



Microsoft

Exam Questions AZ-204

Developing Solutions for Microsoft Azure

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Exam Topic 1)

You need to resolve the Shipping web site error.

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: AllowedOrigins

A CORS request will fail if Access-Control-Allow-Origin is missing. Scenario:

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

```
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin'
header is present on the requested resource. Origin 'http://testwideworldimporters.com/' is
therefore not allowed access.
```

Box 2: http://test-shippingapi.wideworldimporters.com Syntax: Access-Control-Allow-Origin: *

Access-Control-Allow-Origin: <origin> Access-Control-Allow-Origin: null

<origin> Specifies an origin. Only a single origin can be specified. Box 3: AllowedOrigins

Box 4: POST

The only allowed methods are GET, HEAD, and POST. In this case POST is used. "<Corsrule>" "allowedmethods" Failed to load no "Access-control-Origin" header is present References:

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

NEW QUESTION 2

- (Exam Topic 1)

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

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

Answer Area

Option	Value
Tier	<input type="text"/> <ul style="list-style-type: none"> Standard Premium
Profile	<input type="text"/> <ul style="list-style-type: none"> Akamai Microsoft
Optimization	<input type="text"/> <ul style="list-style-type: none"> general web delivery large file download dynamic site acceleration video-on-demand media streaming

- A. Mastered
- B. Not Mastered

Answer: A

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>

NEW QUESTION 3

- (Exam Topic 1)

You need to correct the VM issues.

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

NOTE: Each correct selection is worth one point.

Issue	Tool
Backup and Restore	<input type="text"/> Azure Site Recovery Azure Backup Azure Data Box Azure Migrate
Performance	<input type="text"/> Azure Network Watcher Azure Traffic Manager ExpressRoute Accelerated Networking

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Backup and Restore: Azure Backup

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

In-Place restore of disks in IaaS VMs is a feature of Azure Backup. Performance: Accelerated Networking

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

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

References:

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

NEW QUESTION 4

- (Exam Topic 2)

You need to configure the ContentUploadService deployment.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Add the following markup to line CS23: types: Private
- B. Add the following markup to line CS24: osType: Windows
- C. Add the following markup to line CS24: osType: Linux
- D. Add the following markup to line CS23: types: Public

Answer: A

Explanation:

Scenario: All Internal services must only be accessible from Internal Virtual Networks (VNets) There are three Network Location types – Private, Public and Domain

Reference:

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

NEW QUESTION 5

- (Exam Topic 2)

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

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

NOTE: Each correct selection is worth one point.

```
"optionalClaims": [
  " ",
  {
    "acct": " ",
    "platf": " ",
    "sid": " ",
    "tenant_ctry": " "
  },
  " ",
  {
    "sid": " ",
    "upn": " ",
    "email": " ",
    "enfpolids": " "
  },
  ],
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: sid

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

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

Box 2: email

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

NEW QUESTION 6

- (Exam Topic 2)

You need to ensure that network security policies are met.

How should you configure network security? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Technology	Value
SSL certificate	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> Valid root certificate Self-signed certificate </div> </div>
Proxy type	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> nginx Azure Application Gateway </div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Valid root certificate

Scenario: All websites and services must use SSL from a valid root certificate authority. Box 2: Azure Application Gateway

Scenario:

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

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

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

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

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

Reference:

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

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

NEW QUESTION 7

- (Exam Topic 3)

You need to correct the corporate website error.

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

Actions	Answer Area
Upload the certificate to Azure Key Vault.	
Update line SC05 of Security.cs to include error handling and then redeploy the code.	
Update line SC03 of Security.cs to include a using statement and then re-deploy the code.	⤵
Add the certificate thumbprint to the WEBSITE_LOAD_CERTIFICATES app setting.	⤴
Upload the certificate to source control.	
Import the certificate to Azure App Service.	
Generate a certificate.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Scenario: Corporate website

While testing the site, the following error message displays: CryptographicException: The system cannot find the file specified. Step 1: Generate a certificate

Step 2: Upload the certificate to Azure Key Vault

Scenario: All SSL certificates and credentials must be stored in Azure Key Vault. Step 3: Import the certificate to Azure App Service

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

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

NEW QUESTION 8

- (Exam Topic 3)

You need to ensure that all messages from Azure Event Grid are processed. What should you use?

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

Answer: B

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-compa>

NEW QUESTION 9

- (Exam Topic 3)

You need to configure API Management for authentication.

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

NOTE: Each correct selection is worth one point.

Setting	Value
Policy	<input type="text"/> Check HTTP header Restrict caller IPs Limit call rate by key Validate JWT
Policy section	<input type="text"/> Inbound Outbound

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Box 1: Validate JWT

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

Scenario: User authentication (see step 5 below)

The following steps detail the user authentication process:

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

Box 2: Outbound Reference:

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

NEW QUESTION 10

- (Exam Topic 4)

You need to implement the Log policy.

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

NOTE: Each correct selection is worth one point.

