Number: AI-102
Passing Score: 800
Time Limit: 120
File Version: 13.0

Exam Code: AI-102
Exam Name: Designing and Implementing a Microsoft Azure AI Solution

01 - Implement Computer Vision Solutions

QUESTION 1
DRAG DROP
You are developing a webpage that will use the Video Indexer service to display videos of internal company meetings.
You embed the Player widget and the Cognitive Insights widget into the page.
You need to configure the widgets to meet the following requirements:
Ensure that users can search for keywords.
Display the names and faces of people in the video.
Show captions in the video in English (United States).
How should you complete the URL for each widget? To answer, drag the appropriate values to the correct targets. Each value 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:

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/azure-video-analyzer/video-analyzer-for-media-docs/video-indexer-embed-widgets

QUESTION 2
DRAG DROP

You train a Custom Vision model to identify a company's products by using the Retail domain.
You plan to deploy the model as part of an app for Android phones.
You need to prepare the model for deployment.
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:

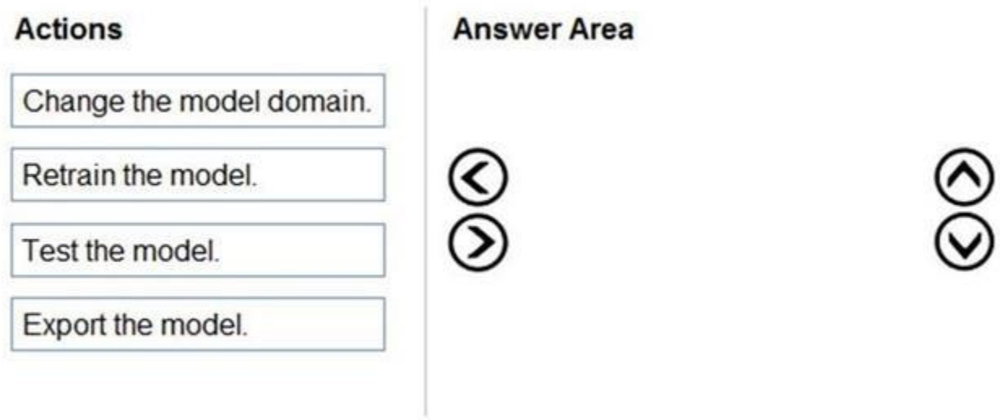

Correct Answer:

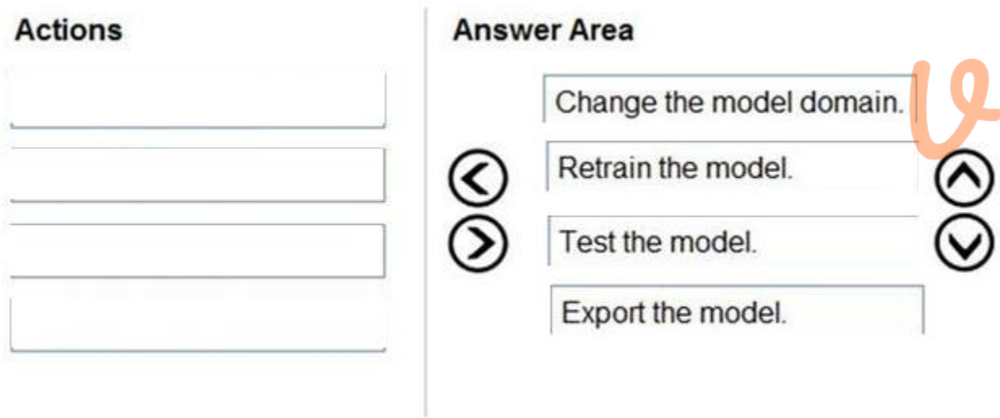

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/export-your-model

QUESTION 3
HOTSPOT
You develop an application that uses the Face API.
You need to add multiple images to a person group.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
Parallel.For(0, PersonCount, async i =>
{
    Guid personId = persons[i].PersonId;
    string personImageDir = $"/path/to/person/{i}/images";
    foreach (string imagePath in Directory.GetFiles(personImageDir, "*.jpg"))
    {
        using (                    t = File.OpenRead(imagePath))
        {
          await faceClient.PersonGroupPerson.
(personGroupId, personId, t);
        }
});
```

File
Stream
Uri
Url

AddFaceFromStreamAsync
AddFaceFromUrlAsync
CreateAsync
GetAsync

Answer Area:

Answer Area

```
Parallel.For(0, PersonCount, async i =>
{
    Guid personId = persons[i].PersonId;
    string personImageDir = $"/path/to/person/{i}/images";
    foreach (string imagePath in Directory.GetFiles(personImageDir, "*.jpg"))
    {
        using ( [ ▼ ] t = File.OpenRead(imagePath))
                  File
                  Stream
                  Uri
                  Url

        {
          await faceClient.PersonGroupPerson. [ ▼ ]
                                                AddFaceFromStreamAsync
                                                AddFaceFromUrlAsync
                                                CreateAsync
                                                GetAsync
(personGroupId, personId, t);
        }
});
```

Section:
Explanation:
Box 1: Stream
The File.OpenRead(String) method opens an existing file for reading.
Example: Open the stream and read it back.
using (FileStream fs = File.OpenRead(path))
Box 2: CreateAsync
Create the persons for the PersonGroup. Persons are created concurrently.
Example:
await faceClient.PersonGroupPerson.CreateAsync(personGroupId, personName);
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-add-faces

QUESTION 4
HOTSPOT
You are developing an application that will use the Computer Vision client library. The application has the following code.

```
public async TaskAnalyzeImage(ComputerVisionClient client, string localImage)
{
	List<VisualFeatureTypes> features = new List<VisualFeatureTypes>()
		{
			VisualFeatureTypes.Description,
			VisualFeatureTypes.Tags,
		};
	using (Stream imageStream = File.OpenRead(localImage))
	{
		try
		{
		  ImageAnalysis results = await client.AnalyzeImageInStreamAsync(imageStream, features);

		  foreach (var caption in results.Description.Captions)
		  {
			Console.WriteLine($"{caption.Text} with confidence {caption.Confidence}");
		  }

		  foreach (var tag in results.Tags)
		  {
			Console.WriteLine($"{tag.Name} {tag.Confidence}");
		  }
		}
		catch (Exception ex)
		{
			Console.WriteLine(ex.Message);
		}
	}
}
```

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
The code will perform face recognition.	○	○
The code will list tags and their associated confidence.	○	○
The code will read a file from the local file system.	○	○

Answer Area:

Answer Area

Statements	Yes	No
The code will perform face recognition.	○	○
The code will list tags and their associated confidence.	○	○
The code will read a file from the local file system.	○	○

Section:
Explanation:
Box 1: No
Box 2: Yes
The ComputerVision.analyzeImageInStreamAsync operation extracts a rich set of visual features based on the image content.
Box 3: No
Images will be read from a stream.
Reference:
https://docs.microsoft.com/en-us/java/api/com.microsoft.azure.cognitiveservices.vision.computervision.computervision.analyzeimageinstreamasync

QUESTION 5
HOTSPOT
You are developing an application to recognize employees' faces by using the Face Recognition API. Images of the faces will be accessible from a URI endpoint.
The application has the following code.

