

02 - Connect to and consume Azure services and third-party services

Case study

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

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

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

To start the case study

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

Background

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

Current environment

Architecture overview

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

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

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

Architecture diagram

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

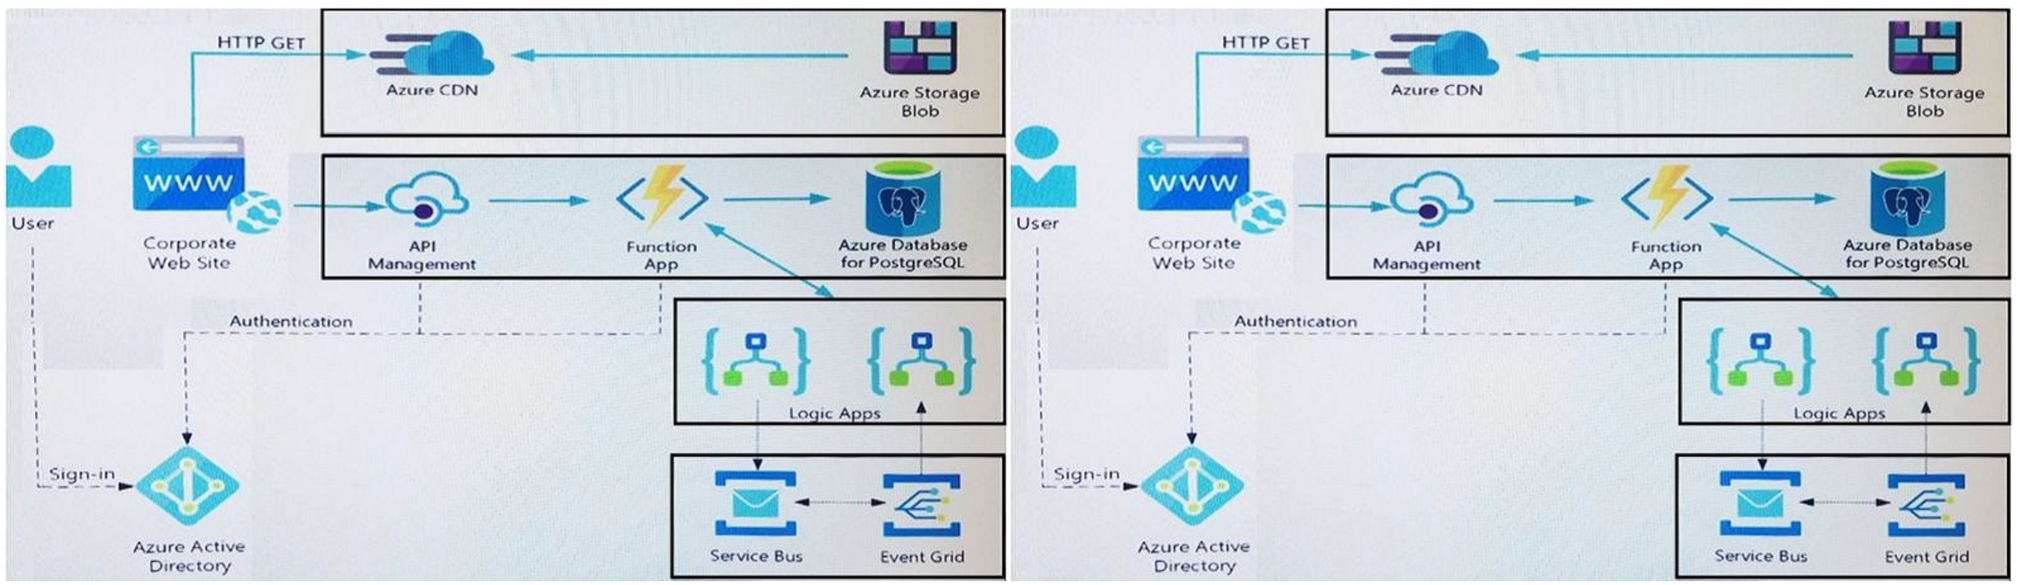

User authentication

The following steps detail the user authentication process:

1. The user selects Sign in in the website.

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

3. The user signs in.

4. Azure AD redirects the user's session back to the web application. The URL includes an access token. 5. The web application calls an API and includes the access token in the authentication header. The application ID is sent as the audience ('aud') claim in the access token. 6. The back-end API validates the access token.

Requirements

Corporate website

Communications and content must be secured by using SSL.

Communications must use HTTPS.

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

Data storage costs must be minimized.

Azure Database for PostgreSQL

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

Azure Key Vault name: cpandlkeyvault

Secret name: PostgreSQLConn

Id: 80df3e46ffcd4f1cb187f79905e9a1e8

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

Azure Service Bus and Azure Event Grid

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

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

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

Security

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

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

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

Compliance

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

Issues

Corporate website

While testing the site, the following error message displays:

CryptographicException: The system cannot find the file specified.

Function app

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

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

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

FunctionAppLogs

| where FunctionName = = "RequestUserApproval"

Logic app

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

'400 Bad Request'

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

Code

Corporate website

Security.cs:

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

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

Function app

RequestUserApproval.cs:

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

QUESTION 1
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

Correct Answer: C
Section:
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

QUESTION 2
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.

Hot Area:

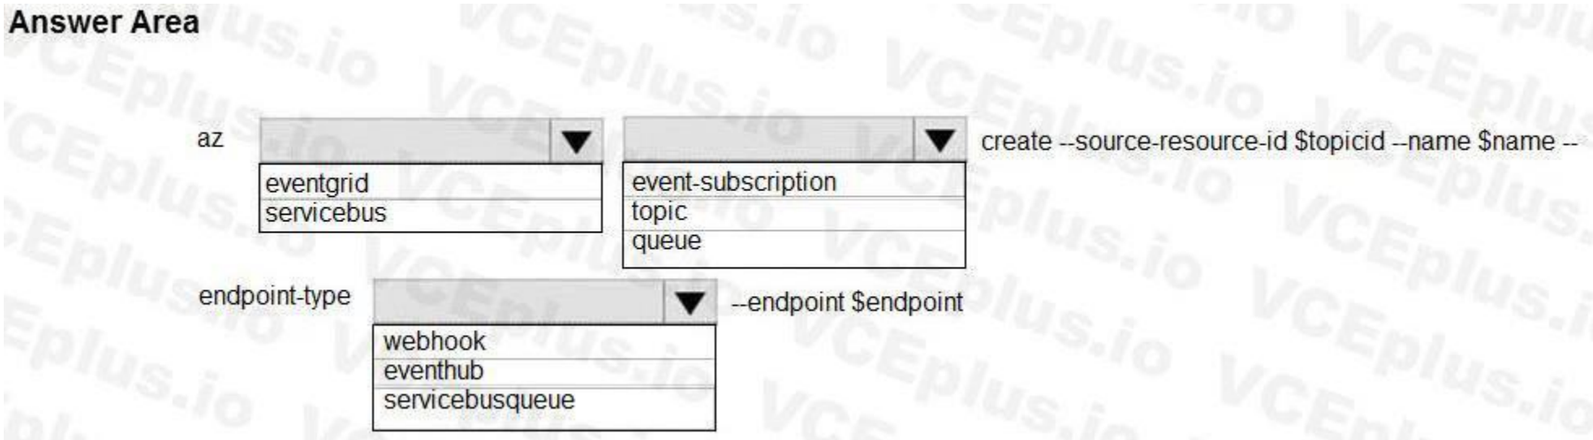

Answer Area:

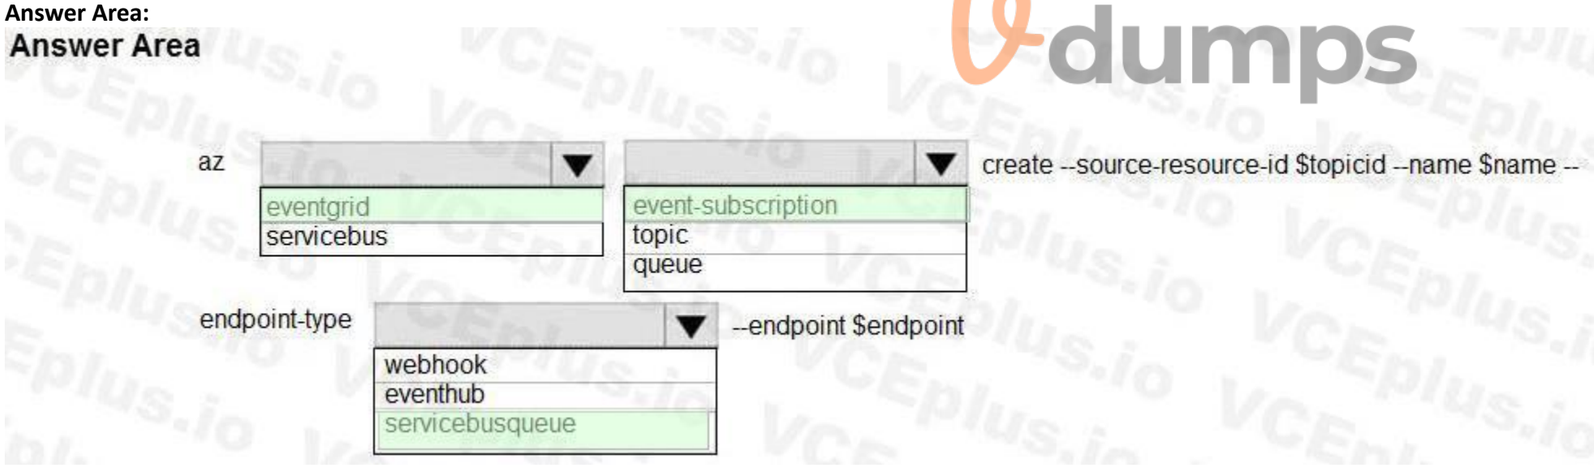

Section:
Explanation:
Box 1: eventgrid
To create event subscription use: az eventgrid event-subscription create
Box 2: event-subscription
Box 3: servicebusqueue
Scenario: Azure Service Bus and Azure Event Grid
Azure Event Grid must use Azure Service Bus for queue-based load leveling.
Events in Azure Event Grid must be routed directly to Service Bus queues for use in buffering.
Events from Azure Service Bus and other Azure services must continue to be routed to Azure Event Grid for processing.
Reference:

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

Mix Questions

QUESTION 1
You are developing a complex workflow by using Azure Durable Functions.
During testing you observe that the results of the workflow differ based on how many instances of the Azure Function are running.
You need to resolve the issue.
What should you do?

A. Ensure that all Orchestrator code is deterministic.
B. Read all state data from the durable function context
C. Configure the Azure Our able f unction to run on an App Service Plan with one instance.
D. Implement the monitor pattern within the workflow.

Correct Answer: A
Section:

QUESTION 2
You are developing an Azure Function App that generates end of day reports (or retail stores. All stores dose at 11 PM each day. Reports must be run one hour after dosing. You configure the function to use a Timer trigger that runs at midnight Customers in the Western United States Pacific Time zone (UTC - 8) report that the Azure Function runs before the stores dose. You need to ensure that the Azure Function runs at midnight in the Pacific Time zone.
What should you do?

A. Configure the Azure Function to run in the West US region.
B. Add an app setting named WEBSITE_TIME_ZONE that uses the value Pacific Standard Time
C. Change the Timer trigger to run at 7 AM
D. Update the Azure Function to a Premium plan.

Correct Answer: A
Section:

QUESTION 3
You are developing an application to manage shipping information for cargo ships. The application will use Azure Cosmos D8 for storage.
The application must run offline when ships are at sea The application must be connected to Azure when ships are in port.
Which Azure Cosmos D8 API should you use for the application?

A. Core
B. MongoDe
C. Cassandra
D. Gremlin

Correct Answer: C
Section:

QUESTION 4
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. Soft delete
B. Change feed
C. Snapshots
D. Object replication
E. Immutability
F. Versioning

Correct Answer: A, B, F
Section:
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 5
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. Create a free tier Azure App Configuration instance with a new Azure AD service principal.
D. Create a standard tier Azure App Configuration instance with an assigned Azure AD managed identity.
E. Store the RSA-HSM key in Azure Cosmos DB. Apply the built-in policies for customer-managed keys and allowed locations.

Correct Answer: A, D
Section:
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 6
You a web application that provides access to legal documents that are stored on Azure Blob Storage with version level immutability policies. Documents are protected with both time-based policies legal hold policies. All time,Ä"based retention policies have AllowProtectedAppendWrites property enabled.
You have a requirement to prevent the user from attempting to perform operations that would fail only a legal is in effect and when all other are expired.
You reed to meet the requirement.
Which two operations you prevent?

A. overwriting existing
B. adding data to documents
C. deleting documents
D. creating document

Correct Answer: A, C

Section:

QUESTION 7
You are developing an Azure Durable Function to manage an online ordering process.
The process must call an external API to gather product discount information.
You need to implement Azure Durable Function.
Which Azure Durable Function types should you use? Each correct answer presents part of the solution
NOTE: Each correct selection is worth ore point

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

Correct Answer: A, B
Section:
Explanation:
https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-featuresoverview

QUESTION 8
You develop a Python application for image rendering that uses GPU resources to optimize rendering processes. You deploy the application to an Azure Container Instances (ACI) Linux container.
The application requires a secret value to be passed when the container is started. The value must only be accessed from within the container.
You need to pass the secret value.
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 environment variable Set the secureValue property to the secret value.
B. Add the secret value to the container image. Use a managed identity.
C. Add the secret value to the application code Set the container startup command.
D. Add the secret value to an Azure Blob storage account. Generate a SAS token.
E. Mount a secret volume containing the secret value in a secrets file.

Correct Answer: A, E
Section:
Explanation:
Objects with secure values are intended to hold sensitive information like passwords or keys for your application. Using secure values for environment variables is both safer and more flexible than including it in your container's image.
Another option is to use secret volumes, described in Mount a secret volume in Azure Container Instances..... https://docs.microsoft.com/en-us/azure/containerinstances/container-instances-environment-variables

QUESTION 9
You develop and deploy a web app to Azure App Service. The Azure App Service uses a Basic plan in a region.
Users report that the web app is responding must capture the complete call stack to help performance issues in code. Call stack data must be correlated across app instances. You must minimize cost and impact to users on the web app.
You need to capture the telemetry.
Which three actions should you perform? Each answer presents part Of the solution
NOTE: Each correct selection is worth point

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

Correct Answer: C, D, F
Section:

QUESTION 10
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

Correct Answer: A
Section:
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.
Incorrect Answers:
Account SAS: Account SAS is specified at the account level. It is secured with the storage account key. User Delegation SAS: A user delegation SAS applies to Blob storage only. Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

QUESTION 11
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. 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.
C. Enable blob versioning for the storage account. Use an Azure Function to process a list of the blob versions per day.
D. Process an Azure Storage blob inventory report by using an Azure Function. Create rule filters on the blob inventory report.
E. Subscribe to blob storage events by using an Azure Function and Azure Event Grid. Filter the events by store location.

Correct Answer: B, E
Section:
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 deadlettering.
Incorrect Answers:
"Enable blob versioning for the storage account. Use an Azure Function to process a list of the blob versions per day": You can enable Blob storage versioning to automatically maintain previous versions of an object. When blob versioning is enabled, you can access earlier versions of a blob to recover your data if it is modified or deleted. Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed https://docs.microsoft.com/enus/azure/storage/blobs/storage-blob-event-overview

QUESTION 12
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.

Hot Area:

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 Area:

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

Section:
Explanation:
Scenario: Retail store locations: Azure Functions must process data immediately when data is uploaded to Blob storage.
Box 1: HTTP
Binding configuration example: https://.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 13
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.

