DUMPSARENA

Developing Solutions for Microsoft Azure

Microsoft AZ-204

Version Demo

Total Demo Questions: 20

Total Premium Questions: 536

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, New Update	284
Topic 2, Case Study 1	2
Topic 3, Case Study 2	2
Topic 4, Case Study 3	2
Topic 5, Case Study 4	3
Topic 6, Case Study 5	2
Topic 7, Case Study 6	2
Topic 8, Case Study 7	5
Topic 9, Case Study 8	2
Topic 10, Case Study 9	6
Topic 11, Case Study 10	2
Topic 12, Case Study 11	2
Topic 13, Case Study 12	2
Topic 14, Case Study 13	4
Topic 15, Case Study 14	3
Topic 16, Case Study 15	2
Topic 17, Case Study 16	2
Topic 18, Case Study 17	2
Topic 19, Case Study 18	3
Topic 20, Mixed Questions	204
Total	536

QUESTION NO: 1

This question is part of a series that presents a consistent scenario. Each question in the series supplies a unique solution to evaluate. Determine whether the given solution achieves the described goals.

You are responsible for developing and deploying a series of ASP.NET web applications on Azure App Service. Your objective is to preserve session state data and HTML output utilizing a storage mechanism with the following requirements:

Share session state among all ASP.NET web applications
Support controlled, concurrent access to the same session state data, allowing multiple readers and a single writer
Preserve complete HTTP responses for concurrent requests

To store this information, the proposed solution is to deploy and configure Azure Cache for Redis and then update the web applications accordingly.

Does this solution fulfill the objective?

A. Yes

B. No

ANSWER: A

Explanation:

Azure Cache for Redis is an optimal solution for achieving the goals outlined. It allows for sharing session state across multiple instances of ASP.NET web applications and supports concurrent read and write operations with session state data. Moreover, Redis facilitates storing complete HTTP responses for concurrent requests. It offers flexibility through its configuration options, including controlling eviction policies and managing access through roles. For more details, refer to the official documentation here: https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching

QUESTION NO: 2

This question belongs to a series of questions sharing the same scenario. Each question in the series contains a unique solution that may or may not meet the requirements stated in the scenario. Some series of questions might have more than one correct solution, while others might not have any correct solutions.

Once you have answered a question in this section, you will NOT be able to revisit it, and therefore, these questions will not be displayed in the review screen.

You are in the process of developing a solution on Azure to collect data from point-of-sale (POS) devices across 2,000 stores worldwide. Each individual device is capable of generating 2 megabytes (MB) of data per day. Stores can have between one to five devices transmitting data.

The collected device data must be stored in Azure Blob Storage and correlated based on a specific device identifier. Additional stores are projected to be established in the future.

Your task is to devise a method to efficiently receive the incoming device data.

Solution: Set up an Azure Event Hub. Configure the machine identifier to serve as the partition key and activate capture.

A. Yes

B. No

ANSWER: A

Explanation:

Azure Event Hubs is an ideal solution for ingesting large amounts of data from various sources, such as the devices mentioned in the scenario. By using the device identifier as a partition key, you can ensure that data from the same device will always be grouped together, which facilitates easier processing and analysis. Enabling capture allows you to persist stream data automatically in Azure Blob Storage, which aligns with your requirement to store device data in Azure Blob Storage. For more detailed programming guidance, refer to the Azure Event Hubs Programming Guide.

QUESTION NO: 3 - (DRAG DROP)

DRAG DROP

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

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

Select and Place:

ANSWER:

Explanation:

FutureOrders: SQLFilter

HighPriortyOrders: CorrelationFilter CorrelationID only

InternationalOrders: SQLFilter

Country NOT USA requires an SQL Filter

HighQuantityOrders: SQLFilter

Need to use relational operators so an SQL Filter is needed.

AllOrders: No Filter

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

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

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

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

QUESTION NO: 4 - (DRAG DROP)

You are a developer for a company that provides a bookings management service in the tourism industry. You are implementing Azure Search for the tour agencies listed in your company's solution.

