QUESTION 122
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

QUESTION 123
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: Deploy and configure an Azure Database for PostgreSQL. Update the web applications.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
The worst solution from a performance and scalability standpoint is to use a database backed session state provider. Instead use Azure Cache for Redis.

QUESTION 124
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: Deploy and configure Azure Cache for Redis. Update the web applications.
Does the solution meet the goal?

A. Yes
B. No

Answer: A
Explanation:
Azure Cache for Redis provides a session state provider that you can use to store your session state inmemory 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.
Reference: https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-aspnet-session-stateprovider

QUESTION 125
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 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 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 Hub. Configure the machine identifier as the partition key and enable capture.
Does the solution meet the goal?

A. Yes
B. No

Answer: A
Explanation:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-programming-guide

QUESTION 126
You are a developing a SaaS application that stores data as key value pairs.
You must make multiple editions of the application available. In the lowest cost edition, the performance must be best-effort, and there is no regional failover. In higher 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: C

QUESTION 127
Hotspot Question
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.

Answer Area

Requirement	Solution
Configure an Azure Storage account	Implement Blob Storage Implement Azure Cosmos DB Implement Storage (general purpose v1) Implement StorageV2 (general purpose v2)
Configure data retention	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

Answer:

Answer Area

Requirement	Solution
Configure an Azure Storage account	Implement Blob Storage Implement Azure Cosmos DB Implement Storage (general purpose v1) Implement StorageV2 (general purpose v2)
Configure data retention	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

Explanation:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers
https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy?toc=/azure/storage/blobs/toc.json

QUESTION 128
You develop and deploy a Java RESTful API to Azure App Service.
You open a browser and navigate to the URL for the API.
You receive the following error message:

```
Failed to load http://api.azurewebsites.net:6000/#/api/Products: No 'Access-
Control-Allow-Origin'header is present on the requested resource.
Origin 'http://localhost:6000' is therefore not allowed access
```

You need to resolve the error.
What should you do?

A. Bind an SSL certificate
B. Enable authentication
C. Enable CORS
D. Map a custom domain
E. Add a CDN

Answer: C
Explanation:
We need to enable Cross-Origin Resource Sharing (CORS).
References:
https://medium.com/@xinganwang/a-practical-guide-to-cors-51e8fd329a1f

QUESTION 129
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

QUESTION 130
You are building a web application that uses the Microsoft identity platform for user authentication.
You are implementing user identification for the web application.
You need to retrieve a claim to uniquely identify a user.
Which claim type should you use?

A. oid
B. aud
C. idp
D. nonce

Answer: A
Explanation:
oid - The object identifier for the user in Azure AD. This value is the immutable and non-reusable identifier of the user. Use this value, not email, as a unique identifier for users; email addresses can change. If you use the Azure AD Graph API in your app, object ID is that value used to query profile information.
Incorrect:
Not B: aud - Who the token was issued for. This will be the application's client ID.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/multitenant-identity/claims

QUESTION 131
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 website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.
You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level. You need to configure authorization.
Solution:
Create a new Azure AD application's manifest, set value of the groupMembershipClaims option to All.
In the website, use the value of the groups claim from the JWI for the user to determine permissions.
Does the solution meet the goal?

A. Yes
B. No

Answer: A
Explanation:
To configure Manifest to include Group Claims in Auth Token
1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:
2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.
3. Locate the "groupMembershipClaims" setting. Set its value to either "SecurityGroup" or "All". To help you decide which:
"SecurityGroup" groups claim will contain the identifiers of all security groups of which the user is a member.
"All" groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member
Now your application will include group claims in your manifest and you can use this fact in your code.
References:
https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/

QUESTION 132
You develop and deploy an Azure App Service web app to a production environment. You enable the Always On setting and the Application Insights site extensions.
You deploy a code update and receive multiple failed requests and exceptions in the web app.
You need to validate the performance and failure counts of the web app in near real time.
Which Application Insights tool should you use?

A. Snapshot Debugger
B. Profiler
C. Smart Detection
D. Live Metrics Stream
E. Application Map

Answer: D
Explanation:

Live Metrics Stream
Deploying the latest build can be an anxious experience. If there are any problems, you want to know about them right away, so that you can back out if necessary. Live Metrics Stream gives you key metrics with a latency of about one second.
With Live Metrics Stream, you can:
* Validate a fix while it's released, by watching performance and failure counts.
* Etc.