Hot Area:

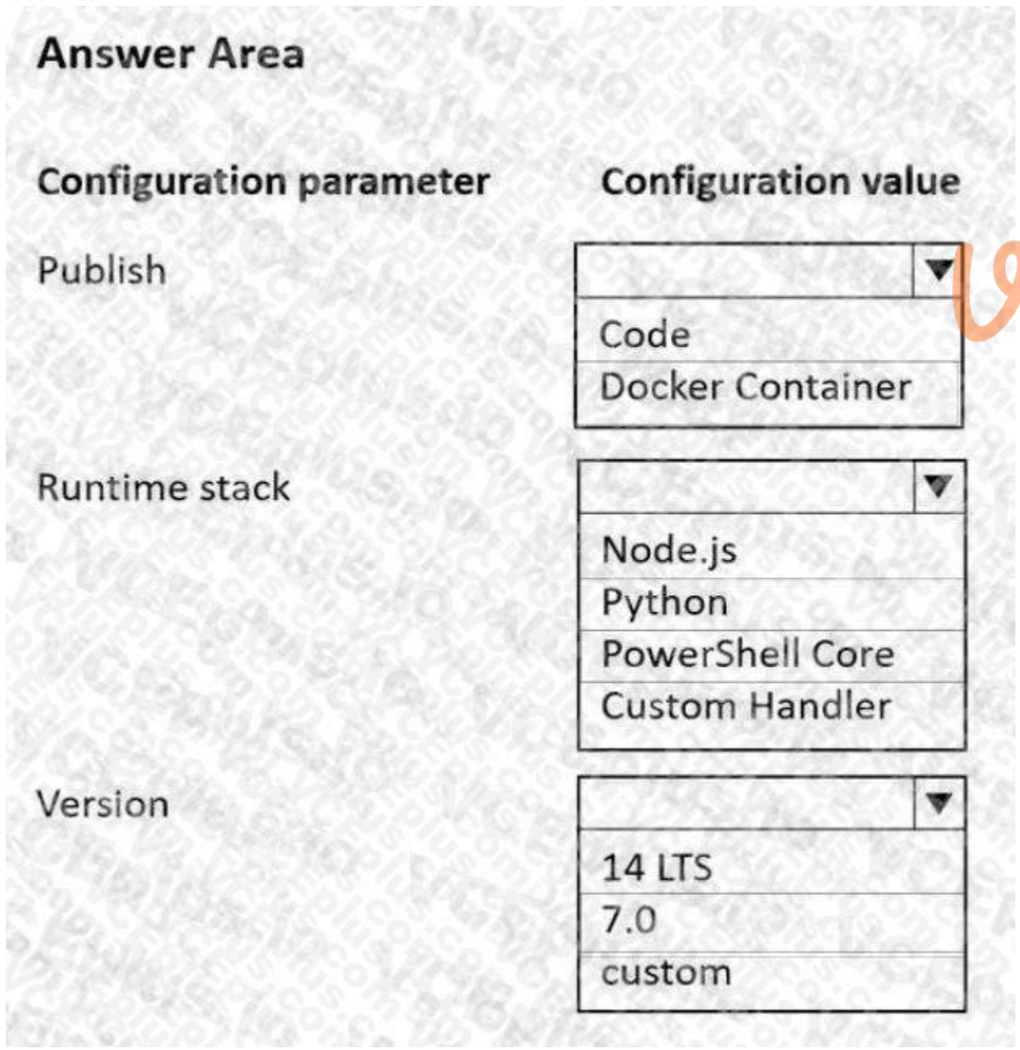

Answer Area:

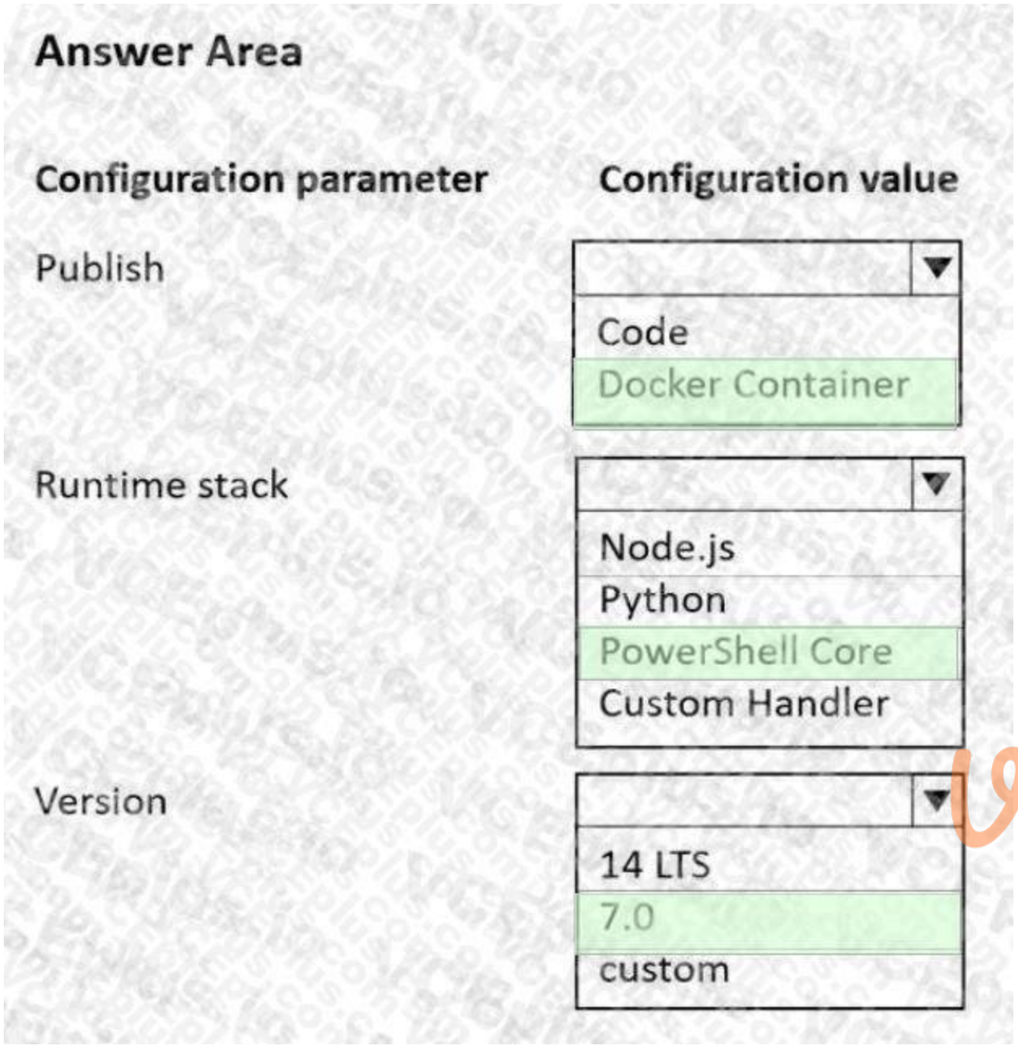

Section:
Explanation:
Box 1: Docker container
A custom handler can be deployed to every Azure Functions hosting option. If your handler requires operating system or platform dependencies (such as a language runtime), you may need to use a custom container. You can create and deploy your code to Azure 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 14
DRAG DROP
You provision virtual machines (VMs) as development environments.

One VM does not start. 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.

Select and Place:

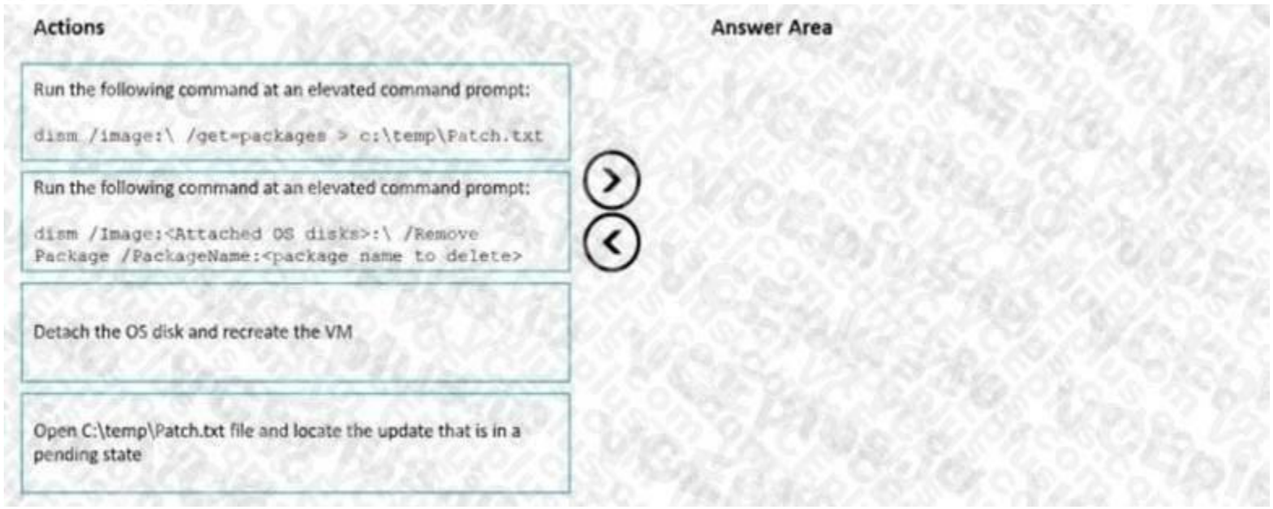

Correct Answer:

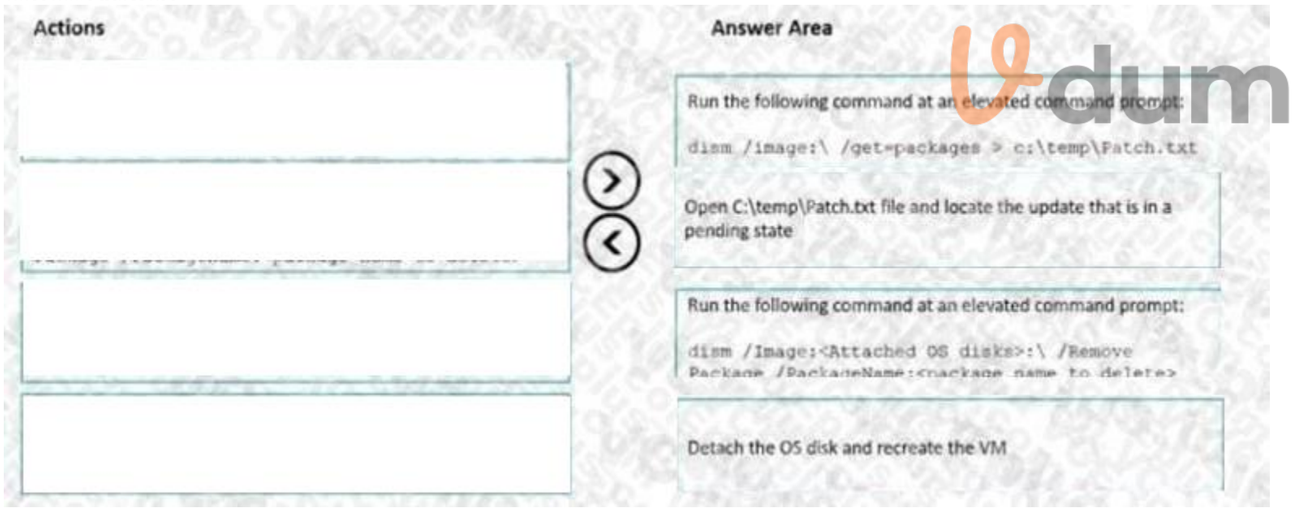

Section:
Explanation:
Remove the update that causes the problem
1. Take a snapshot of the OS disk of the affected VM as a backup.
2. Attach the OS disk to a recovery VM.
3. Once the OS disk is attached on the recovery VM, run diskmgmt.msc to open Disk Management, and ensure the attached disk is ONLINE.
4. (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::\ /get-packages > c:\temp\Patch_level
5. (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.
6. Remove the update that caused the problem:

dism /Image::\ /Remove-Package /PackageName:<>
7. (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 15
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.

Hot Area:

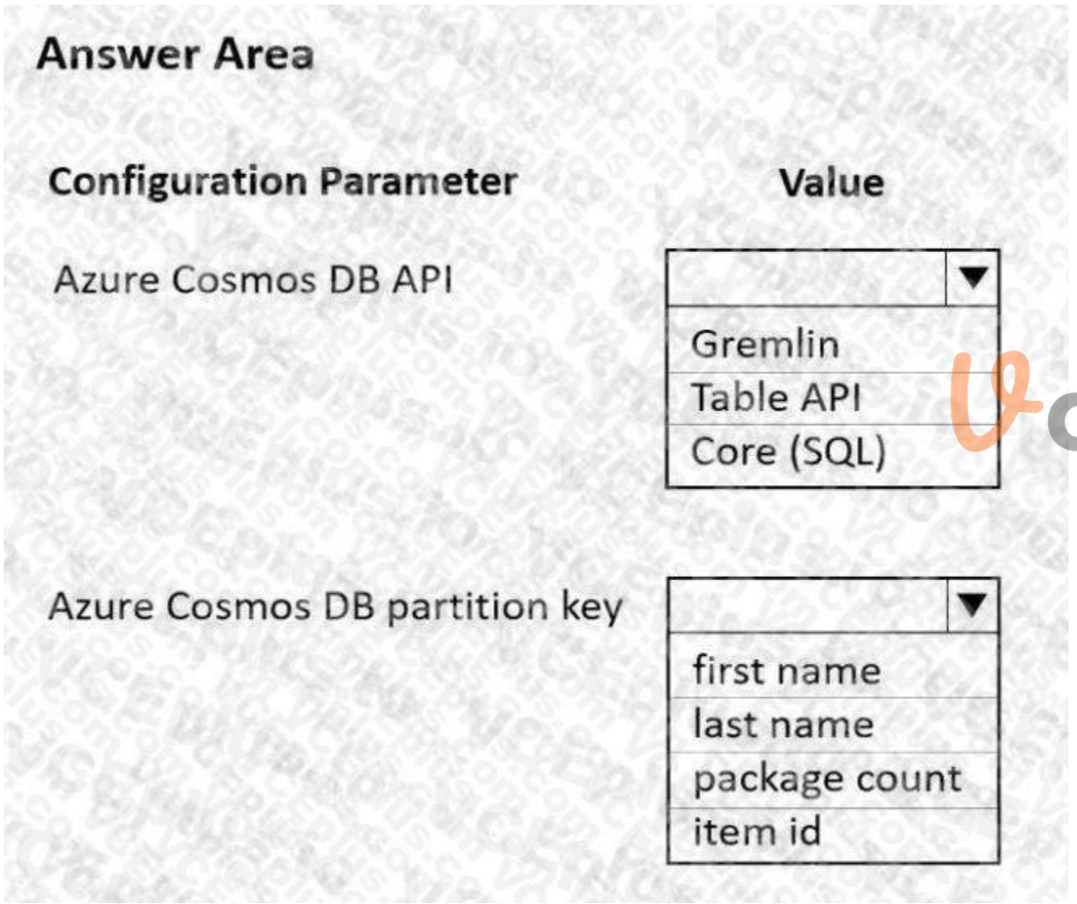

Answer Area:

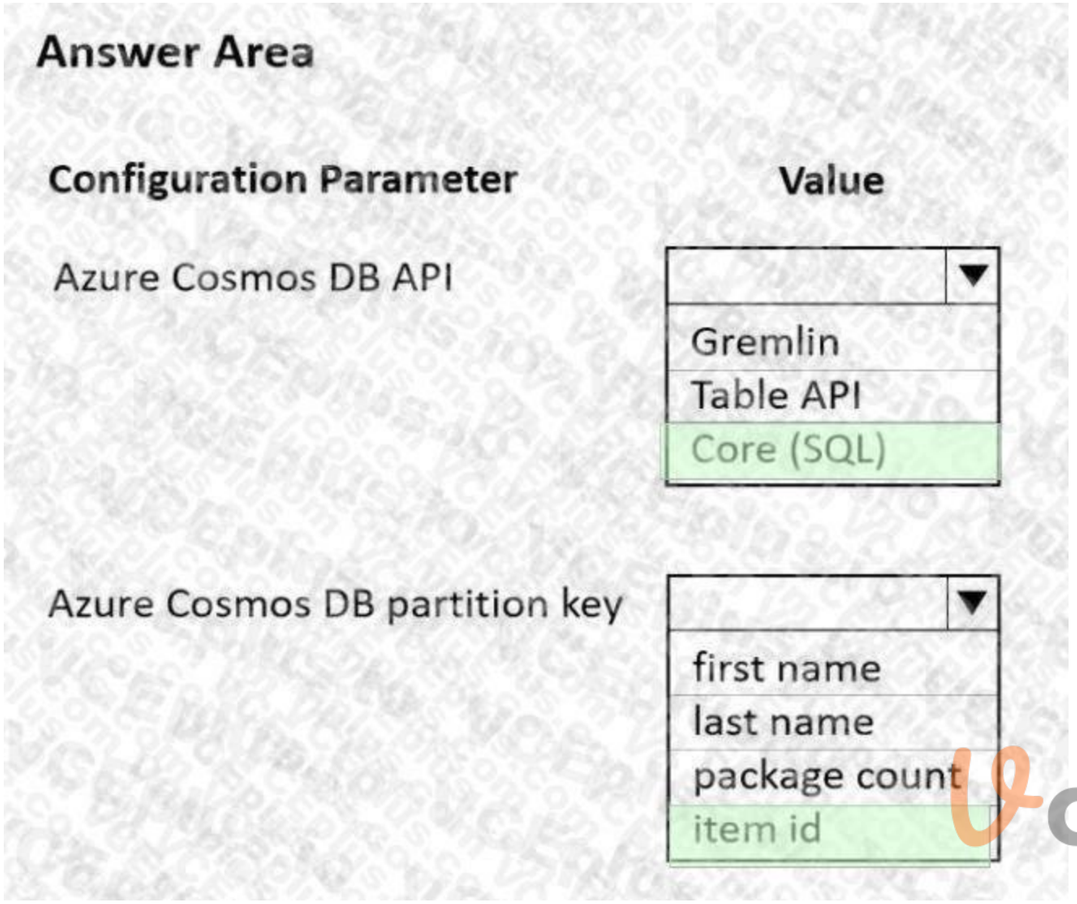

Section:
Explanation:
Box 1: Core (SQL)
Core(SQL) API stores data in document format. It offers the best end-to-end experience as we have full control over the interface, service, and the SDK client libraries. SQL API supports analytics and offers performance isolation between operational and analytical 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 16
You are developing an application to store business-critical data in Azure Blob storage. The application must meet the following requirements:
- Data must not be modified or deleted for a user-specified interval.
- Data must be protected from overwntes and deletes.
- Data must be written once and allowed to be read many times.
You need to protect the data fen the Azure Blob storage account.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

E. Enable the blob change feed for the storage account.

Correct Answer: D, E
Section:

QUESTION 17
You develop Azure Durable Functions to manage vehicle loans.
The loan process includes multiple actions that must be run in a specified order. One of the actions includes a customer credit check process, which may require multiple days to process.
You need to implement Azure Durable Functions for the loan process.
Which Azure Durable Functions type should you use?

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

Correct Answer: A
Section:

QUESTION 18
HOTSPOT
You develop and deploy an Azure App Service web app that connects to Azure Cache for Redis as a content cache. An resources have been deployed to East US 2 region.
The security team requires the from Azure Cache for Redis:
The number of Redis client connections from an associated IP address.
Redis operations completed on the content cache.
The location (region) in which the Azure Cache for Redis instance was accessed.
The audit information must be captured and analyzed by a security team application deployed to Central US region
You need to log information on all client corrections to the cache.
Which configuration values should you use?

