



Microsoft

Exam Questions AZ-204

Developing Solutions for Microsoft Azure

About ExamBible

Your Partner of IT Exam

Found in 1998

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

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

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

NEW QUESTION 1

- (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 2

- (Topic 8)

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

You need to complete the configuration.

Which Azure CLI or PowerShell command should you run?

A)

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

B)

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

C)

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

D)

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

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

Answer: C

NEW QUESTION 3

- (Topic 8)

You develop and deploy a web application to Azure App Service. The application accesses data stored in an Azure Storage account. The account contains several containers with several blobs with large amounts of data. You deploy all Azure resources to a single region.

You need to move the Azure Storage account to the new region. You must copy all data to the new region.

What should you do first?

- A. Export the Azure Storage account Azure Resource Manager template
- B. Initiate a storage account failover
- C. Configure object replication for all blobs
- D. Use the AzCopy command line tool
- E. Create a new Azure Storage account in the current region
- F. Create a new subscription in the current region

Answer: A

Explanation:

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

To get started, export, and then modify a Resource Manager template.

Reference:

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

NEW QUESTION 4

- (Topic 8)

You are developing an Azure Durable Function to manage an online ordering process. The process must call an external API to gather product discount information.

You need to implement Azure Durable Function.

Which Azure Durable Function types should you use? Each correct answer presents part of the solution

NOTE: Each correct selection is worth one point

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

Answer: AB

Explanation:

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview>

NEW QUESTION 5

- (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 to Application Insights.
- D. Add a new diagnostic setting to the Azure Function app
- 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.

D18912E1457D5D1DDCDB40AB3BF70D5D

Reference:

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

NEW QUESTION 6

HOTSPOT - (Topic 8)

You are developing a web application that will use Azure Storage. Older data will be less frequently used than more recent data.

You need to configure data storage for the application. You have the following requirements:

- ? Retain copies of data for five years.
- ? Minimize costs associated with storing data that is over one year old.
- ? Implement Zone Redundant Storage for application data.

What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Requirement	Solution
Configure an Azure Storage account	<div style="border: 1px solid black; padding: 5px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> Implement Blob Storage Implement Azure Cosmos DB Implement Storage (general purpose v1) Implement StorageV2 (general purpose v2) </div>
Configure data retention	<div style="border: 1px solid black; padding: 5px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> Snapshot blobs and move them to the archive tier Set a lifecycle management policy to move blobs to the cool tier Use AzCopy to copy the data to an on-premises device for backup Set a lifecycle management policy to move blobs to the archive tier </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Requirement	Solution
Configure an Azure Storage account	<ul style="list-style-type: none"> Implement Blob Storage Implement Azure Cosmos DB Implement Storage (general purpose v1) Implement StorageV2 (general purpose v2)
Configure data retention	<ul style="list-style-type: none"> Snapshot blobs and move them to the archive tier Set a lifecycle management policy to move blobs to the cool tier Use AzCopy to copy the data to an on-premises device for backup Set a lifecycle management policy to move blobs to the archive tier

NEW QUESTION 7

HOTSPOT - (Topic 8)

You provisioned an Azure Cosmos DB for NoSQL account named account1 with the default consistency level.

You plan to configure the consistency level on a per request basis. The level needs to be set for consistent prefix for read and write operations to account1.

You need to identify the resulting consistency level for read and write operations. Which levels should you configure? To answer, select the appropriate options in the

answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Operation type	Resulting consistency level
Read operations	<ul style="list-style-type: none"> strong session consistent prefix
Write operations	<ul style="list-style-type: none"> strong session consistent prefix

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Operation type	Resulting consistency level
Read operations	<ul style="list-style-type: none"> strong session consistent prefix
Write operations	<ul style="list-style-type: none"> strong session consistent prefix

NEW QUESTION 8

- (Topic 8)

You are developing an ASP.NET Core website that uses Azure FrontDoor. The website is used to build custom weather data sets for researchers. Data sets are downloaded by users as Comma Separated Value (CSV) files. The data is refreshed every 10 hours.

Specific files must be purged from the FrontDoor cache based upon Response Header values.

You need to purge individual assets from the Front Door cache. Which type of cache purge should you use?

- A. single path
- B. wildcard
- C. root domain

Answer: A

Explanation:

These formats are supported in the lists of paths to purge:

? Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example, /pictures/strasbourg.png;

? Wildcard purge: Asterisk (*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /*, for example, /pictures/*.

? Root domain purge: Purge the root of the endpoint with "/" in the path.

Reference:

<https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching>

NEW QUESTION 9

- (Topic 8)

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

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

You need to ensure that the subscription client processes all messages. Which code segment should you use?

- A. `await subscriptionClient.AddRuleAsync(new RuleDescription (RuleDescription.DefaultRuleName, new TrueFilter()));`
- B. `subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName); D18912E1457D5D1DDCCBD40AB3BF70D5D`
- C. `await subscriptionClient.CloseAsync();`
- D. `subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);`

Answer: D

Explanation:

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

`subscriptionClient.RegisterMessageHandler(ReceiveMessagesAsync, messageHandlerOptions);`

Reference:

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

NEW QUESTION 10

DRAG DROP - (Topic 8)


You develop and deploy a Java application to Azure. The application has been instrumented by using the Application Insights SDK.

The telemetry data must be enriched and processed before it is sent to the Application Insights service.

You need to modify the telemetry data.

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


NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 10

DRAG DROP - (Topic 8)

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

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

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

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

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: user_impersonation

Box 2: delegated Example:

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

Box 3: delegated Example

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

NEW QUESTION 14

HOTSPOT - (Topic 8)

You have an Azure Web app that uses Cosmos DB as a data store. You create a CosmosDB container by running the following PowerShell script:

```
$resourceGroupName = "testResourceGroup"
$accountName = "testCosmosAccount"
$databaseName = "testDatabase"
$containerName = "testContainer"
$partitionKeyPath = "/EmployeeId"
$autoscaleMaxThroughput = 5000
New-AzCosmosDBSqlContainer -ResourceGroupName $resourceGroupName -AccountName $accountName -DatabaseName $databaseName -Name $containerName -PartitionKeyKind Hash -PartitionKeyPath $partitionKeyPath -AutoscaleMaxThroughput $autoscaleMaxThroughput
```

You create the following queries that target the container:

```
SELECT * FROM c WHERE c.EmployeeId > '12345'
SELECT * FROM c WHERE c.UserID = '12345'
```

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

NOTE: Each correct selection is worth one point.

	Yes	No
The minimum throughput for the container is 400 R/Us.	<input type="radio"/>	<input type="radio"/>
The first query statement is an in-partition query.	<input type="radio"/>	<input type="radio"/>
The second query statement is a cross-partition query.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

You set the highest, or maximum RU/s Tmax you don't want the system to exceed. The system automatically scales the throughput T such that $0.1 * Tmax \leq T \leq Tmax$.

In this example we have autoscaleMaxThroughput = 5000, so the minimum throughput for the container is 500 R/Us.

Box 2: No

First query: SELECT * FROM c WHERE c.EmployeeId > '12345'

Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition. In order to be an in-partition query, the query must have an equality filter that includes the partition key:

```
SELECT * FROM c WHERE c.DeviceId > 'XMS-0001'
```

Box 3: Yes

Example of In-partition query:

Consider the below query with an equality filter on DeviceId. If we run this query on a container partitioned on DeviceId, this query will filter to a single physical partition.

```
SELECT * FROM c WHERE c.DeviceId = 'XMS-0001'
```

NEW QUESTION 19

- (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 24

- (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 28

- (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 policie
- 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 30

- (Topic 8)

You are developing an application that uses Azure Blob storage.

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

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

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

Answer: B

Explanation:

Change feed support in Azure Blob Storage

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

Reference:

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

NEW QUESTION 34

- (Topic 8)

You are developing a solution that will use a multi-partitioned Azure Cosmos DB database. You plan to use the latest Azure Cosmos DB SDK for development.

The solution must meet the following requirements:

? Send insert and update operations to an Azure Blob storage account.

? Process changes to all partitions immediately.

? Allow parallelization of change processing.

You need to process the Azure Cosmos DB operations.

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

NOTE: Each correct selection is worth one point.

A. Create an Azure App Service API and implement the change feed estimator of the SD

B. Scale the API by using multiple Azure App Service instances.

C. Create a background job in an Azure Kubernetes Service and implement the change feed feature of the SDK.

D. Create an Azure Function to use a trigger for Azure Cosmos D

E. Configure the trigger to connect to the container.

F. Create an Azure Function that uses a FeedIterator object that processes the change feed by using the pull model on the container.

G. Use a FeedRange object to parallelize the processing of the change feed across multiple functions.

Answer: CD

Explanation:

Azure Functions is the simplest option if you are just getting started using the change feed. Due to its simplicity, it is also the recommended option for most change feed use cases. When you create an Azure Functions trigger for Azure Cosmos DB, you select the container to connect, and the Azure Function gets triggered whenever there is a change in the container. Because Azure Functions uses the change feed processor behind the scenes, it automatically parallelizes change processing across your container's partitions.

Note: You can work with change feed using the following options:

? Using change feed with Azure Functions

? Using change feed with change feed processor

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed>

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model> <https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed#azure-functions>

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model#using-feedrange-for-parallelization>

NEW QUESTION 37

- (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 40

- (Topic 8)

Your company purchases an Azure subscription and plans to migrate several on-premises virtual machines to Azure. You need to design the infrastructure required (or the Azure virtual machines solution). What should you include in the design?

A. the number of Azure Storage accounts

B. the settings of the Azure virtual networks

C. the size of the virtual machines

D. the number of Azure regions

Answer: C

NEW QUESTION 41

HOTSPOT - (Topic 8)

You are developing an ASP.NET Core app that includes feature flags which are managed by Azure App Configuration. You create an Azure App Configuration store named AppfeatureflagStore as shown in the exhibit:

Key	Label	State	Description	Last modified
Export	Export	<input checked="" type="checkbox"/> Off <input type="checkbox"/> On	Ability to export data.	6/11/2020, 9:13:26 ... ***

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

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

You went to update the app to use the feature flag.

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

NOTE: Each correct selection is worth one point.

Answer Area

Code section	Value
Controller attribute	<input type="checkbox"/> FeatureGate <input type="checkbox"/> Route <input type="checkbox"/> ServiceFilter <input type="checkbox"/> TypeFilter
Startup method	<input type="checkbox"/> AddAzureAppConfiguration <input type="checkbox"/> AddControllersWithViews <input type="checkbox"/> AddUserSecrets
AppConfig endpoint setting	<input type="checkbox"/> https://appfeatureflagstore.azureconfig.io <input type="checkbox"/> https://appfeatureflagstore.vault.azure.net <input type="checkbox"/> https://export.azureconfig.io <input type="checkbox"/> https://export.vault.azure.net

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: FeatureGate

You can use the FeatureGate attribute to control whether a whole controller class or a specific action is enabled.

Box 2: AddAzureAppConfiguration

The extension method AddAzureAppConfiguration is used to add the Azure App Configuration Provider.

Box 3: https://appfeatureflagstore.azureconfig.io

You need to request the access token with resource=https://<yourstorename>.azureconfig.io


NEW QUESTION 44

HOTSPOT - (Topic 8)

You are developing an application that needs access to an Azure virtual machine (VM). The access lifecycle for the application must be associated with the VM service instance. You need to enable managed identity for the VM.

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

NOTE Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

\$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM Update-AzVM -ResourceGroupName myResourceGroup -VM \$vm -

AssignIdentity:\$SystemAssigned

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

NEW QUESTION 49

- (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 50

HOTSPOT - (Topic 8)

You are developing an ASP.NET Core time sheet application that runs as an Azure Web App. Users of the application enter their time sheet information on the first day of every month.

The application uses a third-party web service to validate data.

The application encounters periodic server errors due to errors that result from calling a third-party web server. Each request to the third-party server has the same chance of failure.

You need to configure an Azure Monitor alert to detect server errors unrelated to the third-party service. You must minimize false-positive alerts.

How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

"type": "Microsoft.Insights/metricAlerts",
"properties": {
  "criteria": {
    "odata.type": "...",
    "allOf": [
      {
        "criterionType": "
        "metricName": "
        "alertSensitivity": "
      }
    ]
  }
}

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: DynamicThresholdCriterion

Box 2: Http5xx

Server errors are in the 5xx range. Client errors are in the 4xx range

Box 3: Low

NEW QUESTION 53

- (Topic 8)

You are 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 cost 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: D

NEW QUESTION 58

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 style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>2</p> <p>4</p> <p>8</p> <p>16</p> </div> </div>
Pricing tier	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Isolated</p> <p>Standard</p> <p>Premium</p> <p>Consumption</p> </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 59

- (Topic 8)

You provide an Azure API Management managed web service to clients. The back end web service implements HTTP Strict Transport Security (HSTS).

Every request to the backend service must include a valid HTTP authorization header. You need to configure the Azure API Management instance with an authentication policy. Which two policies can you use? Each correct answer presents a complete solution NOTE: Each correct selection is worth one point.

- A. Certificate Authentication
- B. Basic Authentication
- C. OAuth Client Credential Grant
- D. Digest Authentication

Answer: AB

NEW QUESTION 60

- (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 62

- (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 64

- (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 SearchIndexClient object to connect to the search index.
- * 2. Create a DataContainer that contains the documents which must be added.
- * 3. Create a DataSource instance and set its Container property to the DataContainer.
- * 4. Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use the following method:

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

References:

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

NEW QUESTION 65

- (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 70

- (Topic 8)

You are a developer for a SaaS company that offers many web services. All web services for the company must meet the following requirements:

? Use API Management to access the services

? Use OpenID Connect for authentication

? Prevent anonymous usage

A recent security audit found that several web services can be called without any authentication.

Which API Management policy should you implement?

- A. jsonp

- B. authentication-certificate
- C. check-header
- D. validate-jwt

Answer: D

Explanation:

Add the validate-jwt policy to validate the OAuth token for every incoming request. Reference:
<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad>

NEW QUESTION 74

HOTSPOT - (Topic 8)

You are developing an application that uses Azure Storage Queues. You have the following code:

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

CloudQueue queue = queueClient.GetQueueReference("appqueue") ;
await queue.CreateIfNotExistsAsync() ;

CloudQueueMessage peekedMessage = await queue.PeekMessageAsync() ;
if (peekedMessage != null)
{
    Console.WriteLine("The peeked message is: {0}", peekedMessage.AsString);
}
CloudQueueMessage message = await queue.GetMessageAsync() ;
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
 NOTE: Each correct selection is worth one point.

Statement	Yes	No
The code configures the lock duration for the queue.	<input type="radio"/>	<input type="radio"/>
The last message read remains in the queue after the code runs.	<input type="radio"/>	<input type="radio"/>
The storage queue remains in the storage account after the code runs.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No
 The QueueDescription.LockDuration property gets or sets the duration of a peek lock; that is, the amount of time that the message is locked for other receivers. The maximum value for LockDuration is 5 minutes; the default value is 1 minute.
 Box 2: Yes
 You can peek at the message in the front of a queue without removing it from the queue by calling the PeekMessage method.
 Box 3: Yes

NEW QUESTION 78

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.

Code segments	Answer Area
raw	<pre>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);</pre>
windows	
windowsphone	
application/xml	
application/json	
application/octet-stream	

- 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 82

- (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 85

- (Topic 8)

You deploy an Azure App Service web app. You create an app registration for the app in Azure Active Directory (Azure AD) and Twitter. the app must authenticate users and must use SSL for all communications. The app must use Twitter as the identity provider. You need to validate the Azure AD request in the app code. What should you validate?

- A. HTTP response code
- B. ID token header
- C. ID token signature
- D. Tenant ID

Answer: B

NEW QUESTION 88

HOTSPOT - (Topic 8)

You develop several Azure Grid to include hundreds of event types, such as billing, inventory, and shipping updates.

Events must be sent to a single endpoint for the Azure Functions app to process. The events must be filtered by event type before processing. You must have authorization and authentication control to partition your tenants to receive the event data.

You need to configure Azure Event Grid.

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

NOTE: Each correct selection is worth one point.

Answer Area

Requirement	Configuration Value
Third-party system endpoint to send events	<ul style="list-style-type: none"> system topic system topic custom topic event domain event subscription
Azure Functions app endpoint to handle filtered events	<ul style="list-style-type: none"> event domain system topic custom topic event domain event subscription

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Requirement	Configuration Value
Third-party system endpoint to send events	<ul style="list-style-type: none"> system topic system topic custom topic event domain event subscription
Azure Functions app endpoint to handle filtered events	<ul style="list-style-type: none"> event domain system topic custom topic event domain event subscription

NEW QUESTION 89

- (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/azurermsql/new-azurermsqldatabasedatamaskingrule?view=azurermps-6.13.0>

NEW QUESTION 94

- (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. Determine whether the solution meets the stated goals.

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead use Azure Cache for Redis.

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

References:

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

NEW QUESTION 95

- (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 Windows VM that is triggered from Azure Service Bus Queue.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 98

- (Topic 8)

You are developing a web application that uses the Microsoft identity platform to authenticate users and resources, The web application calls several REST APIs. The APIs require an access token from the Microsoft identity platform. You need to request a token.

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

A. Application name

B. Application secret

C. Application ID

D. Supported account type

E. Redirect URI/URL

Answer: ABC

NEW QUESTION 101

- (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 102

HOTSPOT - (Topic 8)

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

The solution requires a mobile app that stores user information using Azure Table Storage. You need to develop code that can insert multiple sets of user information.

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

NOTE: Each correct selection is worth one point.

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

op = new

TableOperation
TableBatchOperaton
TableEntity
TableQuery

 ();

table.

ExecuteBatch
Execute
Insert
InsertOrMerge

 (op);

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1, Box 2: TableBatchOperation Create the batch operation.

TableBatchOperation op = new TableBatchOperation();

Box 3: ExecuteBatch

/ Execute the batch operation. table.ExecuteBatch(op);

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

You can perform updates, deletes, and inserts in the same single batch operation. A single batch operation can include up to 100 entities.

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

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

References:

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

NEW QUESTION 105

FILL IN THE BLANK - (Topic 8)

You are developing a web application by using the Azure SDK. The web application accesses data in a zone-redundant BlockBlobStorage storage account. The application must determine whether the data has changed since the application last read the data. Update operations must use the latest data changes when writing data to the storages.....

You need to implement the update operations.

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

NOTE Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Code evaluation	Value
HTTP Header value	Versionid
Conditional header	If-Match

NEW QUESTION 107

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 throughput must be maximized. Latency must be minimized. You need to implement the Azure Event Hub.

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

NOTE:Each correct selection is worth one point.

Setting	Value
Number of partitions	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>3</p> <p>4</p> <p>6</p> <p>12</p> </div> </div>
Partition Key	<div style="border: 1px solid black; padding: 2px;"> <div style="background-color: #ccc; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>Highway</p> <p>Department</p> <p>Timestamp</p> <p>VM name</p> </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 110

- (Topic 8)
 You are developing several microservices to deploy to a Azure Service cluster. The microservices manage data stored in Azure Cosmos DB and Azure Blob storage. The data is secured by using customer-managed keys stored in Aue Key Vault.
 You must automate key rotation for all Key Vault keys and allow for manual key rotation. Keys must rotate every three months. Notifications Of expiring keys must be sent before key expiry.
 You need to configure key rotation and enable key expiry notifications.
 Which two actions should you perform? Each correct answer presents part Of solution. NOTE: Each correct selection is worth

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


Answer: BD

Explanation:

<https://learn.microsoft.com/en-us/azure/key-vault/keys/how-to-configure-key-rotation>


NEW QUESTION 114

HOTSPOT - (Topic 8)
 You are validating the configuration of an Azure Search indexer.
 The service has been configured with an indexer that uses the Import Data option. The index is configured using options as shown in the Index Configuration exhibit. (Click the Index Configuration tab.)



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

CategoryKey	RowKey	Timestamp	InStockCount	ItemDescription	ItemName	LocationRow	LocationShelf	SKU
Food	3	2018-08-25T15:47:26.135Z	12	A box of chocolate candy bars	Choco-bar	3	123421	
Hardware	2	2018-08-25T15:46:08.409Z	2	A bag of bolts	Bolts	1	4	678944
Hardware	1	2018-08-25T15:46:41.402Z	23	A box of nails	Nails	2	1	654365



When users search for items by description, Search explorer returns no records. The Search Explorer exhibit shows the query and results for a test. In the test, a user is trying to search for all items in the table that have a description that contains the word bag. (Click the Search Explorer tab.) You need to resolve the issue.

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

NOTE: Each correct selection is worth one point.

	Yes	No
You can resolve the issue by recreating the search index with the same settings for all fields except ItemDescription. Select the SEARCHABLE option for this field	<input type="radio"/>	<input type="radio"/>
You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field.	<input type="radio"/>	<input type="radio"/>
You can resolve the issue by running the indexer.	<input type="radio"/>	<input type="radio"/>
You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

The ItemDescription field is not searchable.

Box 2: No

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

Box 3: Yes

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

Box 4: No References:

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

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

NEW QUESTION 119

DRAG DROP - (Topic 8)

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

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

The policy fragment must be reused across various scopes and APIs. The policy fragment must be applied to all APIs and run when a calling system invokes any API.

You need to implement the policy fragment.

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

NOTE: Each correct selection is worth one point.

XML elements	Answer Area
name	<pre><policies> <[]> <[] = "APICounts" /> <base /> </[]> . . . </policies></pre>
inbound	
outbound	
set-variable	
fragment-id	
include-fragment	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<https://learn.microsoft.com/en-us/azure/api-management/include-fragment-policy>

NEW QUESTION 124

- (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 127

DRAG DROP - (Topic 8)

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

Property	Description
ShipLocation	the country/region where the order will be shipped
CorrelationId	a priority value for the order
Quantity	a user-defined field that stores the quantity of items in an order
AuditedAt	a user-defined field that records the date an order is audited

The system has the following requirements for subscriptions

Subscription type	Comments
FutureOrders	This subscription is reserved for future use and must not receive any orders.
HighPriorityOrders	Handle all high priority orders and International orders.
InternationalOrders	Handle orders where the country/region is not United States.
HighQuantityOrders	Handle only orders with quantities greater than 100 units.
AllOrders	This subscription is used for auditing purposes. This subscription must receive every single order. AllOrders has an Action defined that updates the AuditedAt property to include the date and time it was received by the subscription.

You need to implement filtering and maximize throughput while evaluating filters.

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

NOTE: Each correct selection is worth one point.

Filter types

Answer Area

- SQLFilter
- CorrelationFilter
- No Filter

Subscription	Filter type
FutureOrders	<input type="text"/>
HighPriorityOrders	<input type="text"/>
InternationalOrders	<input type="text"/>
HighQuantityOrders	<input type="text"/>
AllOrders	<input type="text"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

FutureOrders: SQLFilter HighPriorityOrders: CorrelationFilter

CorrelationID only

InternationalOrders: SQLFilter

Country NOT USA requires an SQL Filter

HighQuantityOrders: SQLFilter

Need to use relational operators so an SQL Filter is needed. AllOrders: No Filter

SQL Filter: SQL Filters - A SqlFilter holds a SQL-like conditional expression that is evaluated in the broker against the arriving messages' user-defined properties and system properties. All system properties must be prefixed with sys. in the conditional expression. The SQL-language subset for filter conditions tests for the existence of properties (EXISTS), as well as for null-values (IS NULL), logical NOT/AND/OR, relational operators, simple numeric arithmetic, and simple text pattern matching with LIKE.

Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties. A common use is to match against the CorrelationId property, but the application can also choose to match against ContentType, Label, MessageId, ReplyTo, ReplyToSessionId, SessionId, To, and any user-defined properties. A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter. For string expressions, the comparison is case-sensitive. When specifying multiple match properties, the filter combines them as a logical AND condition, meaning for the filter to match, all conditions must match.

Boolean filters - The TrueFilter and FalseFilter either cause all arriving messages (true) or none of the arriving messages (false) to be selected for the subscription.

References:

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

NEW QUESTION 131

- (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 by calling Indexes.GetClient on the SearchServiceClient. Here we will use the first method.

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

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

```
var batch = IndexBatch.Upload(hotels);
```

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

```
var batch = IndexBatch.Upload(hotels);
```

```
try  
{  
    indexClient.Documents.Index(batch);  
}
```

References:

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

NEW QUESTION 132

DRAG DROP - (Topic 8)

You develop an application. You plan to host the application on a set of virtual machines (VMs) in Azure.

You need to configure Azure Monitor to collect logs from the application.

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

Actions	Answer Area
Create a Log Analytics workspace.	
Install agents on the VM and VM scale set to be monitored.	
Send console logs.	
Add a VMInsights solution.	
Create an Application Insights resource.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Create a Log Analytics workspace. First create the workspace.

Application Insights

Monitor web app performance and usage

Basics Tags Review + create

Create an Application Insights resource to monitor your live web application. With Application Insights, you have full observability into your application across all components and dependencies of your complex distributed architecture. It includes powerful analytics tools to help you diagnose issues and to understand what users actually do with your app. It's designed to help you continuously improve performance and usability. It works for apps on a wide variety of platforms including .NET, Node.js and Java EE, hosted on-premises, hybrid, or any public cloud. [Learn More](#)

PROJECT DETAILS

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

Subscription *

Resource Group * [Create new](#)

INSTANCE DETAILS

Name *

Region *

Resource Mode * Classic **Workspace-based**

WORKSPACE DETAILS

Subscription *

Log Analytics Workspace *

[Review + create](#) [« Previous](#) [Next : Tags >](#)

Step 2: Add a VMInsights solution.

Before a Log Analytics workspace can be used with VM insights, it must have the VMInsights solution installed.

Step 3: Install agents on the VM and VM scale set to be monitored.

Prior to onboarding agents, you must create and configure a workspace. Install or update the Application Insights Agent as an extension for Azure virtual machines and VM scalet sets.

Step 4: Create an Application Insights resource

Sign in to the Azure portal, and create an Application Insights resource.

Once a workspace-based Application Insights resource has been created, configuring monitoring is relatively straightforward.

NEW QUESTION 134

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 playerId, 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 137

- (Topic 8)

You develop and deploy an Azure App Service web app named App1. You create a new Azure Key Vault named Vault 1. You import several API keys, passwords, certificates, and cryptographic keys into Vault1.

You need to grant App1 access to Vault1 and automatically rotate credentials. Credentials must not be stored in code.

What should you do?

- A. Enable App Service authentication for App
- B. Assign a custom RBAC role to Vault1.
- C. Add a TLS/SSL binding to App1.
- D. Assign a managed identity to App1.
- E. Upload a self-signed client certificate to Vault1. Update App1 to use the client certificate.

Answer: D

NEW QUESTION 142

- (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 144

HOTSPOT - (Topic 8)

You have an App Service plan named aspl based on the Free pricing tier.


You plan to use aspl to implement an Azure Function app with a queue trigger. Your solution must minimize cost.

You need to identify the configuration options that will meet the requirements.

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

NOTE: Each correct selection is worth one point.

Answer Area




- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 148

HOTSPOT - (Topic 8)

You are developing a ticket reservation system for an airline.

The storage solution for the application must meet the following requirements:

- ? Ensure at least 99.99% availability and provide low latency.
- ? Accept reservations event when localized network outages or other unforeseen failures occur.
- ? Process reservations in the exact sequence as reservations are submitted to minimize overbooking or selling the same seat to multiple travelers.
- ? Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window.

You provision a resource group named airlineResourceGroup in the Azure South-Central US region.

You need to provision a SQL SPI Cosmos DB account to support the app.

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

NOTE: Each correct selection is worth one point.

```
resourceGroupName- +airlineResourceGroup'
name- +docdb-airline-reservations'
databaseName- 'docdb-tickets-database'
collectionName- 'docdb-tickets-collection'
consistencyLevel-
```

▼
Strong
Eventual
ConsistentPrefix
BoundedStaleness

```
az cosmosdb create \
--name $name \
```

▼
--enable-virtual-network true\
--enable-automatic-failover true\
--kind 'GlobalDocumentDB' \
--kind 'MongoDB'\

```
--resource group $resourceGroupName \
--max interval 5 \
```

▼
--locations 'southcentralus'
--locations 'eastus'
--locations'southcentralus=0 eastus=1 westus=2'
--locations 'southcentralus=0'

```
--default-consistency-level - $consistencylevel
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: BoundedStaleness

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is, "updates") of an item or by "T" time interval. In other words, when you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (T) by which the reads might lag behind the writes

NEW QUESTION 152

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

Answer Area

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

EXPOSE ./ContosoApp/ /apps/ContosoApp

COPY ./

FROM microsoft/aspnetcore:2.0

WORKDIR /apps/ContosoApp

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: WORKDIR /apps/ContosoApp

Step 2: COPY ./-

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

Step 3: EXPOSE ./ContosoApp/ /app/ContosoApp Step 4: CMD powershell ./setupScript.ps1

ENTRYPOINT ["dotnet", "ContosoApp.dll"]

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

? Call setupScript.ps1 when the container is built.

? Run ContosoApp.dll when the container starts.

References:

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

NEW QUESTION 156

HOTSPOT - (Topic 8)

You are building a website that is used to review restaurants. The website will use an Azure CDN to improve performance and add functionality to requests.

You build and deploy a mobile app for Apple iPhones. Whenever a user accesses the website from an iPhone, the user must be redirected to the app store.

You need to implement an Azure CDN rule that ensures that iPhone users are redirected to the app store.

How should you complete the Azure Resource Manager template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
"conditions": [ {
  "name": "IsDevice",
  "parameters": {
    "@odata.type": "#Microsoft.Azure.Cdn.Models.",
    "operator": "Equal"
    "matchValues": [ "
    {
      "name": "RequestHeader",
      "parameters": {
        "@odata.type": "#Microsoft.Azure.Cdn.Models.",
        "operator": "Contains",
        "selector": "
        "matchValues": [ "
      } }
    ]
  } }
]
```

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Box 1: iOS

Azure AD Conditional Access supports the following device platforms:

? Android

? iOS

? Windows Phone

? Windows

? macOS

Box 2: DeliveryRuleIsDeviceConditionParameters

The DeliveryRuleIsDeviceCondition defines the IsDevice condition for the delivery rule. parameters defines the parameters for the condition.

Box 3: HTTP_USER_AGENT

Box 4: DeliveryRuleRequestHeaderConditionParameters DeliveryRuleRequestHeaderCondition defines the RequestHeader condition for the delivery rule.

parameters defines the parameters for the condition.

Box 5: iOS

The Require approved client app requirement only supports the iOS and Android for device platform condition.

NEW QUESTION 157

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

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Step 1: Upgrade the storage account to GPv2

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

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

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

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

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

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

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

NEW QUESTION 160

- (Topic 8)

You are developing an application to store business-critical data in Azure Blob storage. The application must meet the following requirements:

- Data must not be modified or deleted for a user-specified interval.
- Data must be protected from overwrites and deletes.
- Data must be written once and allowed to be read many times. You need to protect the data in the Azure Blob storage account.

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

- A. Enable version-level immutability support for the storage account.
- B. Create an account shared-access signature (SAS).
- C. Enable point-in-time restore for containers in the storage account.
- D. Create a service shared-access signature (SAS).
- E. Enable the blob change feed for the storage account.
- F. Configure a time-based retention policy for the storage account.

Answer: DF

NEW QUESTION 165

- (Topic 8)

You are developing a web application that uses the Microsoft identity platform to authenticate users and resources. The web application calls several REST APIs. The APIs require an access token from the Microsoft identity platform. You need to request a token.

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

- A. Application secret
- B. Redirect URI/URL
- C. Application name
- D. Supported account type
- E. Application ID

Answer: ABE

NEW QUESTION 169

HOTSPOT - (Topic 8)

You create the following PowerShell script:

```
$source = New-AzScheduledQueryRuleSource -Query 'Heartbeat | where TimeGenerated > ago(1h)' -DataSourceId "contoso"
$schedule = New-AzScheduledQueryRuleSchedule -FrequencyInMinutes 60 -TimeWindowInMinutes 60
$triggerCondition = New-AzScheduledQueryRuleTriggerCondition -ThresholdOperator "LessThan" -Threshold 5
$actionGroup = New-AzScheduledQueryRuleAznsActionGroup -ActionGroup "contoso" -EmailSubject "Custom email subject"
               -CustomWebhookPayload "{ 'alert':'#alertrulename', 'IncludeSearchResults':true }"
$alertingAction = New-AzScheduledQueryRuleAlertingAction -AznsAction $actionGroup -Severity "3" -Trigger $triggerCondition
New-AzScheduledQueryRule -ResourceGroupName "contoso" -Location "eastus" -Action $alertingAction -Enabled $true
-Description "Alert description" -Schedule $schedule -Source $source -Name "Alert Name"
```

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

The AzScheduledQueryRuleSource is Heartbeat, not CPU.

Box 2: Yes

The AzScheduledQueryRuleSource is Heartbeat!

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

Box 3: No

The schedule is 60 minutes, not two hours.

-FrequencyInMinutes: The alert frequency.

-TimeWindowInMinutes: The alert time window

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

NEW QUESTION 174

- (Topic 8)

A company is developing a solution that allows smart refrigerators to send temperature information to a central location. You have an existing Service Bus. The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location.

You need to complete the configuration.

Which Azure CLI or PowerShell command should you run?

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

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

Answer: B

Explanation:

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

Note:

Steps:

Step 1: # Create a resource group resourceGroupName="myResourceGroup"

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

Step 2: # Create a Service Bus messaging namespace with a unique name namespaceName=myNameSpace\$RANDOM

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

Step 3: # Create a Service Bus queue

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

Step 4: # Get the connection string for the namespace

`connectionString=$(az servicebus namespace authorization-rule keys list --resource-group`

`$resourceGroupName --namespace-name $namespaceName --name RootManageSharedAccessKey --query primaryConnectionString --output tsv)`

References:

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

NEW QUESTION 179

DRAG DROP - (Topic 8)

You are Implementing an Azure solution that uses Azure Cosmos DB and the latest Azure Cosmos DB SDK. You add a change feed processor to a new container instance.

You attempt to lead a batch of 100 documents. The process falls when reading one of the documents. The solution must monitor the progress of the change feed processor instance on the new container as the change feed is read. You must prevent the change feed processor from retrying the entire batch when one document cannot be read.

You need to implement the change feed processor to read the documents.

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


Each correct selection is worth one point

Features	Requirement	Feature
Change feed estimator	Monitor the progress of the change feed processor.	Feature
Dead-letter queue	Prevent the change feed processor from retrying the entire batch when one document cannot be read.	Feature
Deployment unit		
Lease container		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 184

- (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 185

- (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 190

- (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 194

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	<input type="text"/> <ul style="list-style-type: none"> Basic Standard Premium
Access control (IAM) level	<input type="text"/> <ul style="list-style-type: none"> Contributor Data Receiver Data Sender Data Owner

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Configuration	Value
Tier	<input type="text"/> <ul style="list-style-type: none"> Basic <u>Standard</u> <u>Premium</u>
Access control (IAM) level	<input type="text"/> <ul style="list-style-type: none"> <u>Contributor</u> <u>Data Receiver</u> Data Sender Data Owner

NEW QUESTION 197

DRAG DROP - (Topic 8)

You develop a web app that uses tier D1 app service plan by using the Web Apps feature of Microsoft Azure App Service.

Spikes in traffic have caused increases in page load times.

You need to ensure that the web app automatically scales when CPU load is about 85 percent and minimize costs.

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

NOTE:More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions	Answer Area
Configure the web app to the Premium App Service tier.	
Configure the web app to the Standard App Service tier.	
Enable autoscaling on the web-app.	⬅️ ⬆️
Add a Scale rule.	➡️ ⬇️
Switch to an Azure App Services consumption plan.	
Configure a Scale condition.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Configure the web app to the Standard App Service Tier
 The Standard tier supports auto-scaling, and we should minimize the cost.
 Step 2: Enable autoscaling on the web app First enable autoscale
 Step 3: Add a scale rule
 Step 4: Add a Scale condition

NEW QUESTION 201

DRAG DROP - (Topic 8)

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

Actions	Answer Area
Use AZCopy to copy the data to the new storage account.	
Deploy the template to create a new storage account in the target region.	
Export a Resource Manager template.	⬅️ ⬆️
Create a new template deployment.	➡️ ⬇️
Modify the template by changing the storage account name and region.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 203

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  --name ContosoWeb --resource-group ContosoDevRG

```

config
 download
 show
 tail


```

 filesystem

```

--web-server-logging
 --docker-container-logging
 --application-logging


```

az  log  --name ContosoWeb --resource-group ContosoDevRG

```

webapp
 acr
 aks

config
 download
 show
 tail

- 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 204

HOTSPOT - (Topic 8)

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

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

You need to implement push notifications.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

▼	hub=
NotificationHubClient	
NotificationHubClientSettings	
NotificationHubJob	
NotificationDetails	

▼	▼
NotificationHubClient	GetInstallation
NotificationHubClientSettings	CreateClientFromConnectionString
NotificationHubJob	CreateOrUpdateInstallation
NotificationDetails	PatchInstallation

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: NotificationHubClient

Box 2: NotificationHubClient

Box 3: CreateClientFromConnectionString

// Initialize the Notification Hub NotificationHubClient hub =

NotificationHubClient.CreateClientFromConnectionString(listenConnString, hubName);

Box 4: SendWindowsNativeNotificationAsync Send the push notification.

var result = await hub.SendWindowsNativeNotificationAsync(windowsToastPayload);

References:

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

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

NEW QUESTION 208

- (Topic 8)

You are developing a mobile app that uses an API which stores geospacial data in Azure Cosmos DB. The app will be used to find restaurants in a particular area and related information including food types, menu information and the optimal route to a selected restaurant from the user's current location.

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

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

Answer: A

NEW QUESTION 211

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.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 216

DRAG DROP - (Topic 8)

You are developing an ASP.NET Core website that can be used to manage photographs which are stored in Azure Blob Storage containers. Users of the website authenticate by using their Azure Active Directory (Azure AD) credentials. You implement role-based access control (RBAC) role permission on the containers that store photographs. You assign users to RBAC role. You need to configure the website's Azure AD Application so that user's permissions can be used with the Azure Blob containers. How should you configure the application? To answer, drag the appropriate setting to the correct location. Each setting may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: user_impersonation

Box 2: delegated Example:

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

Box 3: delegated Example

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

References:

<https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect-aspnetcore/calling-a-web-api-in-an-aspnet-core-web-application-using-azure-ad/>

NEW QUESTION 217

- (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 222

- (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 227

- (Topic 8)

D18912E1457D5D1DDCBD40AB3BF70D5D

You are building a website that uses Azure Blob storage for data storage. You configure Azure Blob storage lifecycle to move all blobs to the archive tier after 30 days.

Customers have requested a service-level agreement (SLA) for viewing data older than 30 days.

You need to document the minimum SLA for data recovery. Which SLA should you use?

A. at least two days

B. between one and 15 hours

C. at least one day

D. between zero and 60 minutes

Answer: B

Explanation:

The archive access tier has the lowest storage cost. But it has higher data retrieval costs compared to the hot and cool tiers. Data in the archive tier can take several hours to retrieve depending on the priority of the rehydration. For small objects, a high priority rehydrate may retrieve the object from archive in under 1 hour.

Reference:

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

NEW QUESTION 232

HOTSPOT - (Topic 8)

You are developing a C++ application that compiles to a native application named process.exe. The application accepts images as input and returns images in one of the following image formats: GIF, PNG, or JPEG.

You must deploy the application as an Azure Function. You need to configure the function and host json files.

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

NOTE:Each correct selection is worth one point.

function.json

```
{
  "type": "http"
  "platform": "gcm"
  "datatype": "stream"
  "path": "process.exe"
}
```

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

host.json

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

```
    "defaultExecutablePath": "process.exe"
```

```
  },
  "enableForwardingHttpRequest": true
  "enableForwardingHttpRequest": false
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

function.json

```
{
  "type": "http"
  "platform": "gcm"
  "datatype": "stream"
  "path": "process.exe"
}
```

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

host.json

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

```
    "defaultExecutablePath": "process.exe"
```

```
  },
  "enableForwardingHttpRequest": true
  "enableForwardingHttpRequest": false
}
```

NEW QUESTION 237

HOTSPOT - (Topic 8)

You have a single page application (SPA) web application that manages information based on data returned by Microsoft Graph from another company's Azure Active Directory (Azure AD) instance.

Users must be able to authenticate and access Microsoft Graph by using their own company's Azure AD instance.

You need to configure the application manifest for the app registration.

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

NOTE: Each correct selection is worth one point.

```

{
  "oauth2AllowImplicitFlow":  ,
  "requiredResourceAccess":  :[{
    "resourceAppId": "00000003-0000-0000-c000-000000000000",
    "resourceAccess": [{
      "id" : "24a6cdd6-fab1-4aaf-91b8-3cc8225e90d0",
      "type": "Scope"
    }
  ]
},
  "signInAudience": 
}

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: true

The `oauth2AllowImplicitFlow` attribute Specifies whether this web app can request OAuth2.0 implicit flow access tokens. The default is false. This flag is used for browser- based apps, like JavaScript single-page apps.

In implicit flow, the app receives tokens directly from the Azure Active Directory (Azure AD) authorize endpoint, without any server-to-server exchange. All authentication logic and session handling is done entirely in the JavaScript client with either a page redirect or a pop-up box.

Box 2: requiredResourceAccess

With dynamic consent, `requiredResourceAccess` drives the admin consent experience and the user consent experience for users who are using static consent. However, this parameter doesn't drive the user consent experience for the general case.

`resourceAppId` is the unique identifier for the resource that the app requires access to. This value should be equal to the `appId` declared on the target resource app.

`resourceAccess` is an array that lists the OAuth2.0 permission scopes and app roles that the app requires from the specified resource. Contains the `id` and `type` values of the specified resources.

Example: "requiredResourceAccess": [

```

{
  "resourceAppId": "00000002-0000-0000-c000-000000000000",
  "resourceAccess": [
    {
      "id": "311a71cc-e848-46a1-bdf8-97ff7156d8e6", "type": "Scope"
    }
  ]
}
],

```

Box 3: AzureADMyOrg

The `signInAudience` attribute specifies what Microsoft accounts are supported for the current application. Supported values are:

AzureADMyOrg - Users with a Microsoft work or school account in my organization's Azure AD tenant (for example, single tenant)

AzureADMultipleOrgs - Users with a Microsoft work or school account in any organization's Azure AD tenant (for example, multi-tenant)

AzureADandPersonalMicrosoftAccount - Users with a personal Microsoft account, or a work or school account in any organization's Azure AD tenant

NEW QUESTION 242

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 247

- (Topic 8)

You develop and deploy an ASP.NET web app to Azure App Service. You use Application Insights telemetry to monitor the app.

You must test the app to ensure that the app is available and responsive from various points around the world and at regular intervals. If the app is not responding, you must send an alert to support staff.

You need to configure a test for the web app.

Which two test types can you use? Each correct answer presents a complete solution.

NOTE:Each correct selection is worth one point.

- A. integration
- B. multi-step web
- C. URL ping
- D. unit
- E. load

Answer: BC

Explanation:

There are three types of availability tests:

? URL ping test: a simple test that you can create in the Azure portal.

? Multi-step web test: A recording of a sequence of web requests, which can be played back to test more complex scenarios. Multi-step web tests are created in Visual Studio Enterprise and uploaded to the portal for execution.

? Custom Track Availability Tests: If you decide to create a custom application to run availability tests, the TrackAvailability() method can be used to send the results to Application Insights.

Reference:

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

NEW QUESTION 251

DRAG DROP - (Topic 8)

An organization plans to deploy Azure storage services.

You need to configure shared access signature (SAS) for granting access to Azure Storage.

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

NOTE:Each correct selection is worth one point.

SAS types	Requirement	SAS type
Account-level	Delegate access to resources in one or more of the storage services	
Service-level	Delegate access to a resource in a single storage service	
User delegation	Secure a resource by using Azure AD credentials	

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

SAS types	Requirement	SAS type
Account-level	Delegate access to resources in one or more of the storage services	Account-level
Service-level	Delegate access to a resource in a single storage service	Service-level
User delegation	Secure a resource by using Azure AD credentials	User delegation

NEW QUESTION 253

HOTSPOT - (Topic 8)

A company is developing a mobile app for field service employees using Azure App Service Mobile Apps as the backend.

The company's network connectivity varies throughout the day. The solution must support offline use and synchronize changes in the background when the app is online app.

You need to implement the solution.

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

NOTE: Each correct selection is worth one point.

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

▼

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

await client.SyncContext.PushAsync();

▼

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

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Box 1: var todoTable = client GetSyncTable<TodoItem>()

To setup offline access, when connecting to your mobile service, use the method GetSyncTable instead of GetTable (example):

IMobileServiceSyncTable todoTable = App.MobileService.GetSyncTable(); / Box 2: await todoTable.PullAsync("allTodoItems",todo.Table.CreateQuery());

Your app should now use IMobileServiceSyncTable (instead of IMobileServiceTable) for

CRUD operations. This will save changes to the local database and also keep a log of the changes. When the app is ready to synchronize its changes with the

Mobile Service, use the methods PushAsync and PullAsync (example):

await App.MobileService.SyncContext.PushAsync(); await todoTable.PullAsync();

References:

<https://azure.microsoft.com/es-es/blog/offline-sync-for-mobile-services/>

NEW QUESTION 257

HOTSPOT - (Topic 8)

You develop two Python scripts to process data.

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

You need to configure the Azure Container Instance.

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

NOTE: Each correct selection is worth one point.

Answer Area

Configuration Setting	Configuration Value
External data volume	<div style="border: 1px solid gray; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid gray;">▼</div> <p>Secret</p> <p>Empty directory</p> <p>Cloned git repo</p> <p>Azure file share</p> </div>
Container restart policy	<div style="border: 1px solid gray; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid gray;">▼</div> <p>Never</p> <p>Always</p> <p>OnFailure</p> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Configuration Setting	Configuration Value
External data volume	<div style="border: 1px solid gray; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid gray;">▼</div> <p>Secret</p> <p>Empty directory</p> <p>Cloned git repo</p> <p style="border: 2px dashed green;">Azure file share</p> </div>
Container restart policy	<div style="border: 1px solid gray; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid gray;">▼</div> <p style="border: 2px dashed green;">Never</p> <p>Always</p> <p>OnFailure</p> </div>

NEW QUESTION 260

HOTSPOT - (Topic 8)

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

Answer Area

```

Delete (Azure.Storage.Blobs.Models.DeleteSnapshotsOption
snapshotsOption = Azure.Storage.Blobs.Models.
    
```

▼

DeleteSnapshotsOption

DeleteIfExists

DeleteSnapshotsOption

WithSnapshot

WithSnapshotCore

▼

OnlySnapshots

IncludeSnapshots

None

OnlySnapshots

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

Delete (Azure.Storage.Blobs.Models.DeleteSnapshotsOption
snapshotsOption = Azure.Storage.Blobs.Models.
    
```

▼

DeleteSnapshotsOption

DeleteIfExists

DeleteSnapshotsOption

WithSnapshot

WithSnapshotCore

▼

OnlySnapshots

IncludeSnapshots

None

OnlySnapshots

NEW QUESTION 265

- (Topic 8)

You are designing a web application to manage user satisfaction surveys. The number of questions that a survey includes is variable. Application users must be able to display results for a survey as quickly as possible. Users must also be able to quickly compute statistical measures including average values across various groupings of answers.

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

- A. Core
- B. Mongo DB
- C. Gremlin
- D. Table API

Answer: D

NEW QUESTION 267

- (Topic 8)

You develop and deploy an Azure App Service web app. The app is deployed to multiple regions and uses Azure Traffic Manager. Application Insights is enabled for the app.

You need to analyze app uptime for each month.

Which two solutions will achieve the goal? Each correct answer presents a complete solution

NOTE: Each correct selection is worth one point

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

Answer: AC

Explanation:

Reference:

<https://azure.microsoft.com/en-us/blog/creating-a-web-test-alert-programmatically-with-application-insights/>

NEW QUESTION 272

DRAG DROP - (Topic 8)

You are developing an Azure-hosted application that must use an on-premises hardware security module (HSM) key.

The key must be transferred to your existing Azure Key Vault by using the Bring Your Own Key (BYOK) process.

You need to securely transfer the key to Azure Key Vault.

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

Actions	Answer Area
Generate a key transfer blob file by using the HSM vendor-provided tool.	
Generate a Key Exchange Key (KEK).	
Create a custom policy definition in Azure Policy.	
Run the <code>az keyvault key import</code> command.	
Run the <code>az keyvault key restore</code> command.	
Retrieve the Key Exchange Key (KEK) public key.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

To perform a key transfer, a user performs following steps:

? Generate KEK.

? Retrieve the public key of the KEK.

? Using HSM vendor provided BYOK tool - Import the KEK into the target HSM and exports the Target Key protected by the KEK.

? Import the protected Target Key to Azure Key Vault.

Step 1: Generate a Key Exchange Key (KEK).

Step 2: Retrieve the Key Exchange Key (KEK) public key.

Step 3: Generate a key transfer blob file by using the HSM vendor-provided tool. Generate key transfer blob using HSM vendor provided BYOK tool

Step 4: Run the `az keyvault key import` command Upload key transfer blob to import HSM-key.

Customer will transfer the Key Transfer Blob (".byok" file) to an online workstation and then run a `az keyvault key import` command to import this blob as a new HSM-backed key into Key Vault.

To import an RSA key use this command: `az keyvault key import`

NEW QUESTION 277

HOTSPOT - (Topic 8)

You are developing an application that monitors data added to an Azure Blob storage account.

You need to process each change made to the storage account.

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

NOTE: Each correct selection is worth one point.

```
cf = ChangeFeedClient("", "")
x = None
while True:
    change_feed = cf.
    for c in change_fee
        ProcessChanges(c)
    x = change_feed.
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
cf = ChangeFeedClient("", "")
x = None
while True:
    change_feed = cf.
    for c in change_fee
        ProcessChanges(c)
    x = change_feed.
```

NEW QUESTION 279

DRAG DROP - (Topic 8)

Your company has several websites that use a company logo image. You use Azure Content Delivery Network (CDN) to store the static image. You need to determine the correct process of how the CDN and the Point of Presence (POP) server will distribute the image and list the items in the correct order. In which order do the actions occur? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

- A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.
- Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the files from cache if the TTL has not expired.
- If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.
- The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: A user requests the image..

A user requests a file (also called an asset) by using a URL with a special domain name, such as <endpoint name>.azureedge.net. This name can be an endpoint hostname or a custom domain. The DNS routes the request to the best performing POP location, which is usually the POP that is geographically closest to the user.

Step 2: If no edge servers in the POP have the..

If no edge servers in the POP have the file in their cache, the POP requests the file from the origin server. The origin server can be an Azure Web App, Azure Cloud Service, Azure Storage account, or any publicly accessible web server.

Step 3: The origin server returns the..

The origin server returns the file to an edge server in the POP.

An edge server in the POP caches the file and returns the file to the original requestor (Alice). The file remains cached on the edge server in the POP until the time-to-live (TTL) specified by its HTTP headers expires. If the origin server didn't specify a TTL, the default TTL is seven days.

Step 4: Subsequent requests for..

Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.

If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.

References:

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

NEW QUESTION 283

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.

The screenshot shows the Logic App configuration interface. On the left, under 'Triggers and Action Blocks', there are four available blocks: 'Insert Entity' (with 'Table' set to 'processing' and 'Entity' set to 'Path'), 'Tier blob' (with 'If blob is older than the defined value, tier it to Cool or Archive tier', 'Blob path' set to 'Path', and 'Blob Tier' set to 'Archive'), 'When there are messages in a queue' (with 'Queue Name' set to 'processing'), and 'Recurrence' (with 'Interval' set to '1' and 'Frequency' set to 'Month'). On the right, the 'Answer Area' shows a workflow with the following steps: 1. 'When there are messages in a queue' (Queue Name: processing). 2. 'If true' condition. 3. 'Set tier age variable' action. 4. 'Set tier age variable' action. 5. 'For each' loop containing 'Scan all blobs in this folder' (with 'Select an output from previous steps' set to 'value'). 6. 'When there are messages in a queue' (Queue Name: processing). 7. 'If true' condition. 8. 'If false' condition. There are 'Add an action' buttons at the bottom of the workflow.

- 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 286

- (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 a DataContainer that contains the documents which must be added.
 * 3. Create a DataSource instance and set its Container property to the DataContainer.
 * 4 Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource.
 Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 291

DRAG DROP - (Topic 8)

You develop software solutions for a mobile delivery service. You are developing a mobile app that users can use to order from a restaurant in their area. The app uses the following workflow:

- ? A driver selects the restaurants from which they will deliver orders.
- ? Orders are sent to all available drivers in an area.
- ? Only orders for the selected restaurants will appear for the driver.
- ? The first driver to accept an order removes it from the list of available orders.

You need to implement an Azure Service Bus solution.

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

Actions	Answer Area
Create a single Service Bus topic.	
Create a Service Bus Namespace for each restaurant for which a driver can receive messages.	
Create a single Service Bus subscription.	
Create a Service Bus subscription for each restaurant for which a driver can receive orders.	
Create a single Service Bus Namespace.	
Create a Service Bus topic for each restaurant for which a driver can receive messages.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Create a single Service Bus Namespace

To begin using Service Bus messaging entities in Azure, you must first create a namespace with a name that is unique across Azure. A namespace provides a scoping container for addressing Service Bus resources within your application.

Box 2: Create a Service Bus Topic for each restaurant for which a driver can receive messages. Create topics.

Box 3: Create a Service Bus subscription for each restaurant for which a driver can receive orders.

NEW QUESTION 294

HOTSPOT - (Topic 8)

You are developing an application to collect the following telemetry data for delivery drivers: first name, last name, package count, item id, and current location coordinates.

The app will store the data in Azure Cosmos DB.

You need to configure Azure Cosmos DB to query the data.

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

NOTE:Each correct selection is worth one point.

Configuration Parameter

Value

Azure Cosmos DB API

	▼
Gremlin	
Table API	
Core (SQL)	

Azure Cosmos DB partition key

	▼
first name	
last name	
package count	
item id	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Core (SQL)

Core(SQL) API stores data in document format. It offers the best end-to-end experience as we have full control over the interface, service, and the SDK client libraries. SQL API supports analytics and offers performance isolation between operational and analytical workloads.

Box 2: item id

item id is a unique identifier and is suitable for the partition key.

NEW QUESTION 297

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 placed in the container referenced by the failedContainerSasUrl parameter.

You need to ensure the files are correctly processed.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

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

            outputFileList.Add(new OutputFile(fileTask.Output,
            new OutputFileDestination(failedContainer),
            new OutputFileUploadOptions(OutputFileUploadCondition. ▼ ))) );
            TaskSuccess
            TaskFailure
            TaskCompletion

            task. ▼ =outputFileList;
            OutputFiles
            FilesToStage
            ResourceFiles
            StageFiles

            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 300

DRAG DROP - (Topic 8)

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

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

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

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

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

NOTE: Each correct selection is worth one point.

Consistency levels

- Strong
- Bounded Staleness
- Consistent Prefix
- Eventual

Answer Area

- Return the most recent patient status.
- Return health monitoring data that is no less than one version behind.
- After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Strong

Strong: Strong consistency offers a linearizability guarantee. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Box 2: Bounded staleness

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is "updates") of an item or by "t" time interval. When you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (t) by which the reads might lag behind the writes

Box 3: Eventual

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

NEW QUESTION 303

- (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 306

- (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 309

- (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 310

.....

Relate Links

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

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

Contact us

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

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