Incorrect:
* Profiler
Azure Application Insights Profiler provides performance traces for applications running in production in Azure. Profiler:
Captures the data automatically at scale without negatively affecting your users.
Helps you identify the "hot" code path spending the most time handling a particular web request.
* Snapshot debugger
When an exception occurs, you can automatically collect a debug snapshot from your live web application. The snapshot shows the state of source code and variables at the moment the exception was thrown. The Snapshot Debugger in Azure Application Insights monitors exception telemetry from your web app. It collects snapshots on your top-throwing exceptions so that you have the information you need to diagnose issues in production.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/live-stream

QUESTION 133
Drag and Drop Question
You are preparing to deploy an Azure virtual machine (VM)-based application. The VMs that run the application have the following requirements:
- When a VM is provisioned the firewall must be automatically configured before it can access Azure resources
- Supporting services must be installed by using an Azure PowerShell script that is stored in Azure Storage
You need to ensure that the requirements are met.
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 bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

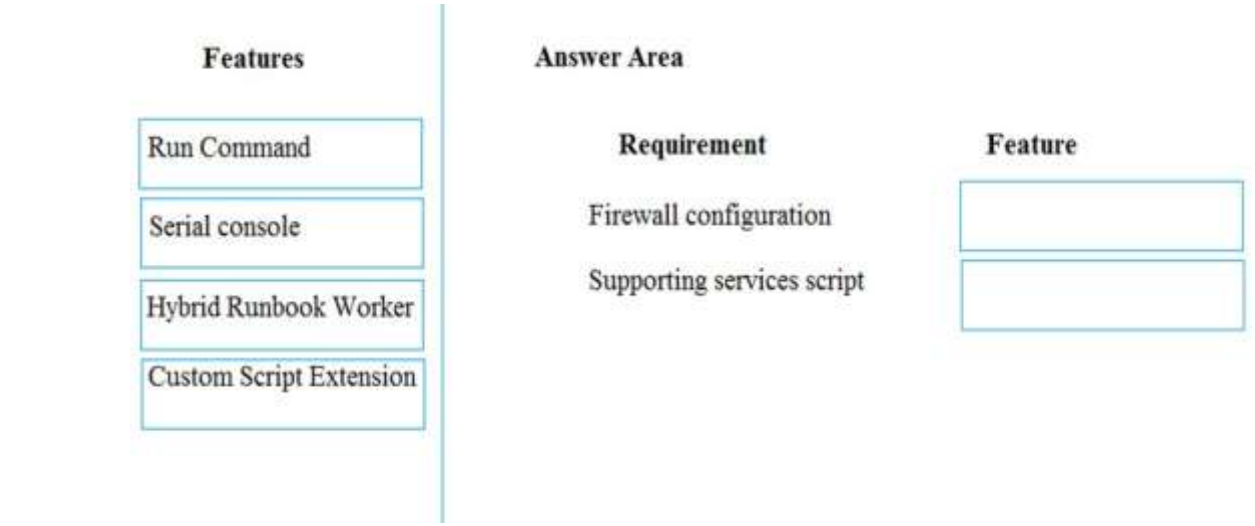

Answer:

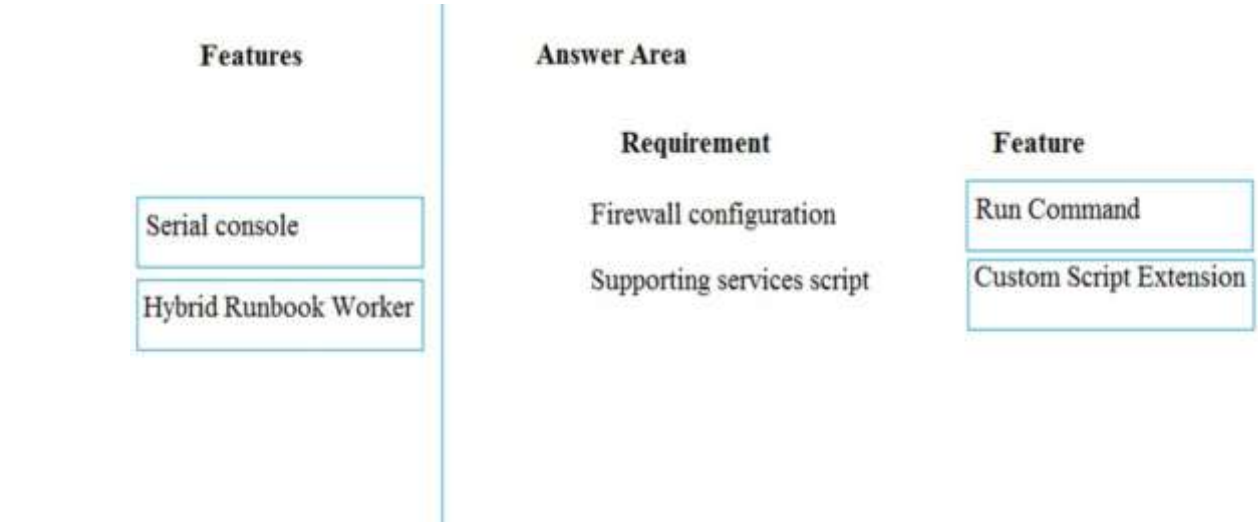

Explanation:
Box 1: Run Command
This capability is useful in all scenarios where you want to run a script within a VM. It's one of the only ways to troubleshoot and remediate a VM that doesn't have the RDP or SSH port open, because of improper network or administrative user configuration.
Box 2: Customer Script Extension
The Custom Script Extension downloads and executes scripts on Azure virtual machines. This extension is useful for post deployment configuration, software installation, or any other configuration or management tasks. Scripts can be downloaded from Azure storage or GitHub, or provided to the Azure portal at extension run time. The Custom Script Extension integrates with Azure Resource Manager templates, and can be run using the Azure CLI, PowerShell, Azure portal, or the Azure Virtual Machine REST API.
Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/custom-script-windows
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-scripts-in-vm

QUESTION 134
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

Answer: D

QUESTION 135
You are developing a user portal for a company.
You need to create a report for the portal that lists information about employees who are subject matter experts for a specific topic. You must ensure that administrators have full control and consent over the data.
Which technology should you use?

A. Microsoft Graph connectors
B. Microosft graph API
C. Microsoft Graph data connect

Answer: C
Explanation:
Data Connect grants a more granular control and consent model: you can manage data, see who is accessing it, and request specific properties of an entity. This enhances the Microsoft Graph model, which grants or denies applications access to entire entities.
Microsoft Graph Data Connect augments Microsoft Graph's transactional model with an intelligent way to access rich data at scale. The data covers how workers communicate, collaborate, and manage their time across all the applications and services in Microsoft 365.
Reference:
https://docs.microsoft.com/en-us/graph/data-connect-concept-overview

QUESTION 136
Drag and Drop Question
You are preparing to deploy an application to an Azure Kubernetes Service (AKS) cluster. The application must only be available from within the VNet that includes the cluster.
You need to deploy the application.
How should you complete the deployment YAML? To answer, drag the appropriate YAML segments to the correct locations. Each YAML 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.

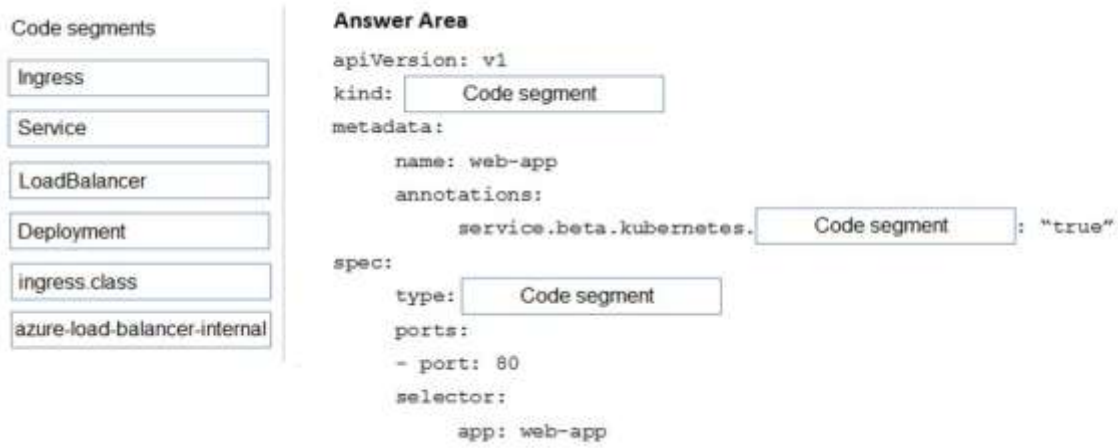

Answer:

Explanation:
To create an internal load balancer, create a service manifest named internal-lb.yaml with the service type LoadBalancer and the azure-load-balancer-internal annotation as shown in the following example:
YAML:
apiVersion: v1
kind: Service
metadata:
name: internal-app
annotations:
service.beta.kubernetes.io/azure-load-balancer-internal: "true" spec:
type: LoadBalancer
ports:
port: 80
selector:
app: internal-app
References:
https://docs.microsoft.com/en-us/azure/aks/internal-lb

QUESTION 137
Hotspot Question
A company is developing a Node.js web app. The web app code is hosted in a GitHub repository located at https://github.com/TailSpinToys/webapp.
The web app must be reviewed before it is moved to production. You must deploy the initial code release to a deployment slot named review.
You need to create the web app and deploy the code.
How should you complete the commands? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Box 1: New-AzResourceGroup

The New-AzResourceGroup cmdlet creates an Azure resource group.