You create the index in Azure Search. You now need to use the Azure Search .NET SDK to import the relevant data into the Azure Search service.

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

Select and Place:

Answer Area

Create a DataSource instance and set its Container property to the DataContainer.	
Create an IndexBatch that contains the documents which must be added.	
Set the DataSources property of the SearchServiceClient.	
Create a SearchIndexClient object to connect to the search index.	
Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.	
Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource.	

ANSWER:

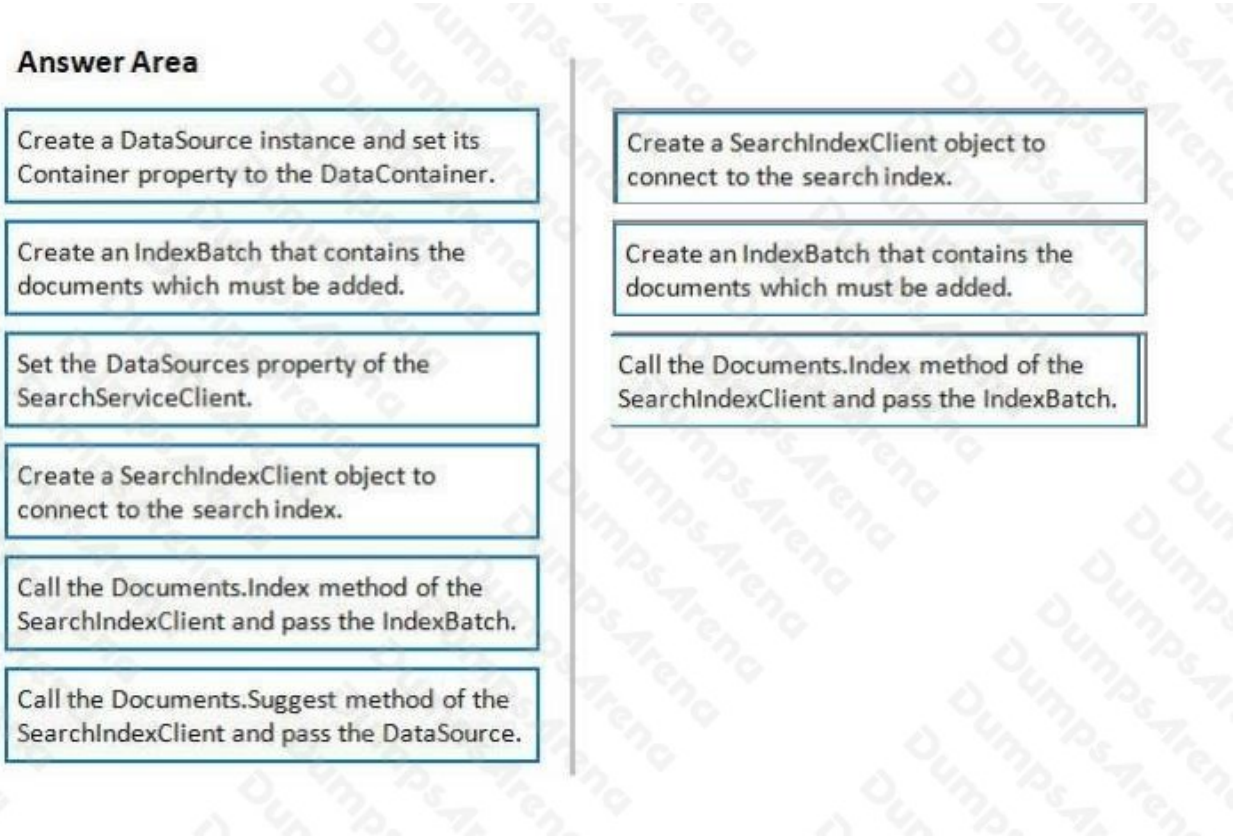

Explanation:

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 wewill use the first method.

2. Create the indexBatch with the documentsSomething 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 indexExample:

var batch = IndexBatch.Upload(hotels);

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

Reference:

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

QUESTION NO: 5

You are developing a web application that processes image analysis on user-uploaded photographs and returns metadata with identified objects. The image analysis process is resource-intensive in terms of both time and computational power. You are planning to use Azure Redis Cache so that uploads do not need to be reprocessed in order to improve efficiency and performance.

In the event of an Azure data center outage, it is critical to minimize any metadata loss.

Which two configurations should you apply to the Azure Redis Cache instance to meet these requirements? Each correct answer contributes to the solution.

Note: Each correct selection is worth one point.

A. Configure Azure Redis with persistence

B. Configure second storage account for persistence

C. Set backup frequency to the minimum value

D. Configure Azure Redis with RDS persistence

ANSWER: A D

Explanation:

To minimize metadata loss in case of an Azure data center outage, it's important to configure persistence for your Azure Redis Cache. This ensures data can be saved and restored when necessary.

https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-premium-persistence

QUESTION NO: 6

You are a developer for a SaaS company that offers numerous web services. All web services provided by the company must adhere to the following requirements:

- Utilize API Management for service access
- Employ OpenID Connect for authentication
- Prohibit anonymous access

A recent security audit has revealed that several web services can be accessed without authentication. Which API Management policy should you implement to address this issue?

A. jsonp

B. authentication-certificate

C. check-header

D. validate-jwt

ANSWER: D

Explanation:

To ensure that any incoming request is authenticated, you should implement the validate-jwt policy. This policy will validate the OAuth token on every request to prevent unauthenticated access, aligning with the requirement to prevent anonymous usage. For more details on configuring API Management policies, please refer to the official documentation: https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies#validate-jwt.

QUESTION NO: 7

You are developing Azure Durable Functions to manage vehicle loans.

In the loan process, there are several actions that must be executed in a specific sequence. One of these actions involves a customer credit check process, which might take several days to complete.

To implement Azure Durable Functions for this loan processing, which type of Azure Durable Function should you use?

A. orchestrator

B. client

C. activity

D. entity

ANSWER: A

Explanation:

The correct Azure Durable Function type to use in this scenario is the "orchestrator" function. Orchestrator functions are designed to coordinate the execution of multiple tasks in a specific order and manage long-running processes, including those with complex chaining and waiting logic. This is particularly suitable for scenarios like vehicle loan processing, where actions might take varying amounts of time to complete, such as customer credit checks.
https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-overview

QUESTION NO: 8

You are developing a web application that uses the Microsoft identity platform to authenticate users and resources. The web application makes calls to several REST APIs, which require an access token from the Microsoft identity platform. You need to perform the necessary steps to request an access token. Which three properties are essential for this process? Each correct answer is part of the overall 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: B C E

Explanation:

To request an access token from the Microsoft identity platform, you typically need to provide the application ID, application secret, and the tenant information associated with the application registration. These elements are crucial in the OAuth 2.0 authentication process. More information can be found in the official documentation: https://learn.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-auth-code-flow

QUESTION NO: 9

You are developing a Software as a Service (SaaS) application that requires storing data in key-value pairs.

For the lowest cost edition of the application, performance must be on a best-effort basis and there should be no regional failover support. Higher-cost editions should enable customers to choose guaranteed levels of performance and support for multiple regions. Cost optimization on Azure is essential.

Which Azure Cosmos DB API should be used for this application?

A. Core

B. MongoDB

C. Cassandra

D. Table API

ANSWER: D

Explanation:

The Table API in Azure Cosmos DB is optimized for cost efficiency and key-value storage, suitable for best-effort performance applications. It is ideal for applications that require flexibility in scaling up to multiple regions with enhanced performance in higher-cost plans. Learn more about Azure Cosmos DB Table API at Azure Cosmos DB Table API.

QUESTION NO: 10

You are developing a REST API using Azure App Service that must be accessed by another Azure App Service web application. This API is required to retrieve and make changes to user profile information stored in Azure Active Directory (Azure AD). Which tools will you use to configure the API for making these updates? Select two tools, as each correct answer forms part of the complete solution.

Note: Each correct selection is worth one point.

A. Microsoft Graph API

B. Microsoft Authentication Library (MSAL)

C. Azure API Management

D. Microsoft Azure Security Center

E. Microsoft Azure Key Vault SDK

ANSWER: A C

Explanation:

A: You can use the Azure AD REST APIs in Microsoft Graph to create unique workflows between Azure AD resources and third-party services.

Enterprise developers use Microsoft Graph to integrate Azure AD identity management and other services to automate administrative workflows, such as employee onboarding (and termination), profile maintenance, license deployment, and more.

C: API Management (APIM) is a way to create consistent and modern API gateways for existing back-end services.

API Management helps organizations publish APIs to external, partner, and internal developers to unlock the potential of their data and services.

Reference:

https://docs.microsoft.com/en-us/graph/azuread-identity-access-management-concept-overview

QUESTION NO: 11

You are in the process of developing an application that will leverage Azure messaging services. Your requirement is to use a publish-subscribe (pub/sub) architecture and avoid the constant need for polling.

Which two services can you use to meet these requirements? Each correct answer presents a complete solution.

Note: Each correct selection is worth one point.

A. Service Bus

B. Event Hub

C. Event Grid

D. Queue

ANSWER: A C

Explanation:

It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka. Reference:

https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber

QUESTION NO: 12 - (HOTSPOT)

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

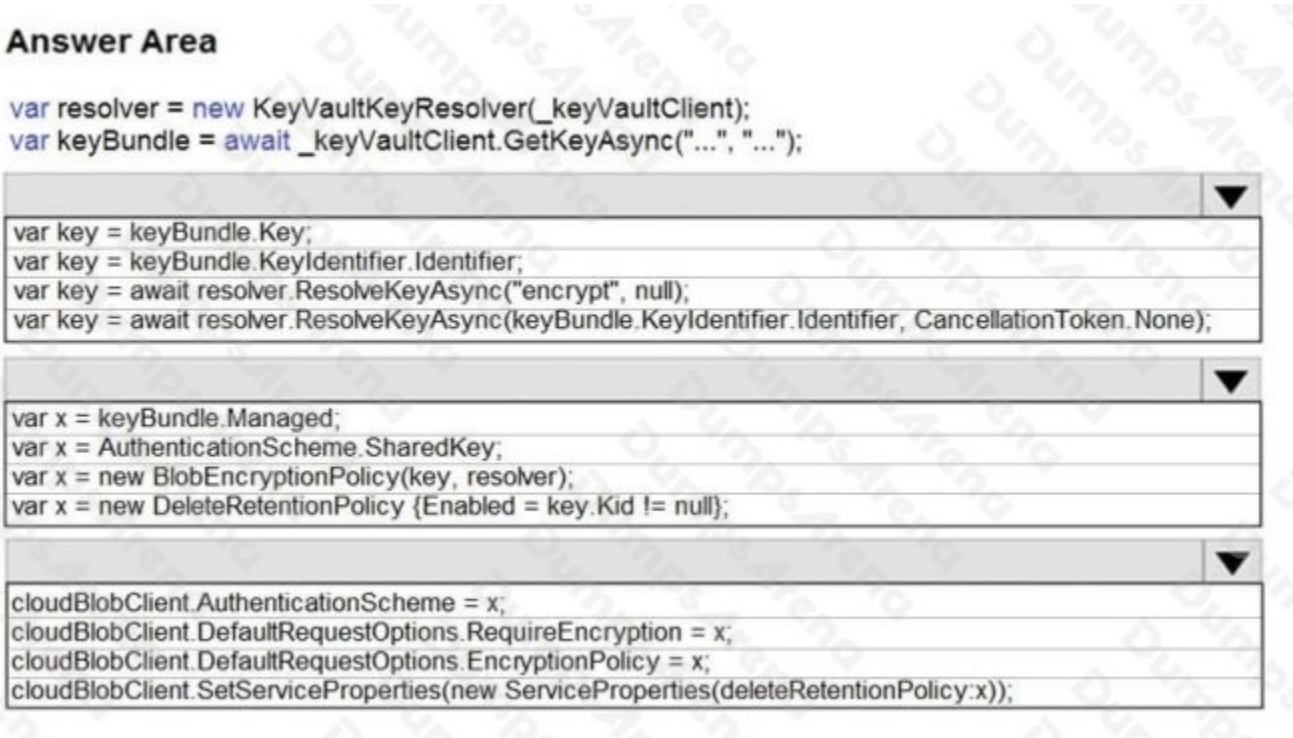

ANSWER:

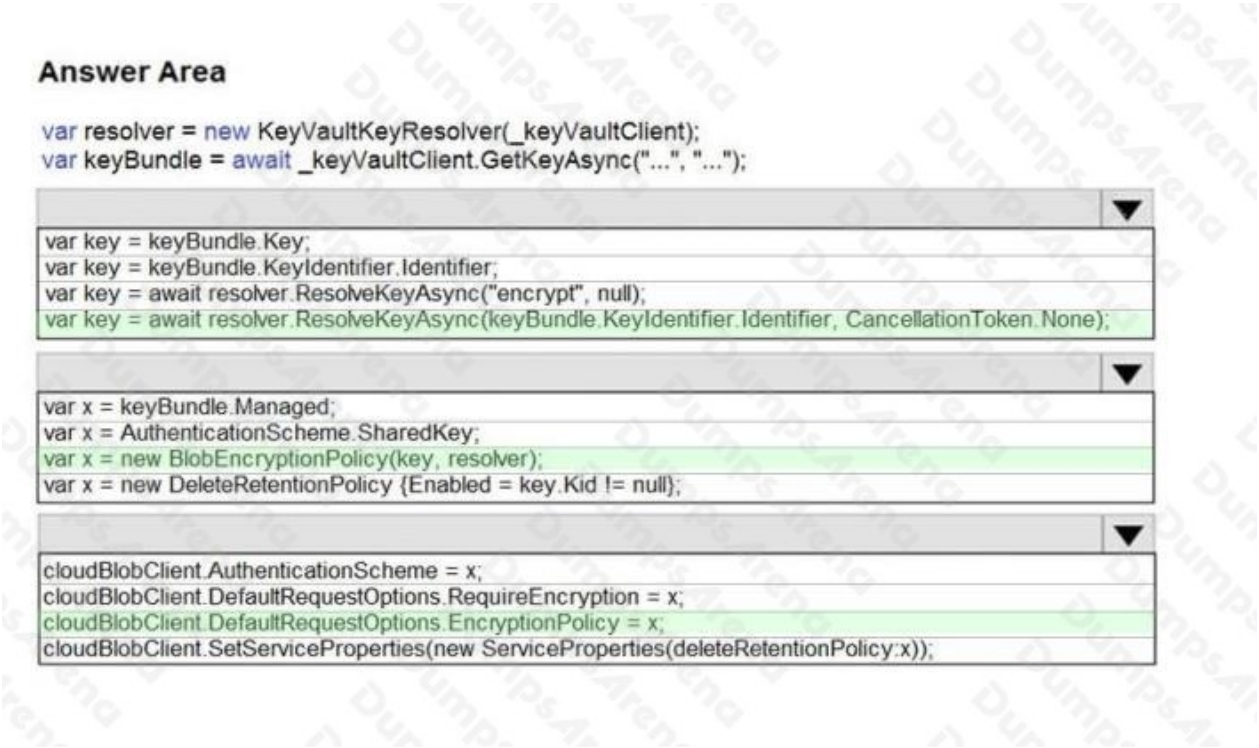

Explanation:

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

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

// We begin with cloudKey1, and a resolver capable of resolving and caching Key Vault secrets. BlobEncryptionPolicy encryptionPolicy = new BlobEncryptionPolicy(cloudKey1, cachingResolver); client.DefaultRequestOptions.EncryptionPolicy = encryptionPolicy; Box 3: cloudblobClient. DefaultRequestOptions.EncryptionPolicy = x;

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

QUESTION NO: 13

You are tasked with creating an Azure Key Vault using PowerShell. Objects that are deleted from the Key Vault should be retained for a duration of 90 days. Which two of the following parameters should be used together to fulfill this requirement? (Select two correct options.)

A. EnabledForDeployment

B. EnablePurgeProtection

C. EnabledForTemplateDeployment

D. EnableSoftDelete

ANSWER: B D

Explanation:

To ensure that deleted objects are retained for a specific period in Azure Key Vault, both **EnableSoftDelete** and **EnablePurgeProtection** parameters are crucial. The **EnableSoftDelete** parameter allows deleted objects to be recoverable for a customized retention period (default is 90 days). The **EnablePurgeProtection** parameter ensures that the deleted objects cannot be purged during the soft delete retention period. More details can be found at https://docs.microsoft.com/en-us/azure/key-vault/general/soft-delete-overview.

QUESTION NO: 14

You have developed and deployed a web application using Azure App Service. This application accesses data stored in an Azure Storage account, which consists of several containers housing large volumes of data. All your Azure resources are currently deployed in a single region.

Your goal is to move the Azure Storage account to a new region, making sure that all existing data is copied to this new location.

What initial step should you take to accomplish this task?

A. Export the Azure Storage account Azure Resource Manager template

B. Initiate a storage account failover

C. Configure object replication for all blobs

D. Use the AzCopy command line tool

E. Create a new Azure Storage account in the current region

F. Create a new subscription in the current region

ANSWER: A

Explanation:

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

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

Reference:

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

QUESTION NO: 15

You are developing a web application hosted on Azure that occasionally goes offline for offline data processing tasks. During these offline periods, multiple Azure Monitor alerts are triggered, resulting in the on-call developer being unnecessarily paged.

To address this issue, it is crucial that the application always records logs indicating when it is offline. However, the on-call developer should not receive pages during these offline processing times.

Which solution should you implement to prevent paging the on-call developer during these offline periods?

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

Explanation:

To prevent unnecessary alerts during defined offline processing times, you can configure an action group in Azure Monitor. Azure Monitor action groups provide a way to customize notifications and automation. Using an action group, you can create an action to suppress alerts temporarily or selectively, ensuring the on-call developer is not paged during expected offline operations.

QUESTION NO: 16

You are developing Azure Web Apps for a company that requires all professional divers to complete a health questionnaire every 15 days, following the commencement of each diving assignment.

Your task involves configuring the Azure Web Apps to automatically scale the instance count up during the periods when divers are filling out these questionnaires and scale down once the process is finished.

Which two autoscaling configuration options can you use to ensure the application scales appropriately according to these requirements? Please note that each correct answer provides a complete solution.

NOTE: Each correct selection earns one point.

A. Predictive autoscaling

B. CPU usage-based autoscaling

C. Recurrence profile

D. Fixed date profile

ANSWER: A C

Explanation:

Azure offers multiple autoscaling options for web applications hosted within its ecosystem.

1. **Predictive autoscaling:** Utilizes machine learning to predict incoming traffic patterns and automatically adjusts resources accordingly, ensuring optimal application performance. More details can be found at Azure Predictive Autoscaling.
2. **Recurrence profile:** Allows automatic adjustment of resources based on scheduled times. This is useful for recurring tasks and predictable patterns of increased load. For additional information, refer to Azure Autoscale - Recurrence Profile.

Deploying a combination of these two approaches can effectively handle the task requirements, accommodating the predictable increase in demand due to questionnaire submissions.