Hot Area:

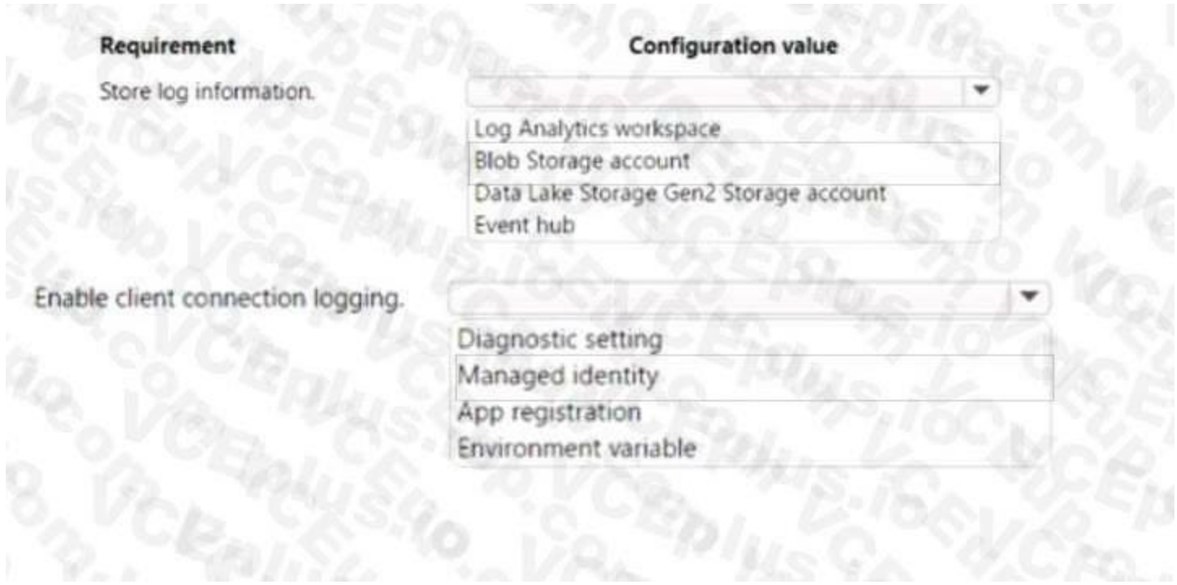

Answer Area:

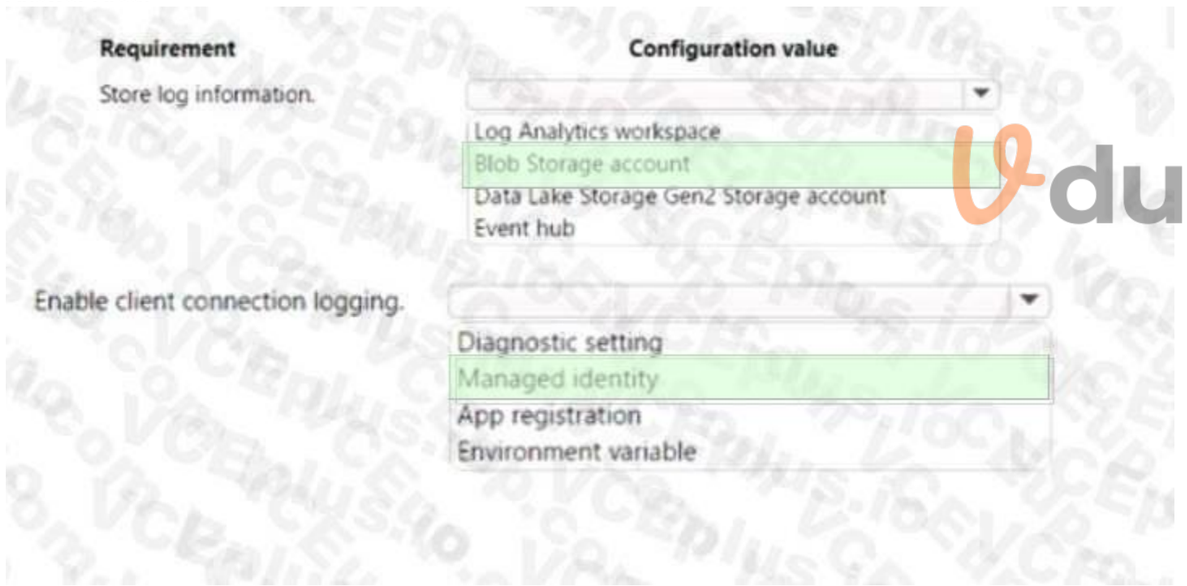

Section:
Explanation:

QUESTION 19
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 as shown in the exhibit:

You must be able to use the feature in the app by using the following markup:

```
<feature name="Export">
  <li class="nav-item">
    <a class="nav-link text-dark" asp-area="" asp-controller="Home" asp-action="Export">Export Data</a>
  </li>
</feature>
```

You need 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.

Hot Area:

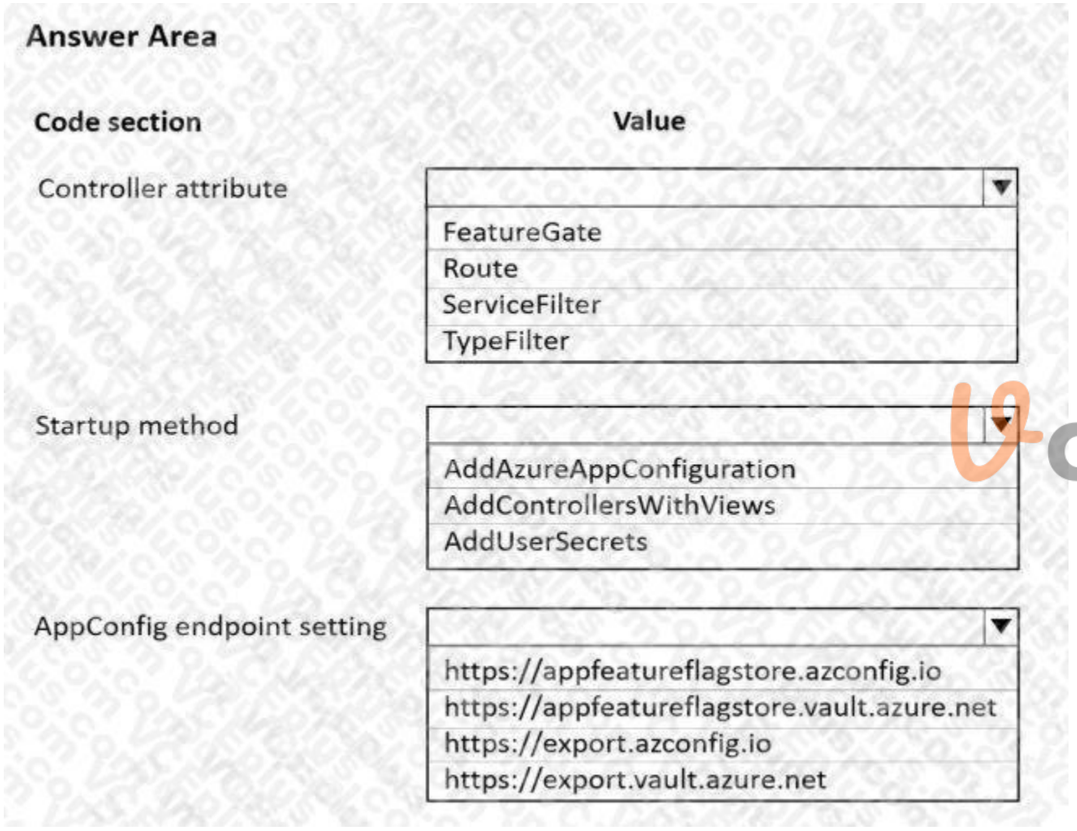

Answer Area:

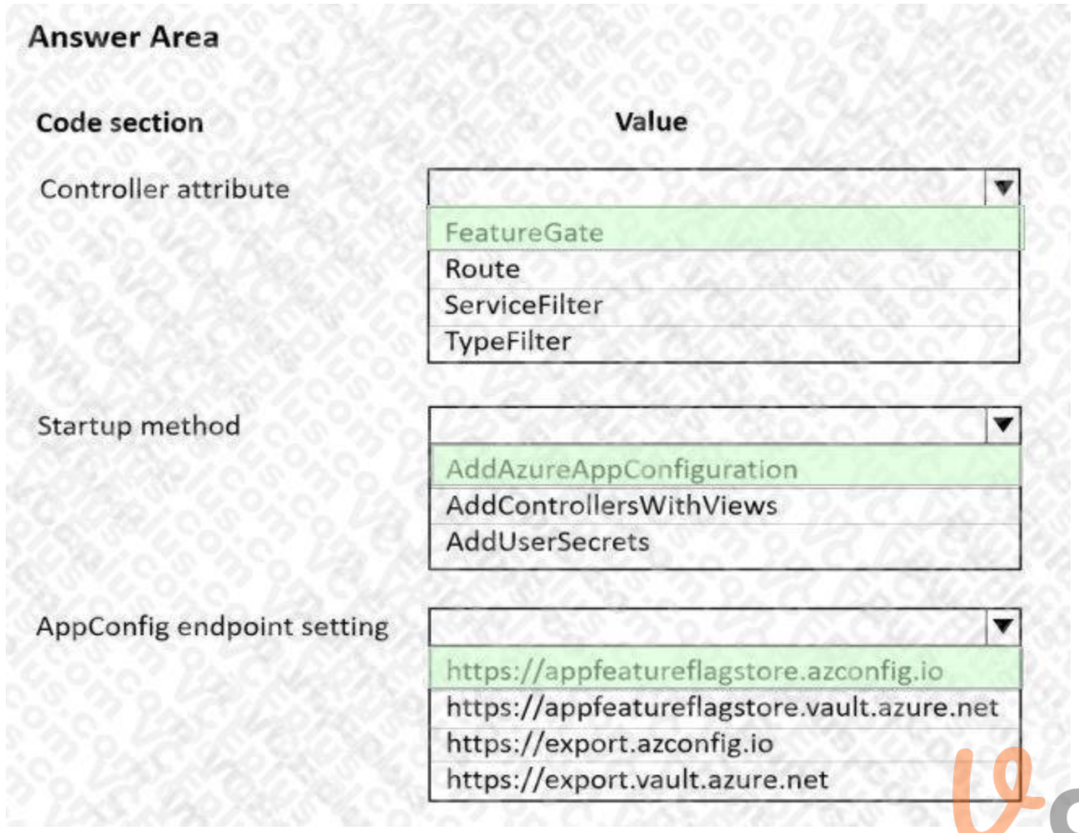

Section:

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.azconfig.io

You need to request the access token with resource=https://.azconfig.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-touse-azure-app-configuration-rest-api

QUESTION 20

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.

Hot Area:

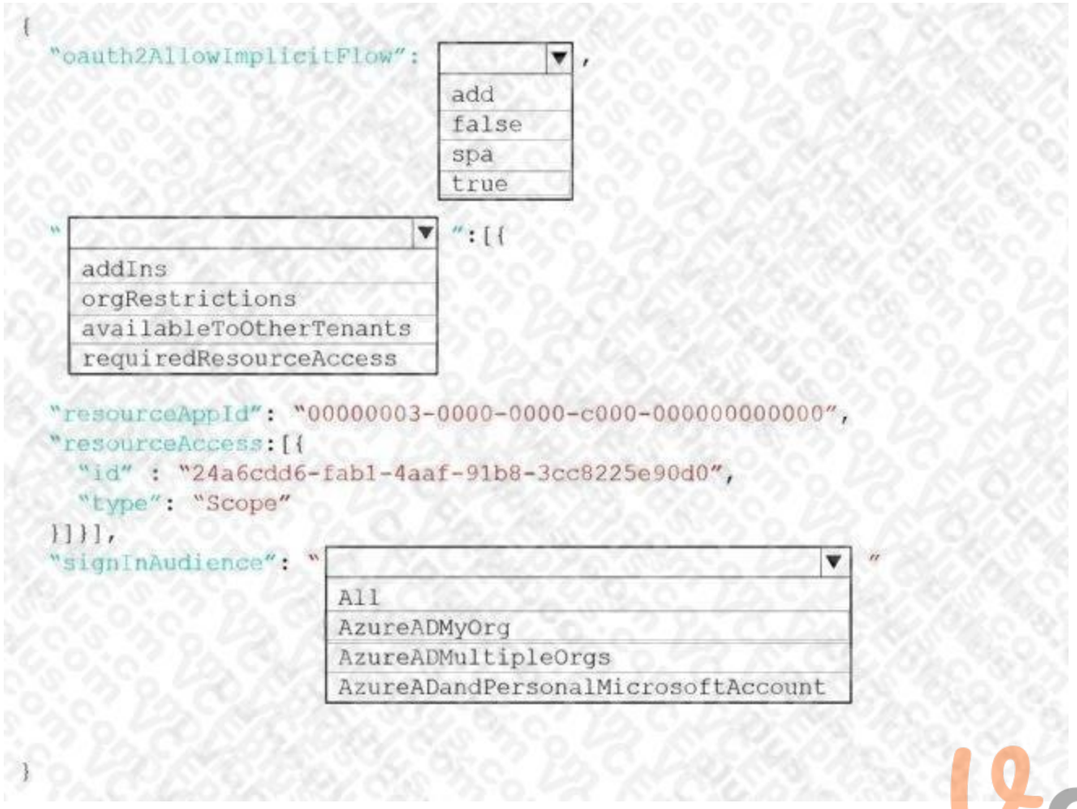

Answer Area:

Section:
Explanation:
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 singlepage 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"
}
]
}],
Incorrect Answers:
The legacy attribute availableToOtherTenants is no longer supported.
The addIns attribute defines custom behavior that a consuming service can use to call an app in specific contexts. For example, applications that can render file streams may set the addIns property for its "FileHandler" functionality. This parameter will let services like Microsoft 365 call the application in the context of a document the user is working on.

Example:
"addIns": [
{
"id": "968A844F-7A47-430C-9163-07AE7C31D407",
"type":" FileHandler",
"properties": [
{
"key": "version",
"value": "2"
}
]
}],
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/enus/azure/active-directory/develop/v2-oauth2-implicit-grant-flow

QUESTION 21
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 Siteto-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.

Hot Area:

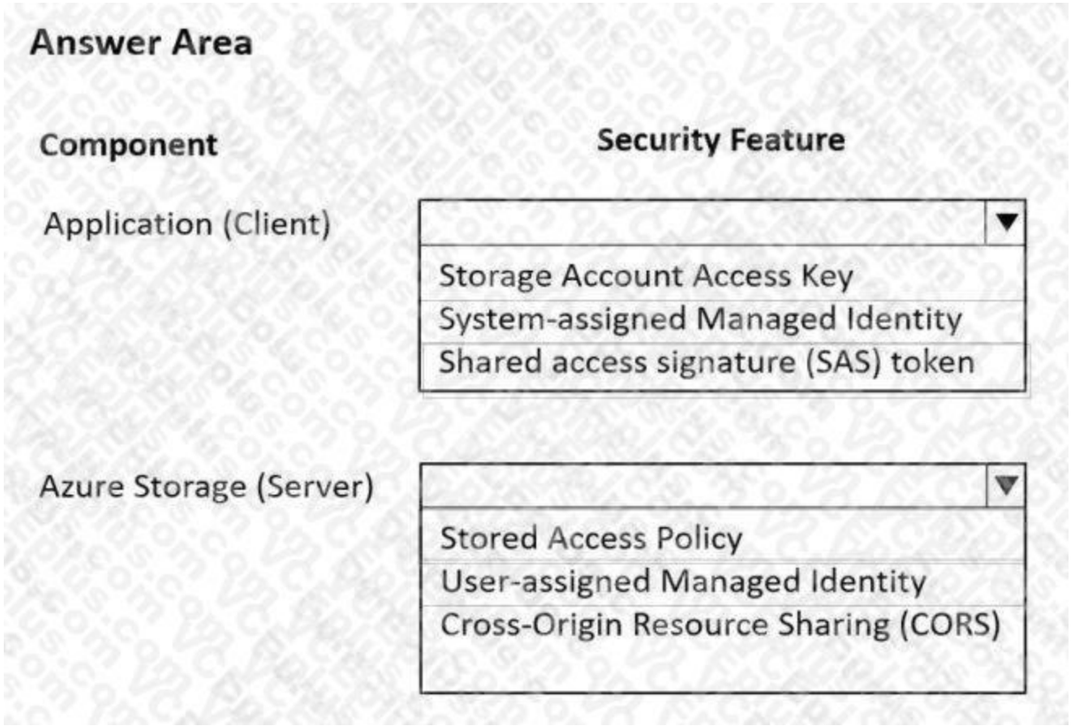

Answer Area:

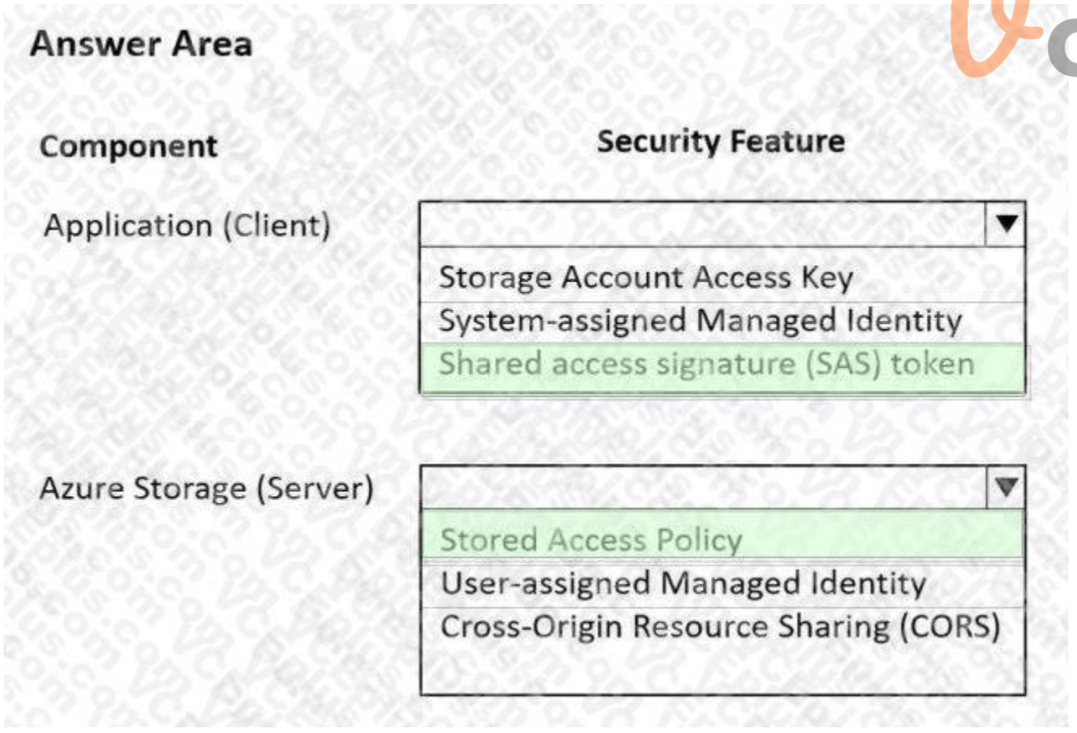

Section:
Explanation:

QUESTION 22
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 SignaIR trigger.
C. Ensure that the function has a retry configured.
D. Modify the trigger to use Consumption mode instead of the App Service plan.