Box 2: New-AzAppServicePlan

The New-AzAppServicePlan cmdlet creates an Azure App Service plan in a given location

Box 3: New-AzWebApp

The New-AzWebApp cmdlet creates an Azure Web App in a given a resource group

Box 4: New-AzWebAppSlot

The New-AzWebAppSlot cmdlet creates an Azure Web App slot.

Reference:

https://docs.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroup?view=azps-2.3.2
https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azappserviceplan?view=azps-2.3.2
https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebapp?view=azps-2.3.2
https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebappslot?view=azps-2.3.2

QUESTION 138
Hotspot Question
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_blob ( [          ] = [       ] )
```

delete_conteiner	False
delete_snapshots	Include
snapshot_blob	Only
snapshots_present	

Answer:

Answer Area

```
delete_blob ( [          ] = [       ] )
```

delete_conteiner	False
delete_snapshots	Include
snapshot_blob	**Only**
snapshots_present	

Explanation:
Box 1: delete_snapshots
# Delete only the snapshot (blob itself is retained)
blob_client.delete_blob(delete_snapshots="only")
Box 2: only
Reference:
https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob_samples_common.py

QUESTION 139
Your company is designing an application named App1 that will use data from Azure SQL Database. App1 will be accessed over the internet by many users. You need to recommend a solution for improving the performance ofApp1.
What should you include in the recommendation?

A. Azure HPC cache
B. ExpressRoute
C. a CON profile
D. Azure Cache for Redis

Answer: D

QUESTION 140
You are designing a multi-tiered application that will be hosted on Azure virtual machines. The virtual machines will run Windows Server. Front-end servers will be accessible from the Internet over port 443. The other servers will NOT be

directly accessible over the internet.
You need to recommend a solution to manage the virtual machines that meets the following requirement:
- Allows the virtual machine to be administered by using Remote Desktop.
- Minimizes the exposure of the virtual machines on the Internet
Which Azure service should you recommend?

A. Azure Bastion
B. Service Endpoint
C. Azure Private Link
D. Azure Front Door

Answer: C

QUESTION 141
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

QUESTION 142
A team has created an Index in the Azure Search service. You have to upload data into the Index. You propose the following steps to carry out from your .Net program
- Create a SearchServiceClient object to connect to the search index.