QUESTION NO: 17

You have developed and deployed an ASP.NET web application to Azure App Service and are using Application Insights telemetry to monitor it. You are required to test the application to ensure its availability and responsiveness from various global locations at regular intervals. In the event of non-responsiveness, an alert must be sent to the support staff. What are the two test types you can configure for the web app to meet these requirements? Each correct choice represents a complete solution.

NOTE: Each correct selection is worth one point.

A. integration

B. multi-step web

C. URL ping

D. unit

E. load

ANSWER: B C

Explanation:

There are three types of availability tests in Azure Monitor: multi-step web tests, URL ping tests, and standard tests. Each of these tests ensures that your application is available from different locations worldwide. For more information, refer to the official documentation: https://docs.microsoft.com/en-us/azure/azure-monitor/app/monitor-web-app-availability.

QUESTION NO: 18 - (HOTSPOT)

HOTSPOT

You create the following PowerShell script:

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

ANSWER:

Answer Area

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

Explanation:

Box 1: No

The AzScheduledQueryRuleSource is Heartbeat, not CPU.

Box 2: Yes

The AzScheduledQueryRuleSource is Heartbeat!

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

Box 3: No

The schedule is 60 minutes, not two hours.

-FrequencyInMinutes: The alert frequency.

-TimeWindowInMinutes: The alert time window

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

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

https://docs.microsoft.com/en-us/powershell/module/az.monitor/newazscheduledqueryruletriggercondition

QUESTION NO: 19

This question is part of a series of questions that share the same scenario. Each question in the series is independent and offers a unique solution which could potentially meet the stated objectives. Some question sets might include multiple correct solutions, while others may not have any correct solutions.

Once you answer a question in this section, you will not be able to revisit it. Consequently, these questions will not appear on the review screen.

You are tasked with developing a website to be hosted as an Azure Web App. User authentication will be facilitated via Azure Active Directory (Azure AD) credentials.

The goal is to assign users to one of the following permission levels on the website: admin, normal, and reader. The user's Azure AD group membership should be instrumental in determining their permission level.

Your task is to implement the appropriate authorization configuration.

Does the proposed solution achieve the goal?

A. Yes

B. No

ANSWER: B

Explanation:

Microsoft Graph is a RESTful web API that allows you to access Microsoft Cloud service resources. To utilize Azure AD groups for permission assignment in an Azure Web App, configure your Azure AD application's manifest to include 'groupMembershipClaims' set to 'All'. This setting ensures that the JWT token for the user contains the 'groups' claim. By accessing the 'groups' claim, you can determine the user's permissions within your web application. For detailed guidance, refer to the official documentation: Azure AD Group Claims Configuration.

QUESTION NO: 20

You are developing a hazard notification system that uses a single signaling server to activate and deactivate audio and visual alarms.

In this system, Azure Service Bus is used to publish alarm signals, while each alarm controller utilizes Azure Service Bus to process these signals as part of a transaction. Recording alarm events for auditing purposes is necessary. Each transaction record should contain details regarding the activated alarm type.

What are two actions you should undertake to ensure an audit trail solution is implemented correctly?

Each correct answer contributes to the solution. Note that each correct choice brings one point.

A. Assign the value of the hazard message SessionID property to the ReplyToSessionId property.

B. Assign the value of the hazard message MessageId property to the DevileryCount property.

C. Assign the value of the hazard message SessionID property to the SequenceNumber property.

D. Assign the value of the hazard message MessageId property to the CorrelationId property.

E. Assign the value of the hazard message SequenceNumber property to the DeliveryCount property.

F. Assign the value of the hazard message MessageId property to the SequenceNumber property.

ANSWER: A D

Explanation:

The correct use of the properties CorrelationId and ReplyToSessionId ensures that messages are properly correlated and can be tracked for audit purposes. The CorrelationId is used to link messages together for the purpose of tracking and context, which is essential for auditing each event back to its source. The ReplyToSessionId helps in demarcating the session for the reply messages, aiding in maintaining a consistency in the transaction chain. For more detailed information, refer to the official Azure Service Bus documentation.