AZ-204 Q&As

Developing Solutions for Microsoft Azure

Pass Microsoft AZ-204 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.certbus.com/az-204.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

QUESTION 1

HOTSPOT

You are building an application that stores sensitive customer data in Azure Blob storage. The data must be encrypted with a key that is unique for each customer.

If the encryption key has been corrupted it must not be used for encryption.

You need to ensure that the blob is encrypted.

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

NOTE: Each correct selection is worth one point.

Hot Area:

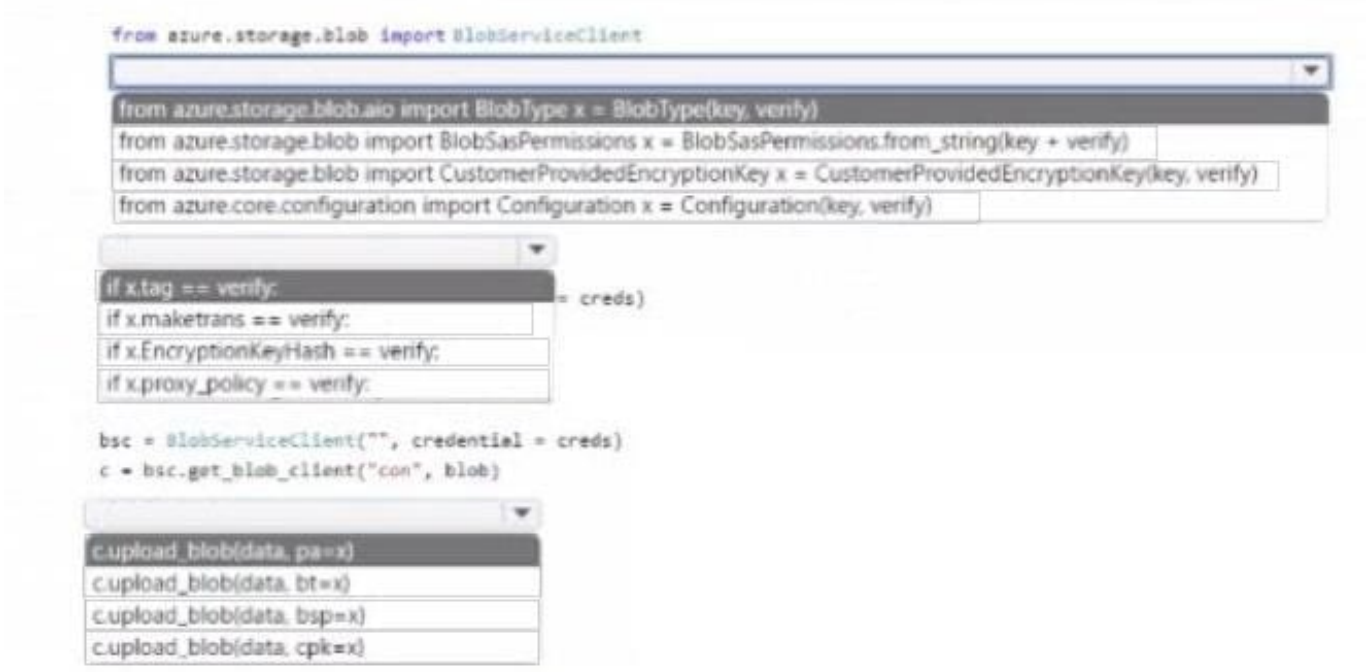

Correct Answer:

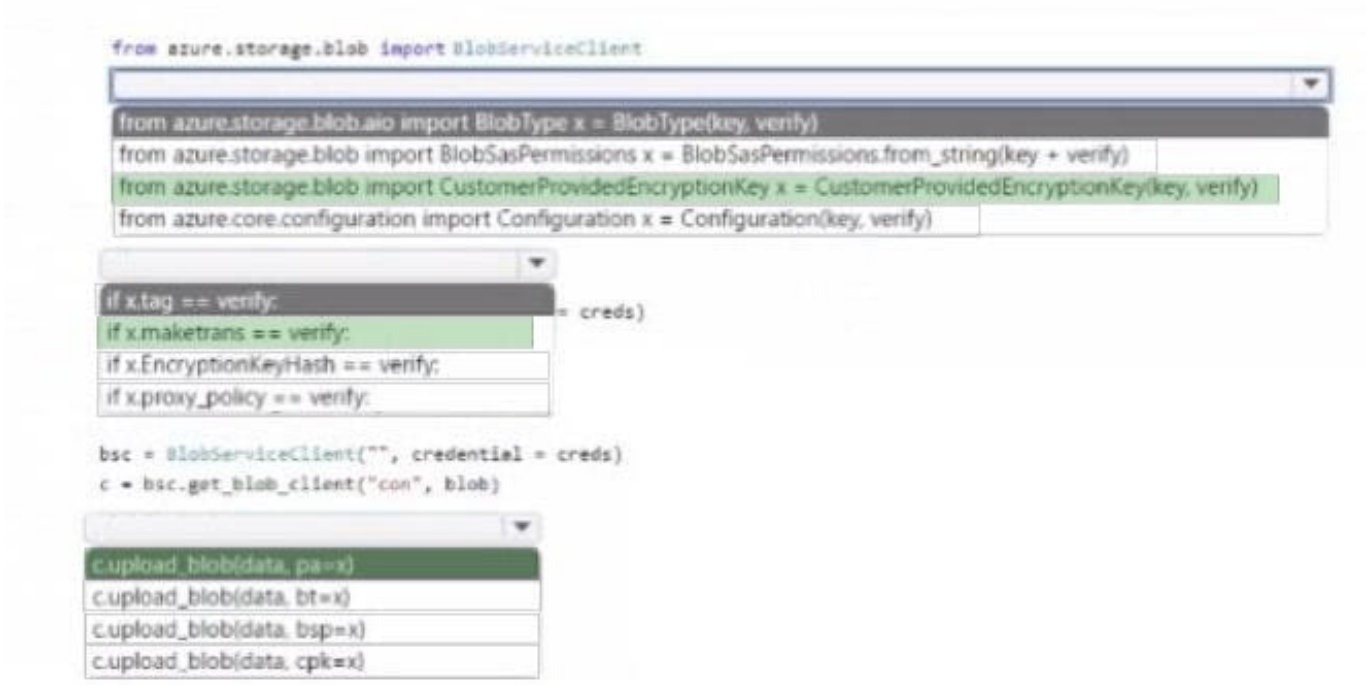

QUESTION 2

DRAG DROP

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

You have the following requirements:

1.

Secret-based authentication mechanisms are not permitted for accessing an Azure Storage account.

2.

Must use only Azure Instance Metadata Service endpoints.

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

between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Code segment 1

```
http://localhost:50342/oauth2/token
```

```
http://169.254.169.254:50432/oauth2/token
```

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

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

Code segment 2

```
XDocument.Parse(payload);
```

```
new MultipartContent(payload);
```

```
new NetworkCredential("Azure", payload);
```

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

Answer Area

```
var url = "  Code segment 1  " ;

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

return  Code segment 2
```

Correct Answer:

Code segment 1

```
http://localhost:50342/oauth2/token
http://169.254.169.254:50432/oauth2/token
http://localhost/metadata/identity/oauth2/token
```

Code segment 2

```
XDocument.Parse(payload);
new MultipartContent(payload);
new NetworkCredential("Azure", payload);
```

Answer Area

```
var url = " http://169.254.169.254/metadata/identity/oauth2/token " ;
var queryString = "...";
var client = new HttpClient();
var response = await client.GetAsync(url + queryString);
var payload = await response.Content.ReadAsStringAsync();
return JsonConvert.DeserializeObject<Dictionary<string, string>>(payload);
```

Azure Instance Metadata Service endpoints "/oauth2/token"

Box 1: http://169.254.169.254/metadata/identity/oauth2/token Sample request using the Azure Instance Metadata Service (IMDS) endpoint (recommended): GET \\'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02- 01andresource=https://management.azure.com/\\' HTTP/1.1 Metadata: true Box 2: JsonConvert.DeserializeObject(payload); Deserialized token response; returning access code.

QUESTION 3

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

You need to analyse app uptime for each month.

Which two solutions will achieve the goal? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Monitor logs

B. Application Insights alerts

C. Azure Monitor metrics

D. Application Insights web tests

Correct Answer: AB

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

QUESTION 4

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

1.

Data must not be modified or deleted for a user-specified interval.

2.

Data must be protected from overwntes and deletes.

3.

Data must be written once and allowed to be read many times.

You need to protect the data fen the Azure Blob storage account.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Configure a time-based retention policy for the storage account.

B. Create an account shared-access signature (SAS).

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

D. Enable version-level immutability support for the storage account.

E. Enable point-in-time restore for containers in the storage account.

F. Create a service shared-access signature (SAS).

Correct Answer: AD

A. Configure a time-based retention policy for the storage account

-

A time-based retention policy stores blob data in a Write-Once, Read-Many (WORM) format for a specified interval. When a time-based retention policy is set, clients can create and read blobs, but can\\'t modify or delete them. After the retention interval has expired, blobs can be deleted but not overwritten. https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-time-based-retention-policy-overview

D.

Before you can apply a time-based retention policy to a blob version, you must enable support for version-level immutability. https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-policy-configure-version-scope?tabs=azure-portal

QUESTION 5

HOTSPOT

You are developing an application that needs access to an Azure virtual machine (VM).

The access lifecycle for the application must be associated with the VM service instance.

You need to enable managed identity for the VM.

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

NOTE: Each correct selection is worth one point.

Hot Area:

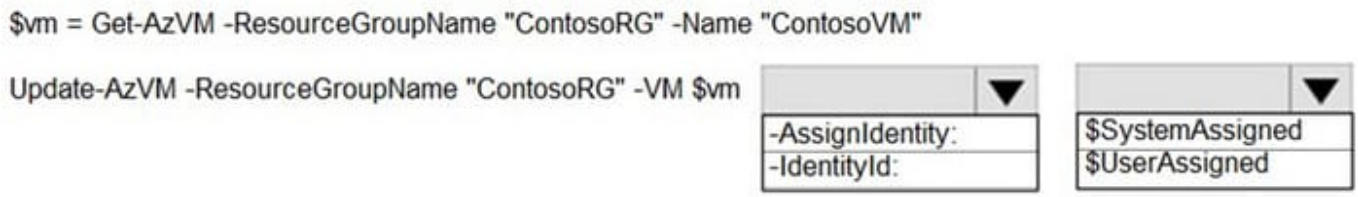

Correct Answer:

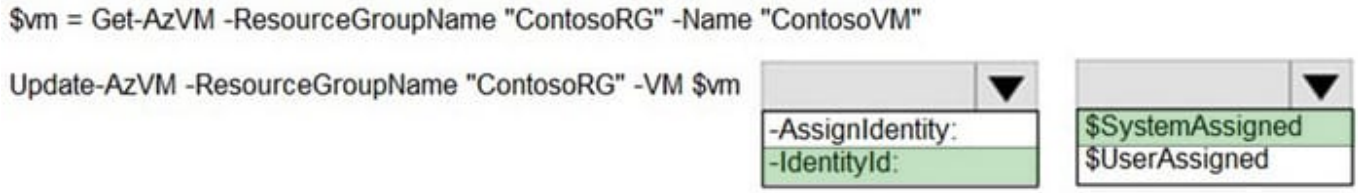

Box 1: -IdentityType

Enable system-assigned managed identity on an existing Azure VM:

To enable a system-assigned managed identity, use the -IdentityType switch on the Update-AzVM cmdlet (see below).

Box 2: $SystemAssigned

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

Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -IdentityType SystemAssigned

Reference:

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

---

QUESTION 6

DRAG DROP

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

Spikes in traffic have caused increases in page load times.

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

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

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

Select and Place:

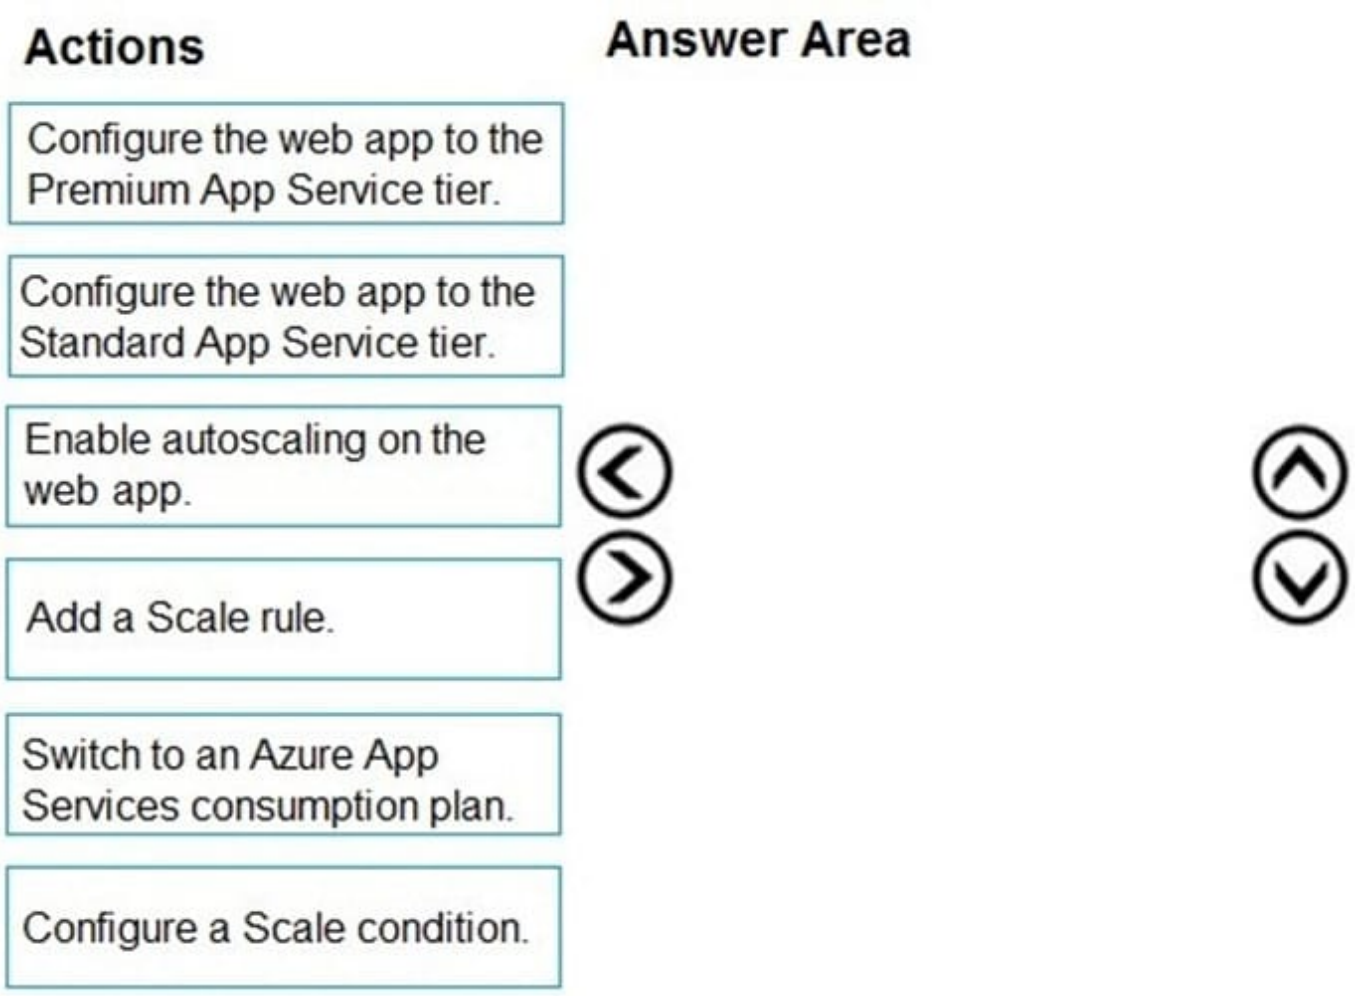

Correct Answer:

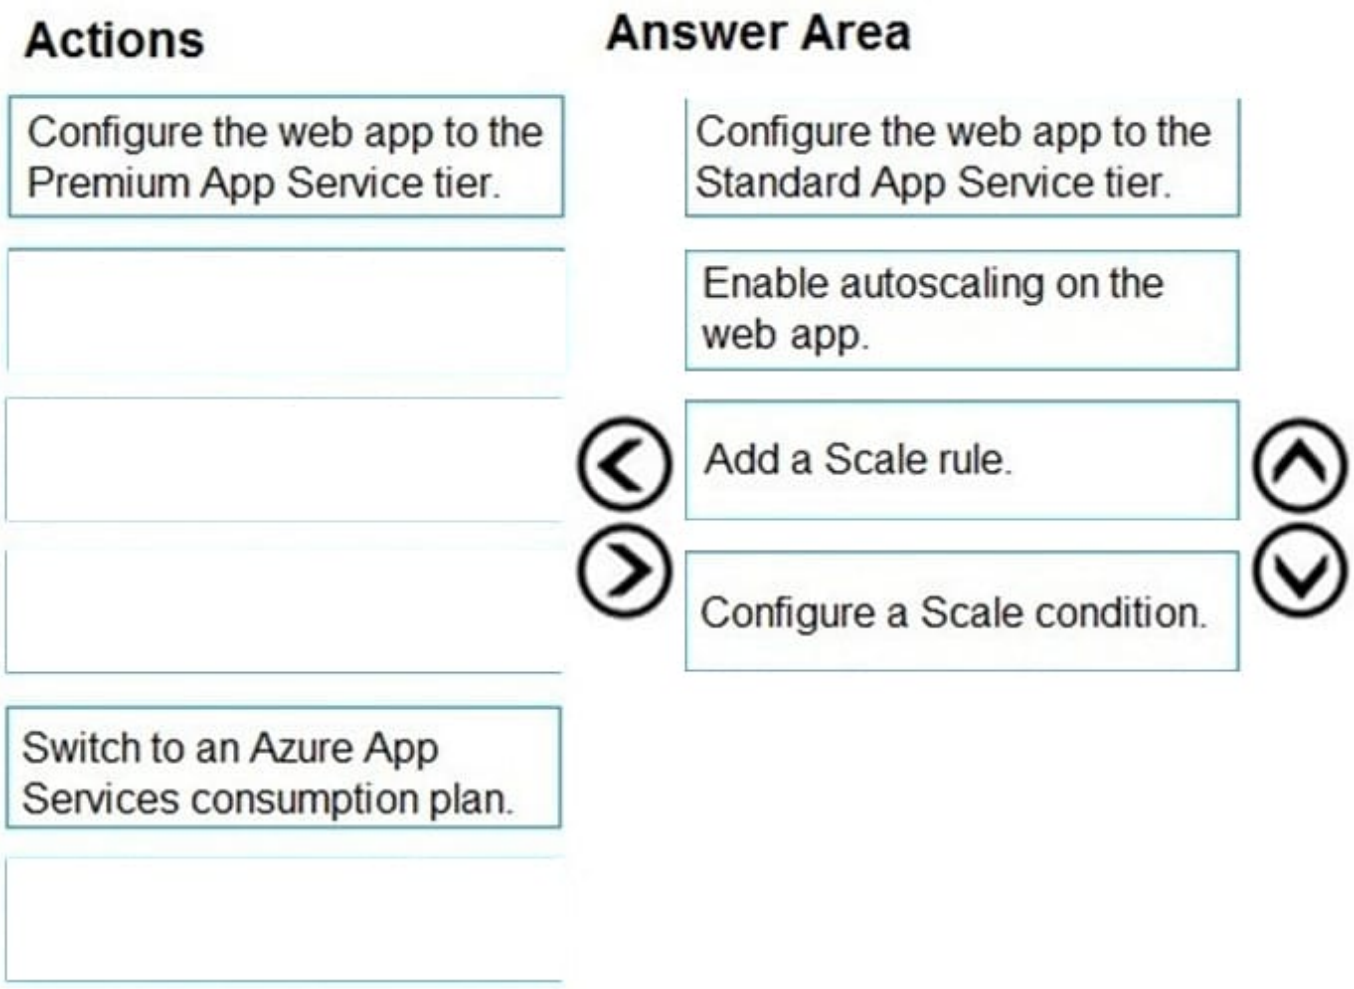

Step 1: Configure the web app to the Standard App Service Tier

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

https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoringautoscale-get-started

https://azure.microsoft.com/en-us/pricing/details/app-service/plans/

QUESTION 7

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:

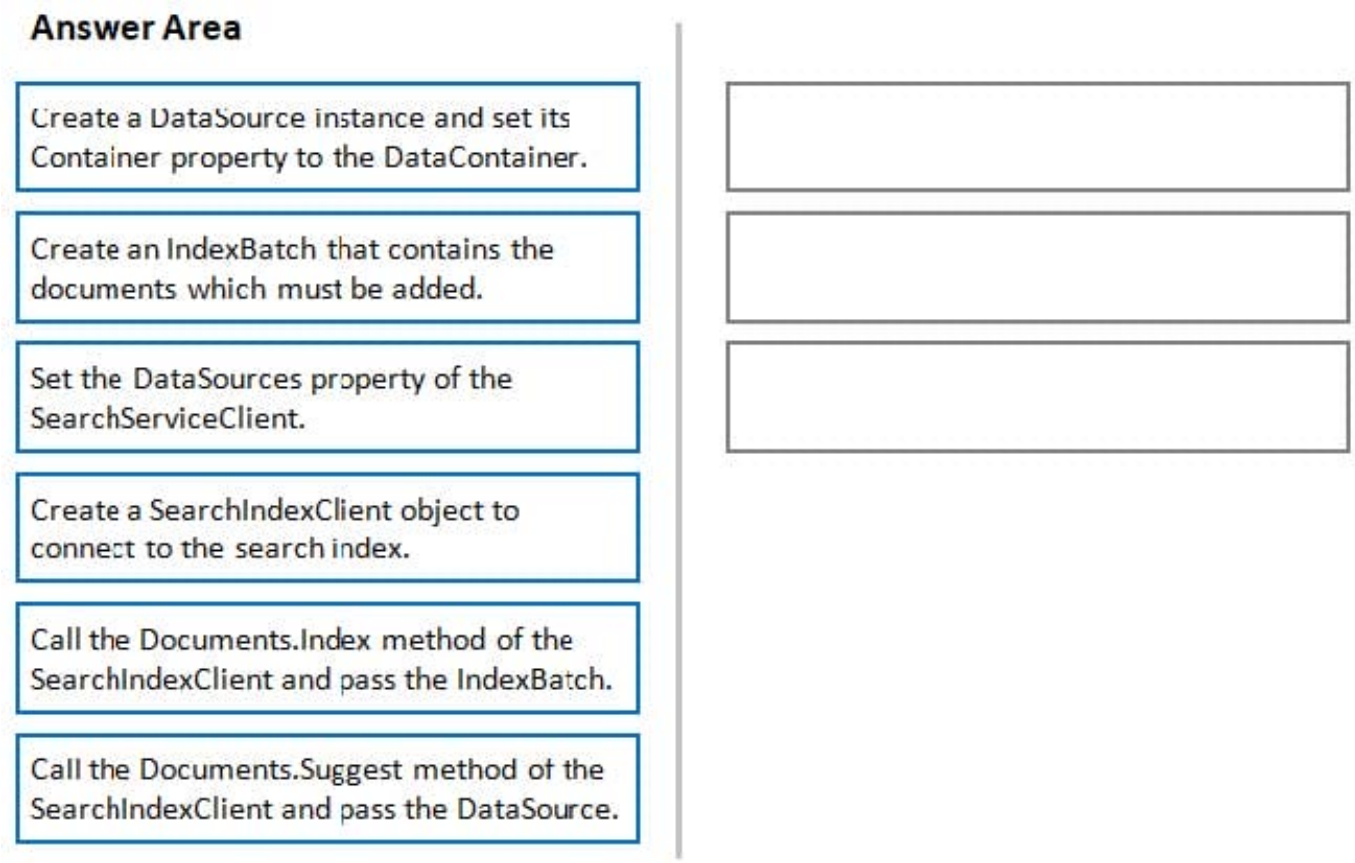

Correct Answer:

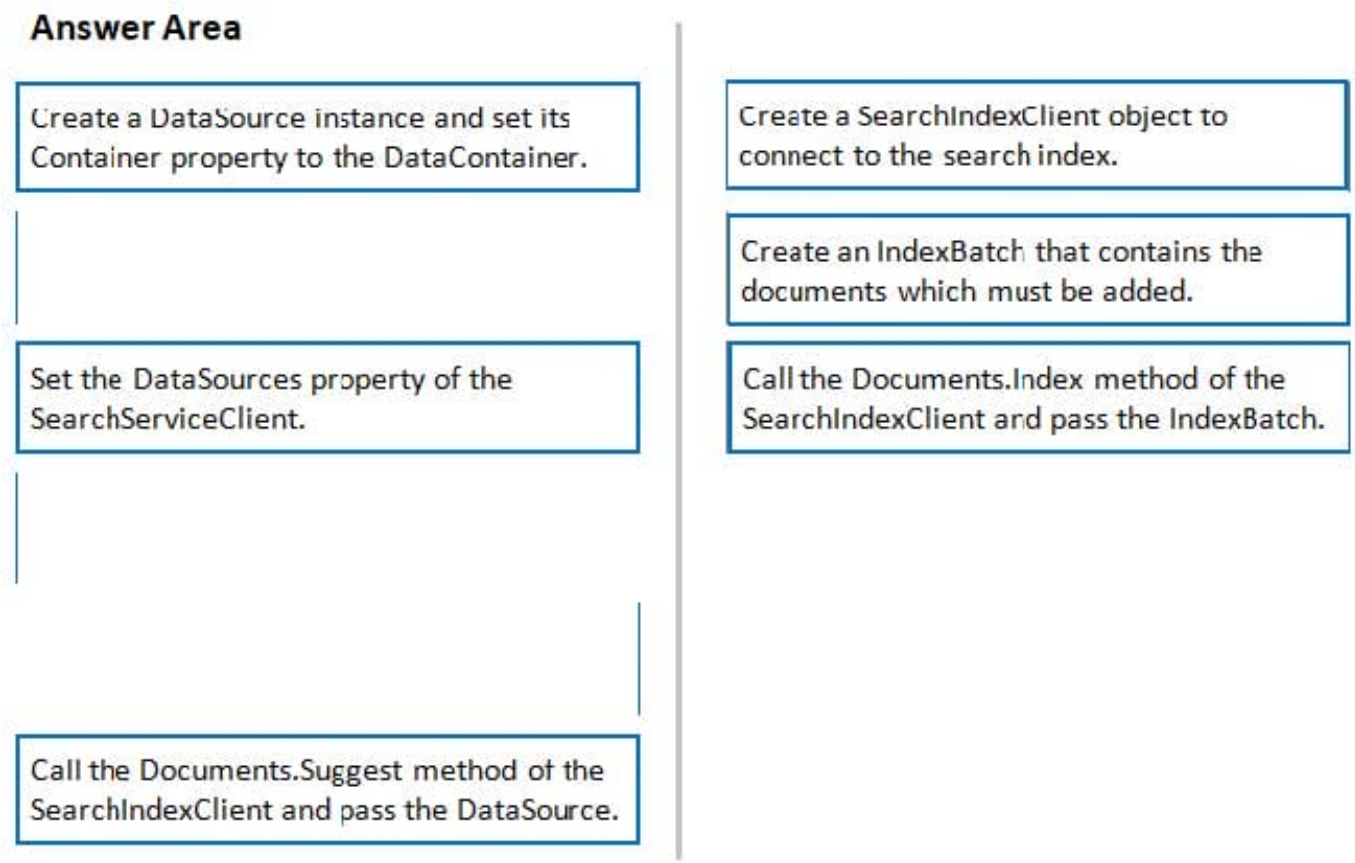

1.

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);

 }
```

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

QUESTION 8

You are developing a web application that runs as an Azure Web App. The web application stores data in Azure SQL Database and stores files in an Azure Storage account. The web application makes HTTP requests to external services as part of normal operations.

The web application is instrumented with Application Insights. The external services are OpenTelemetry compliant.

You need to ensure that the customer ID of the signed in user is associated with all operations throughout the overall system.

What should you do?

A. Add the customer ID for the signed in user to the CorrelationContext in the web application

B. On the current SpanContext, set the TraceId to the customer ID for the signed in user

C. Set the header Ocp-Apim-Trace to the customer ID for the signed in user

D. Create a new SpanContext with the TraceFlags value set to the customer ID for the signed in user

Correct Answer: A

Reference: https://docs.microsoft.com/en-us/azure/azure-monitor/app/correlation

QUESTION 9

HOTSPOT

You are developing an application that includes two Docker containers.

The application must meet the following requirements

The containers must not run as root.

The containers must be deployed to Azure Container Instances by using a YAML file. The containers must share a lifecycle, resources, local network and storage volume. The storage volume must persist through container crashes.

The storage volume must be destroyed on stop or restart of the containers. You need to configure Azure Container

Instances for the application.

Hot Area:

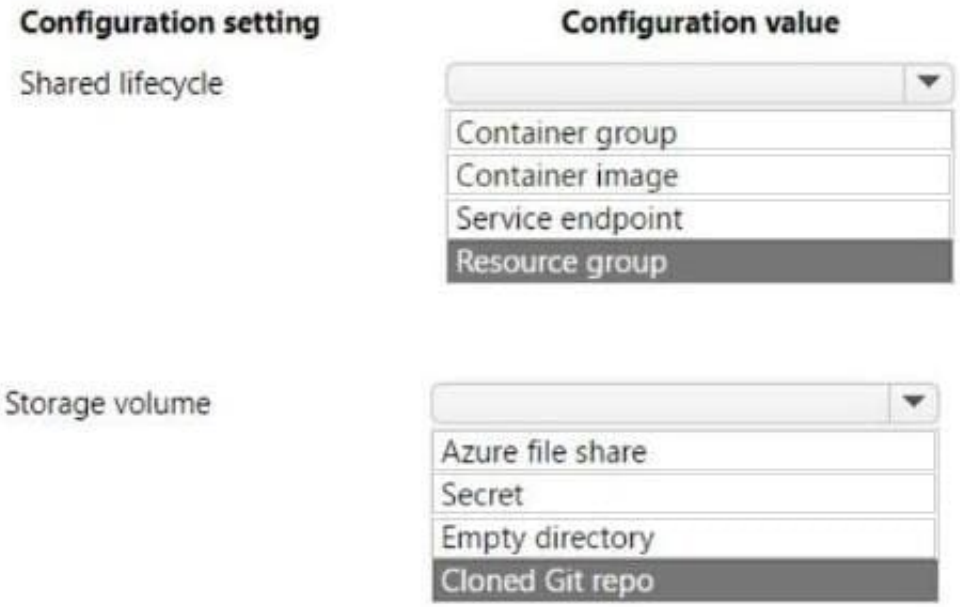

Correct Answer:

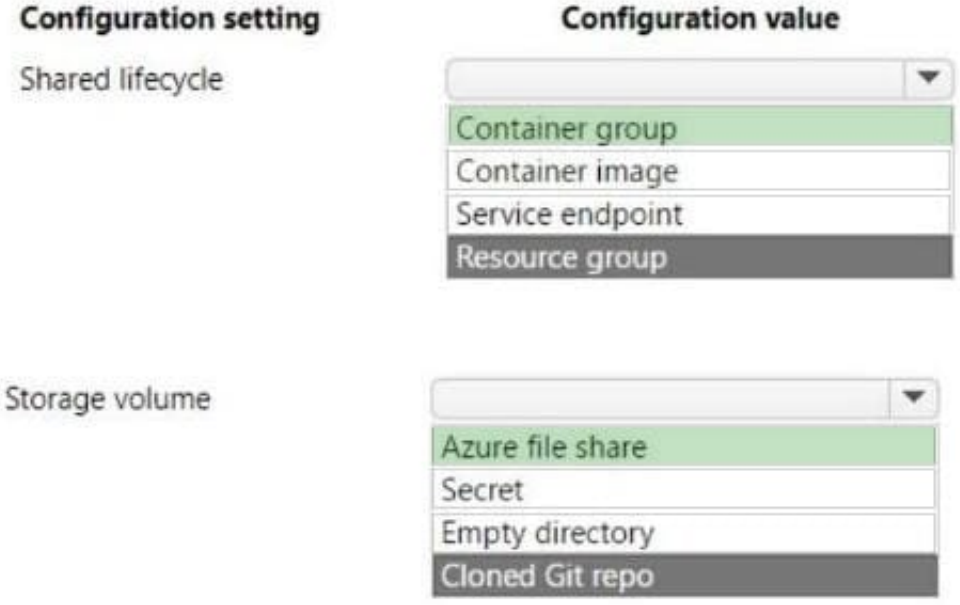

QUESTION 10

HOTSPOT

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

You have the following requirements:

1.

Code must be scalable and allow for growth.

2.

Each record must consist of a playerId, gameId, score, and time played.

3.

When users reach a new high score, the system will save the new score using the SaveScore function below.

4.

Each game is assigned an Id based on the series title.

You plan to store customer information in Azure Cosmos DB. 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

You develop the following code to save scores in the data store. (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 }
```

You develop the following code to query the database. (Line numbers are included for reference only.)

```
01 CloudTableClient tableClient = account.CreateCloudTableClient();
02 CloudTable table = tableClient.GetTableReference ("people");
03 TableQuery < CustomerEntity> query = new TableQuery <CustomerEntity >( )
04 .where( TableQuery.CombineFilters (
05 TableQuery.GenerateFilterCondition (PartitionKey, QueryComparisons.Equal , "Smith"),
06 TableOperators.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.

Hot Area:

Answer Area

Statements	Yes	No
SaveScore will work with Cosmos DB.	○	○
SaveScore will update and replace a record if one already exists with the same playerId and gameId.	○	○
Leader board data for the game will be automatically partitioned using gameId.	○	○
SaveScore will store the values for the gameId and playerId parameters in the database.	○	○

Correct Answer:

Answer Area

Statements	Yes	No
SaveScore will work with Cosmos DB.	○	○
SaveScore will update and replace a record if one already exists with the same playerId and gameId.	○	○
Leader board data for the game will be automatically partitioned using gameId.	○	○
SaveScore will store the values for the gameId and playerId parameters in the database.	○	○

Box 1: Yes

Create a table.

A CloudTableClient object lets you get reference objects for tables and entities. The following code creates a CloudTableClient object and uses it to create a new CloudTable object, which represents a table

// Retrieve storage account from connection-string.

```
CloudStorageAccount storageAccount =

CloudStorageAccount.parse(storageConnectionString);

// Create the table client.

CloudTableClient tableClient = storageAccount.createCloudTableClient();

// Create the table if it doesn\\'t exist.

String tableName = "people";

CloudTable cloudTable = tableClient.getTableReference(tableName);

cloudTable.createIfNotExists();
```

Box 2: No

New records are inserted with TableOperation.insert. Old records are not updated.

To update old records TableOperation.insertOrReplace should be used instead.

Box 3: No

Box 4: Yes

Reference:

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

---

QUESTION 11

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

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

You 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:

1.

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

2.

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

3.

Minimize Azure costs.

You need to implement the messaging solution.

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

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

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

QUESTION 12

You need to migrate on-premises shipping data to Azure. What should you use?

A. Azure Migrate

B. Azure Cosmos DB Data Migration tool (dt.exe)

C. AzCopy

D. Azure Database Migration service

Correct Answer: D

Migrate from on-premises or cloud implementations of MongoDB to Azure Cosmos DB with minimal downtime by using Azure Database Migration Service. Perform resilient migrations of MongoDB data at scale and with high reliability.

Scenario: Data migration from on-premises to Azure must minimize costs and downtime.

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

References:

https://azure.microsoft.com/en-us/updates/mongodb-to-azure-cosmos-db-online-andoffline-migrations-are-now-available/

QUESTION 13

DRAG DROP

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

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

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

Select and Place:

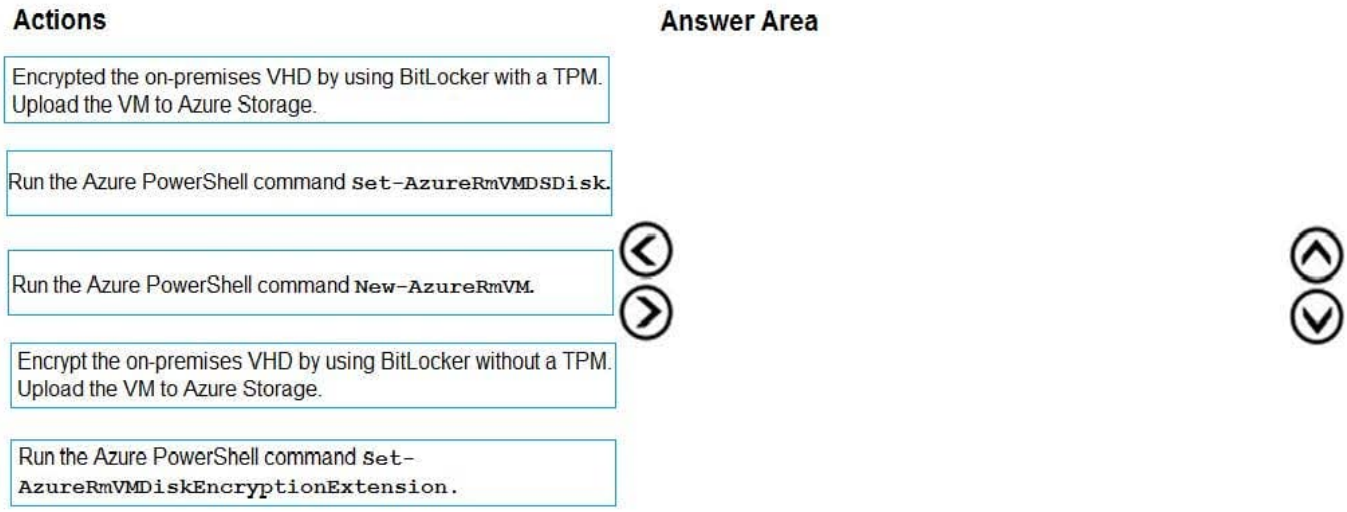

Correct Answer:

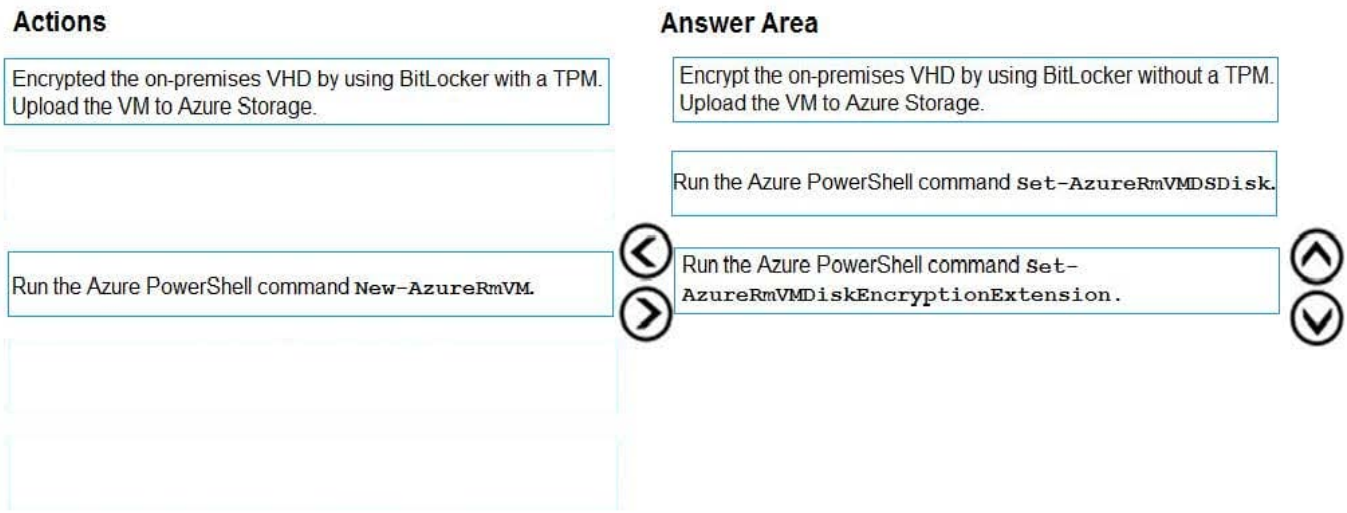

Step 1: Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage

Step 2: Run the Azure PowerShell command Set-AzureRMVMOSDisk

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

Example:

$osDiskName = $vmname+\\'_osDisk\\'

$osDiskCaching = \\'ReadWrite\\'

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

$vm = Set-AzureRmVMOSDisk -VM $vm -VhdUri $osDiskVhdUri -name $osDiskName -Create

Step 3: Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension

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

Incorrect:

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

implementation. A VM does not have a TPM.

References:

https://www.itprotoday.com/iaaspaas/use-existing-vhd-azurerm-vm

---

QUESTION 14

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

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

You are developing a mobile app that uses an Azure SQL Database named Weyland.

The database contains a table names Customers that has a field named email_address.

You want to implement dynamic data masking to hide the data in the email_address field.

Solution: You run the Set-AzSqlDatabaseDataMaskingPolicy -DatabaseName "Weyland" Powershell cmdlet

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

The Set-AzSqlDatabaseDataMaskingPolicy cmdlet sets the masking policy for the database but does not set the masking rule.

Note: Masking is enabled by default.

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/dynamic-data-masking-overview

https://docs.microsoft.com/en-us/powershell/module/az.sql/set-azsqldatabasedatamaskingrule?view=azps- 4.2.0

QUESTION 15

DRAG DROP

You are developing a REST web service. Customers will access the service by using an Azure API Management instance.

The web service does not correctly handle conflicts. Instead of returning an HTTP status code of 409, the service returns a status code of 500. The body of the status message contains only the word conflict.

You need to ensure that conflicts produce the correct response.

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

NOTE: Each correct selection is worth one point.

Select and Place:

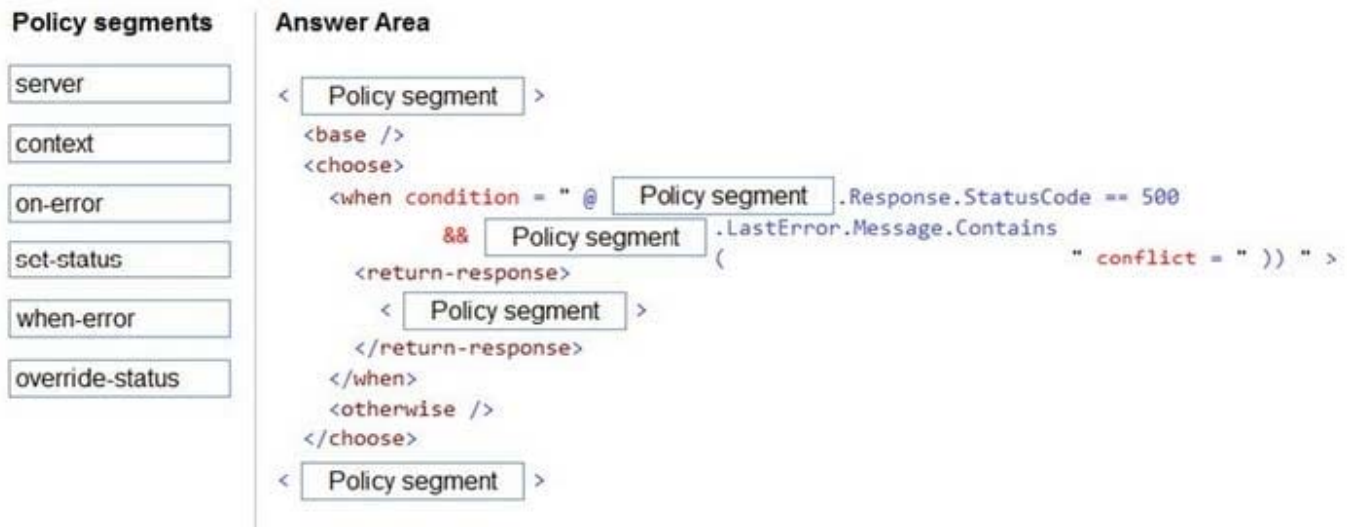

Correct Answer:

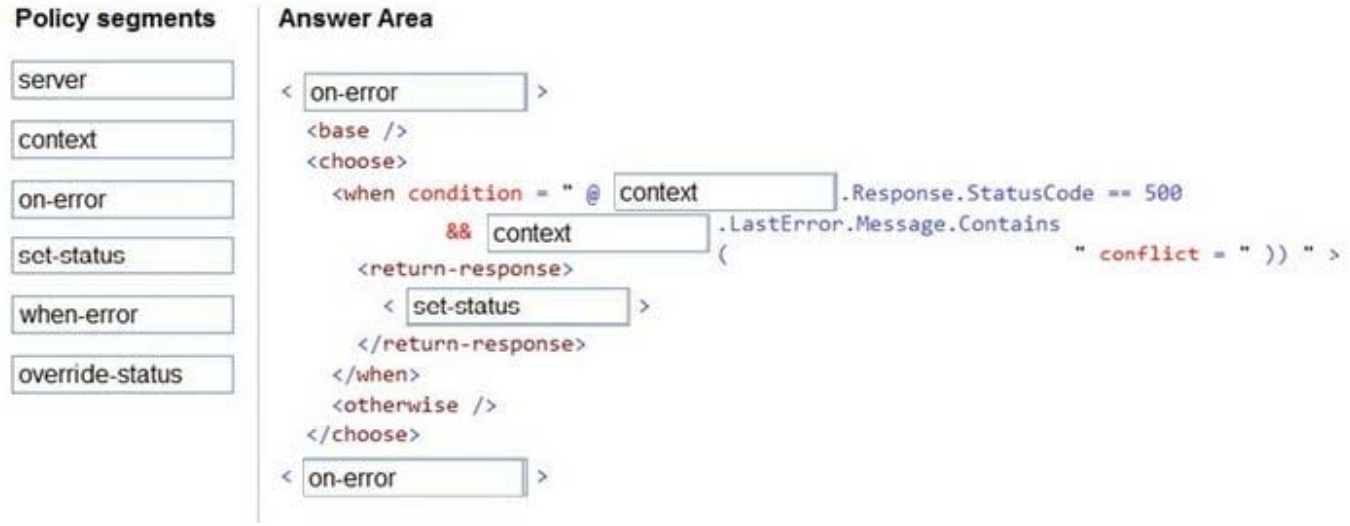

Box 1: on-error

Policies in Azure API Management are divided into inbound, backend, outbound, and on- error.

If there is no on-error section, callers will receive 400 or 500 HTTP response messages if an error condition occurs.

Box 2: context

Box 3: context

Box 4: set-status

The return-response policy aborts pipeline execution and returns either a default or custom response to the caller. Default response is 200 OK with no body. Custom response can be specified via a context variable or policy statements.

Syntax:

Box 5: on-error