- Create a DataContainer that contains the documents which must be added.
- Create a DataSource instance and set its Container property to the DataContainer.
- Set the DataSource property of the SearchServiceClient
Does the list of steps fulfil the requirement?

A. Yes
B. No

Answer: B
Explanation:
https://docs.microsoft.com/en-us/azure/search/search-what-is-azure-search#how-to-use-azure-search

QUESTION 143
A team has created an Index in the Azure Search service. You have to upload data into the Index. You propose the following steps to carry out from your .Net program
- Create a SearchIndexClient object to connect to the search index
- Create an IndexBatch that contains the documents which must be added.
- Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.
Does the list of steps fulfil the requirement?

A. Yes
B. No

Answer: A
Explanation:

https://docs.microsoft.com/en-us/azure/search/search-import-data-dotnet

In order to push documents into your index using the .NET SDK, you will need to:

1. Create a `SearchIndexClient` object to connect to your search index.
2. Create an `IndexBatch` containing the documents to be added, modified, or deleted.
3. Call the `Documents.Index` method of your `SearchIndexClient` to send the `IndexBatch` to your search index.

QUESTION 144
A team has created an Index in the Azure Search service. You have to upload data into the Index. You propose the following steps to carry out from your .Net program
- Create a SearchIndexClient object to connect to the search index.
- Create a DataContainer that contains the documents which must be added.
- Create a DataSource instance and set its Container property to the DataContainer
- Call the Documents.Search method of the SearchIndexClient and pass the DataSource.
Does the list of steps fulfil the requirement?

A. Yes
B. No

Answer: B
Explanation:

In order to push documents into your index using the .NET SDK, you will need to:

1. Create a `SearchIndexClient` object to connect to your search index.
2. Create an `IndexBatch` containing the documents to be added, modified, or deleted.
3. Call the `Documents.Index` method of your `SearchIndexClient` to send the `IndexBatch` to your search index.

https://docs.microsoft.com/en-us/azure/search/search-import-data-dotnet

QUESTION 145
A team is developing container-based applications that need to be deployed to a Kubernetes cluster in Azure. You have to create the cluster and ensure the services are running as desired. Which of the following commands would you execute? Choose 4 answers from the options given below

A. az aks create
B. az group create