```
static async void AddFace(string subscription_key, string personGroupId, string personId, string imageURI)
{
  var client = new HttpClient();
  client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", subscription_key);
  var endpointURI = $"https://westus.api.cognitive.microsoft.com/face/v1.0/persongroups/{personGroupId}/persons/{personId}/persistedFaces";
  HttpResponseMessage response;
  var body = "{ \"url\": \"" + imageURI + "\"}";
  var content = new StringContent(body, Encoding.UTF8, "application/json");
  var response = await client.PutAsync(endpointURI, content);
}
```

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:

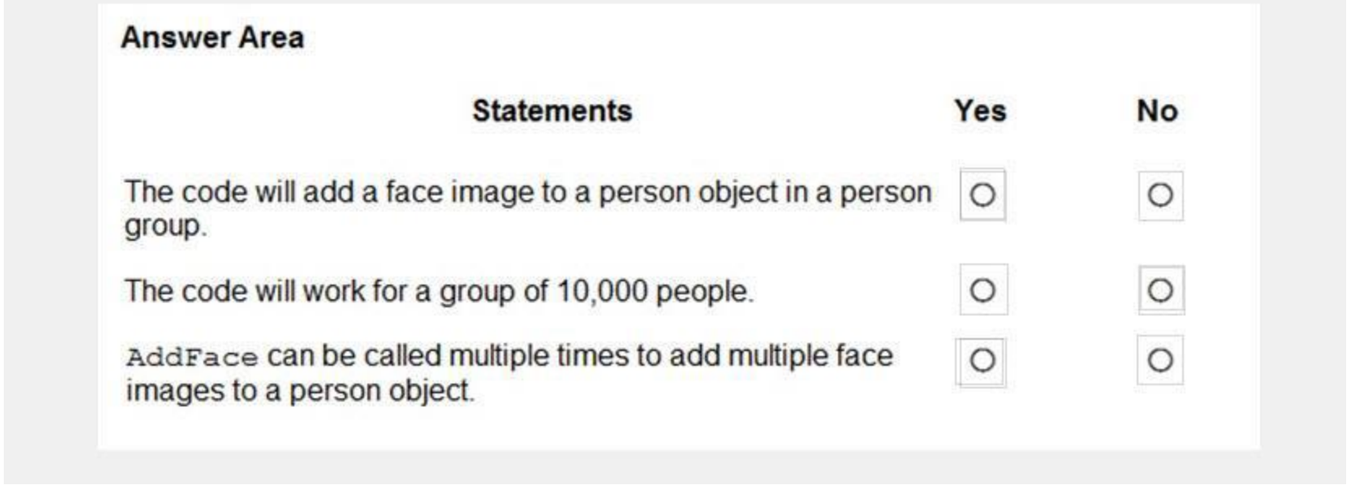

Answer Area:

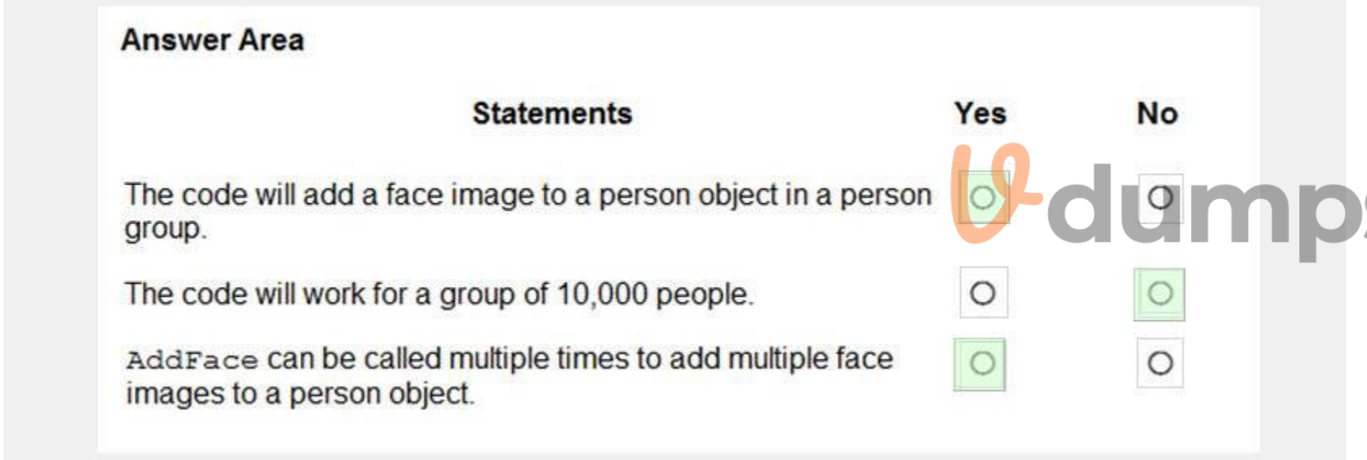

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/use-persondirectory

QUESTION 6
DRAG DROP
You have a Custom Vision resource named acvdev in a development environment.
You have a Custom Vision resource named acvprod in a production environment.
In acvdev, you build an object detection model named obj1 in a project named proj1.
You need to move obj1 to acvprod.
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:

Actions

- Use the `ExportProject` endpoint on acvdev.
- Use the `GetProjects` endpoint on acvdev.
- Use the `ImportProject` endpoint on acvprod.
- Use the `ExportIteration` endpoint on acvdev.
- Use the `GetIterations` endpoint on acvdev.
- Use the `UpdateProject` endpoint on acvprod.

Answer Area

Correct Answer:

Actions

- Use the `ExportIteration` endpoint on acvdev.
- Use the `GetIterations` endpoint on acvdev.
- Use the `UpdateProject` endpoint on acvprod.

Answer Area

- Use the `GetProjects` endpoint on acvdev.
- Use the `ExportProject` endpoint on acvdev.
- Use the `ImportProject` endpoint on acvprod.

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/copy-move-projects

QUESTION 7
DRAG DROP
You are developing an application that will recognize faults in components produced on a factory production line. The components are specific to your business.
You need to use the Custom Vision API to help detect common faults.
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:

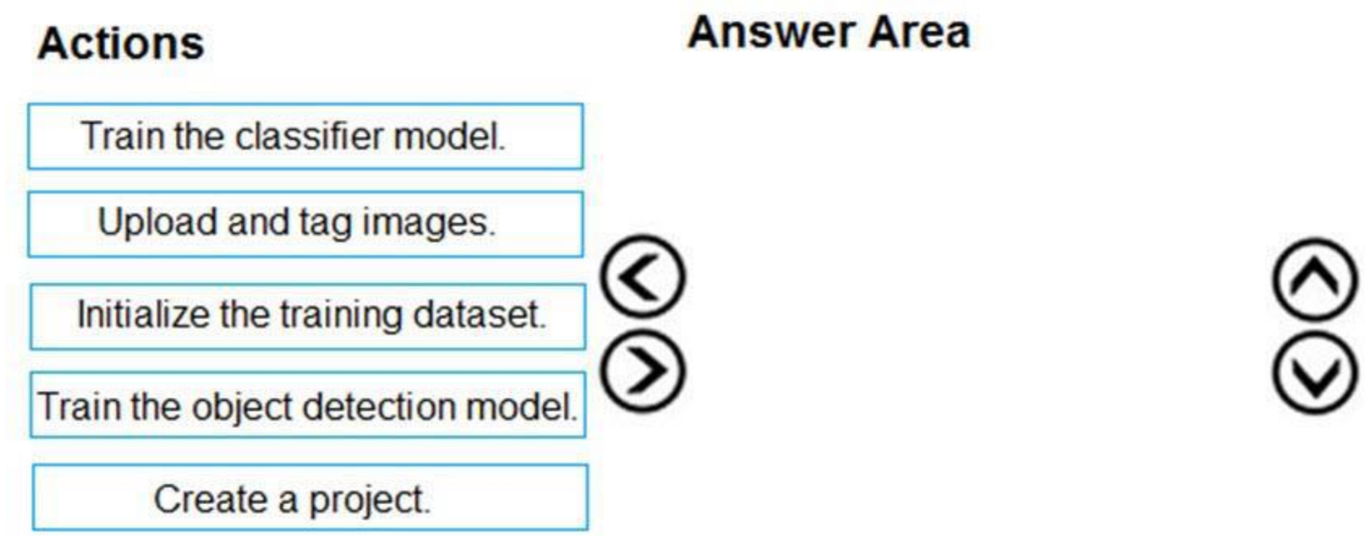

Correct Answer:

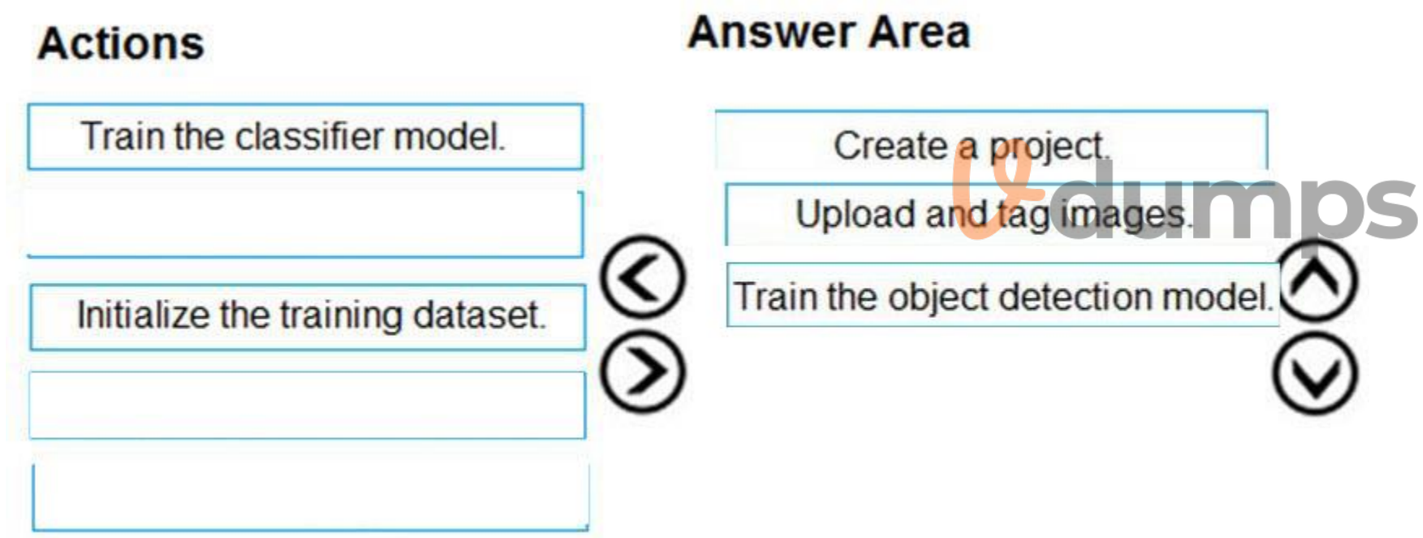

Section:
Explanation:
Step 1: Create a project
Create a new project.
Step 2: Upload and tag the images
Choose training images. Then upload and tag the images.
Step 3: Train the classifier model.
Train the classifier
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/getting-started-build-a-classifier

QUESTION 8
You have a Video Indexer service that is used to provide a search interface over company videos on your company's website. You need to be able to search for videos based on who is present in the video.

What should you do?

A. Create a person model and associate the model to the videos.
B. Create person objects and provide face images for each object
C. Invite the entire staff of the company to Video Indexer.
D. Edit the faces in the videos.
E. Upload names to a language model.

Correct Answer: A
Section:
Explanation:
Video Indexer supports multiple Person models per account. Once a model is created, you can use it by providing the model ID of a specific Person model when uploading/indexing or reindexing a video. Training a newfacefor a video updates the specific custom model that the video was associated with.
Note: Video Indexer supports face detection and celebrity recognition for video content. The celebrity recognition feature covers about one million faces based on commonly requested data source such as IMDB, Wikipedia, and top LinkedIn influencers. Faces that aren't recognized by the celebrity recognition feature are detected but left unnamed. Once you label a face with a name, the face and name get added to your account's Person model. Video Indexer will then recognize this face in your future videos and past videos.
Reference:
https://docs. mi crosoft. com/en-us/azu re/med ia -servi ces/vi deo-i ndexer/customize-pers on-mo del-with-api

QUESTION 9
You use the Custom Vision service to build a classifier.
After training is complete, you need to evaluate the classifier.
Which two metrics are available for review? Each correct answer presents a complete solution. (Choose two.) NOTE: Each correct selection is worth one point.

A. recall
B. F-score
C. weighted accuracy
D. precision
E. area under the curve (AUC)

Correct Answer: A, D
Section:
Explanation:
Custom Vision provides three metrics regarding the performance of your model: precision, recall, and AP.
Reference:
https://www.tallan.com/blog/2020/05/19/azure-custom-vision/

QUESTION 10
You are developing a method that uses the Computer Vision client library. The method will perform optical character recognition (OCR) in images. The method has the following code.

