

Microsoft

Exam Questions AZ-204

Developing Solutions for Microsoft Azure



NEW QUESTION 1

- (Topic 8)

You are developing a road tollway tracking application that sends tracking events by using Azure Event Hubs using premium tier. Each road must have a throttling policy uniquely assigned. You need to configure the event hub to allow for per-road throttling. What should you do?

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

Answer: D

NEW QUESTION 2

- (Topic 8)

You are developing several Azure API Management (APIM) hosted APIs. The APIs have the following requirements: Require a subscription key to access all APIs.

- Include terms of use that subscribers must accept to use the APIs.
- Administrators must review and accept or reject subscription attempts.
- Limit the count of multiple simultaneous subscriptions. You need to implement the APIs.

What should you do? OB.

- A. Create and publish a product.
- B. Configure and apply query string-based versioning.
- C. Configure and apply header-based versioning.
- D. Add a new revision to all API
- E. Make the revisions current and add a change log entr

Answer: B

NEW QUESTION 3

- (Topic 8)

You develop and add several functions to an Azure Function app that uses the latest runtime host. The functions contain several REST API endpoints secured by using SSL. The Azure Function app runs in a Consumption plan.

You must send an alert when any of the function endpoints are unavailable or responding too slowly.

You need to monitor the availability and responsiveness of the functions. What should you do?

- A. Create a URL ping test.
- B. Create a timer triggered function that calls TrackAvailability() and send the results to ApplicationInsights.
- C. Create a timer triggered function that calls GetMetric("Request Size") and send the results toApplication Insights.
- D. Add a new diagnostic setting to the Azure Function ap
- E. Enable the FunctionAppLogs and Send to Log Analytics options.

Answer: B

Explanation:

You can create an Azure Function with TrackAvailability() that will run periodically according to the configuration given in TimerTrigger function with your own business logic. The results of this test will be sent to your Application Insights resource, where you will be able to query for and alert on the availability results data. This allows you to create customized tests similar to what you can do via Availability Monitoring in the portal. Customized tests will allow you to write more complex availability tests than is possible using the portal UI, monitor an app inside of your Azure VNET, change the endpoint address, or create an availability test even if this feature is not available in your region.

D18912E1457D5D1DDCBD40AB3BF70D5D

Reference:

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

NEW QUESTION 4

- (Topic 8)

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

- A. AppServiceEnvironmentPlatformLogs
- B. AppServiceApplogs
- C. AppServiceAuditLogs
- D. AppServiceConsoteLogs

Answer: C

NEW QUESTION 5

DRAG DROP - (Topic 8)

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

Custom handler

Extension bundle

Trigger

Runtime

Policy

Hosting plan

Answer Area

Requirement

Feature

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

Declaratively connect to an Azure Blob Storage account.

Feature

Feature

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Custom handler
Custom handlers can be used to create functions in any language or runtime by running an HTTP server process, for example Go or Rust.
Box 2: Trigger
Functions are invoked by a trigger and can have exactly one. In addition to invoking the function, certain triggers also serve as bindings. You may also define multiple bindings in addition to the trigger. Bindings provide a declarative way to connect data to your code.

NEW QUESTION 6

- (Topic 8)
An organization deploys Azure Cosmos DB.
You need to ensure that the index is updated as items are created, updated, or deleted. What should you do?

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

Answer: B

NEW QUESTION 7

- (Topic 8)
You are building a web application that performs image analysis on user photos and returns metadata containing objects identified. The image is very costly in terms of time and compute resources. You are planning to use Azure Redis Cache so duplicate uploads do not need to be reprocessed.
In case of an Azure data center outage, metadata loss must be kept to a minimum. You need to configure the Azure Redis cache instance.
Which two actions should you perform?

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

Answer: BC

NEW QUESTION 8

DRAG DROP - (Topic 8)
You are developing Azure WebJobs.
You need to recommend a WebJob type for each scenario.
Which WebJob type should you recommend? To answer, drag the appropriate WebJob types to the correct scenarios. Each WebJob type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE:Each correct selection is worth one point.

WebJob types

Triggered

Continuous

Scenario

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

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

Supports remote debugging

WebJob type

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Box 1: Continuous
Continuous runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.
Box 2: Triggered
Triggered runs on a single instance that Azure selects for load balancing.
Box 3: Continuous
Continuous supports remote debugging.
Note:
The following table describes the differences between continuous and triggered WebJobs.

Continuous	Triggered
Starts immediately when the WebJob is created. To keep the job from ending, the program or script typically does its work inside an endless loop. If the job does end, you can restart it.	Starts only when triggered manually or on a schedule.
Runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.	Runs on a single instance that Azure selects for load balancing.
Supports remote debugging.	Doesn't support remote debugging.

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

NEW QUESTION 9

HOTSPOT - (Topic 8)
You are developing a .NET application that communicates with Azure Storage. A message must be stored when the application initializes.
You need to implement the message.
How should you complete the code segment? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse(CloudConfigurationManager.GetSetting("StorageConnectionString"));

CloudQueueClient pVar1 = storageAccount. CreateCloudQueueClient ();
CloudTableClient pVar2 = pVar1. CreateCloudTableClient ();
CloudQueue GetQueueReference ();
CloudTable GetTableReference ();

CloudQueueClient tExistsAsync(); CreateCloudQueueClient ("contoso-storage");
CloudTableClient CreateCloudTableClient ();
CloudQueue GetQueueReference ();
CloudTable GetTableReference ();
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse(CloudConfigurationManager.GetSetting("StorageConnectionString"));

CloudQueueClient pVar1 = storageAccount. CreateCloudQueueClient ();
CloudTableClient pVar2 = pVar1. CreateCloudTableClient ();
CloudQueue GetQueueReference ();
CloudTable GetTableReference ();

CloudQueueClient tExistsAsync(); CreateCloudQueueClient ("contoso-storage");
CloudTableClient CreateCloudTableClient ();
CloudQueue GetQueueReference ();
CloudTable GetTableReference ();
```

NEW QUESTION 10

- (Topic 8)
You develop Azure Durable Functions to manage vehicle loans.
The loan process includes multiple actions that must be run in a specified order. One of the actions includes a customer credit check process, which may require multiple days to process.
You need to implement Azure Durable Functions for the loan process. Which Azure Durable Functions type should you use?

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

Answer: A

NEW QUESTION 10

- (Topic 8)

You are developing an Azure messaging solution.

You need to ensure that the solution that meets the following requirements:

- Provide transactional support
- Provide duplicate detection.
- Store the messages for an unlimited period of time

Which two technologies will meet the requirements? Each correct answer presents a complete solution NOTE Each correct selection is worth one point.

- A. Azure Service Bus Queue
- B. Azure Storage Queue
- C. Azure Service Bus Topic
- D. Azure Event Hub

Answer: AC

Explanation:

The Azure Service Bus Queue and Topic has duplicate detection.

Enabling duplicate detection helps keep track of the application-controlled MessageId of all messages sent into a queue or topic during a specified time window.

Reference:

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

NEW QUESTION 15

- (Topic 8)

You develop a REST API. You implement a user delegation SAS token to communicate with Azure Blob storage.

The token is compromised. You need to revoke the token.

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

NOTE: Each correct selection is worth one point.

- A. Revoke the delegation keys
- B. Delete the stored access policy.
- C. Regenerate the account key.
- D. Remove the role assignment for the security principle.

Answer: AB

Explanation:

A: Revoke a user delegation SAS

To revoke a user delegation SAS from the Azure CLI, call the az storage account revoke- delegation-keys command. This command revokes all of the user delegation keys associated with the specified storage account. Any shared access signatures associated with those keys are invalidated.

B: To revoke a stored access policy, you can either delete it, or rename it by changing the signed identifier.

Changing the signed identifier breaks the associations between any existing signatures and the stored access policy. Deleting or renaming the stored access policy immediately effects all of the shared access signatures associated with it. D18912E1457D5D1DDCBD40AB3BF70D5D

Reference:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/storage/blobs/storage-blob-user-delegationsas-create-cli.md>

<https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy#modifying-or-revoking-a-stored-access-policy>

NEW QUESTION 19

- (Topic 8)

A company maintains multiple web and mobile applications. Each application uses custom in-house identity providers as well as social identity providers.

You need to implement single sign-on (SSO) for all the applications. What should you do?

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

Answer: A

Explanation:

<https://docs.microsoft.com/en-us/azure/active-directory-b2c/custom-policy-reference-ssso>

NEW QUESTION 20

- (Topic 8)

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

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

You are developing an Azure solution to collect point-of-sale (POS) device data from

2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Reference:

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

NEW QUESTION 25

DRAG DROP - (Topic 8)

You develop and deploy an Azure App Service ---- app. The web app accesses data in an Azure SQL database

You must update the web app to store frequently used data in a new Azure Cache for Redis Premium instance.

You need to implement the Azure Cache for Redis features.

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

NOTE: Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 28

- (Topic 8)

You are developing a mobile instant messaging app for a company. The mobile app must meet the following requirements:

- Support offline data sync.
- Update the latest messages during normal sync cycles. You need to implement Offline Data Sync.

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

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

Answer: BE

Explanation:

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

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

References:

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

NEW QUESTION 33

- (Topic 8)

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

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

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

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

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

Solution: Convert the Azure Storage account to a BlobStorage storage account. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Not necessary to convert the account, instead move photo processing to an Azure Function triggered from the blob upload..

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

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

Reference:

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

NEW QUESTION 36

HOTSPOT - (Topic 8)

You are developing an application that runs in several customer Azure Kubernetes Service clusters, within each cluster, a pod runs that collects performance data to be analyzed later, a large amount of data is collected so saving latency must be minimized

The performance data must be stored so that pod restarts do not impact the stored data. Write latency should be minimized.

You need to configure blob storage.

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

```
apiVersion: storage.k8s.io/v1
kind:
metadata: PodStorage
      StorageClass
      PersistentVolume
      PersistentVolumeClaim

name: data-store
provisioner: kubernetes.io,
      azure-disk
      azure-file
      portworx-volume
      scaleio

parameters:
  skuName: Premium_LRS
reclaimPolicy:
      local
      retain
      delete
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
apiVersion: storage.k8s.io/v1
kind:
metadata: PodStorage
      StorageClass
      PersistentVolume
      PersistentVolumeClaim

name: data-store
provisioner: kubernetes.io,
      azure-disk
      azure-file
      portworx-volume
      scaleio

parameters:
  skuName: Premium_LRS
reclaimPolicy:
      local
      retain
      delete
```

NEW QUESTION 37

- (Topic 8)

You are developing an application that allows users to find musicians that are looking for work. The application must store information about musicians, the instruments that they play, and other related data.

The application must also allow users to determine which musicians have played together, including groups of three or more musicians that have performed

together at a specific location.
Which Azure Cosmos D6 API should you use for the application?

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

Answer: B

NEW QUESTION 38

HOTSPOT - (Topic 8)

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

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

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

NOTE: Each correct selection is worth one point.

App service plan setting	Value
Number of VM instances	<div><div></div><div>2</div><div>4</div><div>8</div><div>16</div></div>
Pricing tier	<div><div></div><div>Isolated</div><div>Standard</div><div>Premium</div><div>Consumption</div></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

- (Topic 8)

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

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

You 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 42

- (Topic 8)

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

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

NOTE: Each correct selection is worth one point.

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

Answer: ADE

Explanation:

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

Application Insights manages the ID of a session for you. References:

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

NEW QUESTION 47

- (Topic 8)

You are developing an Azure-based web application. The application goes offline periodically to perform offline data processing. While the application is offline, numerous Azure Monitor alerts fire which result in the on-call developer being paged.

The application must always log when the application is offline for any reason.

You need to ensure that the on-call developer is not paged during offline processing. What should you do?

- A. Add Azure Monitor alert processing rules to suppress notifications.
- B. Create an Azure Monitor Metric Alert.
- C. Build an Azure Monitor action group that suppresses the alerts.
- D. Disable Azure Monitor Service Health Alerts during offline processing.

Answer: C

NEW QUESTION 49

- (Topic 8)

You are developing an e-commerce solution that uses a microservice architecture.

You need to design a communication backplane for communicating transactional messages between various parts of the solution. Messages must be communicated in first-in-first-out (FIFO) order.

What should you use?

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

Answer: C

Explanation:

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

? Your solution requires the queue to provide a guaranteed first-in-first-out (FIFO) ordered delivery.

Reference:

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

NEW QUESTION 52

DRAG DROP - (Topic 8)

You are authoring a set of nested Azure Resource Manager templates to deploy multiple Azure resources.

The templates must be tested before deployment and must follow recommended practices. You need to validate and test the templates before deployment.

Which tools should you use? To answer, drag the appropriate tools to the correct requirements. 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	Requirement	Tool
Parameter file	Determine whether the templates follow recommended practices.	Tool
Template function		
Azure Resource Manager test toolkit	Test and validate changes that templates will make to the environment.	Tool
User-defined function		
What-if operation		
Azure Deployment Manager		

A. Mastered

B. Not Mastered

Answer: A

Explanation:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/test-toolkit>

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/deploy-what-if?tabs=azure-powershell>

NEW QUESTION 54

DRAG DROP - (Topic 8)

You are a developer for a Software as a Service (SaaS) company. You develop solutions that provide the ability to send notifications by using Azure Notification Hubs.

You need to create sample code that customers can use as a reference for how to send raw notifications to Windows Push Notification Services (WNS) devices. The sample code must not use external packages.

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

NOTE: Each correct selection is worth one point.



```
var endpoint = "...";
var payload = "...";
var request = new HttpRequestMessage(HttpMethod.Post, endpoint);
request.Headers.Add("X-WNS-Type", "wns/raw");
request.Headers.Add("ServiceBusNotification-Format", " " Code segment " " );
request.Content = new StringContent(payload, Encoding.UTF8, " " Code segment " " );
var client = new HttpClient();
await client.SendAsync(request);
```

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Box 1: windows Example code:

```
var request = new HttpRequestMessage(method, $"{resourceUri}?api-version=2017-04"); request.Headers.Add("Authorization", createToken(resourceUri, KEY_NAME, KEY_VALUE));
request.Headers.Add("X-WNS-Type", "wns/raw"); request.Headers.Add("ServiceBusNotification-Format", "windows"); return request;
```

Box 2: application/octet-stream

Example code capable of sending a raw notification: string resourceUri =

```
 $"https://{NH_NAMESPACE}.servicebus.windows.net/{HUB_NAME}/messages/"; using (var request = CreateHttpRequest(HttpMethod.Post, resourceUri))
{
    request.Content = new StringContent(content, Encoding.UTF8, "application/octet-stream"); request.Content.Headers.ContentType.CharSet = string.Empty;
    var httpClient = new HttpClient();
    var response = await httpClient.SendAsync(request); Console.WriteLine(response.StatusCode);
}
```

NEW QUESTION 59

- (Topic 8)

You are developing several Azure API Management (APIM) hosted APIs.

You must transform the APIs to hide private backend information and obscure the technology stack used to implement the backend processing.

You need to protect all APIs. What should you do?

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

Answer: A

NEW QUESTION 61

DRAG DROP - (Topic 8)

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

You need to configure the services.

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

NOTE: Each correct selection is worth one point.

Services	Scenario	Service
Logic Apps	Process a queue data item.	
WebJobs		
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 63

DRAG DROP - (Topic 8)

You have an Azure Cosmos DB for NoSQL account.

You plan to develop two apps named App1 and App2 that will use the change feed functionality to track changes to containers.

App1 will use the pull model and App2 will use the push model.

You need to choose the method to track the most recently processed change in App1 and App2.

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

NOTE: Each correct selection is worth one point.

Components	Answer Area						
Lease container	<table><thead><tr><th>App</th><th>Component</th></tr></thead><tbody><tr><td>App1</td><td></td></tr><tr><td>App2</td><td></td></tr></tbody></table>	App	Component	App1		App2	
App		Component					
App1							
App2							
Integrated cache							
Continuation token							

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Components	Answer Area						
Lease container	<table><thead><tr><th>App</th><th>Component</th></tr></thead><tbody><tr><td>App1</td><td>Continuation token</td></tr><tr><td>App2</td><td>Lease container</td></tr></tbody></table>	App	Component	App1	Continuation token	App2	Lease container
App		Component					
App1		Continuation token					
App2	Lease container						
Integrated cache							
Continuation token							

NEW QUESTION 67

- (Topic 8)

You deploy an API to API Management

You must secure all operations on the API by using a client certificate.

You need to secure access to the backend service of the API by using client certificates. Which two security features can you use?

- A. Azure AD token
- B. Self-signed certificate
- C. Certificate Authority (CA) certificate
- D. Triple DES (3DES) cipher
- E. Subscription key

Answer: BC

NEW QUESTION 69

HOTSPOT - (Topic 8)

You are developing a solution to store documents in Azure Blob storage. Customers upload documents to multiple containers. Documents consist of PDF, CSV, Microsoft Office format, and plain text files.

The solution must process millions of documents across hundreds of containers. The solution must meet the following requirements:

- * Document must be categorized by a customer identifier as they are uploaded to the storage account.
- * Allow filtering by the customer identifier.
- * Allow searching of information contained within a document.
- * Minimize costs.

You created and configured a standard general-purpose v2 storage account to support the solution.

You need to implement the solution.

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Solution
Search and filter by customer identifier.	<div> <div></div> <div> Azure Cognitive Search Azure Blob index tags Azure Blob inventory policy Azure Blob metadata </div> </div>
Search information inside documents.	<div> <div></div> <div> Azure Cognitive Search Azure Blob index tags Azure Blob inventory policy Azure Blob metadata </div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Azure Blob Index tags: <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-index-how-to?tabs=azure-portal>

Azure Cognitive Search: Search inside documents

NEW QUESTION 71

- (Topic 8)

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

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

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

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

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

Answer: BE

Explanation:

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

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

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

References:

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

NEW QUESTION 75

DRAG DROP - (Topic 8)

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

You develop the following code:


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

You define a role named SysAdmin.

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

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

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

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

NOTE:Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

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

Box 2:

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

NEW QUESTION 80

- (Topic 8)

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

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

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

Answer: D

Explanation:

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

Note:

Step 1: Create a pool of compute nodes. When you create a pool, you specify the number of compute nodes for the pool, their size, and the operating system.

When each task in your job runs, it's assigned to execute on one of the nodes in your pool.

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

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

NEW QUESTION 84

HOTSPOT - (Topic 8)

You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway.

Data from traffic sensors are stored in Azure Event Hub.

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

Data throughout must be maximized. Latency must be minimized. You need to implement the Azure Event Hub. Which settings should you use? To answer, select the appropriate options in the answer area.
NOTE:Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: 6
The number of partitions is specified at creation and must be between 2 and 32. There are 6 highways.
Box 2: Highway References:
<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features>

NEW QUESTION 87

- (Topic 8)
You are creating an app that will use CosmosDB for data storage. The app will process batches of relational data. You need to select an API for the app. Which API should you use?

- A. MongoDBAPI
- B. Table API
- C. SQL API
- D. Cassandra API

Answer: C

Explanation:

For relational data you will need the SQL API
Reference:
<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>

NEW QUESTION 90

- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Use the Azure Blob Storage change feed to trigger photo processing. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

The change feed is a log of changes that are organized into hourly segments but appended to and updated every few minutes. These segments are created only when there are blob change events that occur in that hour.
Instead catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.
Reference:
<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed> <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

NEW QUESTION 94

- (Topic 8)

Your company has several containers based on the following operating systems:

- Windows Server 2019 Nano Server
- Windows Server 2019 Server Core
- Windows Server 2022 Nano Server
- Windows Server 2022 Server Core
- Linux

You plan to migrate the containers to an Azure Kubernetes cluster. What is the minimum number of node pools that the cluster must have?

- A. 1
- B. 2
- C. 3
- D. 6

Answer: C

NEW QUESTION 98

DRAG DROP - (Topic 8)

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
<div>az group create</div>		
<div>az group update</div>		
<div>az webapp update</div>	⬅	⬆
<div>az webapp create</div>	➡	⬇
<div>az appservice plan create</div>		

- 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 101

DRAG DROP - (Topic 8)

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

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

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

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

NOTE: Each correct selection is worth one point.

Policy types

Inbound

Outbound

Backend

Answer Area

Requirement

Rewrite the request URL to match to the format expected by the web service.

Remove formatting text from responses.

Forward the user ID that is associated with the subscription key for the original request to the back-end service.

Policy type

policy type

policy type

policy type

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Policy types

Inbound

Outbound

Backend

Answer Area

Requirement

Rewrite the request URL to match to the format expected by the web service.

Remove formatting text from responses.

Forward the user ID that is associated with the subscription key for the original request to the back-end service.

Policy type

Outbound

Inbound

Backend

NEW QUESTION 102

HOTSPOT - (Topic 8)

You are developing an Azure Function App by using Visual Studio. The app will process orders input by an Azure Web App. The web app places the order information into Azure Queue Storage.

You need to review the Azure Function App code shown below.

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

NOTE:Each correct selection is worth one point.

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: No

ExpirationTime - The time that the message expires.

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

Box 2: Yes

maxDequeueCount - The number of times to try processing a message before moving it to the poison queue. Default value is 5.

Box 3: Yes

When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them.

By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.

Box 4: Yes References:

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

NEW QUESTION 103

DRAG DROP - (Topic 8)

A company has multiple warehouse. Each warehouse contains IoT temperature devices which deliver temperature data to an Azure Service Bus queue. You need to send email alerts to facility supervisors immediately if the temperature at a warehouse goes above or below specified threshold temperatures. Which five actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Add a logic app trigger that fires when one or more messages arrive in the queue.	
Add a Recurrence trigger that schedules the app to run every 15 minutes.	
Add an action that sends an email to specified personnel if the temperature is outside of those thresholds.	
Add a trigger that reads IoT temperature data from a Service Bus queue.	
Add a logic app action that fires when one or more messages arrive in the queue.	
Add a condition that compares the temperature against the upper and lower thresholds.	
Create a blank Logic app.	
Add an action that reads IoT temperature data from the Service Bus queue.	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Step 1: Create a blank Logic app. Create and configure a Logic App.

Step 2: Add a logical app trigger that fires when one or more messages arrive in the queue. Configure the logic app trigger.

Under Triggers, select When one or more messages arrive in a queue (auto-complete). Step 3: Add an action that reads IoT temperature data from the Service Bus queue

Step 4: Add a condition that compares the temperature against the upper and lower thresholds.

Step 5: Add an action that sends an email to specified personnel if the temperature is outside of those thresholds

NEW QUESTION 107

- (Topic 8)

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 112

- (Topic 8)

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

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

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

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

You have the following requirements:

? Queue size must not grow larger than 80 gigabytes (GB).

? Use first-in-first-out (FIFO) ordering of messages.

? Minimize Azure costs.

You need to implement the messaging solution.

Solution: Use the .Net API to add a message to an Azure Storage Queue from the mobile application. Create an Azure VM that is triggered from Azure Storage Queue events.

Does the solution meet the goal?

- A. Yes
B. No

Answer: B

Explanation:

Don't use a VM, instead create an Azure Function App that uses an Azure Service Bus Queue trigger.

Reference:

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

NEW QUESTION 113

- (Topic 8)

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

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

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

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

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

Answer: AD

Explanation:

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

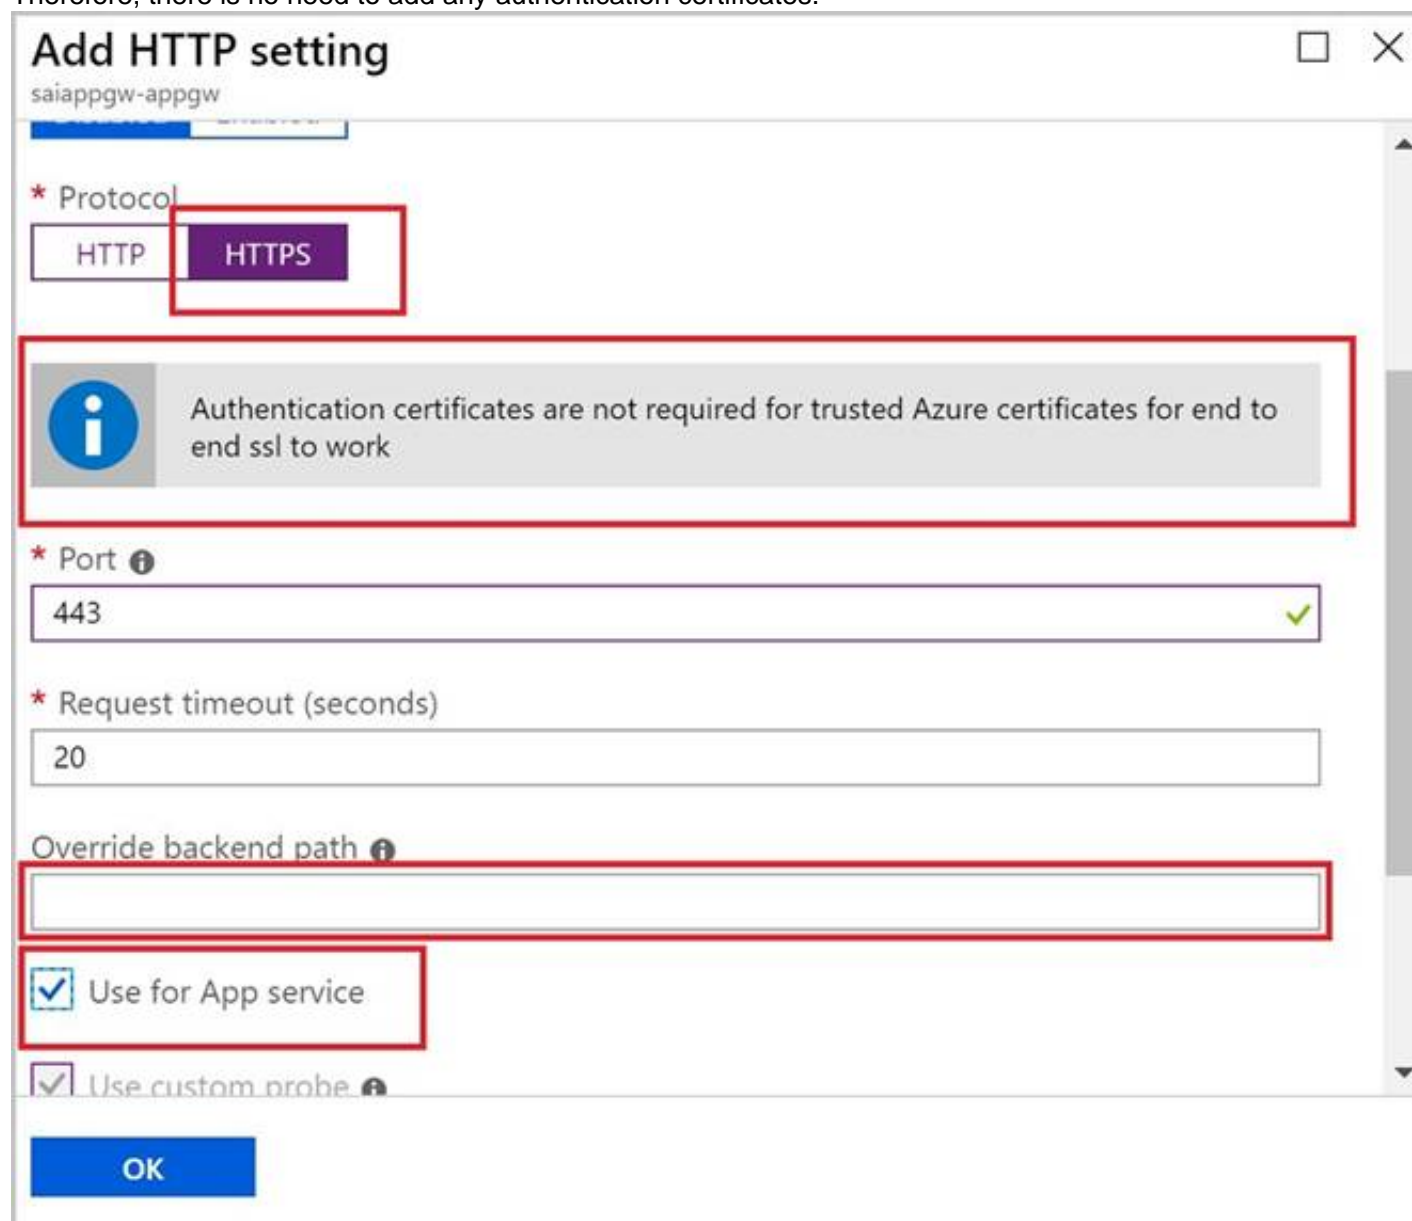
The ability to derive the host name from the IP or FQDN of the back-end pool members.

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

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

In case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway.

Therefore, there is no need to add any authentication certificates.



Reference:

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

NEW QUESTION 115

- (Topic 8)

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

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

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 bycalling 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 116

HOTSPOT - (Topic 8)

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

You have the following requirements:

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

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

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

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

PartitionKey	RowKey	Email
Harp	Walter	wharp@contoso.com
Smith	Steve	ssmith@contoso.com
Smith	Jeff	jsmith@contoso.com

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

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

NOTE:Each correct selection is worth one point.

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

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

Code for CosmosDB, example:

```
// Parse the connection string and return a reference to the storage account. CloudStorageAccount storageAccount = CloudStorageAccount.Parse(
CloudConfigurationManager.GetSetting("StorageConnectionString"));
// Create the table client.
CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
// Retrieve a reference to the table.
CloudTable table = tableClient.GetTableReference("people");
// Create the TableOperation object that inserts the customer entity. TableOperation insertOperation = TableOperation.Insert(customer1);
```

Box 2: No

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

Box 3: No

No partition key is used. Box 4: Yes

References:

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

NEW QUESTION 119

- (Topic 8)

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

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

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

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

Answer: D

Explanation:

To give a managed identity access to an Azure resource, you need to add a role to the target resource for that identity.

Note: To easily authenticate access to other resources that are protected by Azure Active Directory (Azure AD) without having to sign in and provide credentials or secrets, your logic app can use a managed identity (formerly known as Managed Service Identity or MSI). Azure manages this identity for you and helps secure your credentials because you don't have to provide or rotate secrets.

If you set up your logic app to use the system-assigned identity or a manually created, user-assigned identity, the function in your logic app can also use that same identity for authentication.

Reference:

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

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates-for-clients>

NEW QUESTION 124

- (Topic 8)

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

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

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

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

Answer: AC

NEW QUESTION 125

HOTSPOT - (Topic 8)

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

You must deploy the code to a production Azure Function App environment. You need to configure the app for deployment.

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

NOTE:Each correct selection is worth one point.

Configuration parameter	Configuration value
Publish	<div><div></div><div>Code</div><div>Docker Container</div></div>
Runtime stack	<div><div></div><div>Node.js</div><div>Python</div><div>PowerShell Core</div><div>Custom Handler</div></div>
Version	<div><div></div><div>14 LTS</div><div>7.0</div><div>custom</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Docker container

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

Box 2: PowerShell core

When creating a function app in Azure for custom handlers, we recommend you select .NET Core as the stack. A "Custom" stack for custom handlers will be added in the future. PowerShell Core (PSC) is based on the new .NET Core runtime.

Box 3: 7.0

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

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

NEW QUESTION 129

- (Topic 8)

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

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

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

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

Solution: Enable auto swap for the Testing slot. Deploy the app to the Testing slot. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

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

Reference:

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

NEW QUESTION 132

DRAG DROP - (Topic 8)

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

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

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

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

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

Commands

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

EXPOSE ./ContosoApp/ /apps/ContosoApp

COPY /.

FROM microsoft/aspnetcore:2.0

WORKDIR /apps/ContosoApp

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

Answer Area

⬅

➡

⬆

⬆

- 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 ./ContosApp/ /app/ContosoApp Step 4: CMD powershell ./setupScript.ps1
ENTRYPOINT ["dotnet", "ContosoApp.dll"]
You need to create a Dockerfile document that meets the following requirements:
? Call setupScript.ps1 when the container is built.
? Run ContosoApp.dll when the container starts.
References:
<https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker- image>

NEW QUESTION 133

DRAG DROP - (Topic 8)

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

Upgrade the storage account to GPv2

Create a new GPv2 Standard account and set its default access tier level to cool

Change the storage account access tier from hot to cool

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

Answer Area

➡

⬅

⬆

⬆

- 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.

NEW QUESTION 134

- (Topic 8)
A development team is creating a new REST API. The API will store data in Azure Blob storage. You plan to deploy the API to Azure App Service. Developers must access the Azure Blob storage account to develop the API for the next two months. The Azure Blob storage account must not be accessible by the developers after the two-month time period.
You need to grant developers access to the Azure Blob storage account. What should you do?

A. Generate a shared access signature (SAS) for the Azure Blob storage account and provide the SAS to all developers.
B. Create and apply a new lifecycle management policy to include a last accessed date value
C. Apply the policy to the Azure Blob storage account.
D. Provide all developers with the access key for the Azure Blob storage account
E. Update the API to include the Coordinated Universal Time (UTC) timestamp for the request header.
F. Grant all developers access to the Azure Blob storage account by assigning role-based access control (RBAC) roles.

Answer: A

Explanation:

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

NEW QUESTION 136

DRAG DROP - (Topic 8)
You develop a gateway solution for a public facing news API.
The news API back end is implemented as a RESTful service and hosted in an Azure App Service instance.
You need to configure back-end authentication for the API Management service instance. Which target and gateway credential type should you use? To answer, drag the appropriate values to the correct parameters. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

Azure Resource

HTTP(s) endpoint

Basic

Client cert

Configuration parameter	Value
Target	value
Gateway credentials	value

- A. Mastered
B. Not Mastered

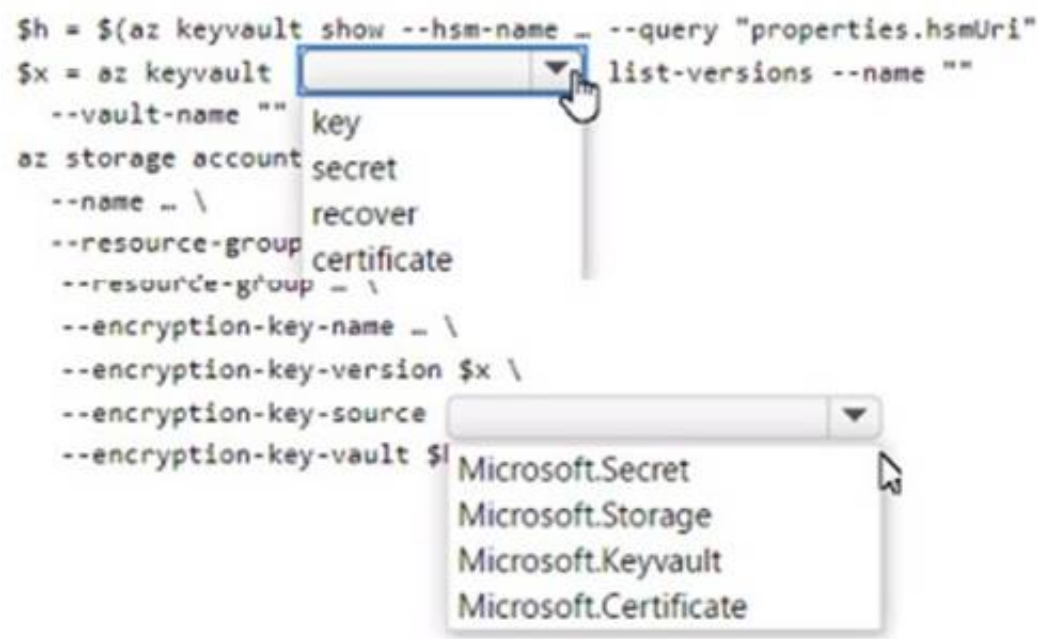
Answer: A

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-backend-entity>

NEW QUESTION 139

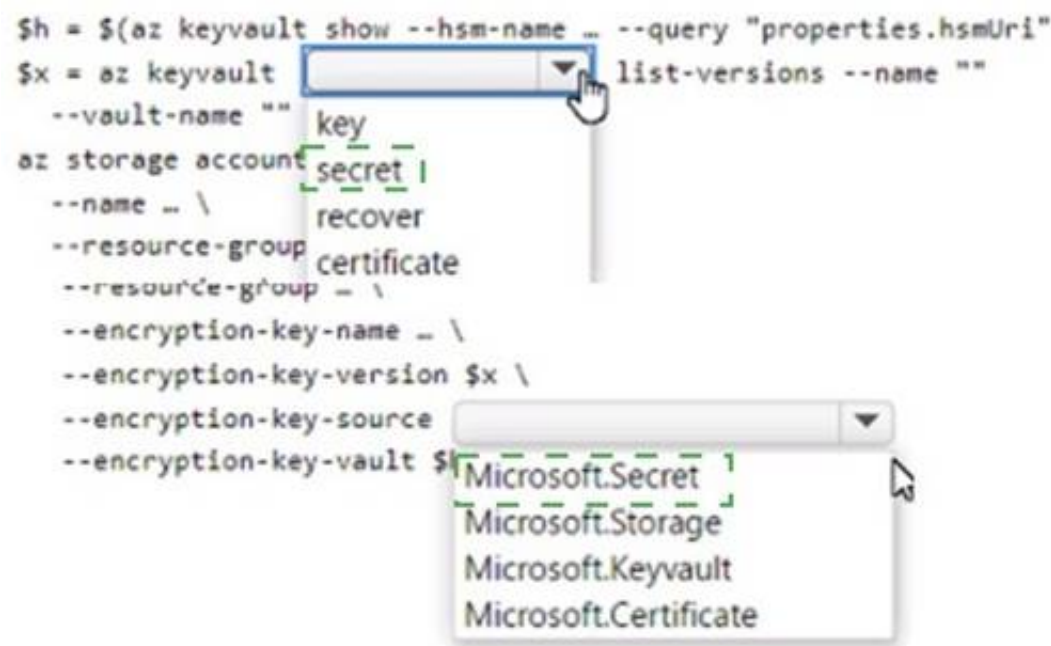
HOTSPOT - (Topic 8)
You are developing an application that uses Azure Storage to store customer data. The data must only be decrypted by the customer and the customer must be provided a script to rotate keys.
You need to provide a script to rotate keys to the customer.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 142

DRAG DROP - (Topic 8)

You are developing an application to securely transfer data between on-premises file systems and Azure Blob storage. The application stores keys, secrets, and certificates in Azure Key Vault. The application uses the Azure Key Vault APIs. The application must allow recovery of an accidental deletion of the key vault or key vault objects. Key vault objects must be retained for 90 days after deletion. You need to protect the key vault and key vault objects. Which Azure Key Vault feature should you use? To answer, drag the appropriate features to the correct actions. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE:Each correct selection is worth one point.

Features	Answer Area	
Access policy	Action	Feature
Purge protection	Enable retention period and accidental deletion.	Feature
Soft delete	Enforce retention period and accidental deletion.	Feature
Shared access signature		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Soft delete
When soft-delete is enabled, resources marked as deleted resources are retained for a specified period (90 days by default). The service further provides a mechanism for recovering the deleted object, essentially undoing the deletion.
Box 2: Purge protection
Purge protection is an optional Key Vault behavior and is not enabled by default. Purge protection can only be enabled once soft-delete is enabled. When purge protection is on, a vault or an object in the deleted state cannot be purged until the retention period has passed. Soft-deleted vaults and objects can

still be recovered, ensuring that the retention policy will be followed.

NEW QUESTION 145

HOTSPOT - (Topic 8)

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	<div><div>▼</div><div>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</div></div>
Actions section	<div><div>▼</div><div>"*/read*", "Microsoft.Support/*" "*/read" "Microsoft.Support/*" "</div></div>

- 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 146

HOTSPOT - (Topic 8)

You develop a containerized application. You plan to deploy the application to a new Azure Container instance by using a third-party continuous integration and continuous delivery (CI/CD) utility. The deployment must be unattended and include all application assets. The third-party utility must only be able to push and pull images from the registry. The authentication must be managed by Azure Active Directory (Azure AD). The solution must use the principle of least privilege. You need to ensure that the third-party utility can access the registry. Which authentication options should you use? To answer, select the appropriate options in the answer area. NOTE:Each correct selection is worth one point.

Authentication	Option
Registry authentication method	<div><div>▼</div><div>Service principal Individual identity Repository-scoped access token Managed identity for Azure resources</div></div>
RBAC role	<div><div>▼</div><div>AcrPull Owner AcrPush Contributor</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Service principal
Applications and container orchestrators can perform unattended, or "headless," authentication by using an Azure Active Directory (Azure AD) service principal. Box 2: AcrPush
AcrPush provides pull/push permissions only and meets the principle of least privilege.

NEW QUESTION 148

- (Topic 8)

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

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

Answer: D

Explanation:

Sampling is an effective way to reduce charges and stay within your monthly quota. You can set sampling manually, either in the portal on the Usage and estimated costs page; or in the ASP.NET SDK in the .config file; or in the Java SDK in the ApplicationInsights.xml file, to also reduce the network traffic. Adaptive sampling is the default for the ASP.NET SDK. Adaptive sampling automatically adjusts to the volume of telemetry that your app sends. It operates automatically in the SDK in your web app so that telemetry traffic on the network is reduced. References: <https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling>

NEW QUESTION 150

HOTSPOT - (Topic 8)

You are developing a web application that uses the Microsoft identify platform for user and resource authentication. The web application calls several REST APIs. You are implementing various authentication and authorization flows for the web application. You need to validate the claims in the authentication token. Which token type should use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Token type
Identify users for the application by using a JWT token that contains claims.	<div><div></div><div>Access</div><div>ID</div><div>Refresh</div><div>SAML</div></div>
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	<div><div>Access</div><div></div></div>
Provide the web application with long-term access to resources on behalf of users without requiring interaction with those users.	<div><div></div><div>Access</div><div>ID</div><div>Refresh</div><div>SAML</div></div>
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	<div><div></div><div>Access</div><div>ID</div><div>Refresh</div><div>SAML</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Requirement	Token type
Identify users for the application by using a JWT token that contains claims.	<div><div></div><div>Access</div><div>ID</div><div>Refresh</div><div>SAML</div></div>
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	<div><div>Access</div><div></div></div>
Provide the web application with long-term access to resources on behalf of users without requiring interaction with those users.	<div><div></div><div>Access</div><div>ID</div><div>Refresh</div><div>SAML</div></div>
Provide XML representations of claims that can be consumed by applications that use WS-Federation.	<div><div></div><div>Access</div><div>ID</div><div>Refresh</div><div>SAML</div></div>

NEW QUESTION 155

- (Topic 8)

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

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

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

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

Reference:

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

NEW QUESTION 159

- (Topic 8)

You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (AAD) for authentication. You need to implement multifactor authentication for the website.

What should you do? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. In Azure AD, create a new conditional access policy.
- B. In Azure AD, enable application proxy.
- C. Configure the website to use Azure AD B2C.
- D. In Azure AD conditional access, enable the baseline policy.
- E. Upgrade to Azure AD Premium.

Answer: AE

Explanation:

References:

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

NEW QUESTION 161

- (Topic 8)

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

After you answer a question in this 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 165

- (Topic 8)

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 170

- (Topic 8)

You have an existing Azure storage account that stores large volumes of data across multiple containers.

You need to copy all data from the existing storage account to a new storage account. The copy process must meet the following requirements:

- ? Automate data movement.
- ? Minimize user input required to perform the operation.
- ? Ensure that the data movement process is recoverable.

What should you use?

- A. AzCopy
- B. Azure Storage Explorer
- C. Azure portal
- D. .NET Storage Client Library

Answer: A

Explanation:

You can copy blobs, directories, and containers between storage accounts by using the AzCopy v10 command-line utility.

The copy operation is synchronous so when the command returns, that indicates that all files have been copied.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy>

NEW QUESTION 172

- (Topic 8)

You are developing several Azure API Management (APIM) hosted APIs.

You must inspect request processing of the APIs in APIM. Requests to APIM by using a REST client must also be included. The request inspection must include the following information:

- requests APIM sent to the API backend and the response it received
 - policies applied to the response before sending back to the caller
 - errors that occurred during the processing of the request and the policies applied to the errors
 - original request APIM received from the caller and the policies applied to the request
- You need to inspect the APIs.

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

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

Answer: ABC

Explanation:

The correct answer is A, B, and C. To inspect request processing of the APIs in APIM, you need to do the following three actions:

? Enable the Allow tracing setting for the subscription used to inspect the API. This setting allows you to trace request processing in APIM using the test console, a REST client, or a client app. You can enable this setting in the portal by selecting Subscriptions and then selecting the subscription you want to use for debugging¹.

? Add the Ocp-Apim-Trace header value to the API call with a value set to true. This header triggers tracing when making requests to APIM using a REST client or a client app. You also need to add the Ocp-Apim-Subscription-Key header value to the key for a subscription that allows access to the API¹.

? Add the Ocp-Apim-Subscription-Key header value to the key for a subscription that allows access to the API. This header authenticates your request and grants you access to the API. You can find the key for your subscription in the portal by selecting Subscriptions and then selecting Show/hide keys¹.

You do not need to create and configure a custom policy for tracing request processing. The trace policy is used to add a custom trace into the request tracing output, Application Insights telemetries, and/or resource logs². It is not required for inspecting the APIs.

NEW QUESTION 175

- (Topic 8)

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

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

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

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

Solution:

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use the following method:

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

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

NEW QUESTION 179

- (Topic 8)

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

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

You develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob. The app continues to time out after four minutes. The app must process the blob data. You need to ensure the app does not time out and processes the blob data.

Solution: Configure the app to use an App Service hosting plan and enable the Always On setting.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.

Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:
<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

NEW QUESTION 182

HOTSPOT - (Topic 8)

You are developing a solution that uses several Azure Service Bus queues. You create an Azure Event Grid subscription for the Azure Service Bus namespace. You use Azure Functions as subscribers to process the messages.

You need to emit events to Azure Event Grid from the queues. You must use principal of least privilege and minimize costs.

Which Azure Service Bus values should you use? TO answer, select the appropriate options in the answer area

Each correct selection is worth ore point

Configuration	Value
Tier	<div><div></div><div>Basic</div><div>Standard</div><div>Premium</div></div>
Access control (IAM) level	<div><div></div><div>Contributor</div><div>Data Receiver</div><div>Data Sender</div><div>Data Owner</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Configuration	Value
Tier	<div><div></div><div>Basic</div><div>Standard</div><div>Premium</div></div>
Access control (IAM) level	<div><div></div><div>Contributor</div><div>Data Receiver</div><div>Data Sender</div><div>Data Owner</div></div>

NEW QUESTION 185

HOTSPOT - (Topic 8)

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

Answer Area

delete_blob (

delete_container

delete_snapshots

snapshot_blob

snapshots_present

 -

False

Include

Only

)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

delete_blob (

delete_container

delete_snapshots

snapshot_blob

snapshots_present

 -

False

Include

Only

)

NEW QUESTION 188

- (Topic 8)

You develop Azure solutions. A .NET application needs to receive a message each time an Azure virtual machine finishes processing data. The messages must NOT persist after being processed by the receiving application. You need to implement the .NET object that will receive the messages. Which object should you use?

- A. QueueClient
- B. SubscriptionClient
- C. TopicClient
- D. CloudQueueClient

Answer: A

Explanation:

A queue allows processing of a message by a single consumer. Need a CloudQueueClient to access the Azure VM. Reference: <https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-queues-topics- subscriptions>

NEW QUESTION 191

HOTSPOT - (Topic 8)

You plan to deploy a web app to App Service on Linux. You create an App Service plan. You create and push a custom Docker image that image that contains the web app to Azure Container Registry. You need to access the console logs generated from inside the container in real-time. How should you complete the Azure CLI command? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

az webapp log

config

download

show

tail

--name ContosoWeb --resource-group ContosoDevRG

filesystem

--web-server-logging

--docker-container-logging

--application-logging

az

webapp

acr

aks

log

config

download

show

tail

--name ContosoWeb --resource-group ContosoDevRG

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: config
To Configure logging for a web app use the command: az webapp log config
Box 2: --docker-container-logging Syntax include:
az webapp log config [--docker-container-logging {filesystem, off}]
Box 3: webapp
To download a web app's log history as a zip file use the command: az webapp log download
Box 4: download References:
<https://docs.microsoft.com/en-us/cli/azure/webapp/log>

NEW QUESTION 193

DRAG DROP - (Topic 8)
You are creating a script that will run a large workload on an Azure Batch pool. Resources will be reused and do not need to be cleaned up after use.
You have the following parameters:

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

You need to write an Azure CLI script that will create the jobs, tasks, and the pool.
In which order should you arrange the commands to develop the solution? To answer, move the appropriate commands from the list of command segments to the answer area and arrange them in the correct order.

Command segments

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

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

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

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

Answer Area

<

>

↑

↓

- 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 197

DRAG DROP - (Topic 8)

You are developing several microservices named serviceA, serviceB, and serviceC. You deploy the microservices to a new Azure Container Apps environment. You have the following requirements.

- The microservices must persist data to storage.
- serviceA must persist data only visible to the current container and the storage must be restricted to the amount of disk space available in the container
- serviceB must persist data for the lifetime of the replica and allow multiple containers in the replica to mount the same storage location.
- serviceC must persist data beyond the lifetime of the replica while allowing multiple containers to access the storage and enable per object permissions.

You need to configure storage for each microservice.

Storage types

Azure Blob Storage

Azure Files storage

Ephemeral volume

Container file system

Answer Area

Microservice	Storage type
serviceA	
serviceB	
serviceC	

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Storage types

Azure Blob Storage

Azure Files storage

Ephemeral volume

Container file system

Answer Area

Microservice	Storage type
serviceA	Ephemeral volume
serviceB	Container file system
serviceC	Azure Files storage

NEW QUESTION 202

- (Topic 8)

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

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

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

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

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

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

Does the solution meet the goal?

- A. Yes
B. No

Answer: B

Explanation:

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

References:

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

NEW QUESTION 203

HOTSPOT - (Topic 8)

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

Images over 500 KB are processed by a different backend that offers a lower tier of service that costs less money. The lower tier of service is denoted by a header named `x-lsrSe- requ«st`. Images over 500 KB must never be processed by backends for smaller images and must always be charged the lower price.

You need to implement API Management policies to ensure that images are processed correctly.

How should you complete the API Management inbound policy? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

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

Hotspot 1: `<set-header name="x-large-request" exists="true" value="true"/>`
Options: delete, skip, append, delete, override

Hotspot 2: `<set-backend-service base-url="@{((large-image-host))}" />`
Options: base-url, dimension, vary-by-header, publish-to-dapr

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

Hotspot 1: `<set-header name="x-large-request" exists="true" value="true"/>`
Options: delete, skip, append, delete, override

Hotspot 2: `<set-backend-service base-url="@{((large-image-host))}" />`
Options: base-url, dimension, vary-by-header, publish-to-dapr

NEW QUESTION 208

- (Topic 8)

You are developing an Azure App Service web app.

The web app must securely store session information in Azure Redis Cache. You need to connect the web app to Azure Redis Cache.

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

Each correct selection is worth one point.

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

Answer: ACD

Explanation:

<https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-web-app-howto>

NEW QUESTION 209

DRAG DROP - (Topic 8)

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

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

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

Actions	Answer area
Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage.	
Run the Azure PowerShell command <code>Set-AzureRmVMDiskEncryptionExtension</code> .	
Run the Azure PowerShell command <code>Set-AzureRmVMOsdisk</code> .	
Encrypt the on-premises VHD by using BitLocker with a TPM. Upload the VM to Azure Storage.	
Run the Azure PowerShell command <code>New-AzureRmVm</code> .	

- 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'
```

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

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

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

Incorrect:

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

References:

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

NEW QUESTION 212

HOTSPOT - (Topic 8)

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

A rule already exists to scale up the App Service when the average queue length of unprocessed and valid queue messages is greater than 1000.

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

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

NOTE: Each correct selection is worth one point.

Answer Area

Scale rule

Metric source

Storage queue

Service Bus queue

Current resource

Storage queue (classic)

Resource type

Service Bus Namespaces

Resource

MessageQueue1103

* Queues

itemqueue

Criteria

* Metric name

Message Count

Active Message Count

* Time grain statistic

1 minute time grain

* Operator

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Service bus queue
You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.
Box 2: ActiveMessage Count
ActiveMessageCount: Messages in the queue or subscription that are in the active state and ready for delivery.
Box 3: Count
Box 4: Less than or equal to
You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.
Box 5: Decrease count by

NEW QUESTION 213

- (Topic 8)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.
When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.
You need to design the process that starts the photo processing.
Solution: Create an Azure Function app that uses the Consumption hosting model and that is triggered from the blob upload.
Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

In the Consumption hosting plan, resources are added dynamically as required by your functions.
Reference:
<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function>

NEW QUESTION 217

- (Topic 8)
Your company is developing an Azure API.
You need to implement authentication for the Azure API. You have the following requirements:
? All API calls must be secure.
? Callers to the API must not send credentials to the API.
Which authentication mechanism should you use?

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

Answer: C

Explanation:

Use the authentication-managed-identity policy to authenticate with a backend service using the managed identity of the API Management service. This policy essentially uses the managed identity to obtain an access token from Azure Active Directory for accessing the specified resource. After successfully obtaining the token, the policy will set the value of the token in the Authorization header using the Bearer scheme.
Reference:
<https://docs.microsoft.com/bs-cyrl-ba/azure/api-management/api-management-authentication-policies>

NEW QUESTION 222

DRAG DROP - (Topic 8)
You develop an Azure solution that uses Cosmos DB.
The current Cosmos DB container must be replicated and must use a partition key that is optimized for queries.
You need to implement a change feed processor solution.
Which change feed processor components should you use? To answer, drag the appropriate components to the correct requirements. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view the content.
NOTE:Each correct selection is worth one point.

Components	Requirement	Component
Host	Store the data from which the change feed is generated.	Component
Delegate	Coordinate processing of the change feed across multiple workers.	Component
Lease container	Use the change feed processor to listen for changes.	Component
Monitored container	Handle each batch of changes.	Component

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: The monitored container

The monitored container has the data from which the change feed is generated. Any inserts and updates to the monitored container are reflected in the change feed of the container.

Box 2: The lease container

The lease container acts as a state storage and coordinates processing the change feed across multiple workers. The lease container can be stored in the same account as the monitored container or in a separate account.

Box 3: The host: A host is an application instance that uses the change feed processor to listen for changes. Multiple instances with the same lease configuration can run in parallel, but each instance should have a different instance name.

Box 4: The delegate

The delegate is the code that defines what you, the developer, want to do with each batch of changes that the change feed processor reads.

NEW QUESTION 226

- (Topic 8)

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

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

You develop Azure solutions.

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

You need to obtain an Azure Resource Manager access token.

Solution: Run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Get an access token using the VM's system-assigned managed identity and use it to call Azure Resource Manager

You will need to use PowerShell in this portion.

? In the portal, navigate to Virtual Machines and go to your Windows virtual machine and in the Overview, click Connect.

? Enter in your Username and Password for which you added when you created the

Windows VM.

? Now that you have created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.

? Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.

Example:

```
$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/'  
-Method GET -Headers @{Metadata="true"}
```

Reference:

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

NEW QUESTION 227

- (Topic 8)

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

What should you do?

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

Answer: A

NEW QUESTION 228

HOTSPOT - (Topic 8)

You are developing an Azure Function app.

The Azure Function app must enable a WebHook to read an image from Azure Blob Storage and create a new Azure Cosmos DB document.

You need to implement the Azure Function app.

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

NOTE: Each correct selection is worth one point.

Answer Area

Trigger	Input binding	Output binding
Blob Storage	Blob Storage	Azure Cosmos DB
HTTP	HTTP	HTTP
Timer	Timer	Timer
Blob Storage	Blob Storage	Blob Storage
Azure Cosmos DB	Azure Cosmos DB	Azure Cosmos DB

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Trigger	Input binding	Output binding
Blob Storage	Blob Storage	Azure Cosmos DB
HTTP	HTTP	HTTP
Timer	Timer	Timer
Blob Storage	Blob Storage	Blob Storage
Azure Cosmos DB	Azure Cosmos DB	Azure Cosmos DB

NEW QUESTION 229

HOTSPOT - (Topic 8)

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

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

You need to configure the API method calls.

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

NOTE: Each correct selection is worth one point.

Code segment	Value
Attribute	<div>Authorize</div> <div>AllowAnonymous</div> <div>AutoValidateAntiforgeryToken</div>
Request Header	<div>X-MS-CLIENT-PRINCIPAL-NAME</div> <div>Proxy-Authorization</div> <div>X-Forwarded-For</div> <div>X-MS-CLIENT-PRINCIPAL-ID</div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Authorize

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

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

X-MS-CLIENT-PRINCIPAL-NAME X-MS-CLIENT-PRINCIPAL-ID

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

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

References:

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

NEW QUESTION 230

- (Topic 8)

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

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

Answer: A

NEW QUESTION 235

DRAG DROP - (Topic 8)

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

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

You need to configure back-end authentication for the API Management service instance. Which target and gateway credential type should you use? To answer, drag the appropriate

values to the correct parameters. Each value may be used once, more than once, or not at

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

NOTE:Each correct selection is worth one point.

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

NEW QUESTION 236

- (Topic 8)

You are a developing a SaaS application that stores data as key value pairs.

You must make multiple editions of the application available. In the lowest cost edition, the performance must be best-effort, and there is no regional failover.

In higher cos! editions customers must be able to select guaranteed performance and support for multiple regions. Azure costs must be minimized.

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

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

Answer: C

NEW QUESTION 238

DRAG DROP - (Topic 8)

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>	⬅
<pre>az webapp create --name \$appName --plan AppServiceLinuxDockerPlan --resource-group fourthCoffeePublicWebResourceGroup</pre>	➡
<pre>az webapp config container set --docker-custom-image-name \$dockerHibContainerPath --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 240

- (Topic 8)

You are developing a medical records document management website. The website is used to store scanned copies of patient intake forms. If the stored intake forms are downloaded from storage by a third party, the content of the forms must not be compromised.

You need to store the intake forms according to the requirements. Solution:

? uk.co.certification.simulator.questionpool.PList@2ffb4e0 Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 241

DRAG DROP - (Topic 8)

You provision virtual machines (VMs) as development environments. One VM does not have host.

The VM is stuck in a Windows update process. You attach the OS disk for the affected VM to a recovery VM.

You need to correct the issue.

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

Actions	Answer area
Open C:\temp\Patch.txt file and locate the update that is in a pending state.	1
Run the following command at an elevated command prompt:	2
dism /Image:<Attached OS disk>:\ /Remove-Package /PackageName:<PACK>	3
Run the following command at an elevated command prompt:	4
dism /Image:\ /get-packages > c:\temp\Patch.txt	
Detach the OS disk and recreate the VM.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Remove the update that causes the problem

? Take a snapshot of the OS disk of the affected VM as a backup.

? Attach the OS disk to a recovery VM.

? Once the OS disk is attached on the recovery VM, run diskmgmt.msc to open Disk Management, and ensure the attached disk is ONLINE.

? (Step 1) Open an elevated command prompt instance (Run as administrator). Run the following command to get the list of the update packages that are on the attached OS disk:

dism /Image:<Attached OS disk>:\ /get-packages > c:\temp\Patch_level

? (Step 2) Open the C:\temp\Patch_level.txt file, and then read it from the bottom up.

Locate the update that's in Install Pending or Uninstall Pending state.

? Remove the update that caused the problem:

dism /Image:<Attached OS disk>:\ /Remove-Package /PackageName:<PACK>

? (Step 4) Detach the OS disk and recreate the VM. Then check whether the issue is resolved.

NEW QUESTION 246

- (Topic 8)

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 249

DRAG DROP - (Topic 8)

A company backs up all manufacturing data to Azure Blob Storage. Admins move blobs from hot storage to archive tier storage every month.

You must automatically move blocks to Archive tier after they have not been accessed for 180 days. The path for any item that is not archived must be placed in an existing queue. This operation must be performed automatically once a month. You set the value of TierAgeInDays to 180.

How should you configure the Logic App? To answer, drag the appropriate triggers or action blocks to the correct trigger or action slots. Each trigger or action block may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Triggers and Action Blocks

Insert Entity

*Table: processing

*Entity: Path

Show advanced options

Tier blob

If blob is older than the defined value, tier it to Cool or Archive tier

*Blob path: Path

*Blob Tier: Archive

When there are messages in a queue

*Queue Name: processing

Show advanced options

Connected to tableStorageAccountConnection. Change connection.

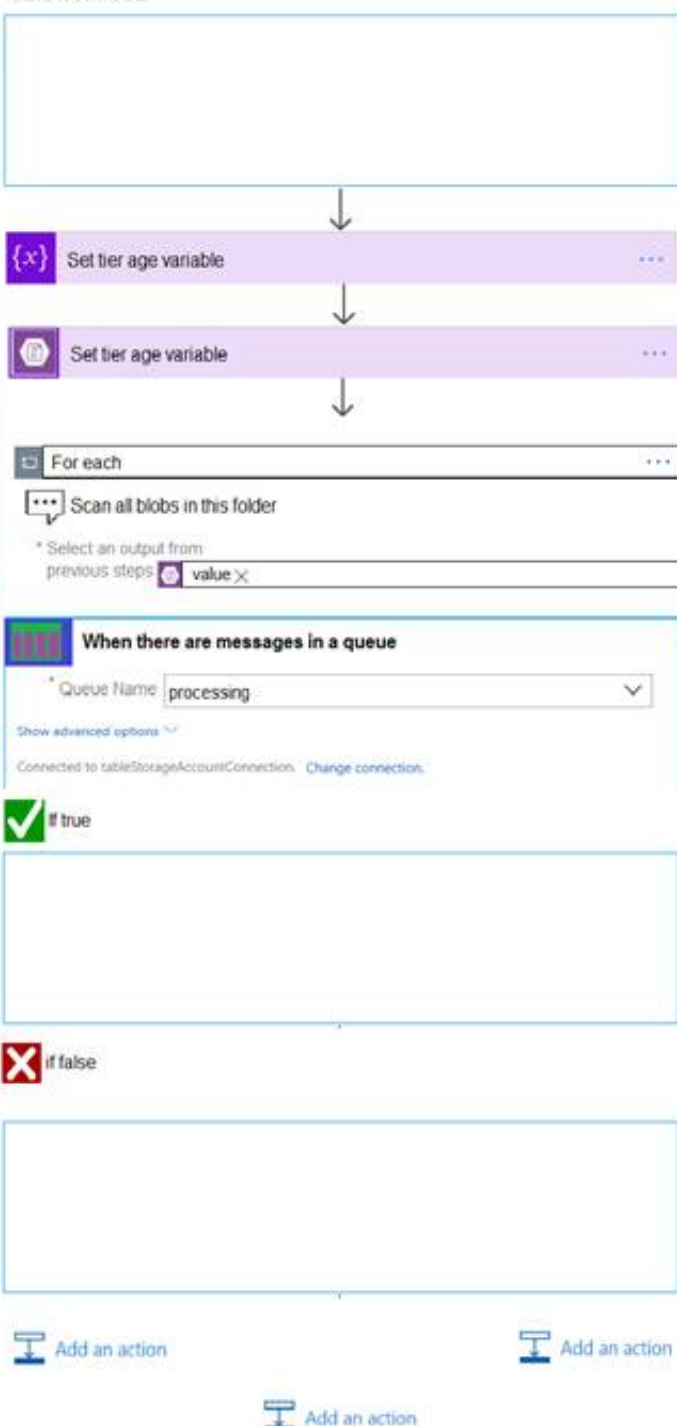
Recurrence

*Interval: 1

*Frequency: Month

Show advanced options

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Recurrence Box 2: Insert Entity

Box 3 (if true): Tier Blob

Box 4: (if false): Leave blank.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations>

NEW QUESTION 252

- (Topic 8)

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

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

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

You create the index in Azure Search.

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

- Solution:
- * 1. Create a SearchServiceClient object to connect to the search index.
 - * 2. Create a DataContainer that contains the documents which must be added.
 - * 3. Create a DataSource instance and set its Container property to the DataContainer.
 - * 4. Set the DataSources property of the SearchServiceClient.
- Does the solution meet the goal?

- A. Yes
B. No

Answer: B

Explanation:

Use the following method:

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

References:

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

NEW QUESTION 253

HOTSPOT - (Topic 8)

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

▼
            

GetJob  

                GetTask  

                EnableJob  

                CreateJob

 ();

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

▼
                    

TaskSuccess  

                        TaskFailure  

                        TaskCompletion

 ));

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

▼
                    

TaskSuccess  

                        TaskFailure  

                        TaskCompletion

 ));

            task.

▼
                

OutputFiles  

                    FilesToStage  

                    ResourceFiles  

                    StageFiles

 =outputFileList;

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

- 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 255

- (Topic 8)

You manage a data processing application that receives requests from an Azure Storage queue.

You need to manage access to the queue. You have the following requirements:

? Provide other applications access to the Azure queue.

? Ensure that you can revoke access to the queue without having to regenerate the storage account keys.

? Specify access at the queue level and not at the storage account level.

Which type of shared access signature (SAS) should you use?

- A. Service SAS with a stored access policy
 B. Account SAS
 C. User Delegation SAS
 D. Service SAS with ad hoc SAS

Answer: A

Explanation:

A service SAS is secured with the storage account key. A service SAS delegates access to a resource in only one of the Azure Storage services: Blob storage, Queue storage, Table storage, or Azure Files.

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

Reference:

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

NEW QUESTION 257

- (Topic 8)

You develop a Python application for image rendering that uses GPU resources to optimize rendering processes. You deploy the application to an Azure Container Instances (ACI) Linux container.

The application requires a secret value to be passed when the container is started. The value must only be accessed from within the container.

You need to pass the secret value.

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

NOTE: Each correct selection is worth one point.

- A. Create an environment variable Set the secureValue property to the secret value.
- B. Add the secret value to the container image
- C. Use a managed identity.
- D. Add the secret value to the application code Set the container startup command.
- E. Add the secret value to an Azure Blob storage account
- F. Generate a SAS token.
- G. Mount a secret volume containing the secret value in a secrets file.

Answer: AE

Explanation:

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

NEW QUESTION 259

- (Topic 8)

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

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

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

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

You need to design the process that starts the photo processing. Solution: Trigger the photo processing from Blob storage events. Does the solution meet the goal?

- A. Yes
- B. NO

Answer: B

Explanation:

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

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

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

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

Reference:

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

NEW QUESTION 263

- (Topic 8)

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

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

You need to select a partition key.

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

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

Answer: AC

Explanation:

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

Another possible strategy to distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in

this way, you can perform parallel write operations across partitions.
Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.
References:
<https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys>

NEW QUESTION 266

HOTSPOT - (Topic 8)

You are authoring a set of nested Azure Resource Manager templates to deploy Azure resources. You author an Azure Resource Manager template named mainTemplate.json that contains the following linked templates: linkedTemplate1.json, linkedTemplate2.json.
You add parameters to a parameters template file named mainTemplate.parameters.json. You save all templates on a local device in the C:\templates\ folder. You have the following requirements:

- Store the templates in Azure for later deployment.
- Enable versioning of the templates.
- Manage access to the templates by using Azure RBAC

You need to store the templates in Azure.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
az ts create \
--template-file "C:\templates\mainTemplate.json" --tags Dept=HumanResource:linkedTemplate1.json linkedTemplate2.json mainTemplate.parameters.json
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
az ts create \
--template-file "C:\templates\mainTemplate.json" --tags Dept=HumanResource:linkedTemplate1.json linkedTemplate2.json mainTemplate.parameters.json
```

NEW QUESTION 268

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AZ-204 Practice Exam Features:

- * AZ-204 Questions and Answers Updated Frequently
- * AZ-204 Practice Questions Verified by Expert Senior Certified Staff
- * AZ-204 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AZ-204 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AZ-204 Practice Test Here](#)