C. kubectl apply
D. az appservice plan create
E. az aks get-credentials

Answer: ABCE
Explanation:
https://docs.microsoft.com/en-us/azure/aks/kubernetes-walkthrough

QUESTION 146
A development team is developing an application. The application will be storing its data in Azure Table storage. Below are the fields that are going to be stored in the table
- Region
- Email address
- Phone number
The following snippet of code needs to be completed that would be used to insert a batch of records.

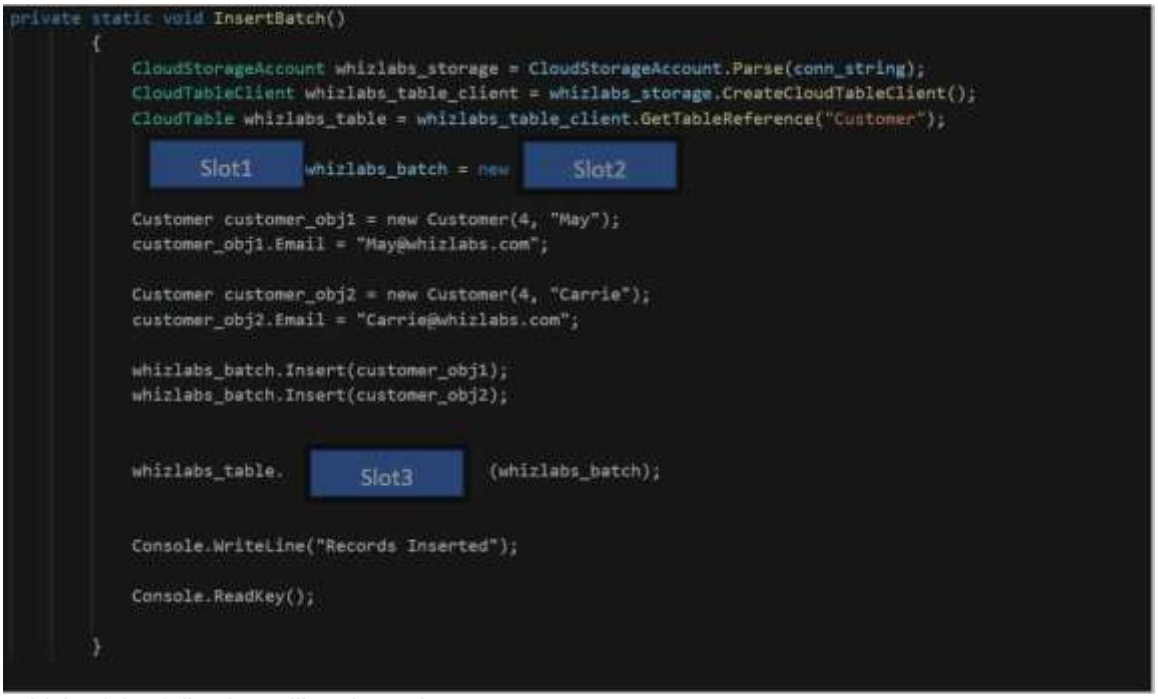

Which of the following will go into Slot1?

A. TableOperation
B. TableBatchOperation
C. TableEntity
D. TableQuery

Answer: B
Explanation:
https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet

QUESTION 147
A development team is developing an application. The application will be storing its data in Azure Table storage. Below are the fields that are going to be stored in the table
- Region
- Email address
- Phone number
The following snippet of code needs to be completed that would be used to insert a batch of records.

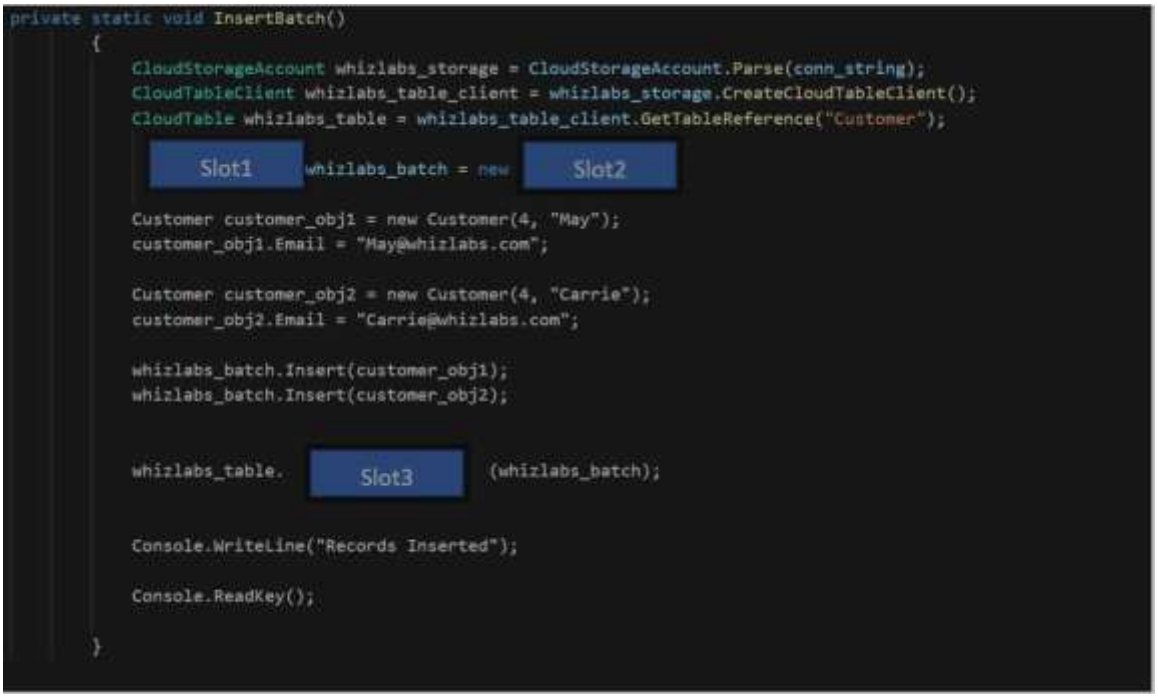

Which of the following will go into Slot2?

A. TableOperation
B. TableBatchOperation
C. TableEntity
D. TableQuery

Answer: B
Explanation:
https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet

QUESTION 148
A development team is developing an application. The application will be storing its data in Azure Table storage. Below are the fields that are going to be stored in the table
- Region
- Email address
- Phone number
The following snippet of code needs to be completed that would be used to insert a batch of records.

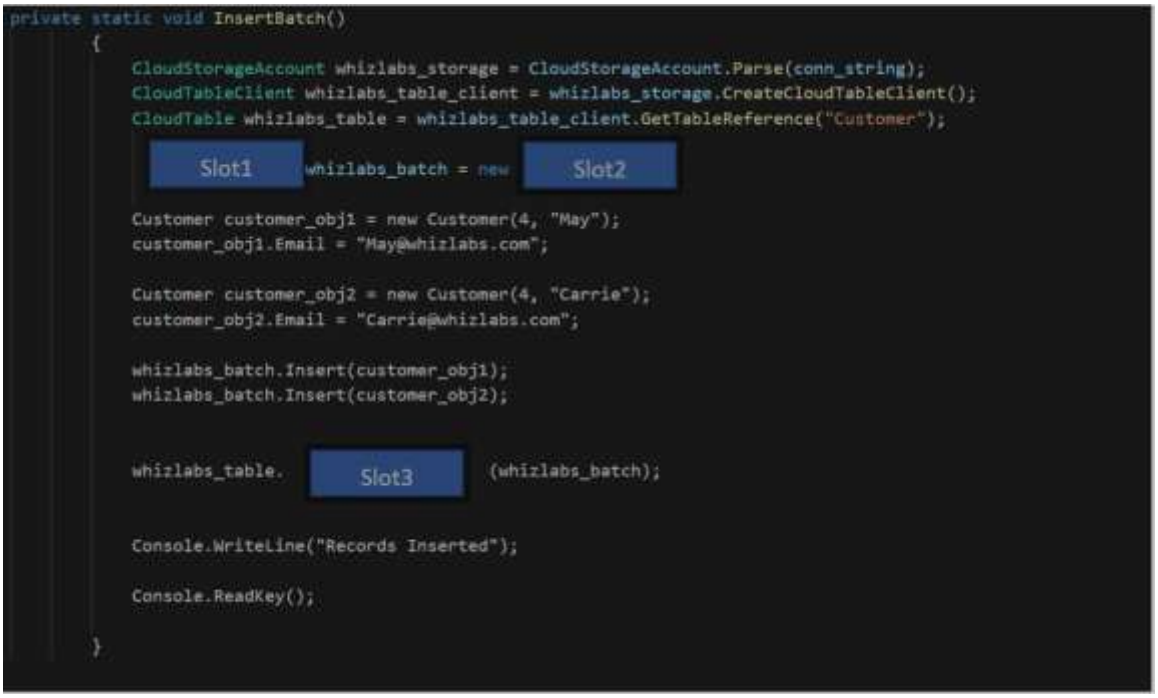

Which of the following will go into Slot3?

A. ExecuteBatch
B. Execute
C. Insert
D. InsertOrMerge

Answer: A
Explanation:
https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet

QUESTION 149
A development team is developing an application. The application will be working with customer data.
The application will also be making use of Azure Redis Cache.
You need to invalidate the cache when the customer data is changed.
You have to complete the below code to comply with the requirement

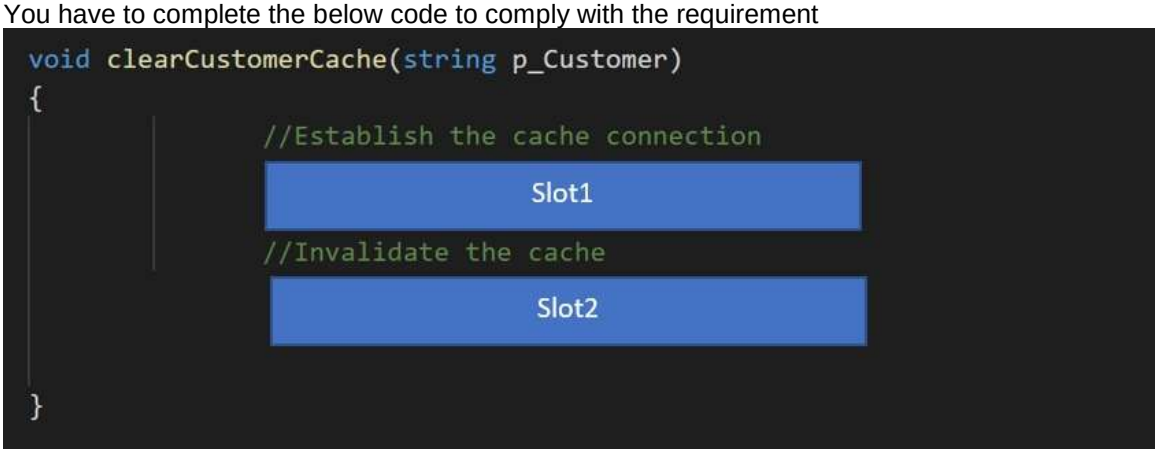

Which of the following will go into Slot1?

A. IDatabase cache=Connection.GetDatabase();
B. IDatabase cache=Connection.GetCache();
C. ICache cache=Connection.GetDatabase();

D. ICache cache=Connection.GetCache();

Answer: A
Explanation:
The right way is to use the IDatabase interface. Also you need to use the GetDatabase() method. This is also mentioned in the Microsoft documentation.