```

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

```

A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Box 1: logdrop

All log files should be saved to a container named logdrop. Box 2: 15

Logs must remain in the container for 15 days. Box 3: UpdateApplicationSettings

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

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

NEW QUESTION 10

- (Exam Topic 4)

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

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

NOTE: Each correct selection is worth one point.

Code segments	Answer Area
topic	<pre> if { @event["data"]["code segment"].ToString() == "code segment" && @event["data"]["code segment"].ToString() == "Microsoft.Web/sites/write" } </pre>
status	
eventType	
Succeeded	
operationName	
resourceProvider	

A. Mastered
 B. Not Mastered

Answer: A

Explanation:

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

Box 1: Status

Box 2: Succeeded

Box 3: operationName

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

Reference:

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

NEW QUESTION 14

- (Exam Topic 4)

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

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

NOTE: Each correct selection is worth one point.



The screenshot shows three identical code snippets in a C# environment. Each snippet consists of a line of code: `public string` followed by a dropdown menu containing the text `(get; set;)`. The dropdown menu is open, showing four options: `id`, `eventType`, `dataVersion`, and `metadataVersion`. The `id` option is highlighted in the first instance.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: id

id is a unique identifier for the event.

Box 2: eventType

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

Box 3: dataVersion

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

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

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

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

Reference:

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

NEW QUESTION 15

- (Exam Topic 5)

You need to ensure disaster recovery requirements are met. What code should you add at line PC16?

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

NOTE: Each correct selection is worth one point.

Values

Answer Area

```

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Box 1: DirectoryTransferContext We transfer all files in the directory.

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

Box 2: ShouldTransferCallbackAsync

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

Box 3: False

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

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

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-data-movement-library> <https://docs.microsoft.com/en-us/dotnet/api/microsoft.windowsazure.storage.datamovement.directorytransfercon>

NEW QUESTION 19

- (Exam Topic 6)

You need to retrieve all order line items from Order.json and sort the data alphabetically by the city. How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

SELECT li.id AS lineitemid, li.price

FROM

JOIN
 IN

ORDER BY

 ASC

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface Description automatically generated

Box 1: orders o

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

Box 2:li

Box 3: o.line_items

Box 4: o.city

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

NEW QUESTION 23

- (Exam Topic 6)

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

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

Actions	Answer Area
Log in to the registry and push image.	
Create an alias of the image with a new build number.	
Create an alias of the image with the fully qualified path to the registry.	⬅️ ⬆️
Download the image to your local computer.	➡️ ⬇️
Build a new application image by using dockerfile.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A picture containing timeline Description automatically generated

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>

NEW QUESTION 24

- (Exam Topic 7)

You are developing applications for a company. You plan to host the applications on Azure App Services. The company has the following requirements:

- > Every five minutes verify that the websites are responsive.
- > Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly.
- > Generate alerts if a website is experiencing issues.
- > If a website fails to load, the system must attempt to reload the site three more times. You need to implement this process with the least amount of effort.

What should you do? D18912E1457D5D1DDCBD40AB3BF70D5D

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

Answer: D

Explanation:

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

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

NEW QUESTION 26

- (Exam Topic 7)

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.

Requirement	Feature
Which pages visited by users most often correlate to a product purchase?	<input type="text"/>
How does load time of the product display page affect a user's decision to purchase a product?	<input type="text"/>
Which events most influence a user's decision to continue to use the application?	<input type="text"/>
Are there places in the application that users often perform repetitive actions?	<input type="text"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box1: Users Box 2: Impact

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

Box 3: Retention

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

Box 4: User flows

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

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

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

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

NEW QUESTION 28

- (Exam Topic 7)

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

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

You need to create the custom role.

To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Item	Value
Powershell command	<input type="text"/> <ul style="list-style-type: none"> Get-AzureRmRoleDefinition-Name "Reader" ConvertTo-Json Out-File C:\SupportRole.json Get-AzureRmRoleDefinition-Name "Operator" ConvertTo-Json Out-File C:\SupportRole.json Set-AzureRmRoleDefinition-Name "Reader" Input-File C:\SupportRole.json Set-AzureRmRoleDefinition Input-File C:\SupportRole.json
Actions section	<input type="text"/> <ul style="list-style-type: none"> "/read", "Microsoft.Support/" "/read" "* *Microsoft.Support/*" "*"

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

The Set-AzureRmRoleDefinition cmdlet updates an existing custom role in Azure Role-Based Access Control. Provide the updated role definition as an input to the command as a JSON file or a PSRoleDefinition object.

The role definition for the updated custom role MUST contain the Id and all other required properties of the role even if they are not updated: DisplayName, Description, Actions, AssignableScope

Box 2: "/read", "Microsoft.Support/" Microsoft.Support/* Create and manage support tickets "Microsoft.Support" role definition azure

NEW QUESTION 29

- (Exam Topic 7)

You are developing an Azure Function app. The app must meet the following requirements:

- > Enable developers to write the functions by using the Rust language.

> Declaratively connect to an Azure Blob Storage account.

You need to implement the app.

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

NOTE: Each correct selection is worth one point.

Features	Answer Area	
	Requirement	Feature
Custom handler	Enable developers to write the functions by using the Rust language.	Feature
Extension bundle		
Trigger	Declaratively connect to an Azure Blob Storage account.	Feature
Runtime		
Policy		
Hosting plan		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, application Description automatically generated

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>

NEW QUESTION 30

- (Exam Topic 7)

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

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

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

You create the index in Azure Search.

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

* 2. Create the indexBatch with the documents Something like:

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

var batch = IndexBatch.Upload(hotels);

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

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

References:

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

NEW QUESTION 33

- (Exam Topic 7)

You develop a web application.

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Register a new application using the Azure portal

- > Sign in to the Azure portal using either a work or school account or a personal Microsoft account.
- > If your account gives you access to more than one tenant, select your account in the upper right corner.