Correct Answer: A, C
Section:

QUESTION 23
HOTSPOT
You develop an application that sells Al generated images based on user input. You recently started a marketing campaign that displays unique ads every second day.
Sales data is stored in Azure Cosmos DB with the date of each sale being stored in a property named 'whenFinished'.
The marketing department requires a view that shows the number of sales for each unique ad.
You need to implement the query for the view.
How should you complete the query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

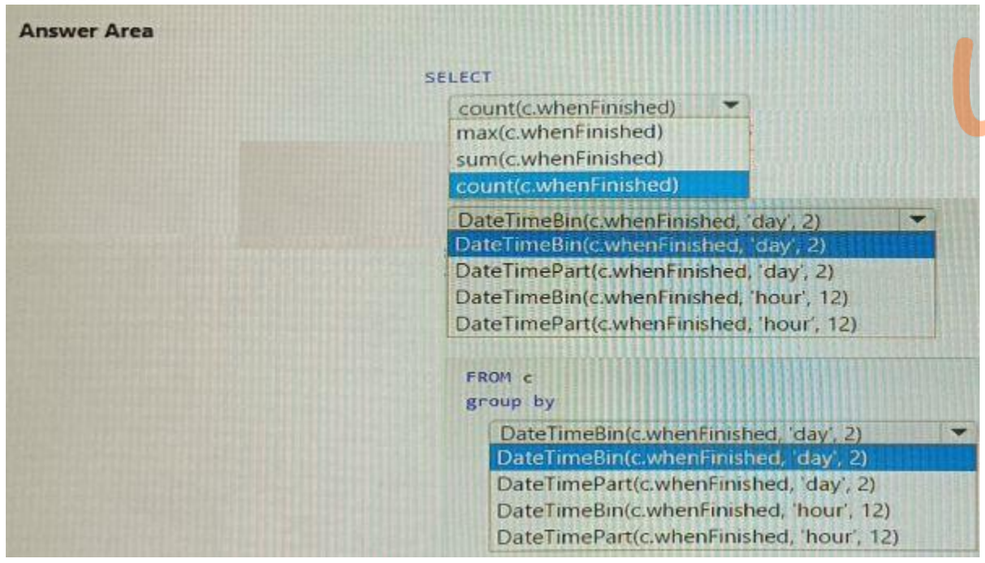

Answer Area:

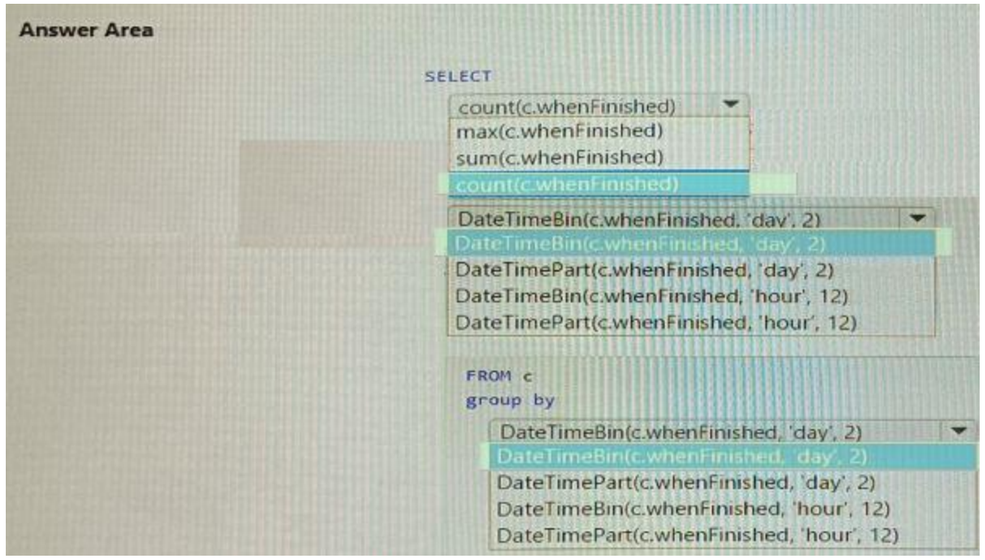

Section:
Explanation:

QUESTION 24
HOTSPOT
You plan to implement an Azure Functions app.
The Azure Functions app has the following requirements:
* Must be triggered by a message placed in an Azure Storage queue.
* Must use the queue name set by an app setting named input-queue.
* Must create an Azure Blob Storage named the same as the content of the message.
You need to identify how to reference the queue and blob name in the function. Just file of the Azure Functions app.
How should you reference the names? To answer, select the appropriate values in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

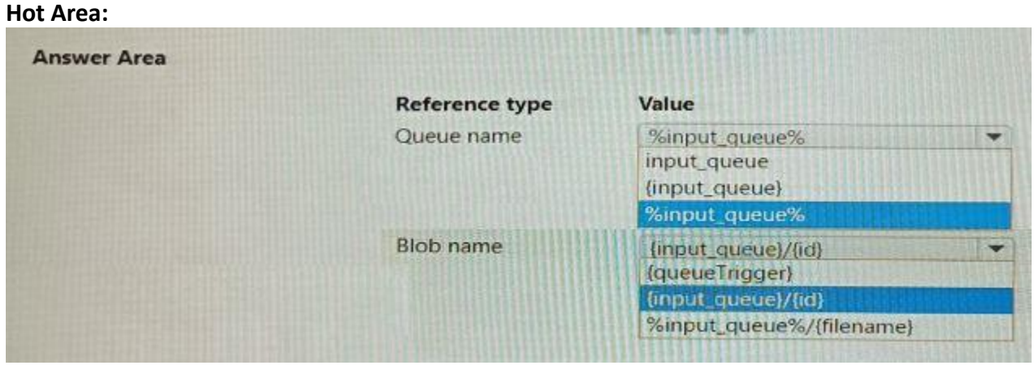

Answer Area:

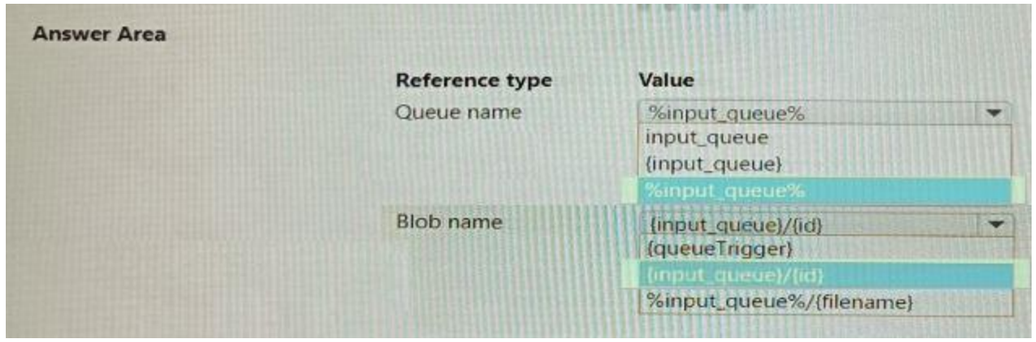

Section:
Explanation:

QUESTION 25
DRAG DROP
You have an Azure Cosmos DB for NoSQL account.
You plan to develop two apps named App1 and App2 that will use the change feed functionality to track changes to containers.
App1 will use the pull model and App2 will use the push model.
You need to choose the method to track the most recently processed change in App1 and App2.
Which component should you use? To answer, drag the appropriate components to the correct apps. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Select and Place:

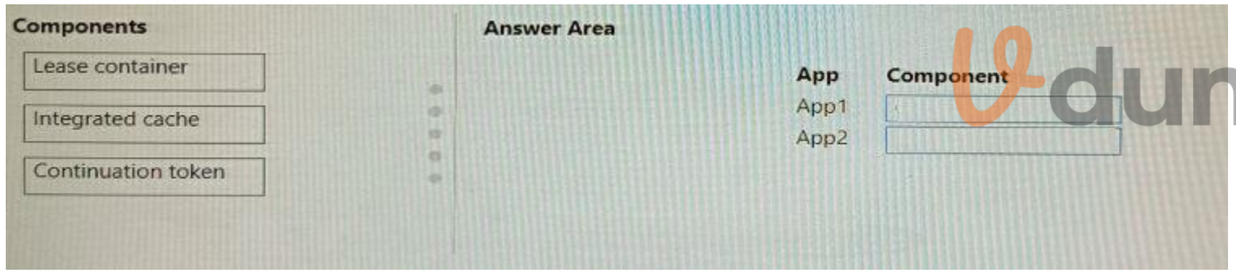

Correct Answer:

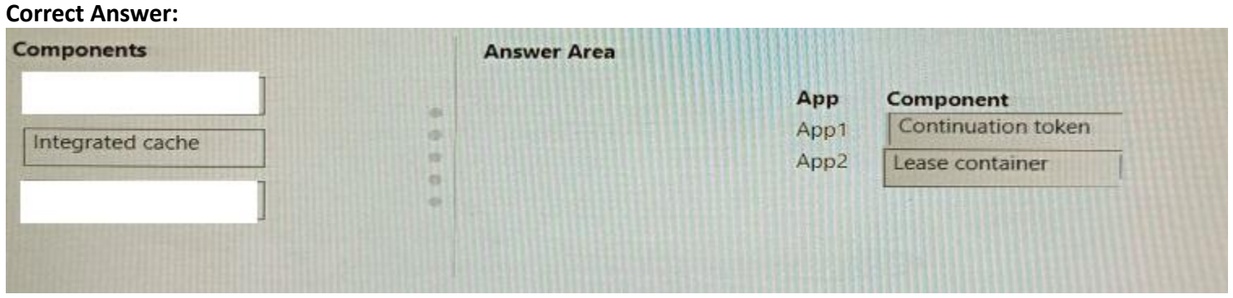

Section:
Explanation:

QUESTION 26
You are a developing a SaaS application that stores data as key value pairs.
You must make multiple editions of the application available. In the lowest cost edition, the performance must be best-effort, and there is no regional failover.
In higher cos! editions customers must be able to select guaranteed performance and support for multiple regions. Azure costs must be minimized.
Which Azure Cosmos OB API should you use for the application?

A. Core
B. MongoDB
C. Cassandra
D. Table API

Correct Answer: D
Section:

QUESTION 27
You are developing an application to store information about the organizational structure for a company.
Users must be able to determine which people report to a particular manager, the office where employees work, and the projects that are assigned to an employee.
Which Azure Cosmos DB API should you use for the application?

A. Core
B. Cassandra
C. Table API
D. Gremlin
E. MongoDB

Correct Answer: E
Section:

QUESTION 28
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.

Correct Answer: C, D
Section:

QUESTION 29
HOTSPOT
You need to implement the corporate website.
How should you configure the solution?

Hot Area:

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

Answer Area:

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

Section:
Explanation:

QUESTION 30
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

Correct Answer: A, B
Section:

QUESTION 31
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.

Correct Answer: D
Section:

QUESTION 32
You are developing several microservices to deploy to a Azure Service cluster. The microservices manage data stored in Azure Cosmos DB and Azure Blob storage. The data is secured by using customer-managed keys stored in Aue Key Vault.
You must automate key rotation for all Key Vault keys and allow for manual key rotation. Keys must rotate every three months. Notifications Of expiring keys must be sent before key expiry.
You need to configure key rotation and enable key expiry notifications.
Which two actions should you perform? Each correct answer presents part Of solution.
NOTE: Each correct selection is worth

A. Create and configure a new Azure Event Grid instance.
B. Create configure a key rotation policy during key creation
C. Create and assign an Azure Key Vault access
D. Configure Azure Key Vault

Correct Answer: B, D
Section:
Explanation:
https://learn.microsoft.com/en-us/azure/key-vault/keys/how-to-configure-key-rotation

QUESTION 33
You develop Azure Web Apps for a commercial diving company. Regulations require that all divers fill out a health questionnaire every 15 days after each diving job starts.
You need to configure the Azure Web Apps so that the instance count scales up when divers are filling out the questionnaire and scales down after they are complete.
You need to configure autoscaling.
What are two possible autoscaling configurations to achieve this goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Predictive autoscaling
B. CPU usage-based autoscaling
C. Recurrence profile
D. Fixed date profile

Correct Answer: A, D
Section:

QUESTION 34
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 name
B. Application secret
C. Application ID

D. Supported account type
E. Redirect URI/URL

Correct Answer: A, B, C
Section:

QUESTION 35
You are developing an Azure App Service web app.
The web app must securely store session information in Azure Redis Cache.
You need to connect the web app to Azure Redis Cache.
Which three Azure Redis Cache properties should you use? Each correct answer presents part of the solution.
Each correct selection is worth one point.

A. SSL port
B. Subscription name
C. Location
D. Host name
E. Access key
F. Subscription id

Correct Answer: A, C, D
Section:
Explanation:
https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-howto

QUESTION 36
You are creating an Azure key vault using PowerShell. Objects deleted from the key vault must be kept for a set period of 90 days.
Which two of the following parameters must be used in conjunction to meet the requirement?
(Choose two.)

A. EnabledForDeployment
B. EnablePurgeProtection
C. EnabledForTemplateDeployment
D. EnableSoftDelete

Correct Answer: B, D
Section:

QUESTION 37
A company maintains multiple web and mobile applications. Each application uses custom in-house identity providers as well as social identity providers.
You need to implement single sign-on (SSO) for all the applications.
What should you do?

A. Use Azure Active Directory B2C (Azure AD B2C) with custom policies. Most Voted
B. Use Azure Active Directory B2B (Azure AD B2B) and enable external collaboration.
C. Use Azure Active Directory B2C (Azure AD B2C) with user flows.
D. Use Azure Active Directory B2B (Azure AD B2B).

Correct Answer: A
Section:
Explanation:
https://docs.microsoft.com/en-us/azure/active-directory-b2c/custom-policy-reference-sso

QUESTION 38
HOTSPOT
You are developing a solution that uses several Azure Service Bus queues. You create an Azure Event Grid subscription for the Azure Service Bus namespace. You use Azure Functions as subscribers to process the messages.
You need to emit events to Azure Event Grid from the queues. You must use principal of least privilege and minimize costs.
Which Azure Service Bus values should you use? TO answer, select the appropriate options in the answer area
Each correct selection is worth ore point

Hot Area:

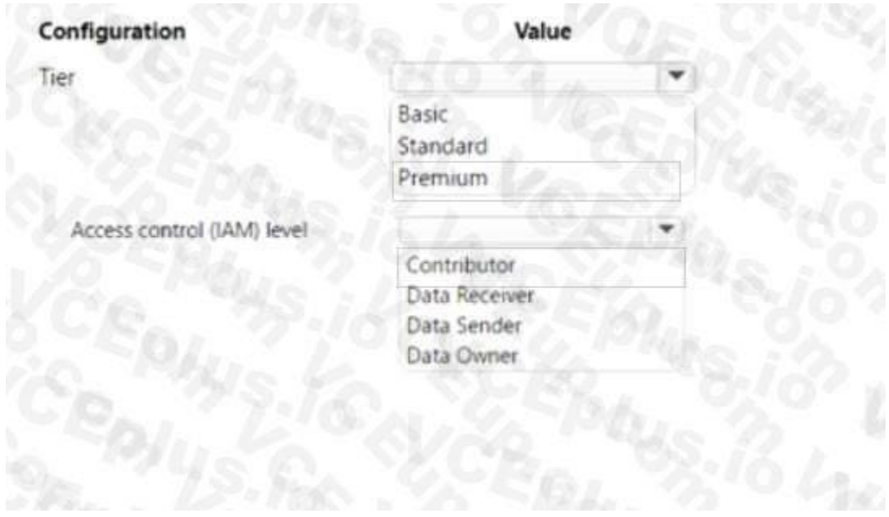

Answer Area:

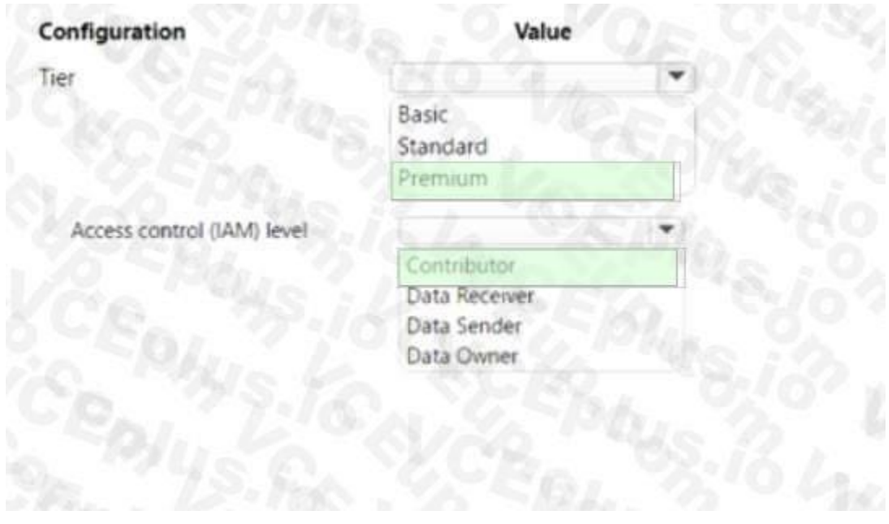

Section:
Explanation:

QUESTION 39
HOTSPOT
You are developing an application that uses Azure Storage to store customer data. The data must only be decrypted by the customer and the customer must be provided a script to rotate keys.
You need to provide a script to rotate keys to the customer.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

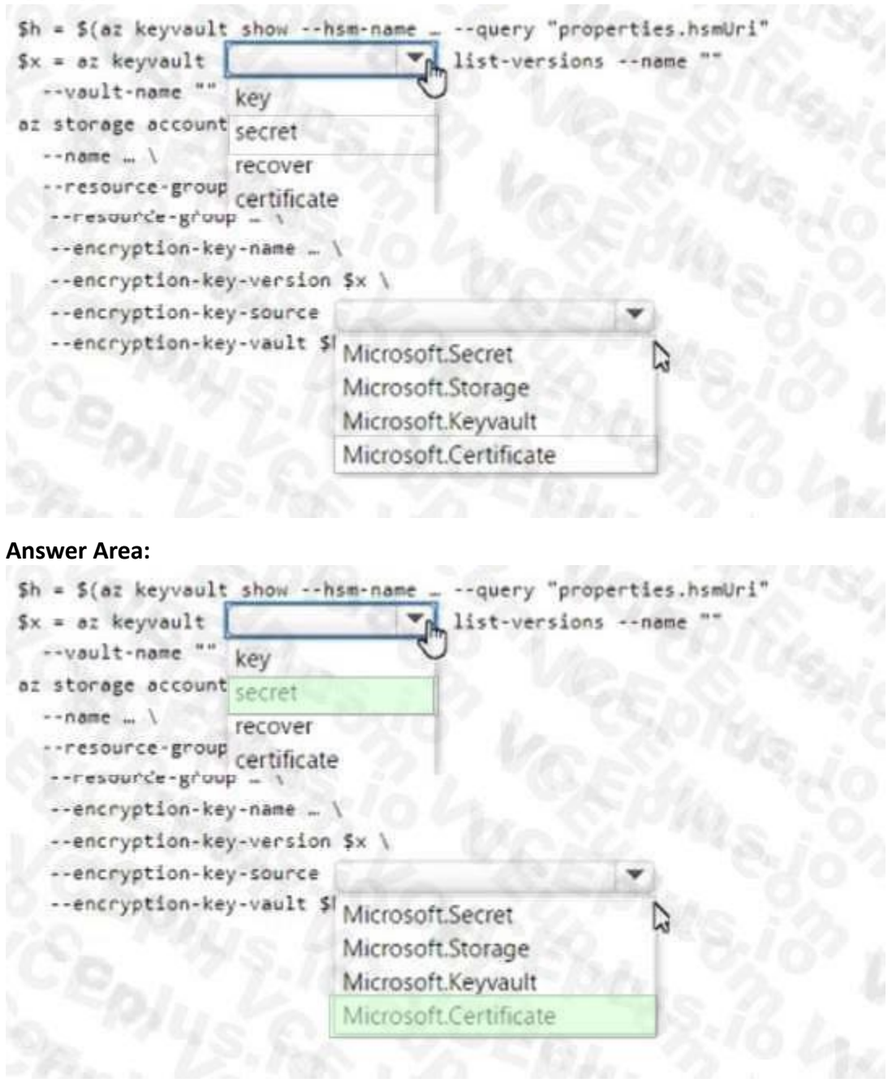

Answer Area:

Section:
Explanation:

QUESTION 40
HOTSPOT
You are developing an application that monitors data added to an Azure Blob storage account.
You need to process each change made to the storage account.

How should you complete the code segment? TO answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

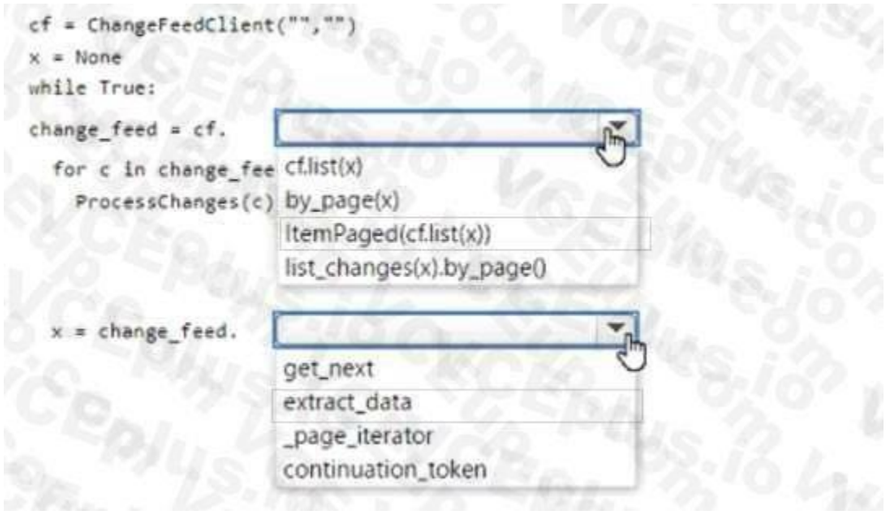

Answer Area:

Section:
Explanation:

QUESTION 41
HOTSPOT
You are developing a solution by using the Azure Event Hubs SDK. You create a standard Azure Event Hub with 16 partitions. You implement eight event processor clients.
You must balance the load dynamically when an event processor client fails. When an event processor client fails, another event processor must continue processing from the exact point at which the failure occurred. All events must be aggregate and upload to an Azure Blob storage account
You need to implement event processing recovery for the solution.
Which SDK features should you use? To answer, select the appropriate options in the answer are a.
Each correct selection is worth one point.

Hot Area:

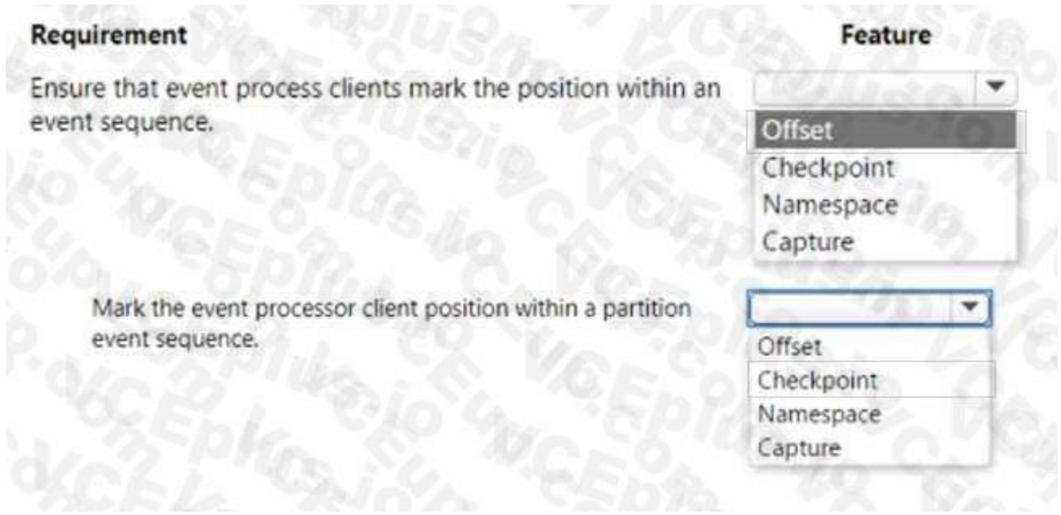

Answer Area:

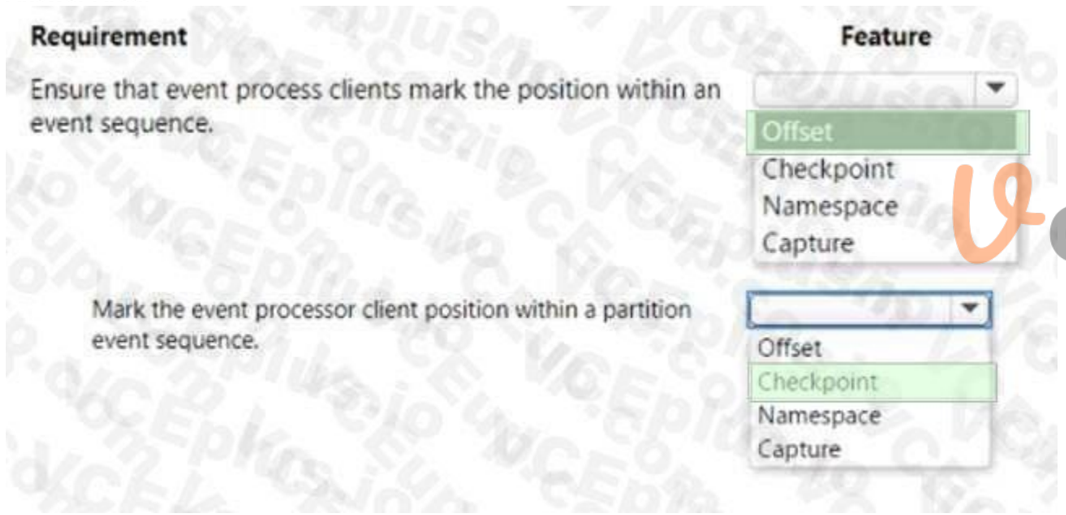

Section:
Explanation:

QUESTION 42
HOTSPOT
You are developing a C+ + application that compiles to a native application named process.exe. The application accepts images as input and returns images in one of the following image formats: GIF, PNG, or JPEG.
You must deploy the application as an Azure Function. You need to configure the function and host json files.
How should you complete the json files? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

function.json

```
{
  [ ▼ ]
    "type": "http"
    "platform": "gcm"
    "datatype": "stream"
    "path": "process.exe"

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

host.json

```
[ ▼ ]
    "customHandler": { "description": {
    "languageWorker": { "path": {
    "extensions": {"worker": {
    "extensionBundle": {
      "defaultExecutablePath": "process.exe"
},
[ ▼ ]
    "enableForwardingHttpRequest": true
    "enableForwardingHttpRequest": false
}
```

Answer Area:

```
function.json
{
  [ ▼ ]
    "direction": "out",
    "name" : "result"
}
host.json
  [ ▼ ]
    "defaultExecutablePath": "process.exe"
},
  [ ▼ ]
}
```

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

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

- "enableForwardingHttpRequest": true
- "enableForwardingHttpRequest": false

Section:
Explanation:

QUESTION 43
DRAG DROP
You develop and deploy a Java application to Azure. The application has been instrumented by using the Application Insights SDK.
The telemetry data must be enriched and processed before it is sent to the Application Insights service.
You need to modify the telemetry data.
Which Application Insights SDK features should you use? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

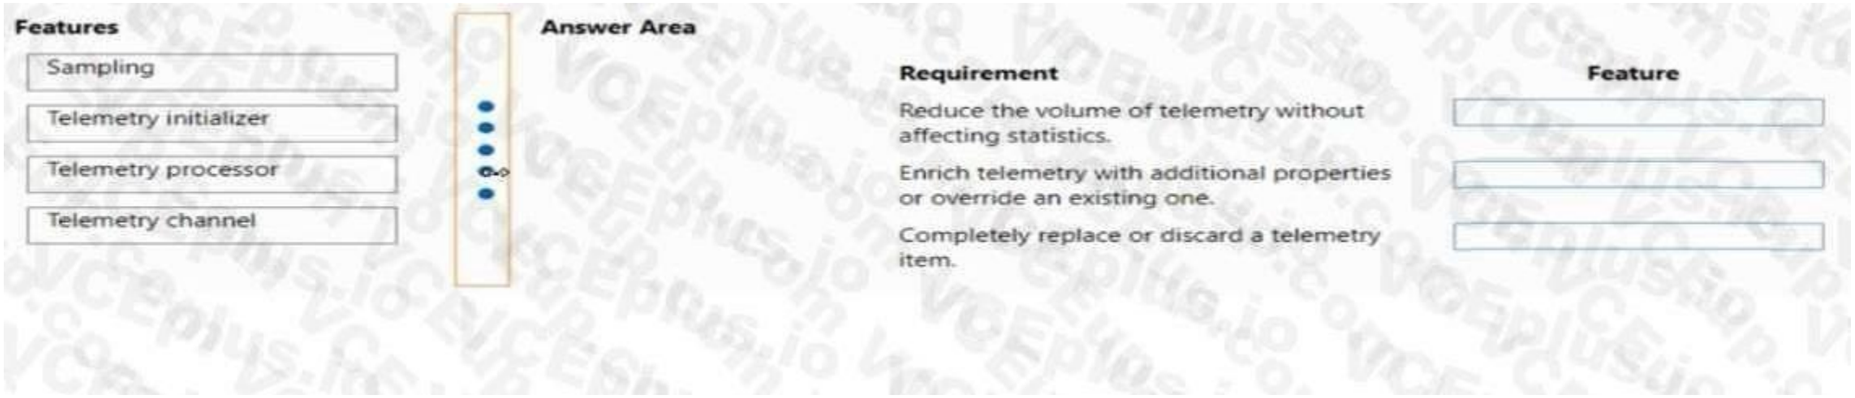

Correct Answer:

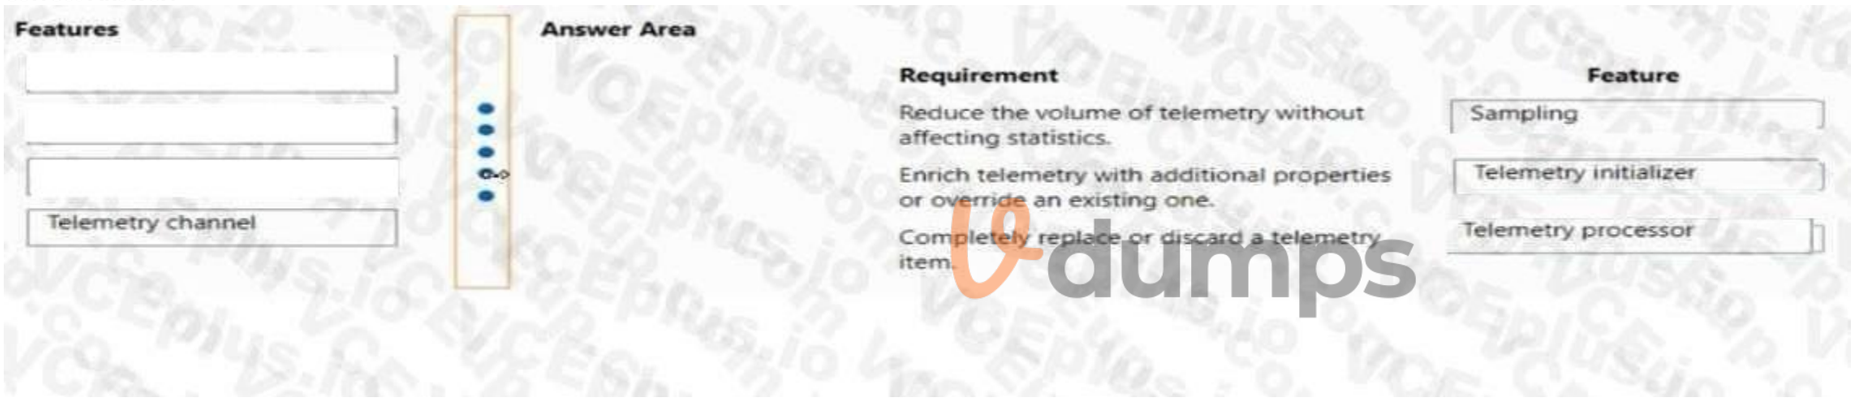

Section:
Explanation:

QUESTION 44
You are developing an Azure Function that calls external APIs by providing an access token for the API. The access token is stored in a secret named token in an Azure Key Vault named mykeyvault.
You need to ensure the Azure Function can access to the token. Which value should you store in the Azure Function App configuration?

A.
```
KeyVault:mykeyvault;Secret:token
```

B.
```
App:Settings:Secret:mykeyvault:token
```

C.
```
AZUREKVCONNSTR_ https://mykeyvault.vault.azure.net/secrets/token/
```

D.
```
@Microsoft.KeyVault(SecretUri=https://mykeyvault.vault.azure.net/secrets/token/)
```

Correct Answer: D
Section:

QUESTION 45
HOTSPOT
You are developing an Azure Function app.
The Azure Function app must enable a WebHook to read an image from Azure Blob Storage and create a new Azure Cosmos DB document.
You need to implement the Azure Function app.
Which configuration should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

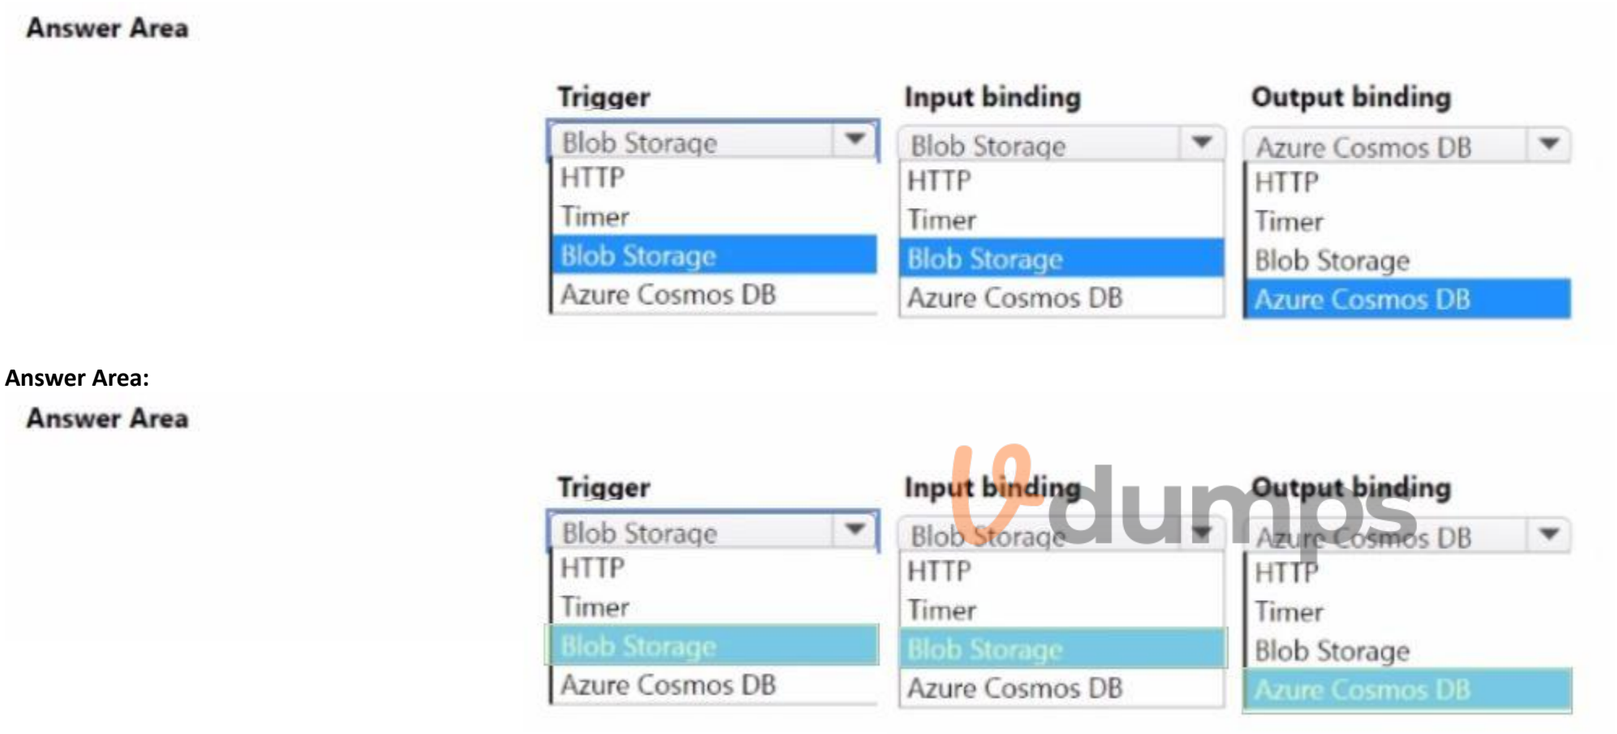

Answer Area:

Section:
Explanation:

QUESTION 46
DRAG DROP
You develop and deploy several APIs to Azure API Management.
You create the following policy fragment named APICounts:

```
<fragment>
    <emit-metric value="1" namespace="custom-metrics">
        <dimension name="User ID" />
        <dimension name="Operation ID" />
        <dimension name="API ID" />
        <dimension name="Client IP" value="@(context.Request.IpAddress)" />
    </emit-metric>
</fragment>
```

The policy fragment must be reused across various scopes and APIs. The policy fragment must be applied to all APIs and run when a calling system invokes any API.
You need to implement the policy fragment.
How should you complete the policy segment? To answer, drag the appropriate XML elements to the correct targets. Each XML element may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Select and Place:

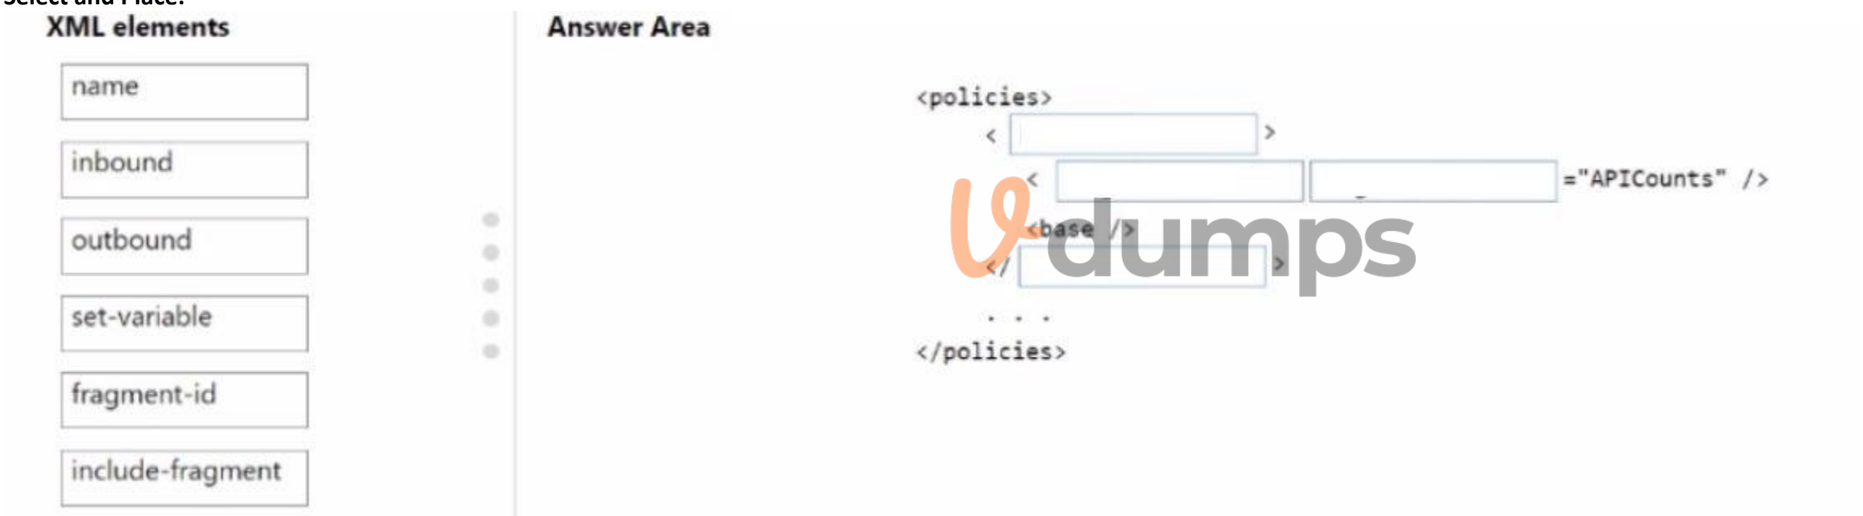

Correct Answer:

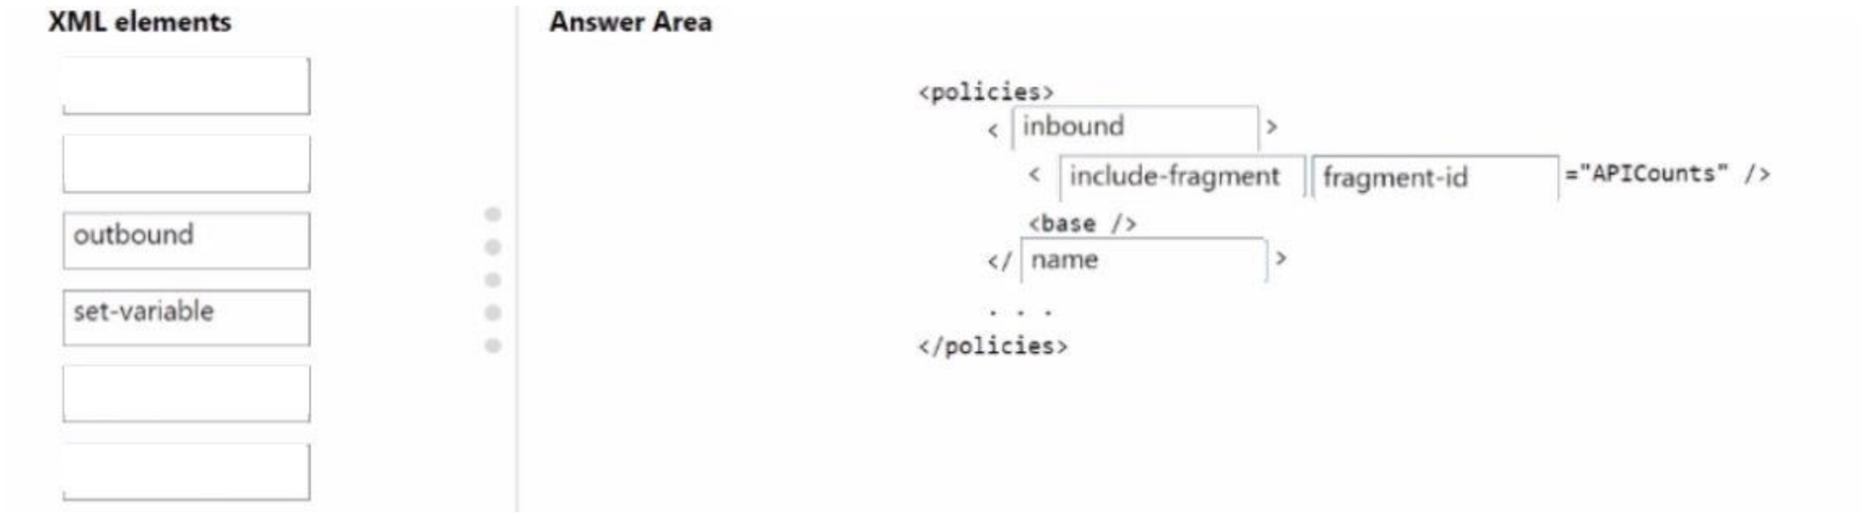

Section:
Explanation:

QUESTION 47
HOTSPOT
You develop an image upload service that is exposed using Azure API Management. Images are analyzed after upload for automatic tagging.
Images over 500 KB arc processed by a different backend that offers a lower tier of service that costs less money. The lower tier of service is denoted by a header named x-lsrSe-requ´st. Images over 500 KB must never be processed by backends for smaller images and must always be charged the lower price.
You need to implement API Management policies to ensure that images are processed correctly.
How should you complete the API Management inbound policy? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

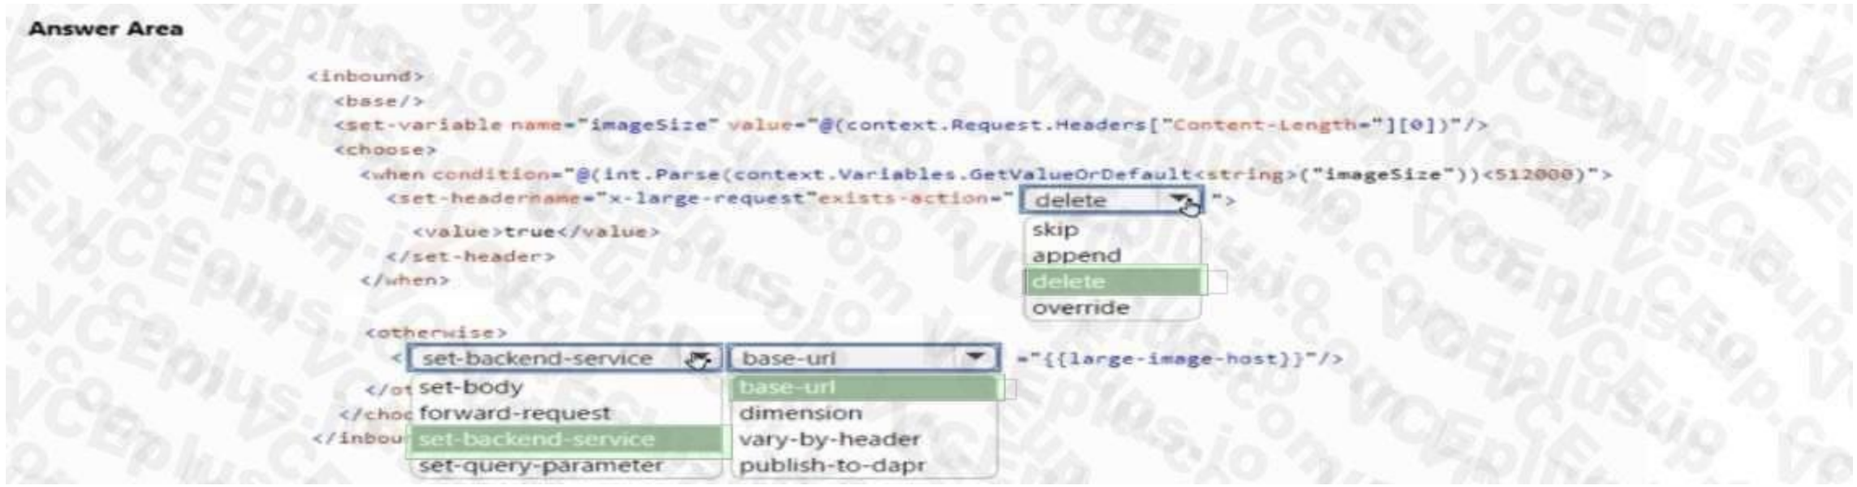

Section:
Explanation:

QUESTION 48
HOTSPOT
You develop several Azure Grid to include hundreds of event types, such as billing, inventory, and shipping updates.
Events must be sent to a single endpoint for the Azure Functions app to process. The events must be filtered by event type before processing. You must have authorization and authentication control to partition your tenants to receive the event data.
You need to configure Azure Event Grid.
Which configuration should you use? To answer, select the appropriate values in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

Section:
Explanation:

QUESTION 49
HOTSPOT
You are authoring a set of nested Azure Resource Manager templates to deploy Azure resources. You author an Azure Resource Manager template named mainTemplate.json that contains the following linked templates: linkedTemplatel.json, linkedTemplate2json.
You add parameters to a parameters template file named mainTemplate.parameters.json. You save all templates on a local device in the C:\templates\ folder.
You have the following requirements:
• Store the templates in Azure for later deployment.
• Enable versioning of the templates.
• Manage access to the templates by using Azure RBAC
You need to store the templates in Azure.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

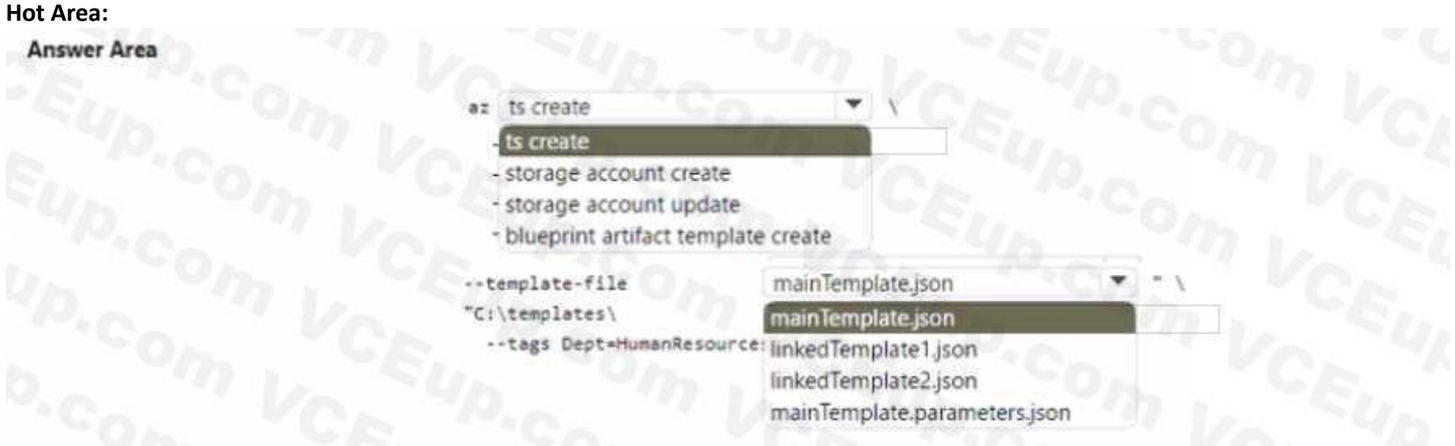

Answer Area:

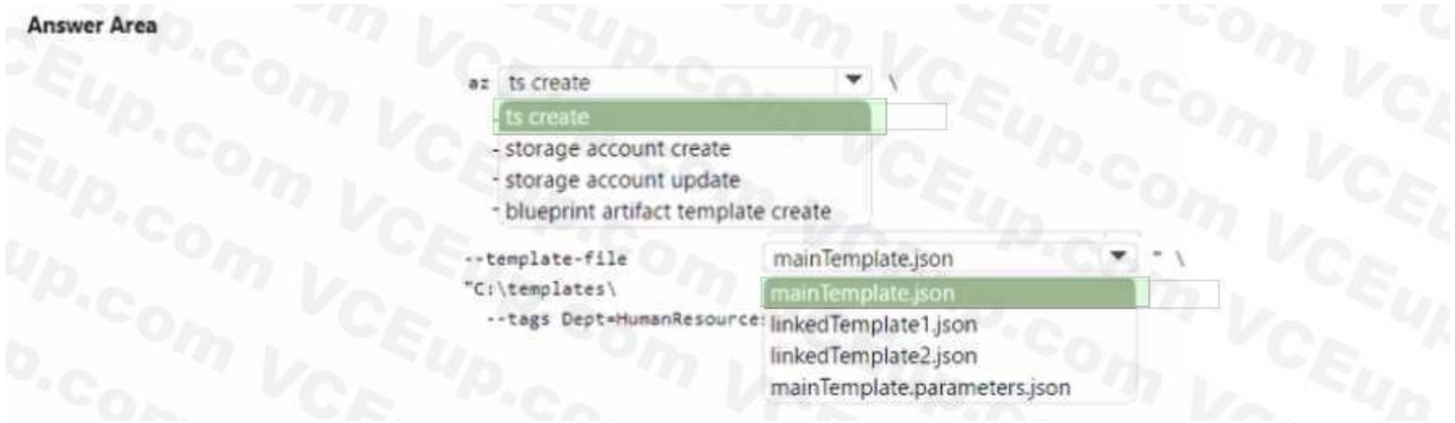

Section:
Explanation:

QUESTION 50
HOTSPOT
You implement an Azure solution to include Azure Cosmos DB. the latest Azure Cosmos DB SDK, and the Azure Cosmos DB for NoSQL API. You also implement a change feed processor on a new container instance by using the Azure
Functions trigger for Azure Cosmos DB.
A large batch of documents continues to fail when reading one of the documents in the batch. The same batch of documents is continuously retried by the triggered function and a new batch of documents must be read.
You need to implement the change feed processor to read the documents.
Which feature should you implement? To answer, select the appropriate features in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

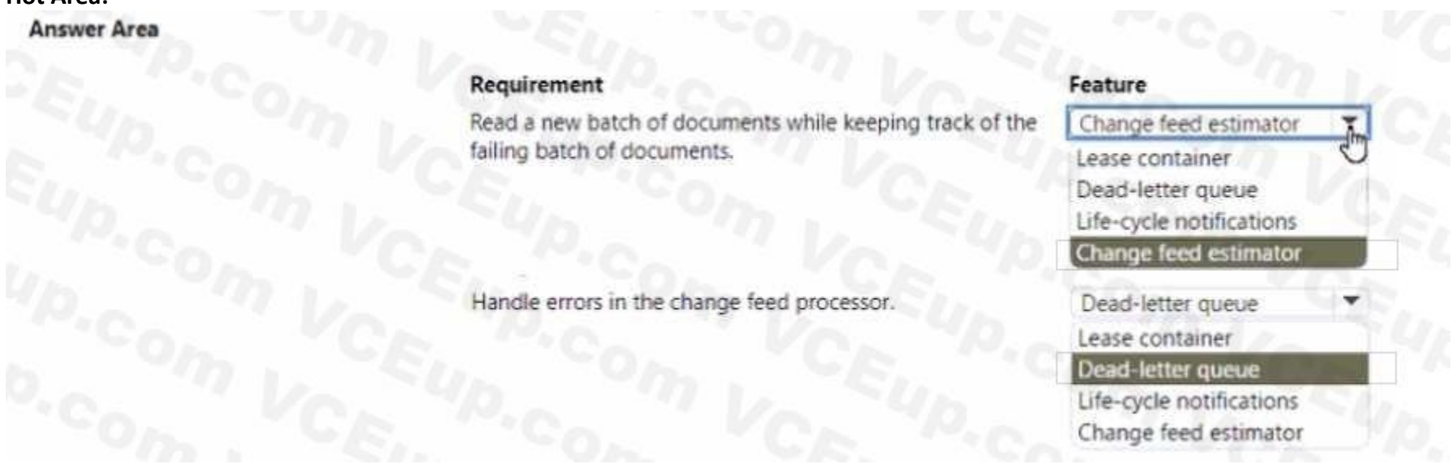

Answer Area:

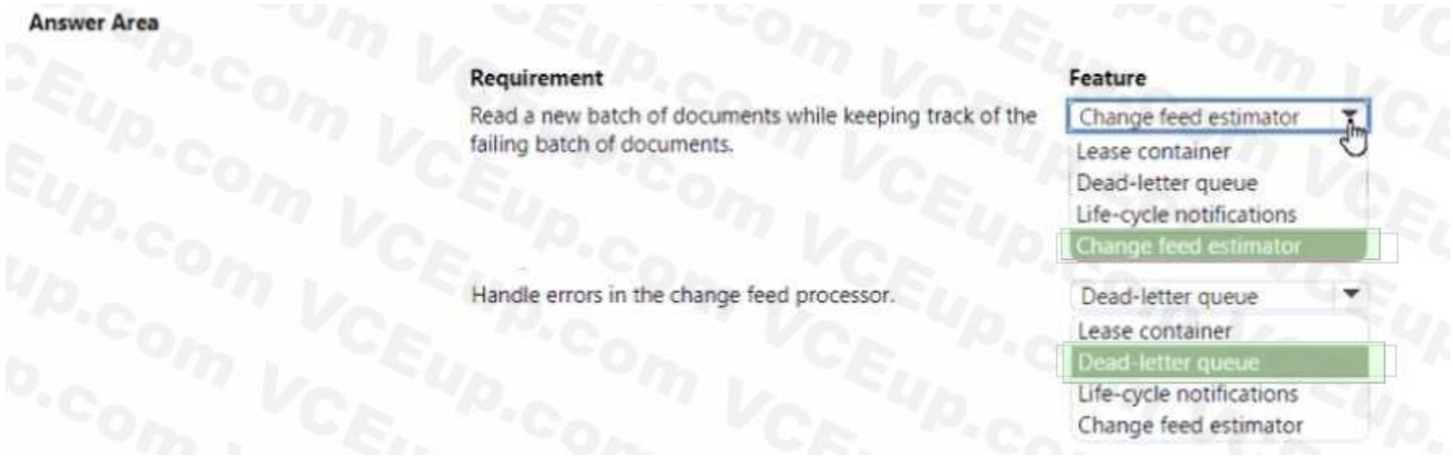

Section:
Explanation:

QUESTION 51
HOTSPOT
You are developing a content management application for technical manuals. The application is deployed as an Azure Static Web app.
Authenticated users can view pages under /manuals but only contributors can access the page /manuals/new html.
You need to configure the routing for the web app.
How should you complete the configuration? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

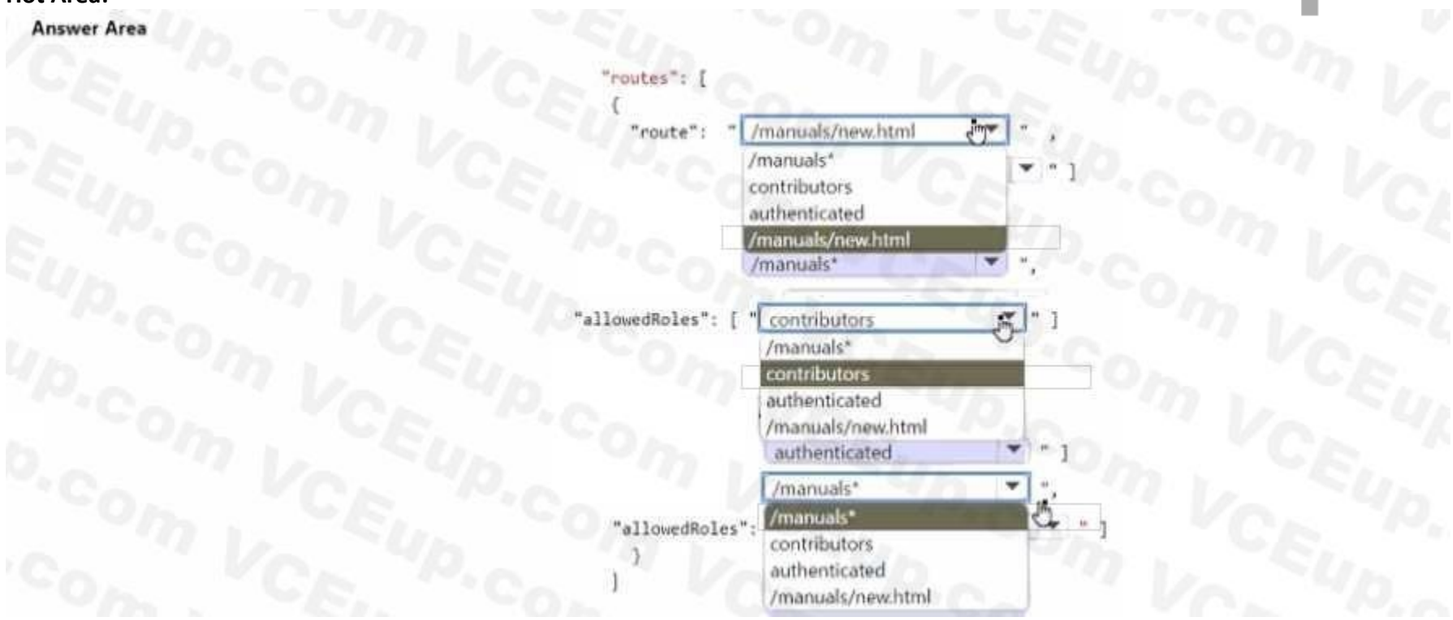

Answer Area:

Section:
Explanation:

QUESTION 52
The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location.
You need to complete the configuration.
Which Azure CLI or PowerShell command should you run?

A.
```
New-AzureRmResourceGroup
  -Name fridge-rg
  -Location fridge-loc
```

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

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

D.
```
New-AzureRmServiceBusNamespace
  -ResourceGroupName fridge-rg
  -NamespaceName fridge-ns
  -Location fridge-loc
```

Correct Answer: C
Section:

QUESTION 53
You are updating an application that stores data on Azure and uses Azure Cosmos DB for storage. The application stores data in multiple documents associated with a single username.
The application requires the ability to update multiple documents for a username in a single ACID operation.
You need to configure Azure Cosmos DB.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Configure Azure Cosmos DB to use the Azure Cosmos DB for Apache Gremlin API.
B. Configure Azure Cosmos DB to use the Azure Cosmos DB for MongoDB API.
C. Create a collection sharded on username to store documents.
D. Create an unsharded collection to store documents.

Correct Answer: B, D
Section:

QUESTION 54
HOTSPOT
An organization deploys a blob storage account. Users take multiple snapshots of the blob storage account over time. You need to delete all snapshots of the blob storage account. You must not delete the blob storage account itself. How should you complete the code segment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

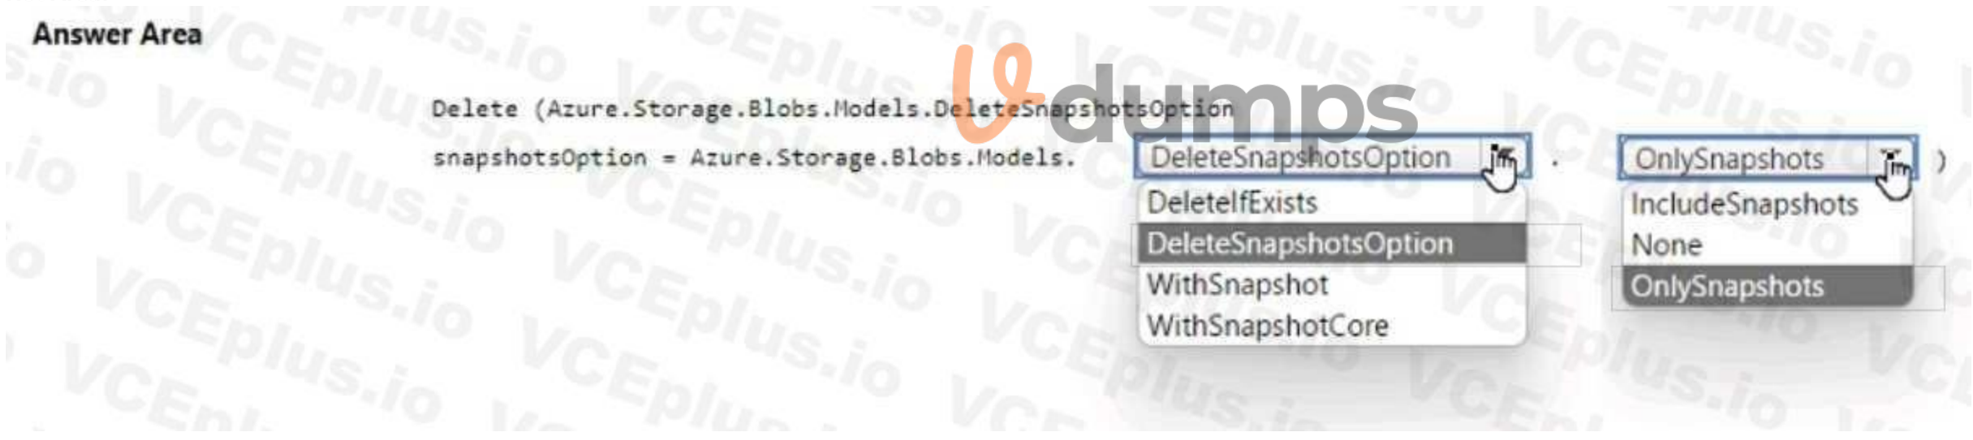

Answer Area:

Answer Area

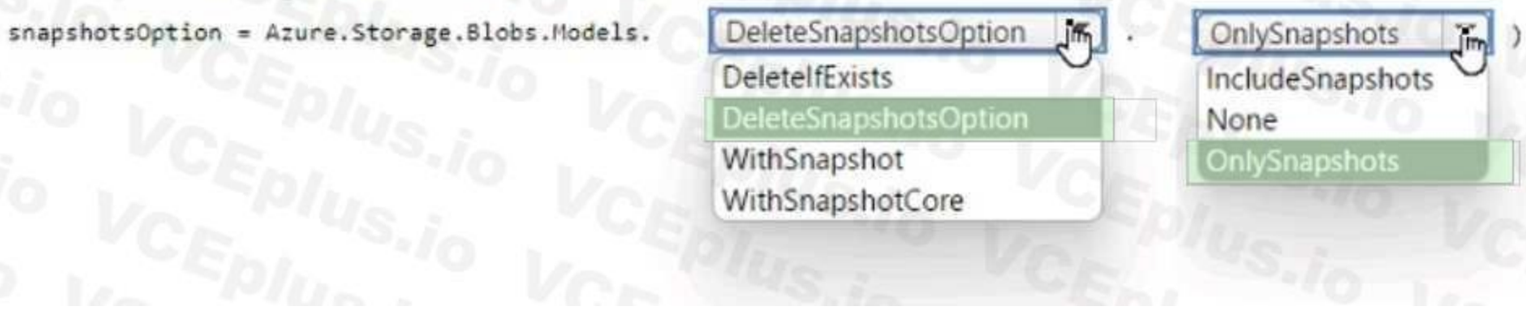

Section:
Explanation:

QUESTION 55
You are developing a Java application to be deployed in Azure. The application stores sensitive data in Azure Cosmos DB. You need to configure Always Encrypted to encrypt the sensitive data inside the application. What should you do first?

A. Create a customer-managed key (CMK) and store the key in a new Azure Key Vault instance.
B. Create an Azure AD managed identity and assign the identity to a new Azure Key Vault instance.
C. Create a data encryption key (DEK) by using the Azure Cosmos DB SDK and store the key in Azure Cosmos DB.
D. Create a new container to include an encryption policy with the JSON properties to be encrypted.

Correct Answer: A
Section:

QUESTION 56
An organization deploys Azure Cosmos DB.
You need to ensure that the index is updated as items are created, updated, or deleted.
What should you do?

A. Set the value of the EnableScanlnQuery option to True.
B. Set the indexing mode to Consistent.
C. Set the indexing mode to Lazy.
D. Set the value of the automatic property of the indexing policy to False.

Correct Answer: B
Section:

QUESTION 57
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 on the review screen.
You are implementing an application by using Azure Event Grid to push near-real-time information to customers.
You have the following requirements:
* You must send events to thousands of customers that include hundreds of various event types.

* The events must be filtered by event type before processing.
* Authentication and authorization must be handled by using Microsoft Entra ID.
* The events must be published to a single endpoint.
You need to implement Azure Event Grid.
Solution: Publish events to a partner topic. Create an event subscription for each customer.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: B
Section:

QUESTION 58
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 on the review screen.
You are implementing an application by using Azure Event Grid to push near-real-time information to customers.
You have the following requirements:
* You must send events to thousands of customers that include hundreds of various event types.
* The events must be filtered by event type before processing.
* Authentication and authorization must be handled by using Microsoft Entra ID.
* The events must be published to a single endpoint
You need to implement Azure Event Grid.
Solution: Publish events to a system topic. Create an event subscription for each customer.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: A
Section:

QUESTION 59
HOTSPOT
You have an App Service plan named aspl based on the Free pricing tier.
You plan to use aspl to implement an Azure Function app with a queue trigger. Your solution must minimize cost.
You need to identify the configuration options that will meet the requirements.
Which value should you configure? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

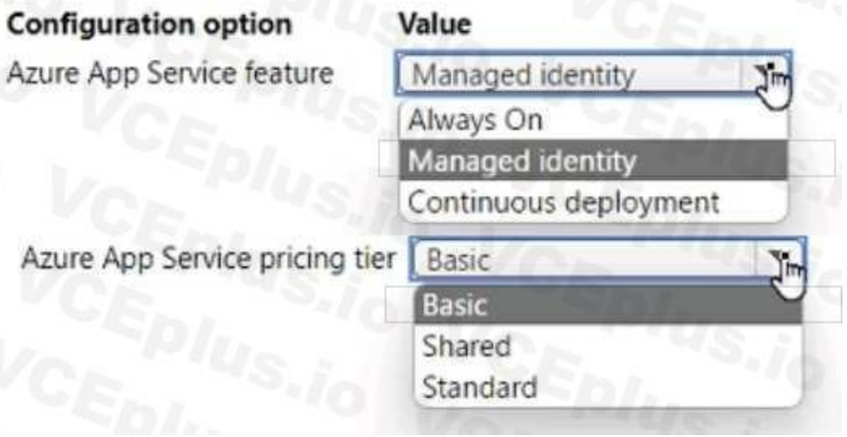

Answer Area:

Answer Area

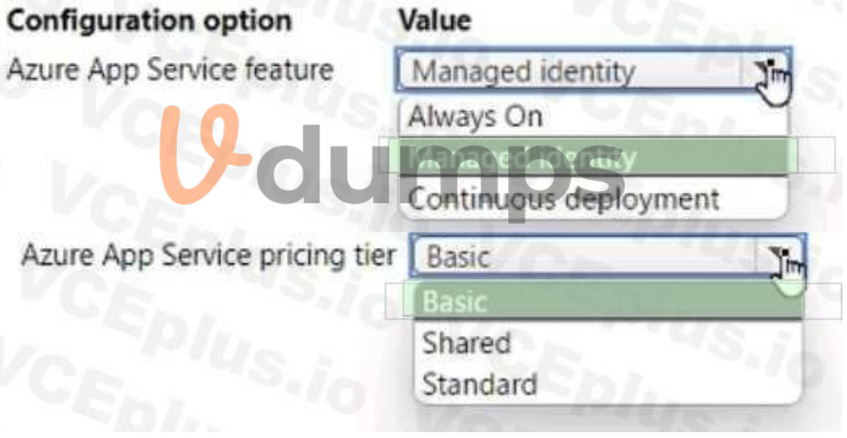

Section:
Explanation:

QUESTION 60
You are developing several Azure API Management (APIM) hosted APIs.
You must transform the APIs to hide private backend information and obscure the technology stack used to implement the backend processing.
You need to protect all APIs.
What should you do?

A. Configure and apply a new inbound policy scoped to a product.