```
public static async Task ReadFileUrl(ComputerVisionClient client, string urlFile)
{
        const int numberOfCharsInOperationId = 36;

        var txtHeaders = await client.ReadAsync(urlFile, language: "en");

        string opLocation = textHeaders.OperationLocation;
        string operationId = opLocation.Substring(opLocation.Length -
numberOfCharsInOperationId);

        ReadOperationResult results;

        results = await client.GetReadResultAsync(Guid.Parse(operationId));

        var textUrlFileResults = results.AnalyzeResult.ReadResults;
        foreach (ReadResult page in textUrlFileResults)
        {
                foreach (Line line in page.Lines)
                {
                        Console.WriteLine(line.Text);
                }
        }
}
```

During testing, you discover that the call to the GetReadResultAsync method occurs before the read operation is complete.
You need to prevent the GetReadResultAsync method from proceeding until the read operation is complete.
Which two actions should you perform? Each correct answer presents part of the solution. (Choose two.)
NOTE: Each correct selection is worth one point.

A. Remove the Guid.Parse (operationid) parameter.
B. Add code to verify the results.Status value.
C. Add code to verify the status of the txtHeaders. status value.
D. Wrap the call to GetReadResultAsync within a loop that contains a delay.

Correct Answer: B, D
Section:
Explanation:
Example code :
do
{
results = await client.GetReadResultAsync(Guid.Parse(operationId));
}
while ((results.Status == OperationStatusCodes.Running ||
results.Status == OperationStatusCodes.NotStarted));
Reference:
https://github.com/Azure-Samples/cognitive-services-quickstart-code/blob/master/dotnet/ComputerVision/ComputerVisionQuickstart.cs

QUESTION 11
DRAG DROP
You are developing a call to the Face API. The call must find similar faces from an existing list named employeefaces. The employeefaces list contains 60,000 images.
How should you complete the body of the HTTP request? To answer, drag the appropriate values to the correct targets. Each value 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:

Values

"faceListId"

"LargeFaceListId"

"matchFace"

"matchPerson"

Answer Area

```
{
"faceId": "18c51a87-3a69-47a8-aedc-a54745f708a1",
                    : "employeefaces",
"maxNumOfCandidatesReturned": 1,
"mode":
}
```

Correct Answer:

Values

"faceListId"

"matchPerson"

Answer Area

```
{
"faceId": "18c51a87-3a69-47a8-aedc-a54745f708a1",
"LargeFaceListId": "employeefaces",
"maxNumOfCandidatesReturned": 1,
"mode": "matchFace"
}
```

Section:
Explanation:
Box 1: LargeFaceListID
LargeFaceList: Add a face to a specified large face list, up to 1,000,000 faces.
Note: Given query face's faceId, to search the similar-looking faces from a faceId array, a face list or a large face list. A "faceListId" is created by FaceList - Create containing persistedFaceIds that will not expire. And a "largeFaceListId" is created by LargeFaceList - Create containing persistedFaceIds that will also not expire.
Incorrect Answers:
Not "faceListId": Add a face to a specified face list, up to 1,000 faces.
Box 2: matchFace
Find similar has two working modes, "matchPerson" and "matchFace". "matchPerson" is the default mode that it tries to find faces of the same person as possible by using internal same-person thresholds. It is useful to find a known person's other photos. Note that an empty list will be returned if no faces pass the internal thresholds. "matchFace" mode ignores same-person thresholds and returns ranked similar faces anyway, even the similarity is low. It can be used in the cases like searching celebrity-looking faces.
Reference:
https://docs.microsoft.com/en-us/rest/api/faceapi/face/findsimilar

QUESTION 12
DRAG DROP
You are developing a photo application that will find photos of a person based on a sample image by using the Face API.
You need to create a POST request to find the photos.
How should you complete the request? To answer, drag the appropriate values to the correct targets. Each value 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:

Values

- detect
- findsimilars
- group
- identify
- matchFace
- matchPerson
- verify

Answer Area

```
POST {Endpoint}/face/v1.0/ [          ]
Request Body

{
  "faceId": "c5c24a82-6845-4031-9d5d-978df9175426",
  "largeFaceListId": "sample_list",
  "maxNumOfCandidatesReturned": 10,
  "mode": "[          ]"
}
```

Correct Answer:

Values

- (blank)
- findsimilars
- group
- identify
- matchFace
- (blank)
- verify

Answer Area

```
POST {Endpoint}/face/v1.0/ [detect]
Request Body

{
  "faceId": "c5c24a82-6845-4031-9d5d-978df9175426",
  "largeFaceListId": "sample_list",
  "maxNumOfCandidatesReturned": 10,
  "mode": "[matchPerson]"
}
```

Section:
Explanation:

Box 1: detect
Face - Detect With Url: Detect human faces in an image, return face rectangles, and optionally with faceIds, landmarks, and attributes.
POST {Endpoint}/face/v1.0/detect
Box 2: matchPerson
Find similar has two working modes, "matchPerson" and "matchFace". "matchPerson" is the default mode that it tries to find faces of the same person as possible by using internal same-person thresholds. It is useful to find a known person's other photos. Note that an empty list will be returned if no faces pass the internal thresholds. "matchFace" mode ignores same-person thresholds and returns ranked similar faces anyway, even the similarity is low. It can be used in the cases like searching celebrity-looking faces.
Reference:
https://docs.microsoft.com/en-us/rest/api/faceapi/face/detectwithurl
https://docs.microsoft.com/en-us/rest/api/faceapi/face/findsimilar

QUESTION 13
HOTSPOT
You develop a test method to verify the results retrieved from a call to the Computer Vision API. The call is used to analyze the existence of company logos in images. The call returns a collection of brands named brands.
You have the following code segment.

```
foreach (var brand in brands)
{
    if (brand.Confidence >= .75)
        Console.WriteLine($"Logo of {brand.Name} between {brand.Rectangle.X},
{brand.Rectangle.Y} and {brand.Rectangle.W}, {brand.Rectangle.H}");
}
```

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:

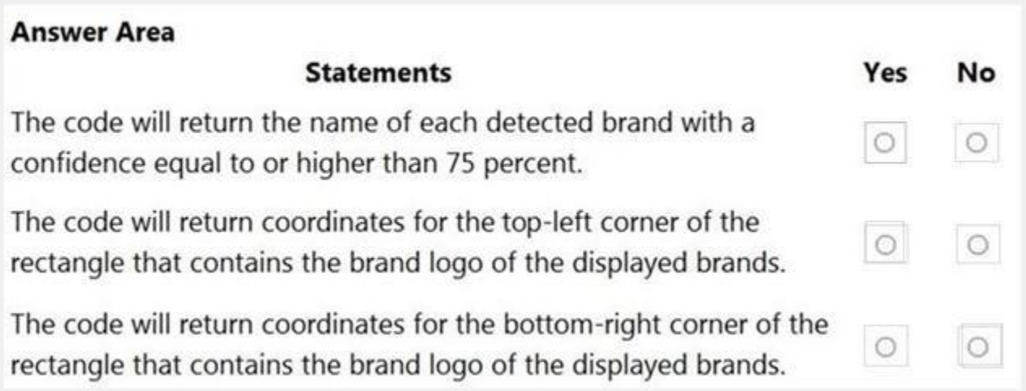

Answer Area:

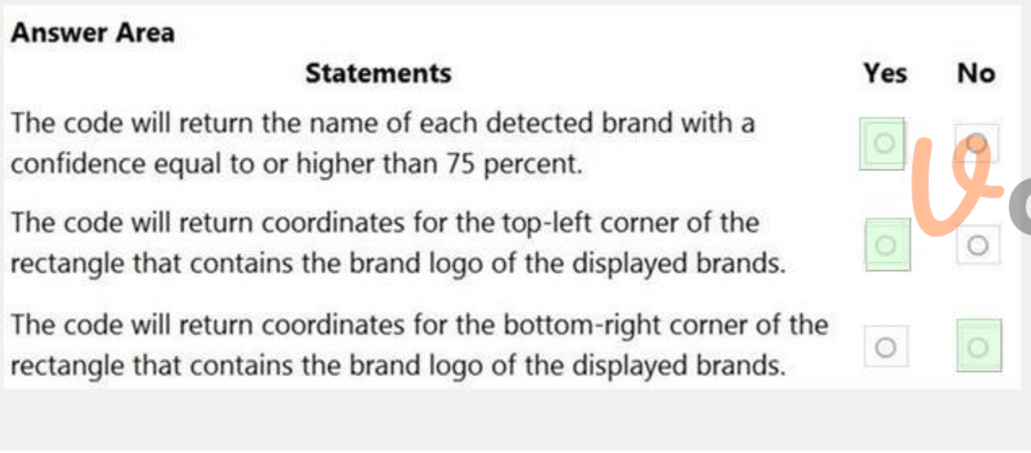

Section:
Explanation:
Box 1: Yes
Box 2: Yes
Coordinates of a rectangle in the API refer to the top left corner.
Box 3: No
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-brand-detection

QUESTION 14
HOTSPOT
You have a Computer Vision resource named contoso1 that is hosted in the West US Azure region.
You need to use contoso1 to make a different size of a product photo by using the smart cropping feature.
How should you complete the API URL? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

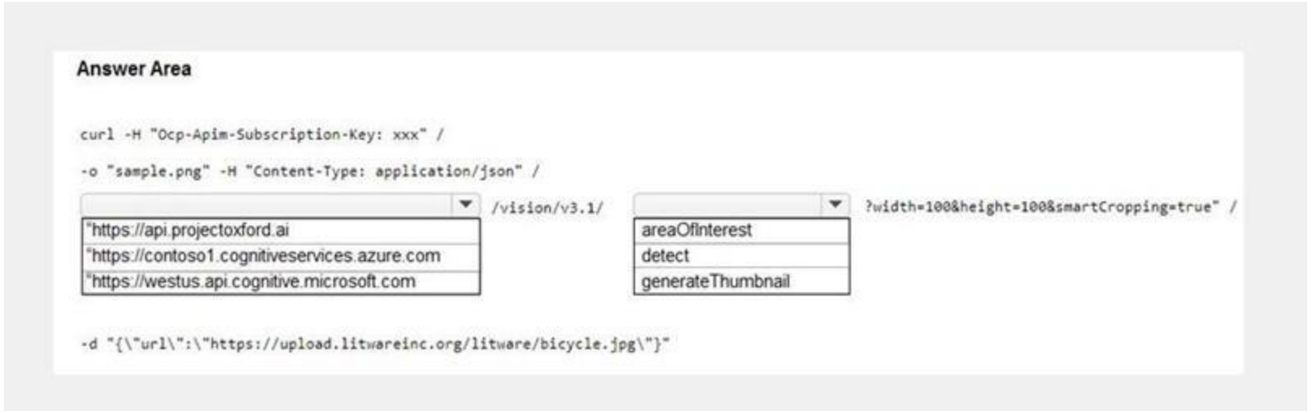

Answer Area:

Section:
Explanation:
Reference:
https://westus.dev.cognitive.microsoft.com/docs/services/computer-vision-v3-2/operations/56f91f2e778daf14a499f21b
https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-generating-thumbnails#examples

QUESTION 15
HOTSPOT
You are building a model that will be used in an iOS app.
You have images of cats and dogs. Each image contains either a cat or a dog.
You need to use the Custom Vision service to detect whether the images is of a cat or a dog.
How should you configure the project in the Custom Vision portal? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Project Types:
- Classification
- Object Detection

Classification Types:
- Multiclass (Single tag per image)
- Multilabel (Multiple tags per image)

Domains:
- Audit
- Food
- General
- General (compact)
- Landmarks
- Landmarks (compact)
- Retail
- Retail (compact)

Answer Area:

Answer Area

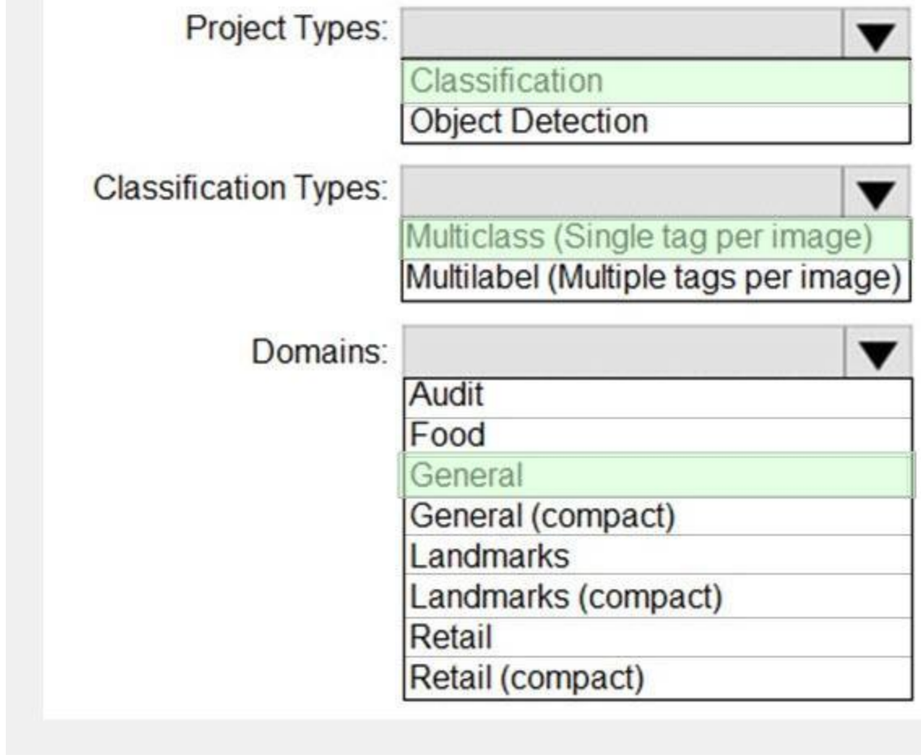

Section:
Explanation:
Box 1: Classification
Incorrect Answers:
An object detection project is for detecting which objects, if any, from a set of candidates are present in an image.
Box 2: Multiclass
A multiclass classification project is for classifying images into a set of tags, or target labels. An image can be assigned to one tag only.
Incorrect Answers:
A multilabel classification project is similar, but each image can have multiple tags assigned to it.
Box 3: General
General: Optimized for a broad range of image classification tasks. If none of the other specific domains are appropriate, or if you're unsure of which domain to choose, select one of the General domains.
Reference:
https://cran.r-project.org/web/packages/AzureVision/vignettes/customvision.html

02 - Implement Computer Vision Solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

A company named Wide World Importers is developing an e-commerce platform.

You are working with a solutions architect to design and implement the features of the e-commerce platform. The platform will use microservices and a serverless environment built on Azure.

Wide World Importers has a customer base that includes English, Spanish, and Portuguese speakers.

Existing Environment

Applications

Wide World Importers has an App Service plan that contains the web apps shown in the following table.

Name	Description
Product Management	An app used by employees to create and manage products. The app and the expected inputs from the employees are in English.
Inventory Tracking	An app used by employees to manage inventory when dispatching orders, receiving refunds, and receiving consignments from suppliers.

Azure Resources

You have the following resources:

▪ An Azure Active Directory (Azure AD) tenant

An Azure Active Directory (Azure AD) tenant

- The tenant supports internal authentication.
- All employees belong to a group named AllUsers.
- Senior managers belong to a group named LeadershipTeam.

An Azure Functions resource

- A function app posts to Azure Event Grid when stock levels of a product change between OK, Low Stock, and Out of Stock. The function app uses the Azure Cosmos DB change feed. An Azure Cosmos DB account
- The account uses the Core (SQL) API.
- The account stores data for the Product Management app and the Inventory Tracking app.

▪ An Azure Storage account

An Azure Storage account

- The account contains blob containers for assets related to products.
- The assets include images, videos, and PDFs.

An Azure Cognitive Services resource named wwics

A Video Indexer resource named wwivi

Requirements

Business Goals

Wide World Importers wants to leverage AI technologies to differentiate itself from its competitors.

Planned Changes

Wide World Importers plans to start the following projects:

A product creation project: Help employees create accessible and multilingual product entries, while expediting product entry creation. A smart e-commerce project: Implement an Azure Cognitive Search solution to display products for customers to browse. A shopping on-the-go project: Build a chatbot that can be integrated into smart speakers to support customers.

Business Requirements

Wide World Importers identifies the following business requirements for all the projects:

Provide a multilingual customer experience that supports English, Spanish, and Portuguese.

Whenever possible, scale based on transaction volumes to ensure consistent performance.

Minimize costs.

Governance and Security Requirements

Wide World Importers identifies the following governance and security requirements:

Data storage and processing must occur in datacenters located in the United States.

Azure Cognitive Services must be inaccessible directly from the internet.

Accessibility Requirements

Wide World Importers identifies the following accessibility requirements:

All images must have relevant alt text.

All videos must have transcripts that are associated to the video and included in product descriptions. Product descriptions, transcripts, and alt text must be available in English, Spanish, and Portuguese.
Product Creation Requirements
Wide World Importers identifies the following requirements for improving the Product Management app:
Minimize how long it takes for employees to create products and add assets.
Remove the need for manual translations.
Smart E-Commerce Requirements
Wide World Importers identifies the following requirements for the smart e-commerce project:
Ensure that the Cognitive Search solution meets a Service Level Agreement (SLA) of 99.9% availability for searches and index writes. Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products. Support autocompletion and autosuggestion based on all product name variants.
Store all raw insight data that was generated, so the data can be processed later.
Update the stock level field in the product index immediately upon changes.
Update the product index hourly.
Shopping On-the-Go Requirements
Wide World Importers identifies the following requirements for the shopping on-the-go chatbot:
Answer common questions.
Support interactions in English, Spanish, and Portuguese.
Replace an existing FAQ process so that all Q&A is managed from a central location.
Provide all employees with the ability to edit Q&As. Only senior managers must be able to publish updates. Support purchases by providing information about relevant products to customers. Product displays must include images and warnings when stock levels are low or out of stock.
Product JSON Sample
You have the following JSON sample for a product.

```
{
    "sku": "b1",
    "name": {
        "en":"Bicycle",
        "es": "Bicicleta",
        "pt": "Bicicleta"
    },
    "stocklevel": "Out of Stock",
    "description": {
        "en":"Bicycle",
        "es": "Bicicleta",
        "pt": "Bicicleta"
    },
    "image":
{"uri": "https://upload.worldwideimporters.org/bicycle.jpg",
            "alttext": {
                "en":"Bicycle",
                "es": "Bicicleta",
                "pt": "Bicicleta"
        }
       },
    "createdUtc": "2020-02-14T06:08:39Z",
    "language": "en"

}
```

QUESTION 1
HOTSPOT
You need to develop code to upload images for the product creation project. The solution must meet the accessibility requirements.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

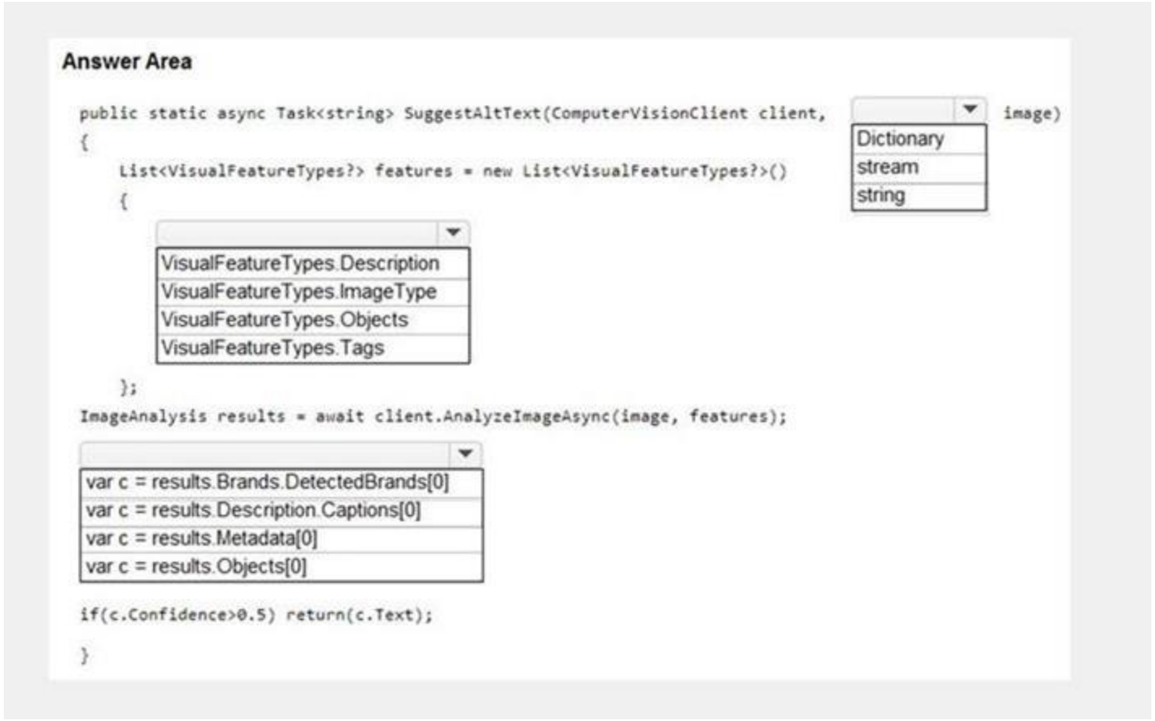

Answer Area:

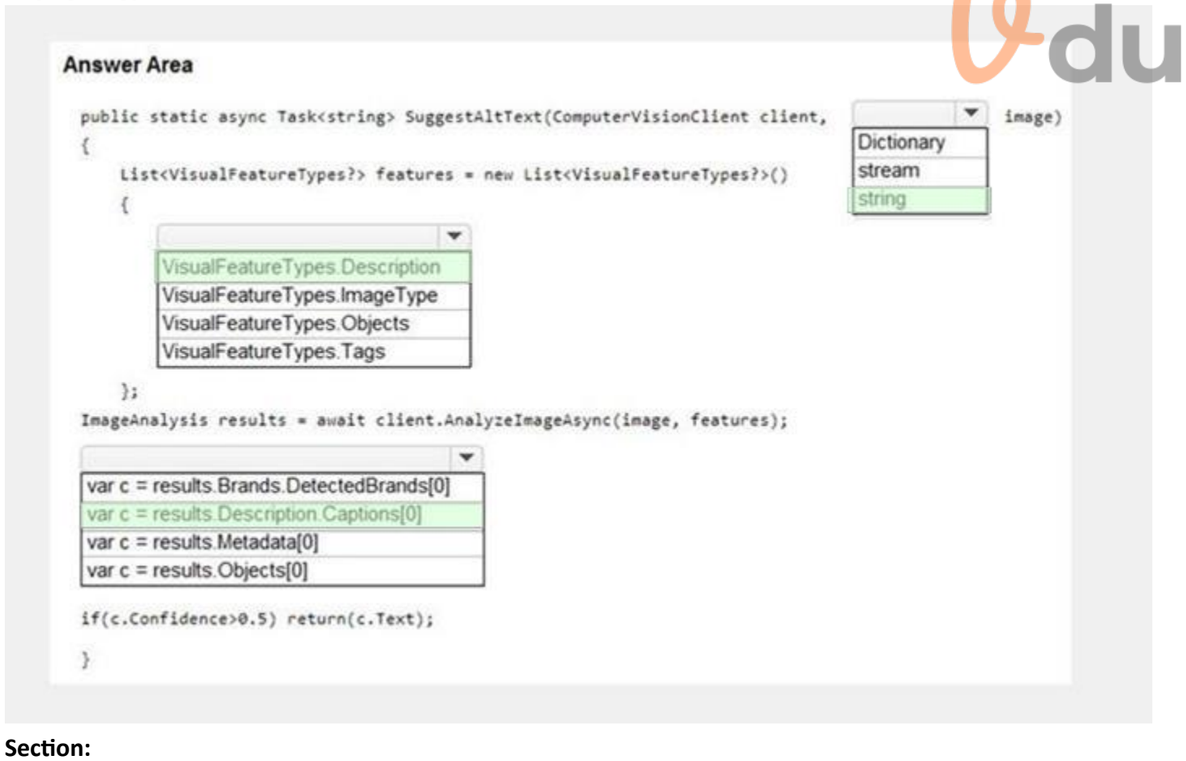

Section:
Explanation:
Reference:

https://github.com/Azure-Samples/cognitive-services-dotnet-sdk-samples/blob/master/documentation-samples/quickstarts/ComputerVision/Program.cs

QUESTION 2
DRAG DROP
You are planning the product creation project.
You need to recommend a process for analyzing videos.
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. (Choose four.)

Select and Place:

Actions	Answer Area
Index the video by using the Video Indexer API.	
Upload the video to blob storage.	
Analyze the video by using the Computer Vision API.	
Extract the transcript from Microsoft Stream.	
Send the transcript to the Language Understanding API as an utterance.	
Extract the transcript from the Video Indexer API.	
Translate the transcript by using the Translator API.	
Upload the video to file storage.	

Correct Answer:

Actions	Answer Area
	Upload the video to blob storage.
	Index the video by using the Video Indexer API.
Analyze the video by using the Computer Vision API.	Extract the transcript from the Video Indexer API.
Extract the transcript from Microsoft Stream.	Translate the transcript by using the Translator API.
Send the transcript to the Language Understanding API as an utterance.	
Upload the video to file storage.	

Section:
Explanation:
Scenario: All videos must have transcripts that are associated to the video and included in product descriptions. Product descriptions, transcripts, and alt text must be available in English, Spanish, and Portuguese.
Step 1: Upload the video to blob storage
Given a video or audio file, the file is first dropped into a Blob Storage. T
Step 2: Index the video by using the Video Indexer API.
When a video is indexed, Video Indexer produces the JSON content that contains details of the specified video insights. The insights include: transcripts, OCRs, faces, topics, blocks, etc.
Step 3: Extract the transcript from the Video Indexer API.
Step 4: Translate the transcript by using the Translator API.
Reference:
https://azure.microsoft.com/en-us/blog/get-video-insights-in-even-more-languages/
https://docs.microsoft.com/en-us/azure/media-services/video-indexer/video-indexer-output-json-v2

03 - Implement Computer Vision Solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study. At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs.

When you are ready to answer a question, click the Question button to return to the question. Overview
General Overview
Contoso, Ltd. is an international accounting company that has offices in France, Portugal, and the United Kingdom. Contoso has a professional services department that contains the roles shown in the following table.

Name	Position	Office
Accountant	Manager	United Kingdom, France, Portugal
Accountant	Consultant	United Kingdom, France, Portugal
Customer Service	Manager	United Kingdom
Customer Service	Agent	United Kingdom
Bookkeeper	Manager	United Kingdom, France, Portugal
Bookkeeper	Consultant	United Kingdom, France, Portugal

Existing environment
Infrastructure
Contoso has the following subscriptions:
Azure
Microsoft 365
Microsoft Dynamics 365
Azure Active (Azure AD) Directory
Contoso has Azure Active Directory groups for securing role-based access. The company uses the following group naming conventions:
[Country]-[Level]-[Role]
[Level]-[Role]
Intellectual Property
Contoso has the intellectual property shown in the following table.

Content	Format	Language	Content store	Domain
Weekly webinars	Video	English	Azure Blob storage	Vid.contoso.com
Blogs	Text	English, French, Portuguese	WordPress	Pt-blog.contoso.com Blog.contoso.com Fr-blog.contoso.com
Wikis	Text	English, French, Portuguese	Azure Cosmos DB	Internal.contoso.com/wiki
Monthly conference recording	Video	English	SharePoint Online	Contoso.sharepoint.com
Frequently asked questions (FAQs)	Text	English	SharePoint Online	Contoso.sharepoint.com

Text-based content is provided only in one language and is not translated.
Requirements
Planned Projects
Contoso plans to develop the following:
A document processing workflow to extract information automatically from PDFs and images of financial documents A customer-support chatbot that will answer questions by using FAQs
A searchable knowledgebase of all the intellectual property
Technical Requirements
Contoso identifies the following technical requirements:
All content must be approved before being published.
All planned projects must support English, French, and Portuguese.
All content must be secured by using role-based access control (RBAC).
RBAC role assignments must use the principle of least privilege.

RBAC roles must be assigned only to Azure Active Directory groups.
AI solution responses must have a confidence score that is equal to or greater than 70 percent.
When the response confidence score of an AI response is lower than 70 percent the response must be improved by human input. Chatbot Requirements
Contoso identifies the following requirements for the chatbot:
Provide customers with answers to the FAQs.
Ensure that the customers can chat to a customer service agent.
Ensure that the members of a group named Management-Accountants can approve the FAQs.
Ensure that the members of a group named Consultant-Accountants can create and amend the FAQs.
Ensure that the members of a group named the Agent-CustomerServices can browse the FAQs.
Ensure that access to the customer service agents is managed by using Omnichannel for Customer Service. When the response confidence score is low, ensure that the chatbot can provide other response options to the customers. Document Processing Requirements
Contoso identifies the following requirements for document processing:
The document processing solution must be able to process standardized financial documents that have the following characteristics: - Contain fewer than 20 pages.
- Be formatted as PDF or JPEG files.
- Have a distinct standard for each office.
The document processing solution must be able to extract tables and text from the financial documents. The document processing solution must be able to extract information from receipt images.
Members of a group named Management-Bookkeeper must define how to extract tables from the financial documents. Members of a group named Consultant-Bookkeeper must be able to process the financial documents. Knowledgebase Requirements
Contoso identifies the following requirements for the knowledgebase:
Supports searches for equivalent terms
Can transcribe jargon with high accuracy
Can search content in different formats, including video
Provides relevant links to external resources for further research

QUESTION 1
You are developing the knowledgebase by using Azure Cognitive Search.
You need to process wiki content to meet the technical requirements.
What should you include in the solution?

A. an indexer for Azure Blob storage attached to a skillset that contains the language detection skill and the text translation skill
B. an indexer for Azure Blob storage attached to a skillset that contains the language detection skill
C. an indexer for Azure Cosmos DB attached to a skillset that contains the document extraction skill and the text translation skill
D. an indexer for Azure Cosmos DB attached to a skillset that contains the language detection skill and the text translation skill

Correct Answer: C
Section:
Explanation:
The wiki contains text in English, French and Portuguese.
Scenario: All planned projects must support English, French, and Portuguese.
The Document Extraction skill extracts content from a file within the enrichment pipeline. This allows you to take advantage of the document extraction step that normally happens before the skillset execution with files that may be generated by other skills.
Note: The Translator Text API will be used to determine the from language. The Language detection skill is not required. Incorrect Answers:
Not A, not B: The wiki is stored in Azure Cosmos DB.
Reference:
https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-document-extraction https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-text-translation

QUESTION 2
You are developing the knowledgebase by using Azure Cognitive Search.
You need to meet the knowledgebase requirements for searching equivalent terms.
What should you include in the solution?

A. synonym map
B. a suggester
C. a custom analyzer
D. a built-in key phrase extraction skill

Correct Answer: A
Section:
Explanation:
Within a search service, synonym maps are a global resource that associate equivalent terms, expanding the scope of a query without the user having to actually provide the term. For example, assuming "dog", "canine", and "puppy" are mapped synonyms, a query on "canine" will match on a document containing "dog".
Create synonyms: A synonym map is an asset that can be created once and used by many indexes.
Reference: https://docs.microsoft.com/en-us/azure/search/search-synonyms

QUESTION 3
You are developing the document processing workflow.
You need to identify which API endpoints to use to extract text from the financial documents. The solution must meet the document processing requirements. Which two API endpoints should you identify? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. /vision/v3.1/read/analyzeResults
B. /formrecognizer/v2.0/custom/models/{modelId}/analyze
C. /formrecognizer/v2.0/prebuilt/receipt/analyze
D. /vision/v3.1/describe
E. /vision/v3.1/read/analyze

Correct Answer: C, E
Section:
Explanation:
C: Analyze Receipt - Get Analyze Receipt Result.
Query the status and retrieve the result of an Analyze Receipt operation.
Request URL: https://{endpoint}/formrecognizer/v2.0-preview/prebuilt/receipt/analyzeResults/{resultId} E: POST {Endpoint}/vision/v3.1/read/analyze
Use this interface to get the result of a Read operation, employing the state-of-the-art Optical Character Recognition (OCR) algorithms optimized for text-heavy documents.
Scenario: Contoso plans to develop a document processing workflow to extract information automatically from PDFs and images of financial documents The document processing solution must be able to process standardized financial documents that have the following characteristics: - Contain fewer than 20 pages.
- Be formatted as PDF or JPEG files.
- Have a distinct standard for each office.
*The document processing solution must be able to extract tables and text from the financial documents. The document processing solution must be able to extract information from receipt images.
Reference: https://westus2.dev.cognitive.microsoft.com/docs/services/form-recognizer-api-v2preview/operations/GetAnalyzeReceiptResult https://docs.microsoft.com/en-us/rest/api/computervision/3.1/read/read

QUESTION 4
HOTSPOT
You are developing the knowledgebase by using Azure Cognitive Search.
You need to build a skill that will be used by indexers.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Answer Area:

Answer Area

```
{
  "@odata.type": "#Microsoft.Skills.Text.EntityRecognitionSkill",
  [ "categories": [],
    "categories": [ "Email", "Persons", "Organizations"],
    "categories": [ "Locations", "Persons", "Organizations"], ]
  "defaultLanguageCode": "en",
  "includeTypelessEntities": true,
  "minimumPrecision": 0.7,
  "inputs": [
    { "name": "text",
      "source": "/document/content"}
  ],
  "outputs": [
    {"name": "persons", "targetName": "people"},
    {"name": "locations", "targetName": "locations"},
    {"name": "organizations", "targetName": "organizations"},
    [ {"name": "entities"}
      {"name": "categories"}
      {"name": "namedEntities"} ]
  ]
}
```

Section:

Explanation:

Box 1: "categories": ["Locations", "Persons", "Organizations"], Locations, Persons, Organizations are in the outputs. Scenario: Contoso plans to develop a searchable knowledgebase of all the intellectual property Note: The categories parameter is an array of categories that should be extracted. Possible category types: "Person", "Location", "Organization", "Quantity", "Datetime", "URL", "Email". If no category is provided, all types are returned.

Box 2: {"name": " entities"}

The include wikis, so should include entities in the outputs.

Note: entities is an array of complex types that contains rich information about the entities extracted from text, with the following fields name (the actual entity name. This represents a "normalized" form) wikipediaId wikipediaLanguage

wikipediaUrl (a link to Wikipedia page for the entity) etc.

Reference: https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-entity-recognition

01 - Implement Natural Language Processing Solutions

QUESTION 1

You are building an Azure Weblob that will create knowledge bases from an array of URLs.

You instantiate a QnAMakerciient object that has the relevant API keys and assign the object to a variable named client. You need to develop a method to create the knowledge bases.

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

A. Create a list of FileDTO objects that represents data from the WebJob.
B. Call the client .Knowledqebase.CreateAsync method.
C. Create a list of QnADTO objects that represents data from the WebJob.
D. Create a createKbDTO object.

Correct Answer: A, C
Section:
Explanation:
Reference: https://docs.microsoft.eom/en-us/rest/api/cognitiveservices-qnamaker/qnamaker4.0/knowledqebase/create

QUESTION 2
You are building a natural language model.
You need to enable active learning.
What should you do?

A. Add show-all-intents=true to the prediction endpoint query.
B. Enable speech priming.
C. Add iog=true to the prediction endpoint query.
D. Enable sentiment analysis.

Correct Answer: C
Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-how-to-review-endpoint-utterances#log-user-queries-to-enable-active-learning

QUESTION 3
DRAG DROP
You are building a retail chatbot that will use a QnA Maker service.
You upload an internal support document to train the model. The document contains the following question: "What is your warranty period?"
Users report that the chatbot returns the default QnA Maker answer when they ask the following question: "How long is the warranty coverage?"
The chatbot returns the correct answer when the users ask the following question: 'What is your warranty period?"
Both questions should return the same answer.
You need to increase the accuracy of the chatbot responses.
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. (Choose three.)

Select and Place:

Actions

- Add a new question and answer (QnA) pair.
- Retrain the model.
- Add additional questions to the document.
- Republish the model.
- Add alternative phrasing to the question and answer (QnA) pair.

Answer Area

Correct Answer:

Actions

- Add a new question and answer (QnA) pair.
- Add additional questions to the document.

Answer Area

- Add alternative phrasing to the question and answer (QnA) pair.
- Retrain the model.
- Republish the model.

Section:

Explanation:

Step 1: Add alternative phrasing to the question and answer (QnA) pair.

Add alternate questions to an existing QnA pair to improve the likelihood of a match to a user query.

Step 2: Retrain the model.

Periodically select Save and train after making edits to avoid losing changes.

Step 3: Republish the model

Note: A knowledge base consists of question and answer (QnA) pairs. Each pair has one answer and a pair contains all the information associated with that answer.

Reference:

https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/edit-knowledge-base

QUESTION 4

HOTSPOT
You are developing a service that records lectures given in English (United Kingdom).
You have a method named AppendToTranscriptFile that takes translated text and a language identifier.
You need to develop code that will provide transcripts of the lectures to attendees in their respective language. The supported languages are English, French, Spanish, and German.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
static async Task TranslateSpeechAsync()
{
    var config =SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");

    var lang = new List<string> [ ]
                                  {"en-GB")
                                  {"fr", "de", "es")
                                  {"French", "Spanish", "German"}
                                  {languages}

    config.SpeechRecognitionLanguage = "en-GB";
    lang.ForEach(config.AddTargetLanguage);

    using var audioConfig = AudioConfig.FromDefaultMicrophoneInput();
    using var recognizer = new [ ] (config, audioConfig);
                                  IntentRecognizer
                                  SpeakerRecognizer
                                  SpeechSynthesizer
                                  TranslationRecognizer

    var result = await recognizer.RecognizeOnceAsync();
    if (result.Reason == ResultReason.TranslatedSpeech)
```


Answer Area:

Answer Area

```
static async Task TranslateSpeechAsync()
{
    var config =SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");

    var lang = new List<string> [ ]
                                  {"en-GB")
                                  {"fr", "de", "es")
                                  {"French", "Spanish", "German"}
                                  {languages}

    config.SpeechRecognitionLanguage = "en-GB";
    lang.ForEach(config.AddTargetLanguage);

    using var audioConfig = AudioConfig.FromDefaultMicrophoneInput();
    using var recognizer = new [ ] (config, audioConfig);
                                  IntentRecognizer
                                  SpeakerRecognizer
                                  SpeechSynthesizer
                                  TranslationRecognizer

    var result = await recognizer.RecognizeOnceAsync();
    if (result.Reason == ResultReason.TranslatedSpeech)
```

Section:
Explanation:
Box 1: {"fr", "de", "es"}
A common task of speech translation is to specify target translation languages, at least one is required but multiples are supported. The following code snippet sets both French and German as translation language targets.

```
static async Task TranslateSpeechAsync()
{
var translationConfig =
SpeechTranslationConfig.FromSubscription(SPEECH__SUBSCRIPTION__KEY, SPEECH__SERVICE__REGION);
translationConfig.SpeechRecognitionLanguage = "it-IT";
// Translate to languages. See, https://aka.ms/speech/sttt-languages
translationConfig.AddTargetLanguage("fr");
translationConfig.AddTargetLanguage("de");
}
```

Box 2: TranslationRecognizer
After you've created a SpeechTranslationConfig, the next step is to initialize a TranslationRecognizer.
Example code:

```
static async Task TranslateSpeechAsync()
{
var translationConfig =
SpeechTranslationConfig.FromSubscription(SPEECH__SUBSCRIPTION__KEY, SPEECH__SERVICE__REGION);
var fromLanguage = "en-US";
var toLanguages = new List<string> { "it", "fr", "de" };
translationConfig.SpeechRecognitionLanguage = fromLanguage;
toLanguages.ForEach(translationConfig.AddTargetLanguage);
using var recognizer = new TranslationRecognizer(translationConfig);
}
```

QUESTION 5
You are building a Language Understanding model for an e-commerce platform.
You need to construct an entity to capture billing addresses.
Which entity type should you use for the billing address?

A. machine learned
B. Regex
C. geographyV2
D. Pattern, any
E. list

Correct Answer: B
Section:
Explanation:
A regular expression entity extracts an entity based on a regular expression pattern you provide. It ignores case and ignores cultural variant. Regular expression is best for structured text or a predefined sequence of alphanumeric values that are expected in a certain format. For example:

Entity	Regular expression	Example
Flight Number	flight [A-Z]{2} [0-9]{4}	flight AS 1234
Credit Card Number	[0-9]{16}	5478789865437632

Incorrect Answers:

C: The prebuilt geographyV2 entity detects places. Because this entity is already trained, you do not need to add example utterances containing GeographyV2 to the application intents. GeographyV2 entity is supported in English culture.
The geographical locations have subtypes:

Subtype	Purpose
`poi`	point of interest
`city`	name of city
`countryRegion`	name of country or region
`continent`	name of continent
`state`	name of state or province

D: Pattern.any is a variable-length placeholder used only in a pattern's template utterance to mark where the entity begins and ends.
E: A list entity represents a fixed, closed set of related words along with their synonyms. You can use list entities to recognize multiple synonyms or variations and extract a normalized output for them. Use the recommend option to see suggestions for new words based on the current list.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-entity-types

QUESTION 6
You need to upload speech samples to a Speech Studio project for use in training.
How should you upload the samples?

A. Combine the speech samples into a single audio file in the wma format and upload the file.
B. Upload a zip file that contains a collection of audio files in the wav format and a corresponding text transcript file.
C. Upload individual audio files in the FLAC format and manually upload a corresponding transcript in Microsoft Word format.
D. Upload individual audio files in the .wma format.

Correct Answer: B
Section:
Explanation:
To upload your data, navigate to the Speech Studio . From the portal, click Upload data to launch the wizard and create your first dataset. You'll be asked to select a speech data type for your dataset, before allowing you to upload your data.
The default audio streaming format is WAV
Use this table to ensure that your audio files are formatted correctly for use with Custom Speech:

Property	Value
File format	RIFF (WAV)
Sample rate	8,000 Hz or 16,000 Hz
Channels	1 (mono)
Maximum length per audio	2 hours
Sample format	PCM, 16-bit
Archive format	.zip
Maximum archive size	2 GB

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/how-to-custom-speech-test-and-train

QUESTION 7
You are developing a method for an application that uses the Translator API.
The method will receive the content of a webpage, and then translate the content into Greek (el). The result will also contain a transliteration that uses the Roman alphabet.
You need to create the URI for the call to the Translator API.
You have the following URI.
https://api.cognitive.microsofttranslator.com/translate?api-version=3.0
Which three additional query parameters should you include in the URI? Each correct answer presents part of the solution. (Choose three.)
NOTE: Each correct selection is worth one point.

A. toScript=Cyrl
B. from=el
C. textType=html
D. to=el
E. textType=plain
F. toScript=Latn

Correct Answer: A, D, F
Section:
Explanation:

QUESTION 8
You have a chatbot that was built by using the Microsoft Bot Framework.
You need to debug the chatbot endpoint remotely.
Which two tools should you install on a local computer? Each correct answer presents part of the solution. (Choose two.) NOTE: Each correct selection is worth one point.

A. Fiddler
B. Bot Framework Composer
C. Bot Framework Emulator

D. Bot Framework CLI
E. ngrok
F. nginx

Correct Answer: C, E
Section:
Explanation:
E xp I an ati o n/Ref e re n ce:
Bot Framework Emulator is a desktop application that allows bot developers to test and debug bots, either locally or remotely. ngrok is a cross-platform application that "allows you to expose a web server running on your local machine to the internet." Essentially, what we'll be doing is using ngrok to forward messages from external channels on the web directly to our local machine to allow debugging, as opposed to the standard messaging endpoint configured in the Azure portal.
Reference:
https://docs.microsoftcom/en-us/azure/bot-service/bot-service-debuq-emulator

QUESTION 9
You need to measure the public perception of your brand on social media by using natural language processing. Which Azure Cognitive Services service should you use?

A. Text Analytics
B. Content Moderator
C. Computer Vision
D. Form Recognizer

Correct Answer: A
Section:
Explanation:
Text Analytics Cognitive Service could be used to quickly determine the public perception for a specific topic, event or brand.
Example: A NodeJS app which pulls Tweets from Twitter using the Twitter API based on a specified search term. Then pass these onto Text Analytics for sentiment scoring before storing the data and building a visualisation in PowerBI. The Architecture looked something like this:

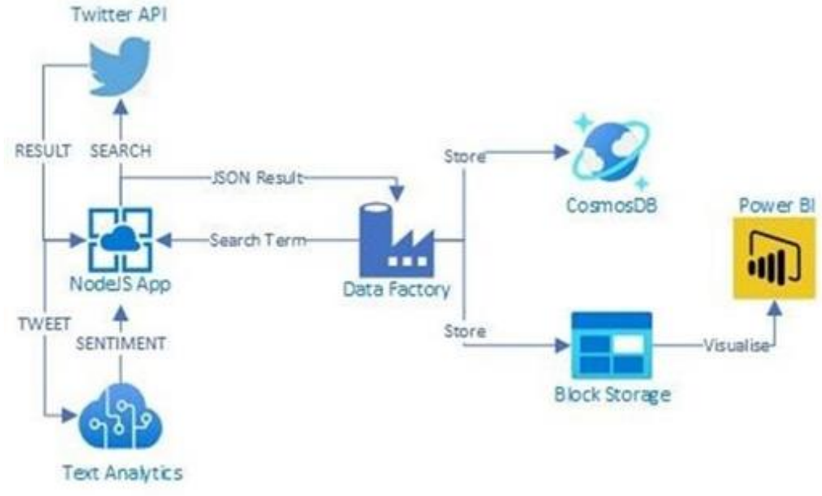

Reference:
https://www.linkedin.com/pulse/measuring-public-perception-azure-cognitive-services-steve-dalai
https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/overview

QUESTION 10

You are building a Language Understanding model for an e-commerce chatbot. Users can speak or type their billing address when prompted by the chatbot. You need to construct an entity to capture billing addresses.
Which entity type should you use?

A. machine learned
B. Regex
C. list
D. Pattern.any

Correct Answer: B
Section:
Explanation:
Reference:
https://docs. microsoft. com/en-us/azure/coqnitive-services/luis/lui s-concept-entity-types

QUESTION 11
DRAG DROP
You train a Custom Vision model used in a mobile app.
You receive 1,000 new images that do not have any associated data.
You need to use the images to retrain the model. The solution must minimize how long it takes to retrain the model.
Which three actions should you perform in the Custom Vision portal? 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:

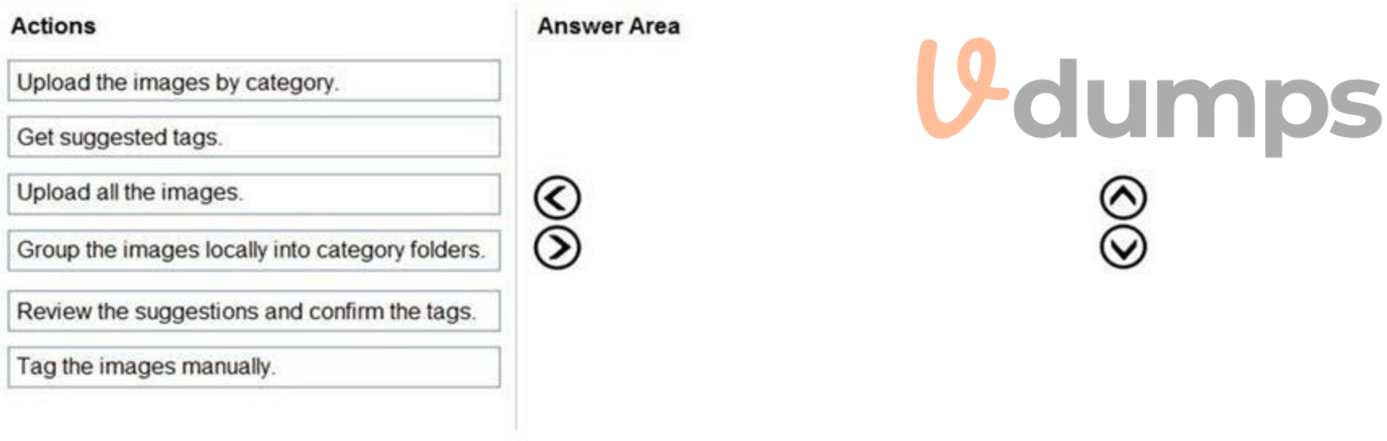

Correct Answer:

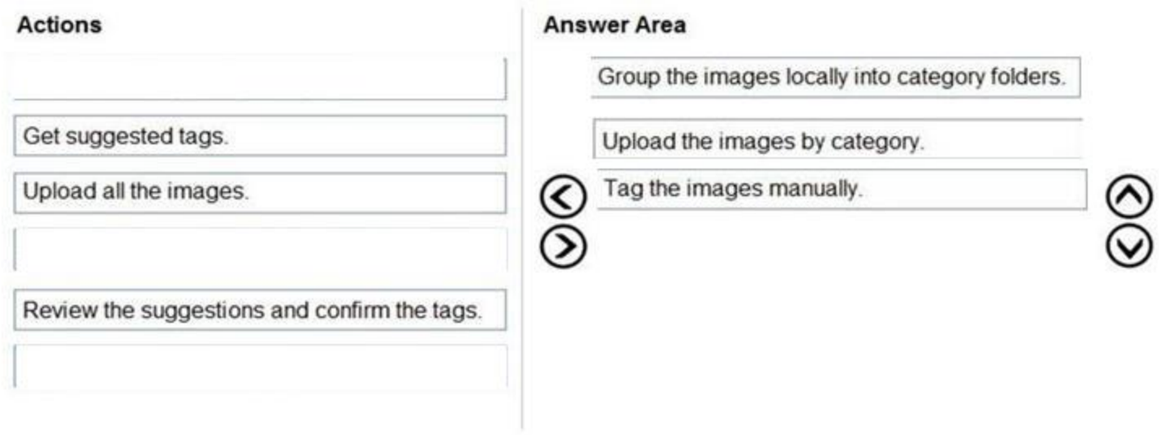

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/getting-started-build-a-classifier

QUESTION 12
HOTSPOT
You are developing an application that includes language translation.
The application will translate text retrieved by using a function named getTextToBeTranslated. The text can be in one of many languages. The content of the text must remain within the Americas Azure geography.
You need to develop code to translate the text to a single language.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
. . .
var endpoint = [ ▼ ] ;
```

- "https://api.cognitive.microsofttranslator.com/translate";
- "https://api.cognitive.microsofttranslator.com/transliterate";
- "https://api-apc.cognitive.microsofttranslator.com/detect";
- "https://api-nam.cognitive.microsofttranslator.com/detect";
- "https://api-nam.cognitive.microsofttranslator.com/translate";

```
var apiKey = "FF956C68B83B21B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);
[ ▼ ]
```

- var uri = endpoint + "?from=en";
- var uri = endpoint + "?suggestedFrom=en";
- var uri = endpoint + "?to=en";

```
HttpResponseMessage response;
var content = new StringContent(body, Encoding.UTF8, "application/json");
var response = await client.PutAsync(uri, content);
. . .
```

Answer Area:

Answer Area

```
. . .
var endpoint = [ ▼ ] ;
```

- "https://api.cognitive.microsofttranslator.com/translate";
- "https://api.cognitive.microsofttranslator.com/transliterate";
- "https://api-apc.cognitive.microsofttranslator.com/detect";
- "https://api-nam.cognitive.microsofttranslator.com/detect";
- "https://api-nam.cognitive.microsofttranslator.com/translate";

```
var apiKey = "FF956C68B83B21B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);
[ ▼ ]
```

- var uri = endpoint + "?from=en";
- var uri = endpoint + "?suggestedFrom=en";
- var uri = endpoint + "?to=en";

```
HttpResponseMessage response;
var content = new StringContent(body, Encoding.UTF8, "application/json");
var response = await client.PutAsync(uri, content);
. . .
```

Section:
Explanation:

QUESTION 13
HOTSPOT
You run the following command.