Set your portal session to the Azure AD tenant that you want.

- > Search for and select Azure Active Directory. Under Manage, select App registrations.
- > Select New registration. (Step 1)
- > In Register an application, enter a meaningful application name to display to users.
- > Specify who can use the application. Select the Azure AD instance. (Step 2)
- > Under Redirect URI (optional), select the type of app you're building: Web or Public client (mobile & desktop). Then enter the redirect URI, or reply URL, for your application. (Step 3)
- > When finished, select Register.

NEW QUESTION 37

- (Exam Topic 7)

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

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

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

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

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

NOTE: Each correct selection is worth one point.

Consistency levels	Answer Area
Strong	Return the most recent patient status. <input type="text"/>
Bounded Staleness	Return health monitoring data that is no less than one version behind. <input type="text"/>
Consistent Prefix	After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges <input type="text"/>
Eventual	

- A. Mastered
- B. Not Mastered

Answer: A

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

Eventual: There's no ordering guarantee for reads. In the absence of any further writes, the replicas eventually converge.

NEW QUESTION 38

- (Exam Topic 7)

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.

Answer Area

App service plan setting	Value
Number of VM instances	<div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">2</div> <div style="border: 1px solid black; padding: 2px;">4</div> <div style="border: 1px solid black; padding: 2px;">8</div> <div style="border: 1px solid black; padding: 2px;">16</div>
Pricing tier	<div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">Isolated</div> <div style="border: 1px solid black; padding: 2px;">Standard</div> <div style="border: 1px solid black; padding: 2px;">Premium</div> <div style="border: 1px solid black; padding: 2px;">Consumption</div>

- A. Mastered
- B. Not Mastered

Answer: A

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).

References:

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

NEW QUESTION 40

- (Exam Topic 7)

You are developing an ASP.NET Core Web API web service that uses Azure Application Insights to monitor performance and track events.

You need to enable logging and ensure that log messages can be correlated to events tracked by Application Insights.

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

NOTE: Each correct selection is worth one point.

Code segments	Answer Area
IncludeEventId	<pre> public class Startup { ... public void ConfigureServices (IServiceCollection services) { services.AddOptions< >(). Configure(o => o. = true); services.AddMvc(); } public void Configure (IApplicationBuilder app, IHostingEnvironment env, ILoggerFactory loggerFactory) { loggerFactory.AddApplicationInsights(app. ,LogLevel.Trace); app.UseMvc(); } } </pre>
ServerFeatures	
LoggerFilterOptions	
ApplicationServices	
ApplicationInsightsLoggerOptions	
TrackExceptionsAsExceptionTelemetry	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: ApplicationInsightsLoggerOptions

If you want to include the EventId and EventName properties, then add the following to the ConfigureServices method:

services.AddOptions<ApplicationInsightsLoggerOptions>().Configure(o => o.IncludeEventId = true);

Box 2: IncludeEventID

Box 3: ApplicationServices

In Asp.Net core apps it turns out that trace logs do not show up in Application Insights out of the box. We need to add the following code snippet to our Configure method in Startup.cs:

loggerFactory.AddApplicationInsights(app.ApplicationServices, logLevel);

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

NEW QUESTION 45

- (Exam Topic 7)

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
$actionGroup = New-AzScheduledQueryRuleAznsActionGroup -ActionGroup "contoso" -EmailSubject "Custom email subject"
               -CustomWebhookPayload "{ 'alert':'#alertrulename', 'IncludeSearchResults':true }"
$alertingAction = New-AzScheduledQueryRuleAlertingAction -AznsAction $actionGroup -Severity "3" -Trigger $triggerCondition
New-AzScheduledQueryRule -ResourceGroupName "contoso" -Location "eastus" -Action $alertingAction -Enabled $true
               -Description "Alert description" -Schedule $schedule -Source $source -Name "Alert Name"
```

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	<input type="radio"/>	<input type="radio"/>
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	<input type="radio"/>	<input type="radio"/>
The log alert is scheduled to run every two hours.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated

Box 1: No

The AzScheduledQueryRuleSource is Heartbeat, not CPU.

Box 2: Yes

The AzScheduledQueryRuleSource is Heartbeat!

Note: New-AzScheduledQueryRuleTriggerCondition creates an object of type Trigger Condition. This object is to be passed to the command that creates Alerting Action object.

Box 3: No

The schedule is 60 minutes, not two hours.

-FrequencyInMinutes: The alert frequency.

-TimeWindowInMinutes: The alert time window

The New-AzAscheduledQueryRuleSchedule command creates an object of type Schedule. This object is to be passed to the command that creates Log Alert Rule.

Reference:

<https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryrule> <https://docs.microsoft.com/en-us/powershell/module/az.monitor/new-azscheduledqueryruletriggercondition>

NEW QUESTION 49

- (Exam Topic 7)

You are developing an application that use an Azure blob named data to store application data. The application creates blob snapshots to allow application state to be reverted to an earlier state. The Azure storage account has soft deleted enabled.

The system performs the following operations in order:

- The blob is updated
- Snapshot 1 is created.
- Snapshot 2 is created.
- Snapshot 1 is deleted.

A system error then deletes the data blob and all snapshots. You need to determine which application states can be restored.

What is the restorability of the application data? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Application State	Restorability
Data blob	<input type="text"/> Can be restored Cannot be restored
Snapshot 1	<input type="text"/> Can be restored Cannot be restored
Snapshot 2	<input type="text"/> Can be restored Cannot be restored

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Can be restored

When enabled, soft delete enables you to save and recover your data when blobs or blob snapshots are deleted. This protection extends to blob data that is erased as the result of an overwrite.

Box 2: Cannot be restored It has been deleted.

Box 3: Can be restored It has not been deleted. References:

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

NEW QUESTION 54

- (Exam Topic 7)

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

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

You must use a storage mechanism with the following requirements:

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

You need to store the information.

Solution: Enable Application Request Routing (ARR). Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 56

- (Exam Topic 7)

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

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

Answer: D

Explanation:

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

Example: Create a Backend Proxy In-Memory Object

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

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

```
AzureRmApiManagementBackendProxy -Url "http://12.168.1.1:8080"
```

```
-ProxyCredential $proxyCreds
```

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

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

```
-Proxy $credential -Description "backend with proxy server"
```

Creates a Backend Proxy Object and sets up Backend

NEW QUESTION 59

- (Exam Topic 7)

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

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

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

The solution must allow dynamic creation and management of all Azure resources within the AKS cluster. You need to configure an AKS cluster for use with the Azure APIs.

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

When you run modern, microservices-based applications in Kubernetes, you often want to control which components can communicate with each other. The principle of least privilege should be applied to how traffic can flow between pods in an Azure Kubernetes Service (AKS) cluster. Let's say you likely want to block traffic directly to back-end applications. The Network Policy feature in Kubernetes lets you define rules for ingress and egress traffic between pods in a cluster.

References:

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

NEW QUESTION 62

- (Exam Topic 7)

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

The Azure Logic app must securely access the Azure Blob storage account. Azure AD resources must remain if the Azure Logic app is deleted.

You need to secure the Azure Logic app. What should you do?

- A. Create an Azure AD custom role and assign role-based access controls.
- B. Create an Azure AD custom role and assign the role to the Azure Blob storage account.
- C. Create an Azure Key Vault and issue a client certificate.
- D. Create a user-assigned managed identity and assign role-based access controls.
- E. Create a system-assigned managed identity and issue a client certificate.

Answer: C

NEW QUESTION 63

- (Exam Topic 7)

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

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

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

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

```
public class PlayerEntity : TableEntity
{
    public PlayerEntity()
    {
    }
    public PlayerEntity(string region, string email)
    {
        PartitionKey =  ;
        RowKey =  ;
    }
    public string Phone { get; set; }
}
public class Player
{
    protected PlayerEntity player;
    async void GetPlayer(string cs,  table, string pk, string rk)
    {
        
        TEntity query = TEntity.Retrieve<PlayerEntity>(pk, rk);
        TableOperation query = TableOperation.Retrieve<PlayerEntity>(pk, rk);
        TableResult query = TableQuery.Retrieve<PlayerEntity>(pk, rk);
        TableResultSegment query = TableResult.Retrieve<PlayerEntity>(pk, rk);
        
        TEntity data = await table.ExecuteAsync(query);
        TableOperation data = await table.ExeucteAsync(query);
        TableQuery data = await table.ExecuteAsync(query);
        TableResult data = await table.ExecuteAsync(query);
        player = data.Result as PlayerEntity;
    }
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: region

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

The core of any table's design is based on its scalability, the queries used to access it, and storage operation requirements. The PartitionKey values you choose will dictate how a table will be partitioned and the type of queries that can be used. Storage operations, in particular inserts, can also affect your choice of PartitionKey values.

Box 2: email

Not phone number some players may not have a phone number. Box 3: CloudTable

Box 4 : TableOperation query =.. Box 5: TableResult

References:

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

NEW QUESTION 68

- (Exam Topic 7)

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

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

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

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

Azure CLI Commands	Answer Area
<input type="text" value="az group create"/>	
<input type="text" value="az group update"/>	
<input type="text" value="az webapp update"/>	⬅
<input type="text" value="az webapp create"/>	➡
<input type="text" value="az appservice plan create"/>	⬆
	⬇

- A. Mastered
- B. Not Mastered

Answer: A

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.

References:

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

NEW QUESTION 71

- (Exam Topic 7)

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

The same connection can be used by multiple concurrent threads. Redis supports both read and write operations.

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

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

Reference:

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

NEW QUESTION 72

- (Exam Topic 7)

You are developing an application that uses Azure Blob storage.

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

You need to process the transaction logs asynchronously. What should you do?

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

Answer: B

Explanation:

Change feed support in Azure Blob Storage

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

Reference:

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

NEW QUESTION 77

- (Exam Topic 7)

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

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 81

- (Exam Topic 7)

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

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

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

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

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

NOTE: Each correct selection is worth one point

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Inbound.

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

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

Box 2: Inbound
A cache-lookup-value policy Example:

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

Box 3: Outbound

A cache-store-value policy. Example:

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

Box 4: Outbound

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

```
<outbound>
<!-- Update response body with user profile-->
<find-and-replace from="$userprofile$"
to="@((string)context.Variables["userprofile"])" />
<base />
</outbound>
```

Reference:

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

NEW QUESTION 82

- (Exam Topic 7)

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

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

You need to design the approach to loading application secrets. What should you do?

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

Answer: C

Explanation:

Use Key Vault references for App Service and Azure Functions.

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

Reference:

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

NEW QUESTION 85

- (Exam Topic 7)

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

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

You develop Azure solutions.

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 87

- (Exam Topic 7)

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

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

Which role-based access control should you use?

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

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 89

- (Exam Topic 7)

A company runs an international travel and bookings management service. The company plans to begin offering restaurant bookings. You must develop a solution that uses Azure Search and meets the following requirements:

- Users must be able to search for restaurants by name, description, location, and cuisine.
- Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness.
- All words in descriptions must be included in searches. You need to add annotations to the restaurant class.

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

```
[SerializePropertyNameAsCamelCase]
public class Restaurant
{
    [Key, IsFilterable]
    public int RestaurantId { get; set; }
    [IsSearchable, IsFilterable, IsSortable]
    public string Name { get; set; }
    [IsSearchable, Required]
    public string location { get; set; }
    public string Phone { get; set; }
    [Required]
    [IsSearchable]
    [IsFilterable, IsFacetable, Required]
    [IsFilterable, IsFacetable, IsSortable]
    public string Description { get; set; }
    [IsFilterable, IsSortable, IsSearchable]
    [IsFilterable, IsSortable, IsFacetable]
    [IsFilterable, IsSortable, Key]
    [IsFilterable, IsSortable, IsSearchable, Required]
    public double Rating { get; set; }
    [IsSearchable, IsFilterable, IsFacetable]
    [IsFilterable, IsSortable, Key]
    [IsFilterable, IsSortable, IsSearchable]
    [IsFilterable, IsSortable, Key, Required]
    public List<string> Cuisines { get; set; }
    [IsFilterable, IsSortable, Key, Required]
    [IsSearchable, IsSortable, IsFacetable]
    [IsFilterable, IsSortable, Key, IsSearchable]
    [IsFilterable, IsFacetable]
    public bool FamilyFriendly { get; set; }
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: [IsSearchable.IsFilterable.IsSortable,IsFacetable] Location

Users must be able to search for restaurants by name, description, location, and cuisine.

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 2: [IsSearchable.IsFilterable.IsSortable,Required]

Description

Users must be able to search for restaurants by name, description, location, and cuisine. All words in descriptions must be included in searches.

Box 3: [IsFilterable,IsSortable,IsFaceTable] Rating

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 4: [IsSearchable.IsFilterable,IsFacetable]

Cuisines

Users must be able to search for restaurants by name, description, location, and cuisine.

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. Box 5: [IsFilterable,IsFacetable]

FamilyFriendly

Users must be able to narrow the results further by location, cuisine, rating, and family-friendliness. References:

<https://www.henkboelman.com/azure-search-the-basics/>

NEW QUESTION 92

- (Exam Topic 7)

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

- A. Deploy the website to an App Service that uses the Shared service tie
- B. Configure the App Service plan to automatically scale when the CPU load is high.
- C. Deploy the website to a virtual machin
- D. Configure the virtual machine to automatically scale when the CPU load is high.
- E. Deploy the website to an App Service that uses the Standard service tie
- F. Configure the App Service plan to automatically scale when the CPU load is high.
- G. Deploy the website to a virtual machin
- H. Configure a Scale Set to increase the virtual machine instancecount when the CPU load

Answer: C

Explanation:

Windows Azure Web Sites (WAWS) offers 3 modes: Standard, Free, and Shared.
 Standard mode carries an enterprise-grade SLA (Service Level Agreement) of 99.9% monthly, even for sites with just one instance.
 Standard mode runs on dedicated instances, making it different from the other ways to buy Windows Azure Web Sites.

NEW QUESTION 93

- (Exam Topic 7)

A company is developing a mobile app for field service employees using Azure App Service Mobile Apps as the backend. The company's network connectivity varies throughout the day. The solution must support offline use and synchronize changes in the background when the app is online app.

You need to implement the solution.

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

```
var client = new MobileServiceClient("MOBILE_APP_URL");
var store = new MobileServiceSQLiteStore
(Constants.OfflineDbPath);
store.DefineTable<TodoItem>();
await client.SyncContext.IntitalizeAsync(store);
```

▼

```
var todoTable = client.GetSyncTable<TodoItem>();
var todoTable = client.GetTable<TodoItem>();
var todoTable = client.SyncTable;
var todoTable = client.Table;
```

```
await client.SyncContext.PushAsync();
```

▼

```
await todoTable.PullAsync("allTodoItems",todoTable.CreateQuery());
await todoTable.UpdateAsync();
todoTable.PullAsync("allTodoItems", todoTable.CreateQuery());
todoTable.UpdateAsync();
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: var todoTable = client GetSyncTable<TodoItem>()
 To setup offline access, when connecting to your mobile service, use the method GetSyncTable instead of GetTable (example):
 IMobileServiceSyncTable todoTable = App.MobileService.GetSyncTable(); / Box 2: await todoTable.PullAsync("allTodoItems",todo.Table.CreateQuery());
 Your app should now use IMobileServiceSyncTable (instead of IMobileServiceTable) for CRUD operations. This will save changes to the local database and also keep a log of the changes. When the app is ready to synchronize its changes with the Mobile Service, use the methods PushAsync and PullAsync (example):
 await App.MobileService.SyncContext.PushAsync(); await todoTable.PullAsync();
 References:
<https://azure.microsoft.com/es-es/blog/offline-sync-for-mobile-services/>

NEW QUESTION 96

- (Exam Topic 7)

You are developing a web application that runs as an Azure Web App. The web application stores data in Azure SQL Database and stores files in an Azure Storage account. The web application makes HTTP requests to external services as part of normal operations. The web application is instrumented with Application Insights. The external services are OpenTelemetry compliant. You need to ensure that the customer ID of the signed in user is associated with all operations throughout the overall system. What should you do?

- A. Create a new SpanContext with the TraceRags value set to the customer ID for the signed in user.
- B. On the current SpanContext, set the Traceld to the customer ID for the signed in user.
- C. Add the customer ID for the signed in user to the CorrelationContext in the web application.
- D. Set the header Ocp-Apim-Trace to the customer ID for the signed in user.

Answer: D

NEW QUESTION 101

- (Exam Topic 7)

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

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

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

You have the following requirements:

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

You need to implement the messaging solution.

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

You can create a function that is triggered when messages are submitted to an Azure Storage queue. Reference: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-queue-triggered-function>

NEW QUESTION 106

- (Exam Topic 7)

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

The news API back end is implemented as a RESTful service and hosted in an Azure App Service instance. You need to configure back-end authentication for the API Management service instance.

Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Configuration parameter	Value
Target	value
Gateway credentials	value

Azure Resource
HTTP(s) endpoint
Basic
Client cert

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Resource

Box 2: Client cert

API Management allows to secure access to the back-end service of an API using client certificates. References: <https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-ba>

NEW QUESTION 111

- (Exam Topic 7)

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

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

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

Actions	Answer area
Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage.	
Run the Azure PowerShell command set-AzureRmVMDiskEncryptionExtension.	
Run the Azure PowerShell command set-AzureRmVMSDisk.	
Encrypt the on-premises VHD by using BitLocker with a TPM. Upload the VM to Azure Storage.	
Run the Azure PowerShell command New-AzureRmVM.	

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Step 1: Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage Step 2: Run the Azure PowerShell command Set-AzureRMVMOSDisk

To use an existing disk instead of creating a new disk you can use the Set-AzureRMVMOSDisk command. Example:

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

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

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

```
$vm = Set-AzureRmVMOSDisk -VM $vm -VhdUri $osDiskVhdUri -name $osDiskName -Create Step 3: Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension
```

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

Incorrect:

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

References:

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

NEW QUESTION 112

- (Exam Topic 7)

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

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

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

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

Solution:

- Create a new Azure AD application's manifest, set value of the groupMembershipClaims option to All.

- In the website, use the value of the groups claim from the JWT for the user to determine permissions. Does the solution meet the goal?

A. Yes

B. No

Answer: A

Explanation:

To configure Manifest to include Group Claims in Auth Token

- * 1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:

- * 2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.

- * 3. Locate the "groupMembershipClaims" setting. Set its value to either "SecurityGroup" or "All". To help you decide which:

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

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

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

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

NEW QUESTION 117

- (Exam Topic 7)

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.



A. Mastered

B. Not Mastered

Answer: A

Explanation:

```
$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM
```

```
Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -AssignIdentity:$SystemAssigned https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-power
```

NEW QUESTION 122

- (Exam Topic 7)

You are maintaining an existing application that uses an Azure Blob GPv1 Premium storage account. Data older than three months is rarely used.

Data newer than three months must be available immediately. Data older than a year must be saved but does not need to be available immediately.

You need to configure the account to support a lifecycle management rule that moves blob data to archive storage for data not modified in the last year.

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

Actions	Answer Area
Upgrade the storage account to GPv2	
Create a new GPv2 Standard account and set its default access tier level to cool	⤵
Change the storage account access tier from hot to cool	⤴
Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account	⤴ ⤵

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Upgrade the storage account to GPv2

Object storage data tiering between hot, cool, and archive is supported in Blob Storage and General Purpose v2 (GPv2) accounts. General Purpose v1 (GPv1) accounts don't support tiering.

You can easily convert your existing GPv1 or Blob Storage accounts to GPv2 accounts through the Azure portal.

Step 2: Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account

Step 3: Change the storage account access tier from hot to cool Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Only the hot and cool access tiers can be set at the account level. The archive access tier can only be set at the blob level.

Reference:

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

NEW QUESTION 126

- (Exam Topic 7)

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

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

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

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

You need to configure authorization. Solution:

- > Create a new Azure AD application. In the application's manifest, define application roles that match the required permission levels for the application.
- > Assign the appropriate Azure AD group to each role. In the website, use the value of the roles claim from the JWT for the user to determine permissions.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

To configure Manifest to include Group Claims in Auth Token

> Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to

App registrations to find your application:

- > Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.
- > Locate the "groupMembershipClaims" setting. Set its value to either "SecurityGroup" or "All". To help you decide which:
- > "SecurityGroup" - groups claim will contain the identifiers of all security groups of which the user is a member.
- > "All" - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member

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

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

NEW QUESTION 131

- (Exam Topic 7)

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

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

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

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

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

NOTE: Each correct selection is worth one point.

Settings	Answer Area		
client_id	API	Permission	Type
profile	Azure Storage	Setting	Setting
delegated	Microsoft Graph	User.Read	Setting
application			
user_impersonation			

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: user_impersonation

Box 2: delegated Example:

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

Example

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

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

NEW QUESTION 134

- (Exam Topic 7)

You manage several existing Logic Apps.

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

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Enterprise Integration Pack

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

Box 2: Code View Editor

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

Box 3: Logical Apps Designer

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

References:

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

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

NEW QUESTION 136

- (Exam Topic 7)

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.

Answer Area

```
{
  "oauth2AllowImplicitFlow": true,
  "addIns": {
    "orgRestrictions": {
      "availableToOtherTenants": true,
      "requiredResourceAccess": [
        {
          "resource": "00000000-0000-0000-0000-000000000000",
          "scope": "openid"
        }
      ]
    }
  },
  "type": "Scope"
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
 Answer Area

```
{
  "oauth2AllowImplicitFlow": true,
  "addIns": {
    "orgRestrictions": {
      "availableToOtherTenants": true,
      "requiredResourceAccess": [
        {
          "resource": "00000000-0000-0000-0000-000000000000",
          "scope": "openid"
        }
      ]
    }
  },
  "type": "Scope"
}
```

NEW QUESTION 141

- (Exam Topic 7)

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

You need to configure the services.

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

NOTE: Each correct selection is worth one point.

Services	Scenario	Service
Logic Apps	Process a queue data item.	<input type="text"/>
WebJobs	Manage all code segments from the same DevOps environment.	<input type="text"/>
Flow		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: WebJobs

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

Box 2: Flow

NEW QUESTION 145

- (Exam Topic 7)

You are preparing to deploy an application to an Azure Kubernetes Service (AKS) cluster. The application must only be available from within the VNet that includes the cluster. You need to deploy the application.

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

NOTE: Each correct selection is worth one point.

Code segments

- Ingress
- Service
- LoadBalancer
- Deployment
- ingress.class
- azure-load-balancer-internal

Answer Area

```

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

YAML:

```

apiVersion: v1 kind: Service metadata:
name: internal-app annotations:
service.beta.kubernetes.io/azure-load-balancer-internal: "true" spec:
type: LoadBalancer ports:
- port: 80 selector:
app: internal-app References:
https://docs.microsoft.com/en-us/azure/aks/internal-lb
  
```

NEW QUESTION 149

- (Exam Topic 7)

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

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

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

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

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

Commands

Answer Area

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

⏪
⏩

⏴
⏵

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: WORKDIR /apps/ContosoApp Step 2: COPY ./The Docker document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.

Step 3: EXPOSE ./ContosoApp/ /app/ContosoApp Step 4: CMD powershell ./setupScript.ps1 ENTRYPOINT ["dotnet", "ContosoApp.dll"]

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

- > Call setupScript.ps1 when the container is built.
- > Run ContosoApp.dll when the container starts. References:

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

NEW QUESTION 152

- (Exam Topic 7)

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

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

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

You assign users to RBAC role.

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

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

NOTE: Each correct selection is worth one point.



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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: user_impersonation

Box 2: delegated Example:

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

Box 3: delegated Example

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

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

NEW QUESTION 155

- (Exam Topic 7)

You are implementing an Azure API app that uses built-in authentication and authorization functionality. All app actions must be associated with information about the current user.

You need to retrieve the information about the current user.

What are two possible ways to achieve the goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

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

Answer: AC

Explanation:

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

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

References:

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

NEW QUESTION 158

- (Exam Topic 7)

You are developing a web service that will run on Azure virtual machines that use Azure Storage. You configure all virtual machines to use managed identities.

You have the following requirements:

- Must use only Azure Instance Metadata Service endpoints.

You need to write code to retrieve an access token to access Azure Storage. To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Code segment 1

http://localhost:50342/oauth2/token

http://169.254.169.254:50432/oauth2/token

http://localhost/metadata/identity/oauth2/token

http://169.254.169.254/metadata/identity/oauth2/token

Code segment 2

XDocument.Parse(payload);

new MultipartContent(payload);

new NetworkCredential("Azure", payload);

JsonConvert.DeserializeObject<Dictionary<string, string>>(payload);

Answer Area

```
var url = "Code segment 1";
var queryString = "...";
var client = new HttpClient();
var response = await client.GetAsync(url + queryString);
var payload = await response.Content.ReadAsStringAsync();

return Code segment 2;
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Azure Instance Metadata Service endpoints "/oauth2/token"

Box 1: http://169.254.169.254/metadata/identity/oauth2/token

Sample request using the Azure Instance Metadata Service (IMDS) endpoint (recommended): GET 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01

&resource=https://management.azure.com/' HTTP/1.1 Metadata: true

Box 2: JsonConvert.DeserializeObject<Dictionary<string,string>>(payload); Deserialized token response; returning access code.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token> <https://docs.microsoft.com/en-us/azure/service-fabric/how-to-managed-identity-service-fabric-app-code>

NEW QUESTION 162

- (Exam Topic 7)

You develop a solution that uses Azure Virtual Machines (VMs).

The VMs contain code that must access resources in an Azure resource group. You grant the VM access to the resource group in Resource Manager.

You need to obtain an access token that uses the VMs system-assigned managed identity. Which two actions should you perform? Each correct answer presents part of the solution.

- A. Use PowerShell on a remote machine to make a request to the local managed identity for Azure resources endpoint.
- B. Use PowerShell on the VM to make a request to the local managed identity for Azure resources endpoint.
- C. From the code on the V
- D. call Azure Resource Manager using an access token.
- E. From the code on the V
- F. call Azure Resource Manager using a SAS token.
- G. From the code on the V
- H. generate a user delegation SAS token.

Answer: BC

NEW QUESTION 167

- (Exam Topic 7)

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

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

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

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

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

Solution: Provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.

- A. Yes
- B. No

Answer: A

Explanation:

References:

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

NEW QUESTION 171

- (Exam Topic 7)

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

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

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

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

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

Commands

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

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 173

- (Exam Topic 7)

You plan to deploy a new application to a Linux virtual machine (VM) that is hosted in Azure.

The entire VM must be secured at rest by using industry-standard encryption technology to address organizational security and compliance requirements.

You need to configure Azure Disk Encryption for the VM.

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

NOTE: Each correct selection is worth one point.

Answer Area

```

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

- Box 1: keyvault
- Create an Azure Key Vault with az keyvault create and enable the Key Vault for use with disk encryption. Specify a unique Key Vault name for keyvault_name as follows:
 keyvault_name=myvaultname\$RANDOM az keyvault create \
 --name \$keyvault_name \
 --resource-group \$resourcegroup \
 --location eastus \
 --enabled-for-disk-encryption True
- Box 2: keyvault key
- The Azure platform needs to be granted access to request the cryptographic keys when the VM boots to decrypt the virtual disks. Create a cryptographic key in your Key Vault with az keyvault key create. The following example creates a key named myKey:

```
az keyvault key create \
--vault-name $keyvault_name \
--name myKey \
--protection software Box 3: vm
```

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

```
az vm create \
--resource-group $resourcegroup \
--name myVM \
--image Canonical:UbuntuServer:16.04-LTS:latest \
--admin-username azureuser \
--generate-ssh-keys \ Box 4: vm encryption
```

Encrypt your VM with az vm encryption enable: az vm encryption enable \

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

Note: seems to an error in the question. Should have enable instead of create. Box 5: all Encrypt both data and operating system.

References:

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

NEW QUESTION 178

- (Exam Topic 7)

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

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

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

You need to ensure the files are correctly processed.

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

Answer Area

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

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

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

            task [dropdown] =outputFileList;

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: CreateJob

Box 2: TaskSuccess

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

Incorrect: TaskCompletion: Upload the file(s) after the task process exits, no matter what the exit code was. Box 3: TaskFailure

TaskFailure: Upload the file(s) only after the task process exits with a nonzero exit code. Box 4: OutputFiles

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

References:

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

NEW QUESTION 181

- (Exam Topic 7)

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

You need to create compute nodes for the solution on Azure Batch. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 186

- (Exam Topic 7)

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

Fourth Coffee is migrating this application to Azure.

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

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

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

Azure CLI commands	Answer area
<pre>az webapp config hostname add --webapp-name \$appName --resource-group fourthCoffeePublicWebResourceGroup --hostname \$fqdn</pre>	
<pre>#!/bin/bash appName="FourthCoffeePublicWeb\$random" location "WestUS" dockerHubContainerPath="FourthCoffee/publicweb:v1" fqdn=http://www.fourthcoffee.com>www.fourthcoffee.com</pre>	<div style="text-align: center;"> ⬅ </div>
<pre>az webapp create --name \$appName --plan AppServiceLinuxDockerPlan --resource-group fourthCoffeePublicWebResourceGroup</pre>	<div style="text-align: center;"> ⬆ </div>
<pre>az webapp config container set --docker-custom-image-name \$dockerHubContainerPath --name \$appName --resource-group fourthCoffeePublicWebResourceGroup</pre>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: #bin/bash

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

Step 2: az webapp config hostname add

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

Step 3: az webapp create

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

Step : az webapp config container set

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

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

References:

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

NEW QUESTION 188

- (Exam Topic 7)

You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage containers named Container1 and Container2.

Uploading of videos occurs on an irregular basis.

You need to copy specific blobs from Container1 to Container2 when a new video is uploaded. What should you do?

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

Answer: B

Explanation:

The Start-AzureStorageBlobCopy cmdlet starts to copy a blob. Example 1: Copy a named blob

```
C:\PS>Start-AzureStorageBlobCopy -SrcBlob "ContosoPlanning2015" -DestContainer "ContosoArchives"
```

```
-SrcContainer "ContosoUploads"
```

This command starts the copy operation of the blob named ContosoPlanning2015 from the container named ContosoUploads to the container named ContosoArchives.

Reference:

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

NEW QUESTION 191

- (Exam Topic 7)

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

You have the following parameters:

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

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

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

Command segments

Answer Area

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

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

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

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: az batch pool create

Create a new Linux pool with a virtual machine configuration. az batch pool create \

```
--id mypool \
```

```
--vm-size Standard_A1 \
```

```
--target-dedicated 2 \
```

```
--image canonical:ubuntu:16.04-LTS \
```

```
--node-agent-sku-id "batch.node.ubuntu 16.04" Step 2: az batch job create
```

Create a new job to encapsulate the tasks that are added. az batch job create \

```
--id myjob \
```

```
--pool-id mypool
```

Step 3: az batch task create

```
# Add tasks to the job. Here the task is a basic shell command. az batch task create \  
--job-id myjob \  
--task-id task1 \  
--command-line "/bin/bash -c 'printenv AZ_BATCH_TASK_WORKING_DIR'" Step 4: for i in {1..$numberOfJobs} do  
References:  
https://docs.microsoft.com/bs-latn-ba/azure/batch/scripts/batch-cli-sample-run-job
```

NEW QUESTION 195

.....

Relate Links

100% Pass Your AZ-204 Exam with ExamBible Prep Materials

<https://www.exambible.com/AZ-204-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>