B. Configure and apply a new outbound policy scoped to the operation.
C. Configure and apply a new outbound policy scoped to global.
D. Configure and apply a new backend policy scoped to global.

Correct Answer: A
Section:

QUESTION 61
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing a solution that will be deployed to an Azure Kubernetes Service (AKS) cluster. The solution will include a custom VNet, Azure Container Registry images, and an Azure Storage account.
The solution must allow dynamic creation and management of all Azure resources within the AKS cluster.
You need to configure an AKS cluster for use with the Azure APIs.
Solution: Enable the Azure Policy Add-on for Kubernetes to connect the Azure Policy service to the GateKeeper admission controller for the AKS cluster. Apply a built-in policy to the cluster.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
Instead create an AKS cluster that supports network policy. Create and apply a network to allow traffic only from within a defined namespace
https://docs.microsoft.com/en-us/azure/aks/use-network-policies

QUESTION 62
HOTSPOT
You provisioned an Azure Cosmos DB for NoSQL account named account1 with the default consistency level.
You plan to configure the consistency level on a per request basis The level needs to be set for consistent prefix for read and write operations to account1.
You need to identify the resulting consistency level for read and write operations.
Which levels should you configure? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Operation type	Resulting consistency level
Read operations	strong session consistent prefix
Write operations	strong session consistent prefix

Answer Area:

Answer Area

Operation type	Resulting consistency level
Read operations	strong session consistent prefix
Write operations	strong session consistent prefix

Section:

Explanation:

QUESTION 63
You are developing several Azure API Management (APIM) hosted APIs.
You must inspect request processing of the APIs in APIM. Requests to APIM by using a REST client must also be included. The request inspection must include the following information:
* requests APIM sent to the API backend and the response it received
* policies applied to the response before sending back to the caller
* errors that occurred during the processing of the request and the policies applied to the errors
* original request APIM received from the caller and the policies applied to the request
You need to inspect the APIs.
Which three actions should you do? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Enable the Allow tracing setting for the subscription used to inspect the API.
B. Add the Ocp-Apim-Trace header value to the API call with a value set to true
C. Add the Ocp-Apim-Subscription-Key header value to the key for a subscription that allows access to the API.
D. Create and configure a custom policy. Apply the policy to the outbound policy section with an API scope.
E. Create and configure a custom policy. Apply the policy to the inbound policy section with a global scope.

Correct Answer: A, B, C
Section:
Explanation:
The correct answer is A, B, and C. To inspect request processing of the APIs in APIM, you need to do the following three actions:
Enable the Allow tracing setting for the subscription used to inspect the API. This setting allows you to trace request processing in APIM using the test console, a REST client, or a client app. You can enable this setting in the portal by selecting Subscriptions and then selecting the subscription you want to use for debugging1.
Add the Ocp-Apim-Trace header value to the API call with a value set to true. This header triggers tracing when making requests to APIM using a REST client or a client app. You also need to add the Ocp-Apim-Subscription-Key header value to the key for a subscription that allows access to the API1.
Add the Ocp-Apim-Subscription-Key header value to the key for a subscription that allows access to the API. This header authenticates your request and grants you access to the API. You can find the key for your subscription in the portal by selecting Subscriptions and then selecting Show/hide keys1.
You do not need to create and configure a custom policy for tracing request processing. The trace policy is used to add a custom trace into the request tracing output, Application Insights telemetries, and/or resource logs2. It is not required for inspecting the APIs.

QUESTION 64
DRAG DROP
You are developing several microservices named service
* The microservices must persist data to storage.
* serviceA must persist data only visible to the current container and the storage must be restricted to the amount of disk space available in the container
* servtceB must persist data for the lifetime of the replica and allow multiple containers in the replica to mount the same storage location.
* serviceC must persist data beyond the lifetime of the replica while allowing multiple containers to access the storage and enable per object permissions.
You need to configure storage for each microservice.

Select and Place:

Correct Answer:

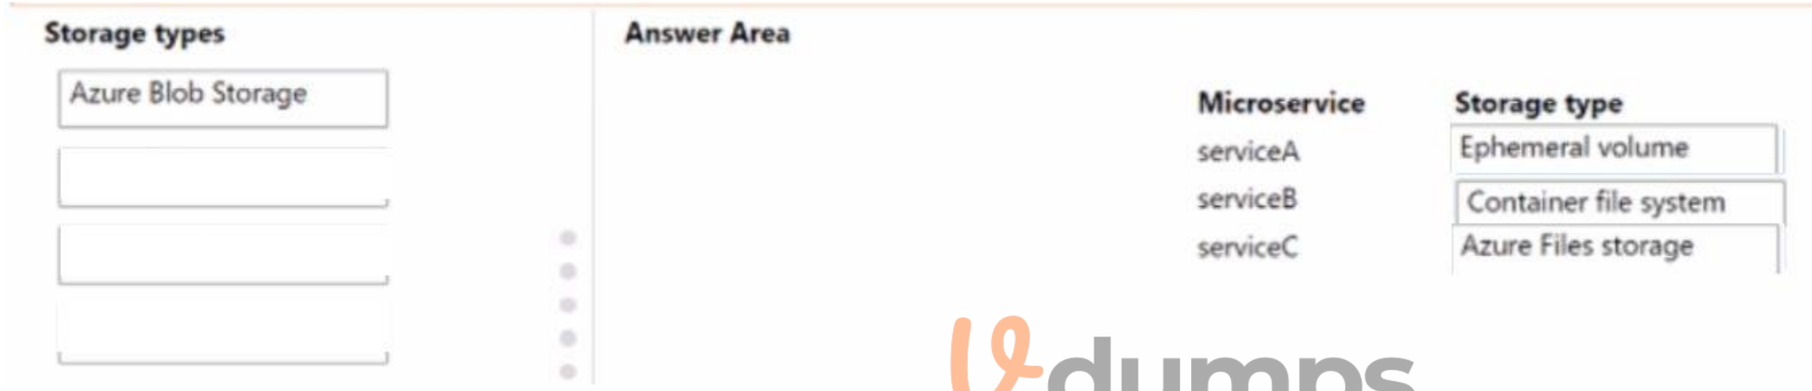

Section:
Explanation:

QUESTION 65
You are building a B2B web application that uses Azure B2B collaboration for authentication Paying customers authenticate to Azure B2B using federation
The application allows users to sign up for trial accounts using any email address
When a user converts to a paying customer, the data associated with the teal should be kept, but the user must authenticate using federation
You need to update the user in Azure Active Directory (Azure AD) when they convert to a paying customer
Which Graph API parameter is used to change authentication from one-time pass odes to federation?

A. uscrFlowType
B. Status
C. invittdUstr
D. resetRedemption

Correct Answer: B
Section:

QUESTION 66
HOTSPOT
You develop two Python scripts to process data.
The Python scripts must be deployed to two, separate Linux containers running in an Azure Container Instance container group. The containers must access external data by using the Server Message Block (SMB) protocol.
Containers in the container group must run only once
You need to configure the Azure Container Instance.

Which configuration value should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Configuration Setting	Configuration Value
External data volume	Secret Empty directory Cloned git repo Azure file share
Container restart policy	Never Always OnFailure

Answer Area:

Answer Area

Configuration Setting	Configuration Value
External data volume	Secret Empty directory Cloned git repo Azure file share
Container restart policy	Never Always OnFailure

Section:
Explanation:

QUESTION 67
You are developing a road tollway tracking application that sends tracking events by using Azure Event Hubs using premium tier.
Each road must have a throttling policy uniquely assigned.

You need to configure the event hub to allow for per-road throttling.
What should you do?

A. Ensure each road has a unique connection string.
B. Use a unique consumer group for each road
C. Use a unique application group for each road
D. Ensure each road stores events in a different partition.

Correct Answer: D
Section:

QUESTION 68
DRAG DROP
An organization has web apps hosted in Azure.
The organization wants to track events and telemetry data in the web apps by using Application Insights.
You need to configure the web apps for Application Insights.
Which three actions should you perform in sequence? To answer, move the appropnate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

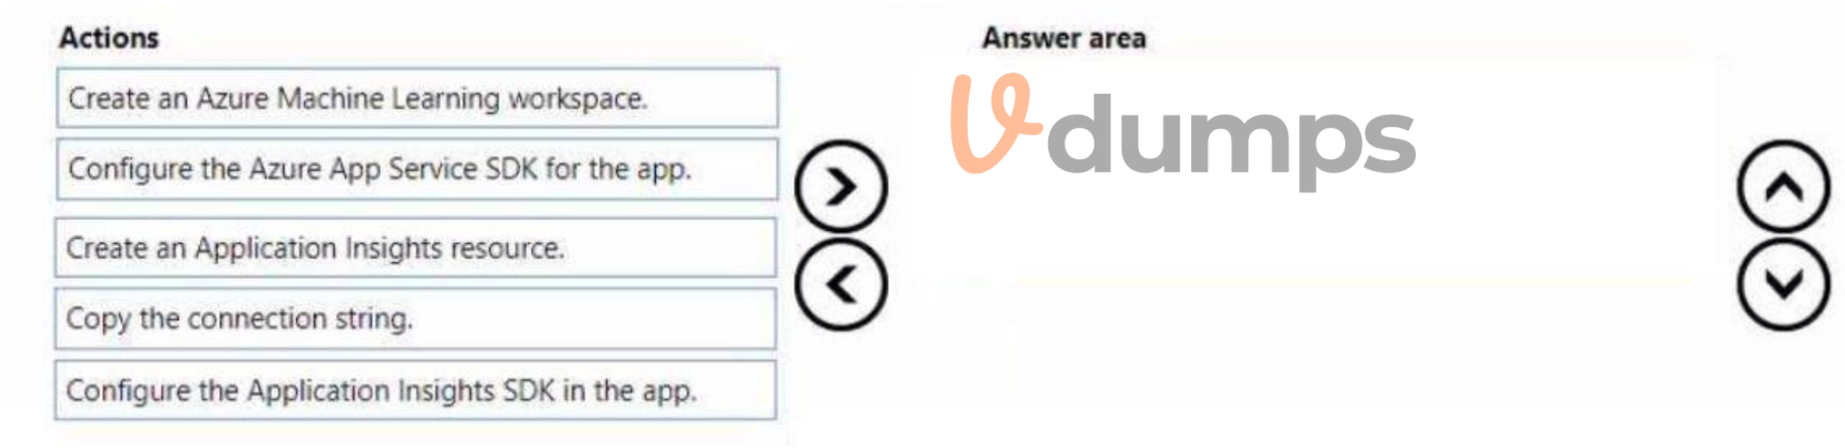

Correct Answer:

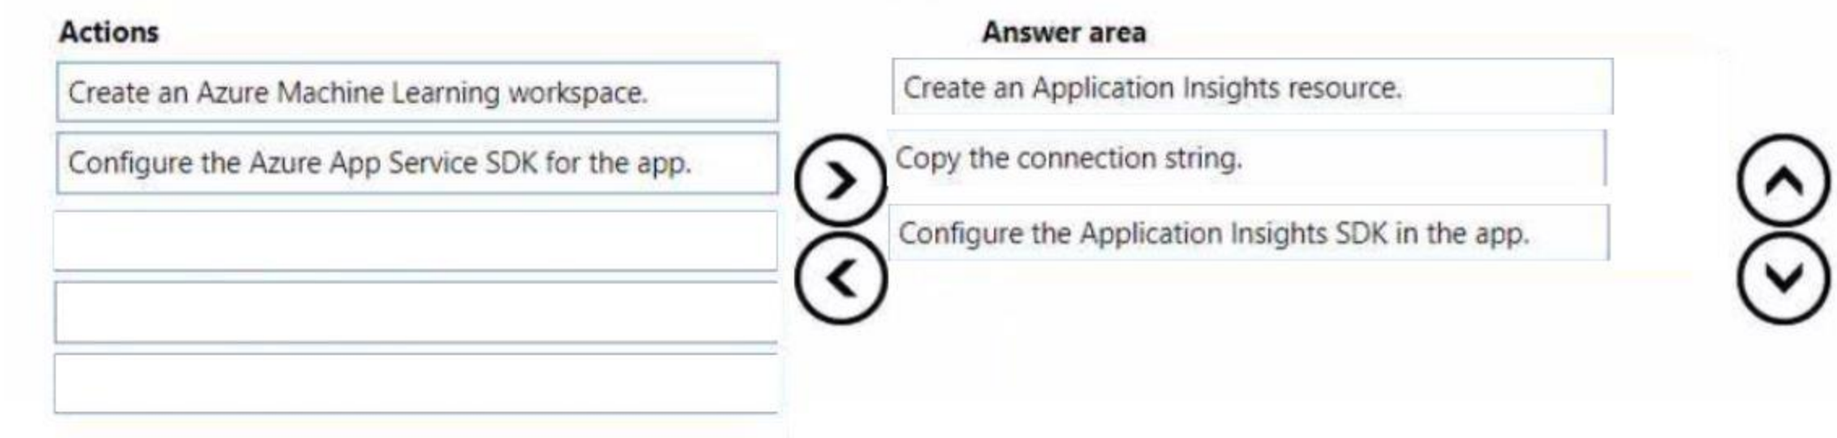

Section:
Explanation:
Create an Application Insights resource
Copy the instrumentation key
Install the SDK in your app

QUESTION 69
A company uses an Azure Blob Storage for archiving.
The company requires that data in the Blob Storage is only in the archive tier.
You need to ensure data copied to the Blob Storage is moved to the archive tier.
What should you do?

A. Use a Put Block List operation with a request header of x-ms-immutability-policy-mode.
B. Create a lifecycle policy with an action of tierToArchive and configure daysAftenModificationGreaterThan for 0.
C. Use a Put Blob operation with a request header of x-ms-immutability-policy-unrjl-date.
D. Create a lifecycle policy with an action of tierToArchive and configure a filter for bloblndexMatch.

Correct Answer: B
Section:

QUESTION 70
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 role Assign 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 OB keys from the function app.

Correct Answer: B, C, E
Section:

QUESTION 71
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.

Correct Answer: A, C
Section:

QUESTION 72
You develop an Azure App Service web app and deploy to a production environment. You enable Application Insights for the web app.
The web app is throwing multiple exceptions in the environment.
You need to examine the state of the source code and variables when the exceptions are thrown.
Which Application Insights feature should you configure?

A. Smart detection
B. Snapshot Debugger
C. Standard test
D. Profiler

Correct Answer: B
Section:

QUESTION 73
HOTSPOT
You have the following data lifecycle management policy:

```
{
    "rules": [
        {
            "enabled": true,
            "name": "Policy1",
            "type": "Lifecycle",
            "definition": {
                "actions": {
                    "baseBlob": {
                        "tierToArchive": {
                            "daysAfterModificationGreaterThan": 0
                        }
                    }
                },
                "filters": {
                    "blobIndexMatch": [
                        {
                            "name": "Customer",
                            "op": "==",
                            "value": "Adatum"
                        }
                    ]
                }
            }
        }
    ]
}
```

You plan to implement an Azure Blob Storage account and apply to it Policy 1. The solution should maximize resiliency and performance.
You need to configure the account to support the policy.
Which redundancy option and storage account type should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

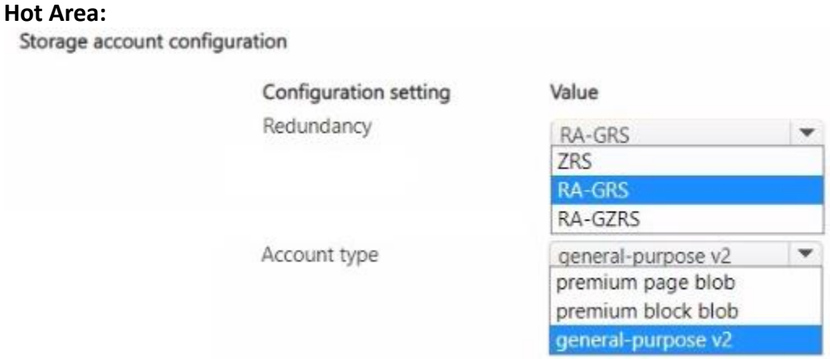

Answer Area:

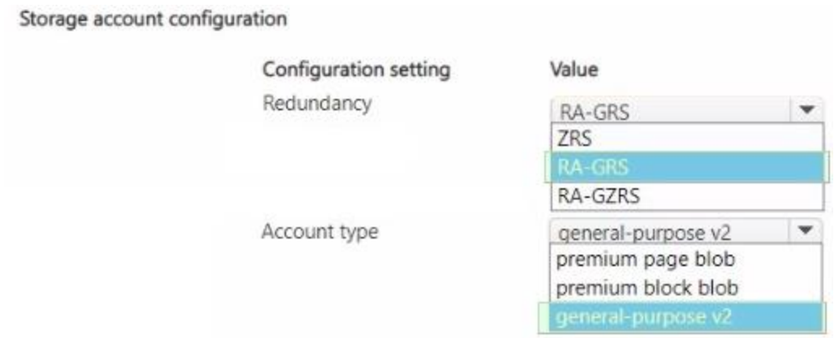

Section:
Explanation:

QUESTION 74
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear on the review screen.
You have an Azure App Service web app named WebApp1 and an Azure Functions app named Function 1. WebApp1 is associated with an Application Insights instance named appinsights1.
You configure a web test and a corresponding alert for WebApp1 in appinsights1. Each alert triggers a delivery of email to your mailbox.
You need to ensure that each alert also triggers execution of Function1.
Solution: Configure an Application Insights funnel.
Does the solution meet the goal?

A. Yes
B. No

Correct Answer: A
Section:

QUESTION 75
You are developing a web application that uses the Microsoft identity platform for user and resource authentication. The web application calls several REST APIs.
A REST API call must read the user's calendar. The web application requires permission to send an email as the user.
You need to authorize the web application and the API.
Which parameter should you use?

A. code_challenge
B. tenant
C. scope
D. clientjd
E. state

Correct Answer: C
Section: