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.pass4itsure.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

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.

Correct Answer: B

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

QUESTION 2

HOTSPOT

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.

Hot Area:

Answer Area

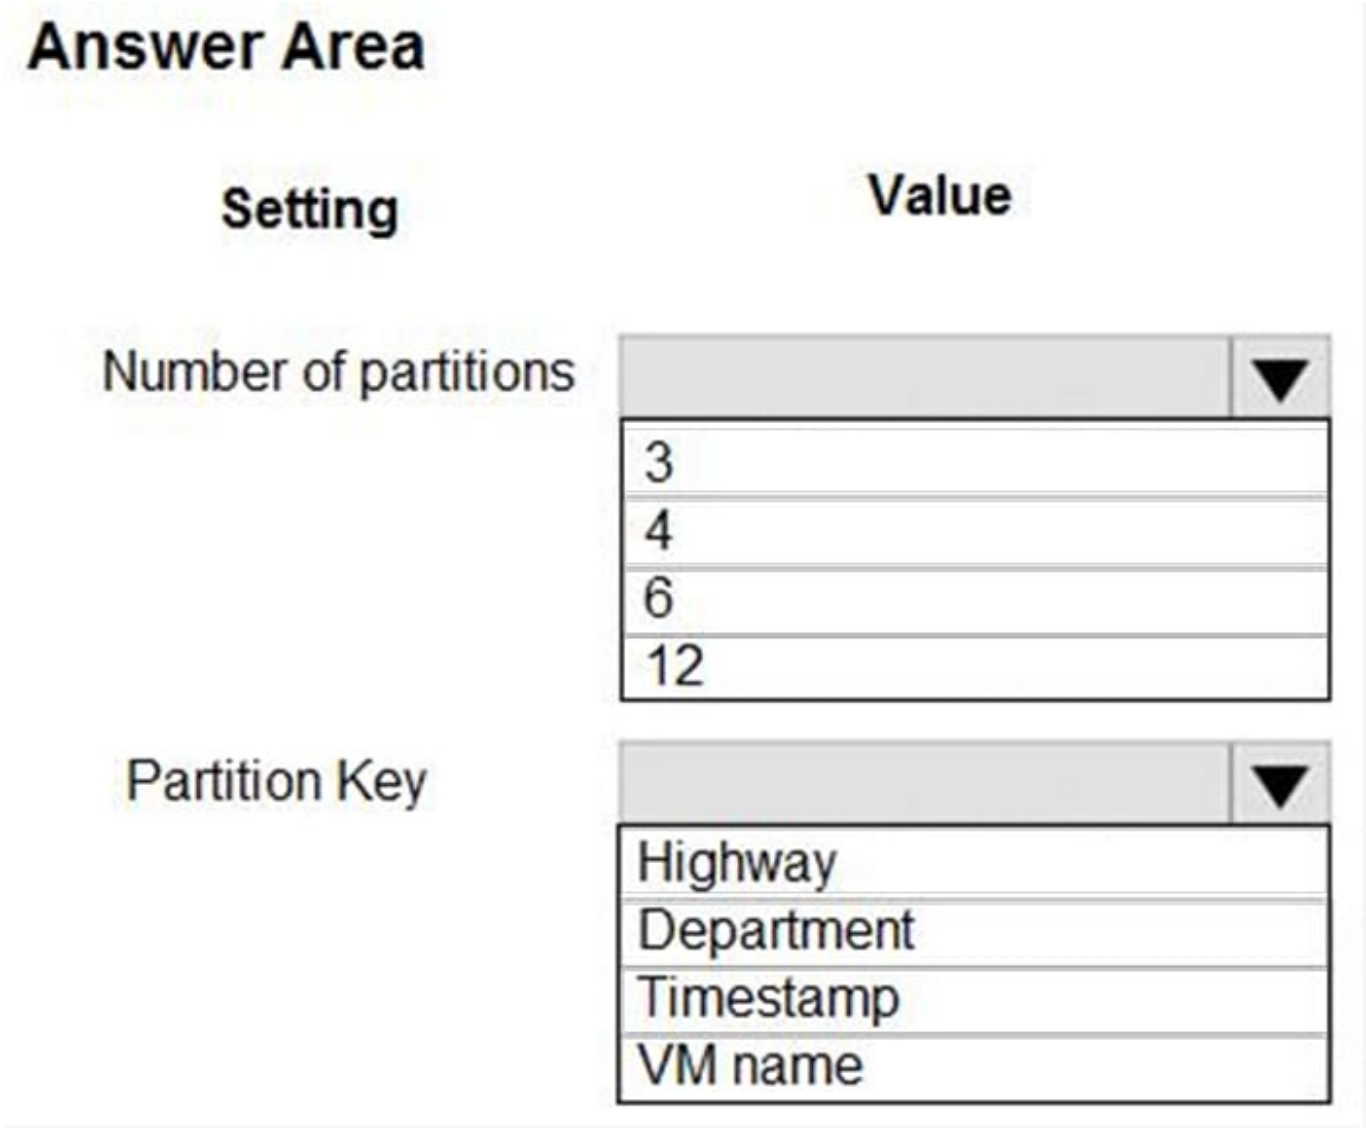

Correct Answer:

Answer Area

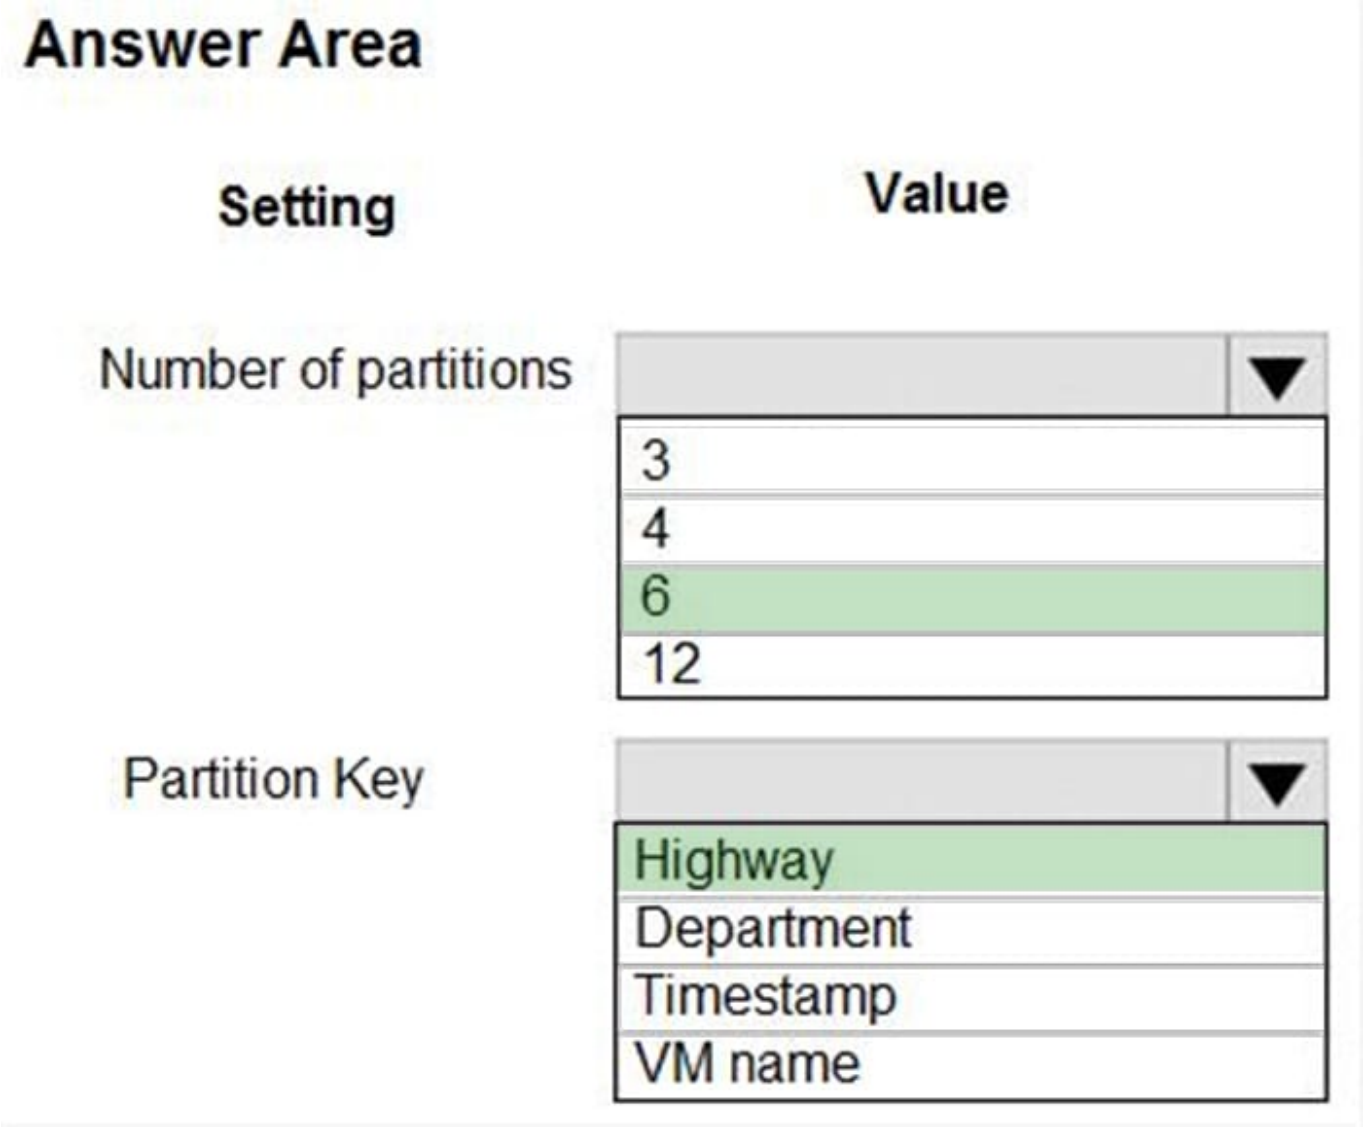

Box 1: 6

The number of partitions is specified at creation and must be between 2 and 32.