```
static void Main(string[] args)
{
    // Connection refers to a property that returns a ConnectionMultiplexer
    // as shown in the previous example.
    IDatabase cache = lazyConnection.Value.GetDatabase();
```

Since this is clearly given in the Microsoft documentation, all other options are incorrect.
https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-dotnet-how-to-use-azure-redis-cache

QUESTION 150
A development team is developing an application. The application will be working with customer data. The application will also be making use of Azure Redis Cache. You need to invalidate the cache when the customer data is changed. You have to complete the below code to comply with the requirement

```
void clearCustomerCache(string p_Customer)
{
        //Establish the cache connection
        Slot1
        //Invalidate the cache
        Slot2
}
```

Which of the following will go into Slot2?

A. cache.KeyDelete(p_Customer);
B. cache.ValueDelete(p_Customer);
C. cache.StringGet(p_Customer);
D. cache.StringSet(p_Customer);

Answer: A
Explanation:
Since you have to invalidate the cache, you have to delete the Key itself
Option B is incorrect since you need to work with keys and not the values
Option C is incorrect this is used to get the string value
Option D is incorrect this is used to set the string value
https://docs.microsoft.com/en-us/azure/azure-cache-for-redis/cache-dotnet-how-to-use-azure-redis-cache

QUESTION 151
A development team is developing an application that works with Azure Table storage.

Column	
FirstName	Partition Key
LastName	RowKey
Email	Property

Below are some of the rows in the table

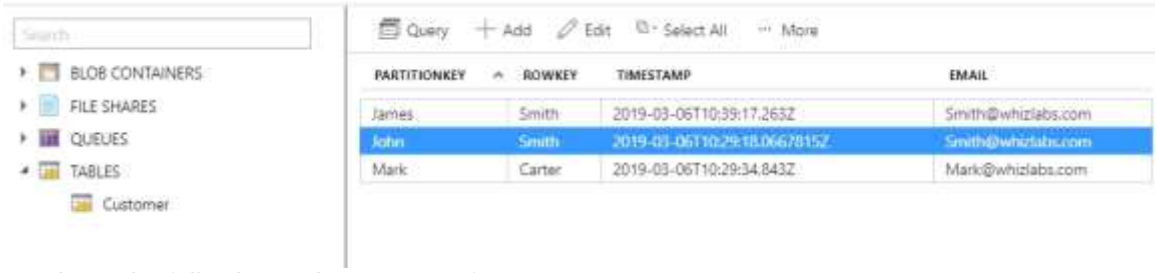

You have the following code statement from a C# program

```
TableQuery<CustomerEntity> rangeQuery = new TableQuery<CustomerEntity>().Where(
    TableQuery.CombineFilters(
        TableQuery.GenerateFilterCondition("PartitionKey", QueryComparisons.Equal, "James"),
        TableOperators.And,
        TableQuery.GenerateFilterCondition("RowKey", QueryComparisons.Equal, "Smith@whizlabs.com")));
```

Would this return all the entities where the RowKey is Smith@whizlabs.com?

A. Yes
B. No

Answer: B
Explanation:
https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet

QUESTION 152
A development team is developing an application that works with Azure Table storage.

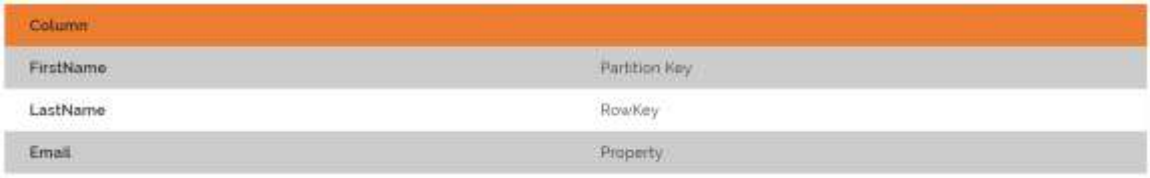

Below are some of the rows in the table

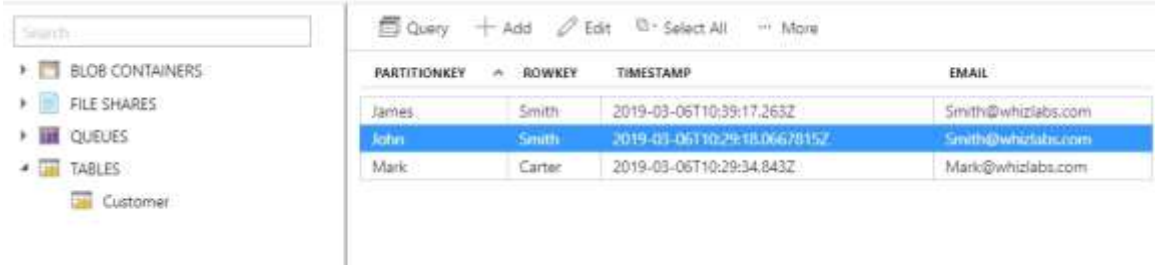

Is the below.Net Code query

```
TableQuery<CustomerEntity> rangeQuery = new TableQuery<CustomerEntity>().Where(
    TableQuery.CombineFilters(
        TableQuery.GenerateFilterCondition("PartitionKey", QueryComparisons.Equal, "James"),
        TableOperators.And,
        TableQuery.GenerateFilterCondition("RowKey", QueryComparisons.Equal, "Smith")));
```

Same as executing the below REST API call along with a valid Shared Access Signature
https://whizlabsstore.table.core.windows.net/Customer(PartitionKey='James',RowKey='Smith')

A. Yes
B. No

Answer: A
Explanation:

https://docs.microsoft.com/en-us/rest/api/storageservices/querying-tables-and-entities

QUESTION 153
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: Enable Application Request Routing (ARR)
Does the solution meet the goal?

A. Yes
B. No

Answer: B
Explanation:
Instead deploy and configure Azure Cache for Redis. Update the web applications.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching#managing-concurrency-in-a-cache