There are 6 highways.

Box 2: Highway

Reference:

https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features

QUESTION 3

DRAG DROP

You have an application that uses Azure Blob storage.

You need to update the metadata of the blobs.

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

Select and Place:

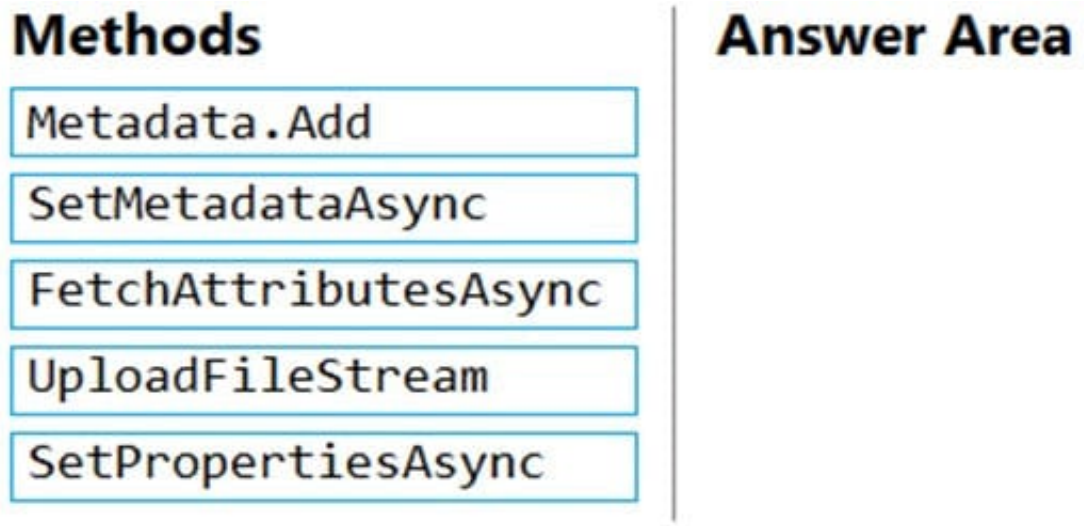

Correct Answer:

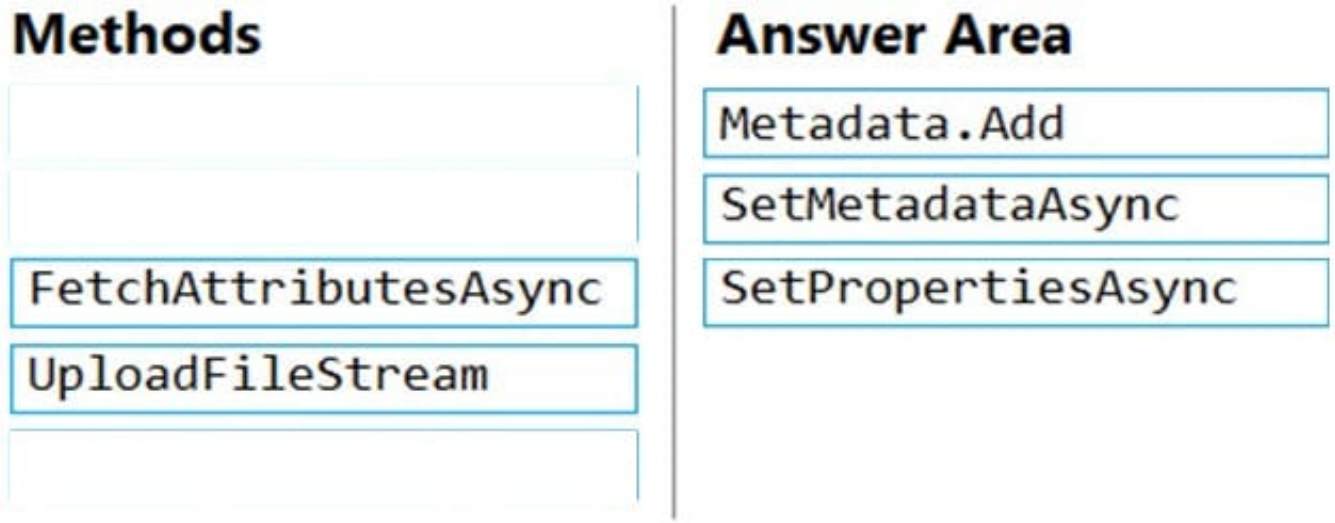

Metadata.Add example: // Add metadata to the dictionary by calling the Add method metadata.Add("docType", "textDocuments");

SetMetadataAsync example: // Set the blob\\'s metadata. await blob.SetMetadataAsync(metadata);

// Set the blob\\'s properties. await blob.SetPropertiesAsync();

Reference: https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-properties-metadata

QUESTION 4

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

1.

Every five minutes verify that the websites are responsive.

2.

Verify that the websites respond within a specified time threshold. Dependent requests such as images and JavaScript files must load properly.

3.

Generate alerts if a website is experiencing issues.

4.

If a website fails to load, the system must attempt to reload the site three more times.

You need to implement this process with the least amount of effort.

What should you do?

A. Create a Selenium web test and configure it to run from your workstation as a scheduled task.

B. Set up a URL ping test to query the home page.

C. Create an Azure function to query the home page.

D. Create a multi-step web test to query the home page.

E. Create a Custom Track Availability Test to query the home page.

Correct Answer: D

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

A: Selenium is an umbrella project for a range of tools and libraries that enable and support the automation of web browsers.

It provides extensions to emulate user interaction with browsers, a distribution server for scaling browser allocation, and the infrastructure for implementations of the W3C WebDriver specification that lets you write interchangeable code for all major web browsers.

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

QUESTION 5

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:

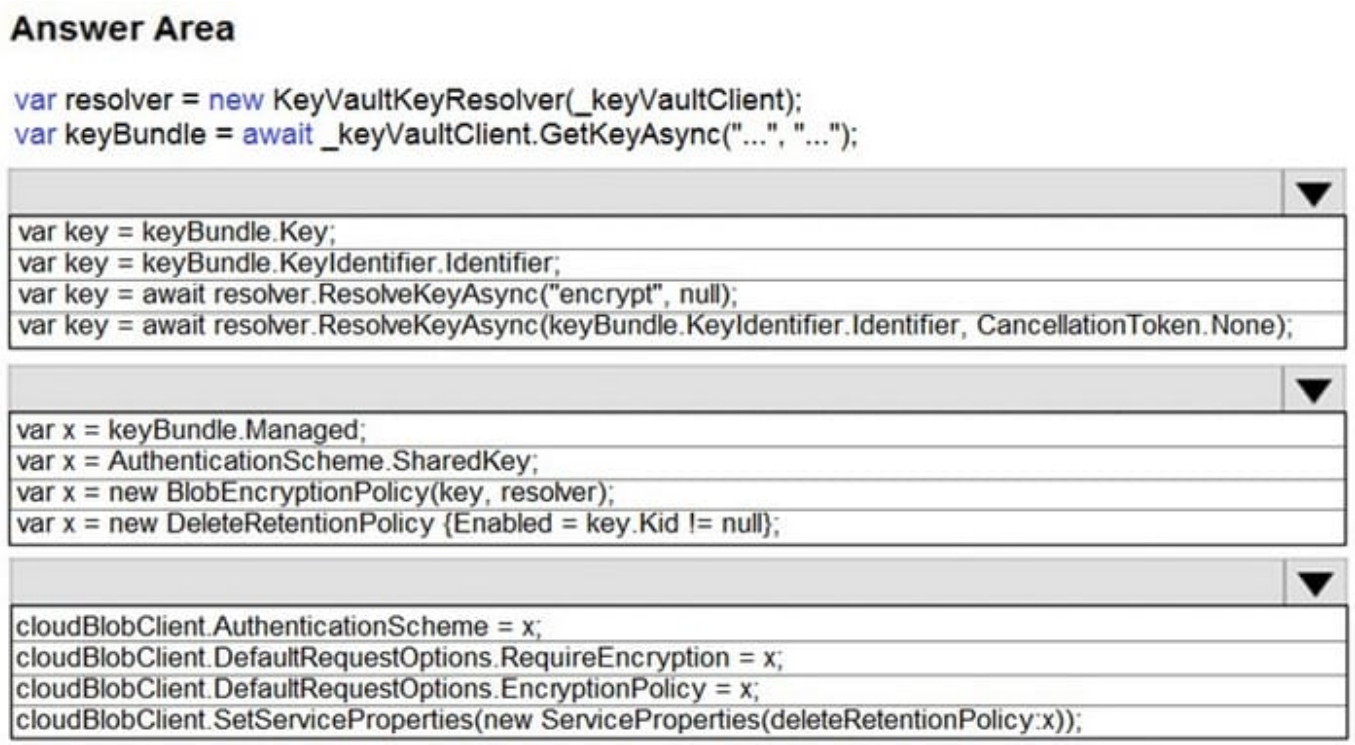

Correct Answer:

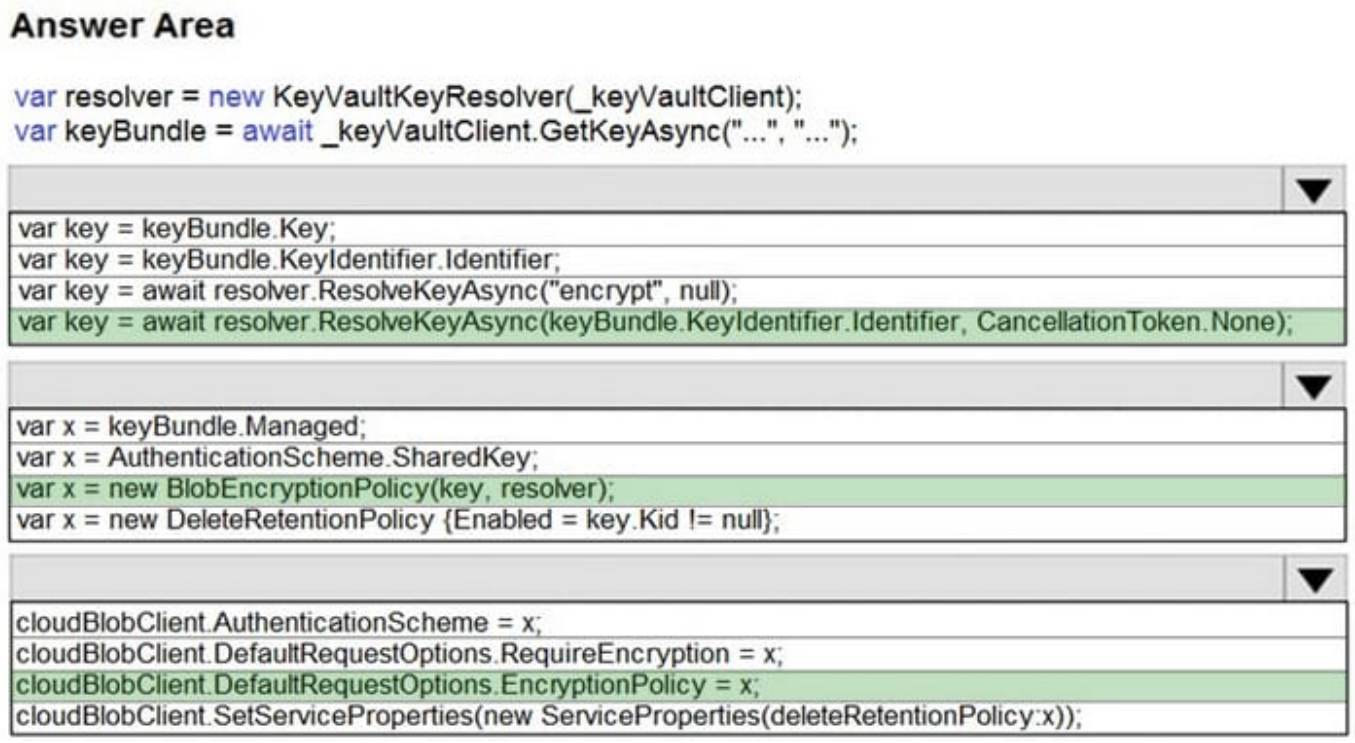

Box 1: var key = await Resolver.ResolveKeyAsyn(keyBundle,KeyIdentifier.CancellationToken.None);

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 6

HOTSPOT

You need to ensure that validation testing is triggered per the requirements.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
var event = getEvent();
if (event.eventType === '                    '
                         ImagePushed
                         RepositoryItem
                         ImageDeployed
                         RepositoryUpdated
 && event.data.target.                       === 'contentanalysisservice'
                         aci
                         image
                         service
                         repository
 && event.                    .contains('contosoimages'))
                         topic
                         service
                         repository
                         imageCollection
{
  startValidationTesting();
}
```

Correct Answer:

Answer Area

```
var event = getEvent();
if (event.eventType === '                    '
                         ImagePushed
                         RepositoryItem
                         ImageDeployed
                         RepositoryUpdated
 && event.data.target.                       === 'contentanalysisservice'
                         aci
                         image
                         service
                         repository
 && event.                                   .contains('contosoimages'))
                         topic
                         service
                         repository
                         imageCollection
{
  startValidationTesting();
}
```

Box 1: RepositoryUpdated

When a new version of the ContentAnalysisService is available the previous seven days of content must be processed with the new version to verify that the new version does not significantly deviate from the old version.

Box 2: service

Box 3: imageCollection

Reference: https://docs.microsoft.com/en-us/azure/devops/notifications/oob-supported-event-types

QUESTION 7

You need to reduce read latency for the retail store solution.

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

NOTE: Each correct selection is worth one point.

A. Create a new composite index for the store location data queries in Azure Cosmos DB. Modify the queries to support parameterized SQL and update the Azure function app to call the new Queries.

B. Configure Azure Cosmos DB consistency to strong consistency Increase the RUs for the container supporting store location data.

C. Provision an Azure Cosmos OB dedicated gateway, update blob storage to use the new dedicated gateway endpoint.

D. Configure Azure Cosmos DB consistency to session consistency. Cache session tokens in a new Azure Redis cache instance after every write. Update reads to use the session token stored in Azure Redis.

E. Provision an Azure Cosmos DB dedicated gateway Update the Azure Function app connection string to use the new dedicated gateway endpoint.

Correct Answer: AE

Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

Reference:

https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/data-cache-with-redis-cache
https://docs.microsoft.com/en-us/azure/cosmos-db/dedicated-gateway

QUESTION 8

DRAG DROP

You manage several existing Logic Apps.

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

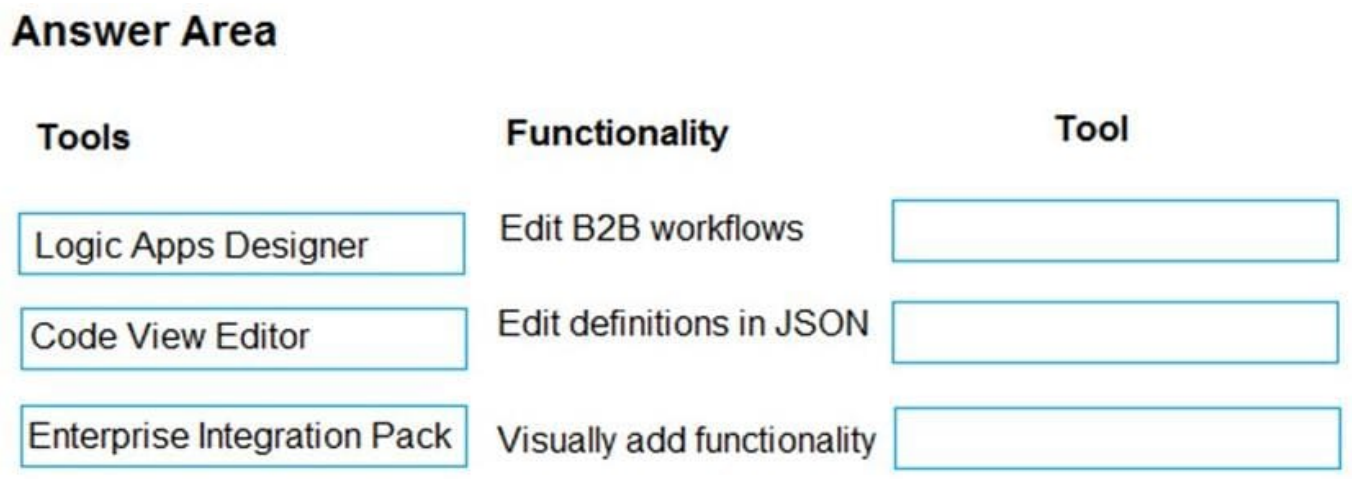

Correct Answer:

Answer Area

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

QUESTION 9

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 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

Correct Answer: A

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

QUESTION 10

DRAG DROP

You are developing a microservices solution. You plan to deploy the solution to a multinode Azure Kubernetes Service (AKS) cluster.

You need to deploy a solution that includes the following features:

1.

reverse proxy capabilities

2.

configurable traffic routing

3.

TLS termination with a custom certificate

Which component 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 content.

NOTE: Each correct selection is worth one point.

Select and Place:

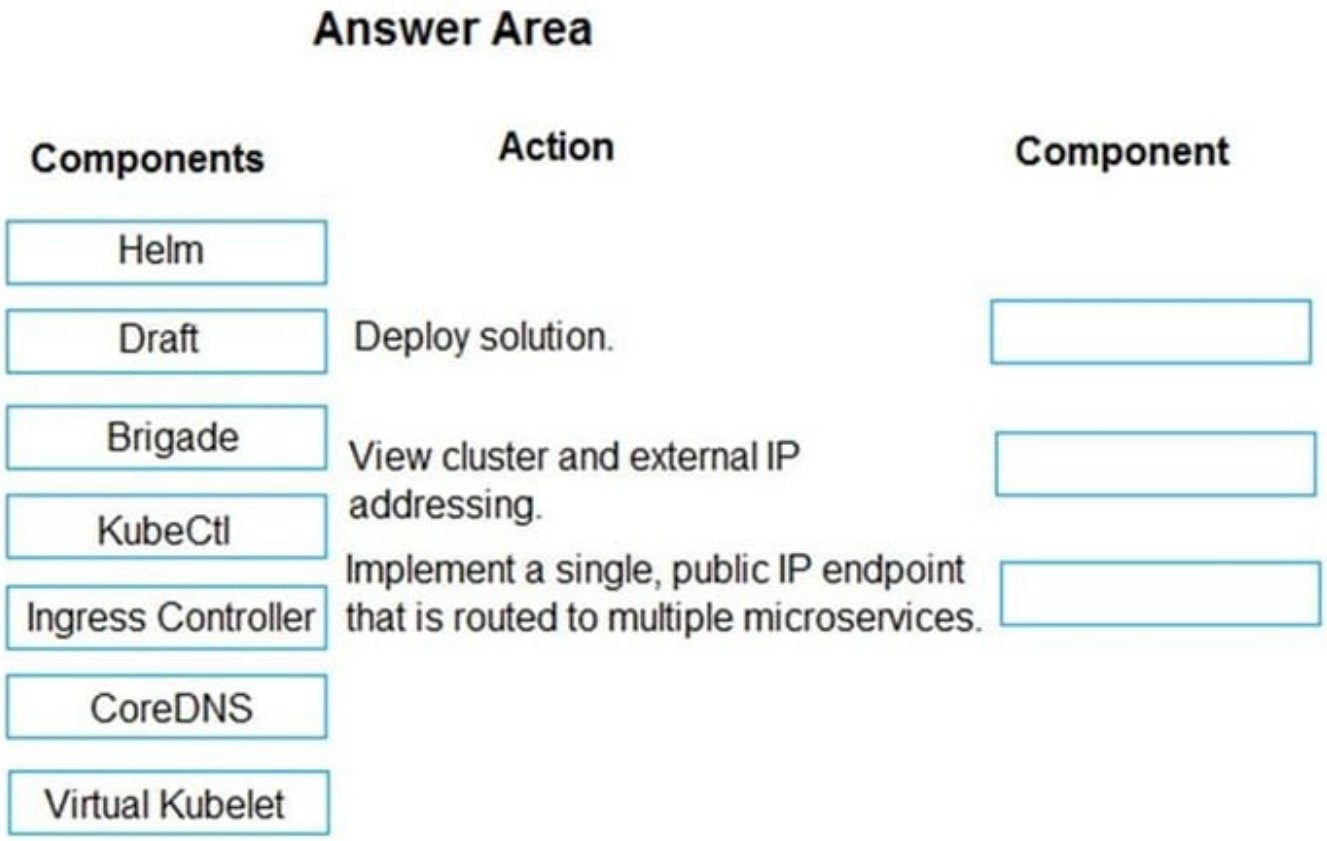

Correct Answer:

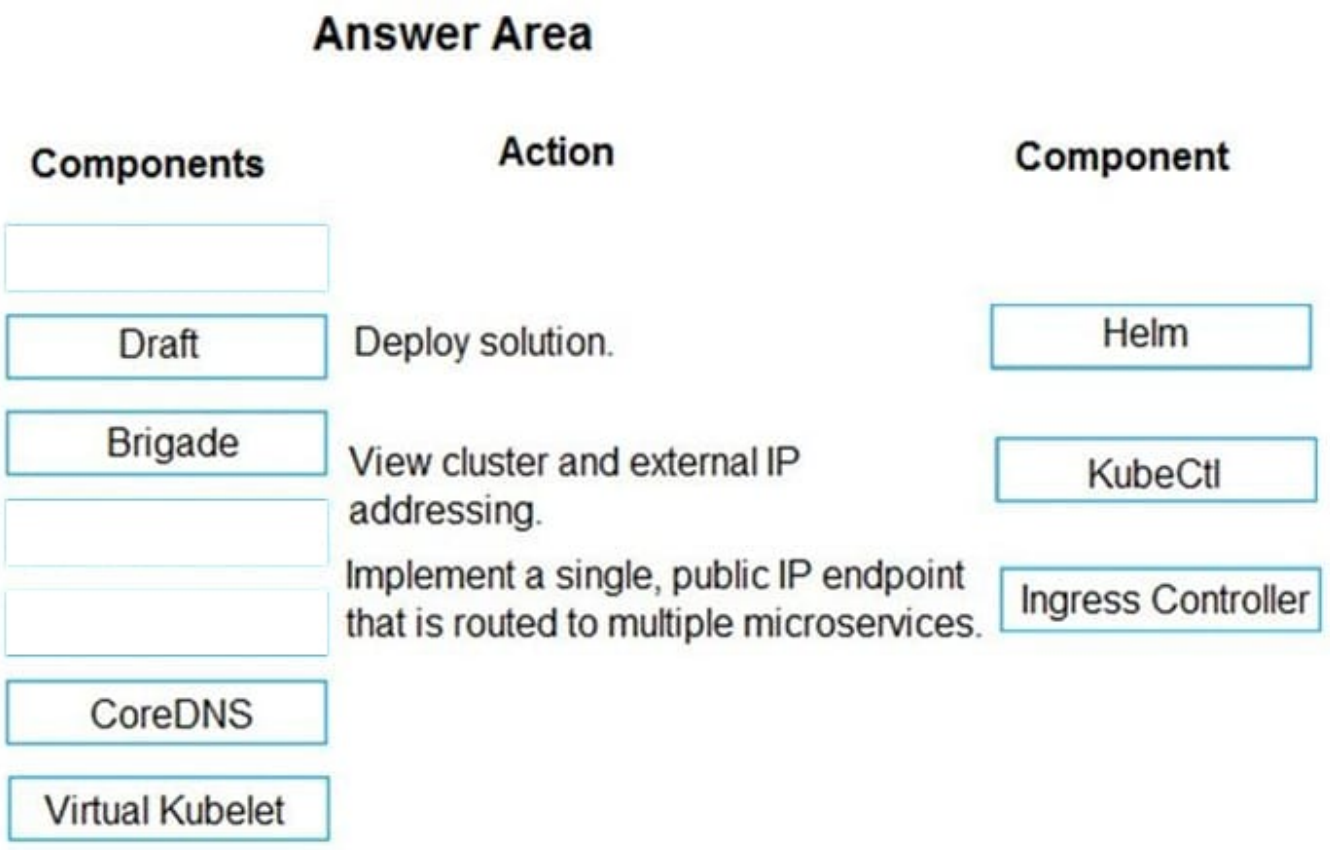

Box 1: Helm

To create the ingress controller, use Helm to install nginx-ingress.

Box 2: kubectl

To find the cluster IP address of a Kubernetes pod, use the kubectl get pod command on your local machine, with the option -o wide .

Box 3: Ingress Controller

An ingress controller is a piece of software that provides reverse proxy, configurable traffic routing, and TLS termination for Kubernetes services. Kubernetes ingress resources are used to configure the ingress rules and routes for individual

Kubernetes services.

Incorrect Answers:

Virtual Kubelet: Virtual Kubelet is an open-source Kubernetes kubelet implementation that masquerades as a kubelet. This allows Kubernetes nodes to be backed by Virtual Kubelet providers such as serverless cloud container platforms.

CoreDNS: CoreDNS is a flexible, extensible DNS server that can serve as the Kubernetes cluster DNS. Like Kubernetes, the CoreDNS project is hosted by the CNCF.

Reference:

https://docs.microsoft.com/bs-cyrl-ba/azure/aks/ingress-basic

https://www.digitalocean.com/community/tutorials/how-to-inspect-kubernetes-networking

QUESTION 11

You develop Azure solutions.

A .NET application needs to receive a message each time an Azure virtual machine finishes processing data. The messages must NOT persist after being processed by the receiving application.

You need to implement the .NET object that will receive the messages.

Which object should you use?

A. QueueClient

B. SubscriptionClient

C. TopicClient

D. CloudQueueClient

Correct Answer: A

A queue allows processing of a message by a single consumer. Need a CloudQueueClient to access the Azure VM.

Incorrect Answers:

B, C: In contrast to queues, topics and subscriptions provide a one-to-many form of communication in a publish and subscribe pattern. It\\'s useful for scaling to large numbers of recipients.

Reference:

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-queues-topics-subscriptions

QUESTION 12

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 06 API should you use for the application?

A. Core

B. Mongo DB

C. Gremlin

D. Table API

Correct Answer: D

QUESTION 13

You are developing an inventory tracking solution. The solution includes an Azure Function app containing multiple functions triggered by Azure Cosmos DB. You plan to deploy the solution to multiple Azure regions.

The solution must meet the following requirements:

Item results from Azure Cosmos DS must return the most recent committed version of an item.

Items written to Azure Cosmos DB must provide ordering guarantees.

You need to configure the consistency level for the Azure Cosmos DB deployments.

Which consistency level should you use?

A. consistent prefix

B. eventual

C. bounded staleness

D. strong

E. session

Correct Answer: D

Explanation:

Strong consistency

Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. 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.

Incorrect:

* bounded staleness (and the other weaker consistency levels)

With Bounded Staleness consistency, reads issued against a non-primary region may not necessarily return the most recent version of the data globally, but are guaranteed to return the most recent version of the data in that region, which will

be within the maximum staleness boundary globally.

Note: Azure Cosmos DB offers five well-defined levels. From strongest to weakest, the levels are:

Strong Bounded staleness Session Consistent prefix Eventual

Reference: https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels

QUESTION 14

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.

Correct Answer: C

You can use alert processing rules to add action groups or remove (suppress) action groups from your fired alerts.
Reference: https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-action-rules

QUESTION 15

You need to resolve the capacity issue. What should you do?

A. Convert the trigger on the Azure Function to an Azure Blob storage trigger

B. Ensure that the consumption plan is configured correctly to allow scaling

C. Move the Azure Function to a dedicated App Service Plan

D. Update the loop starting on line PC09 to process items in parallel

Correct Answer: D

If you want to read the files in parallel, you cannot use forEach. Each of the async callback function calls does return a promise. You can await the array of promises that you\\'ll get with Promise.all. Scenario: Capacity issue: During busy periods, employees report long delays between the time they upload the receipt and when it appears in the web application.

```
PC08        var container = await GetCloudBlobContainer();
PC09        foreach (var fileItem in await ListFiles())
PC10        {
PC11          var file = new CloudFile(fileItem.StorageUri.PrimaryUri);
PC12          var ms = new MemoryStream();
PC13          await file.DownloadToStreamAsync(ms);
PC14          var blob = container.GetBlockBlobReference(fileItem.Uri.ToString());
PC15          await blob.UploadFromStreamAsync(ms);
PC16
PC17        }
```

Reference: https://stackoverflow.com/questions/37576685/using-async-await-with-a-foreach-loop