```
docker run --rm -it -p 5000:5000 --memory 10g --cpus 2 \
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment\
Eula=accept \
Billing={ENDPOINT_URI} \
ApiKey={API_KEY}
```

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

Hot Area:

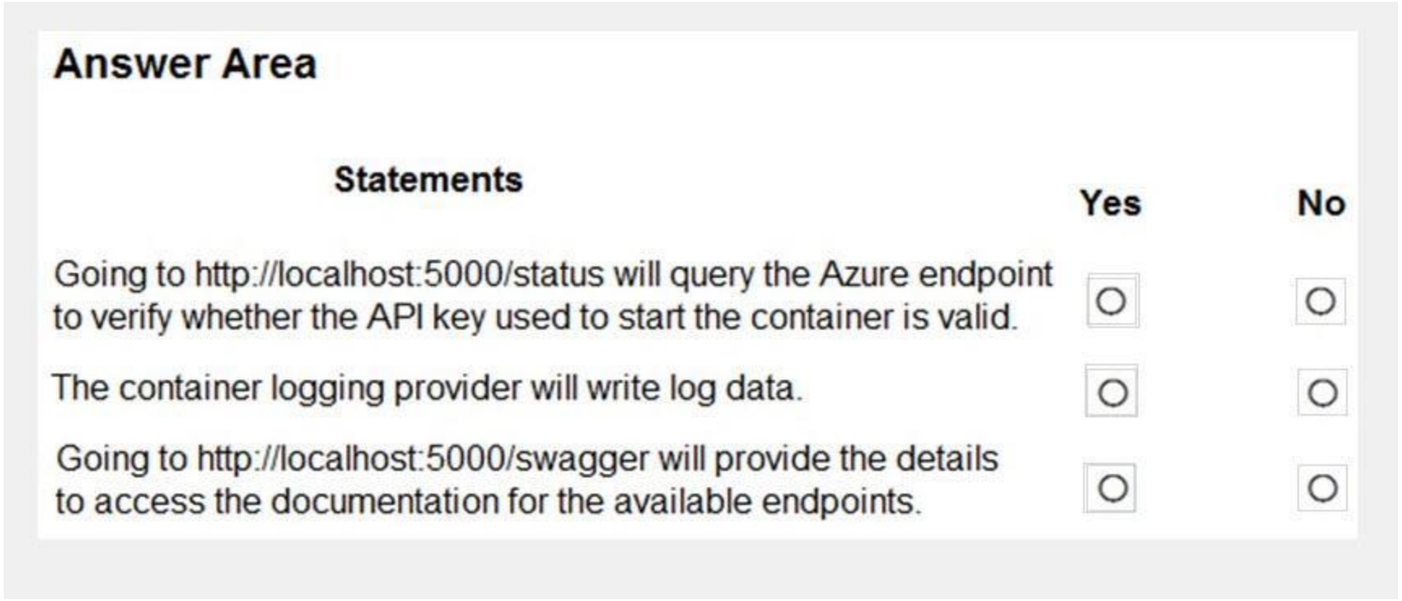

Answer Area:

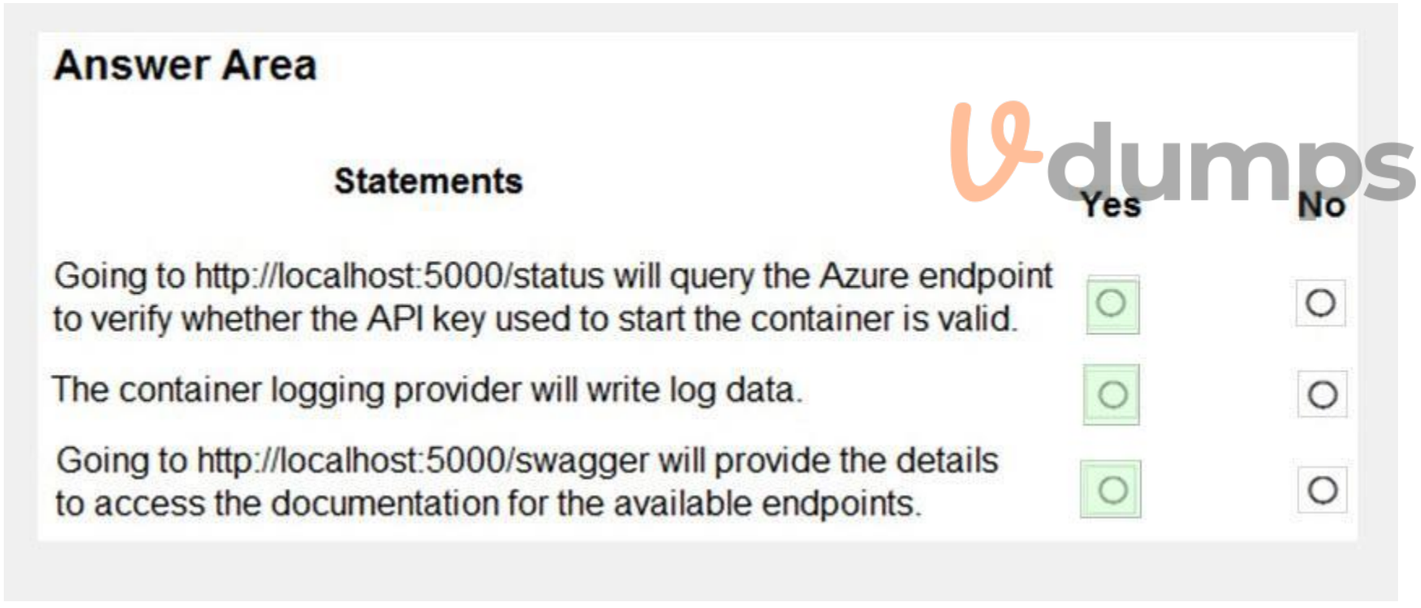

Section:
Explanation:
Box 1: Yes
http://localhost:5000/status : Also requested with GET, this verifies if the api-key used to start the container is valid without causing an endpoint query.
Box 2: Yes
The command saves container and LUIS logs to output mount at C:\output, located on container host
Box 3: Yes
http://localhost:5000/swagger : The container provides a full set of documentation for the endpoints and a Try it out feature. With this feature, you can enter your settings into a web-based HTML form and make the query without having to write any code. After the query returns, an example CURL command is provided to demonstrate the HTTP headers and body format that's required.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-container-howto

QUESTION 14
You are training a Language Understanding model for a user support system.
You create the first intent named GetContactDetails and add 200 examples.
You need to decrease the likelihood of a false positive.
What should you do?

A. Enable active learning.
B. Add a machine learned entity.
C. Add additional examples to the GetContactDetails intent.
D. Add examples to the None intent.

Correct Answer: A
Section:
Explanation:
Active learning is a technique of machine learning in which the machine learned model is used to identify informative new examples to label. In LUIS, active learning refers to adding utterances from the endpoint traffic whose current predictions are unclear to improve your model.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-glossary

QUESTION 15
DRAG DROP
You are building a Language Understanding model for purchasing tickets.
You have the following utterance for an intent named PurchaseAndSendTickets.
Purchase [2 audit business] tickets to [Paris] [next Monday] and send tickets to [email@domain.com]
You need to select the entity types. The solution must use built-in entity types to minimize training data whenever possible.
Which entity type should you use for each label? To answer, drag the appropriate entity types to the correct labels. Each entity 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.

Select and Place:

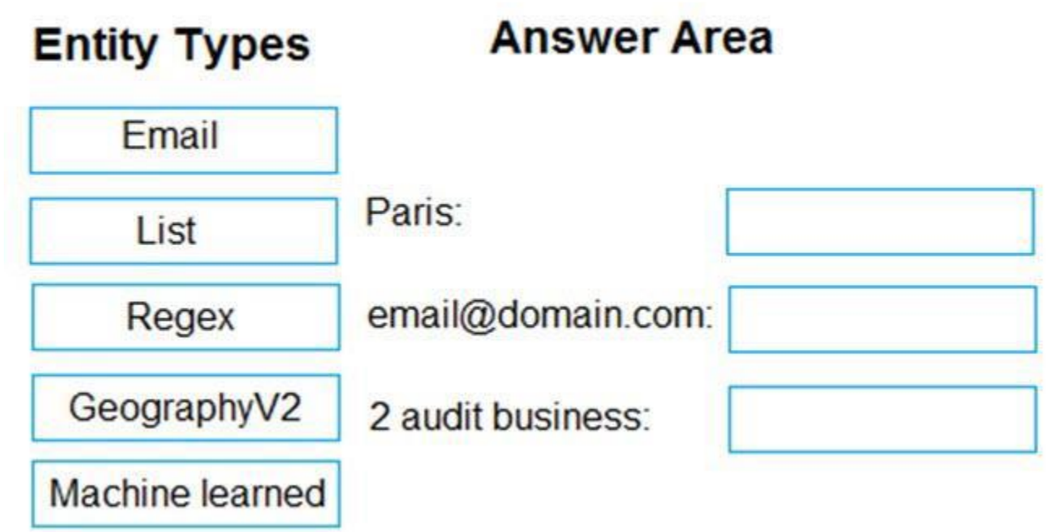

Correct Answer:

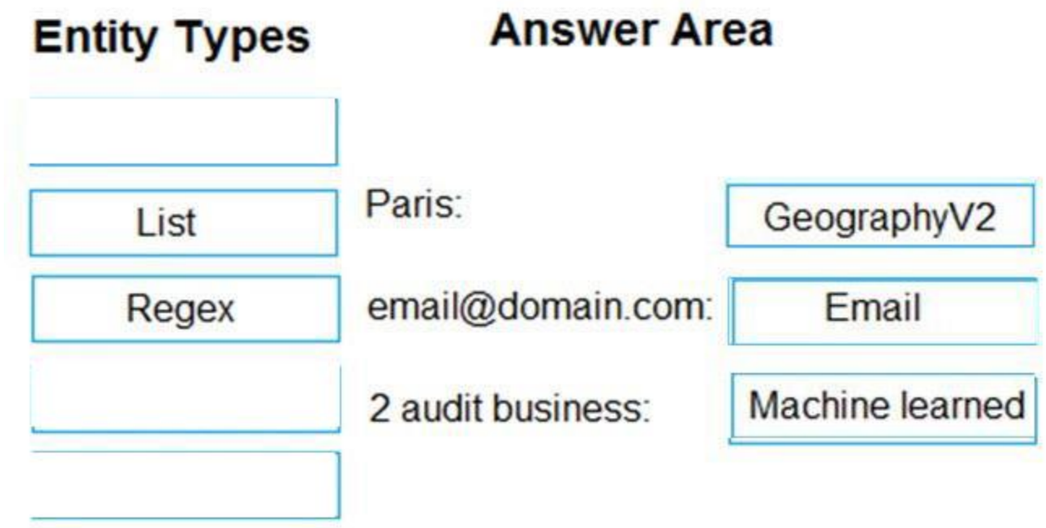

Section:
Explanation:
Box 1: GeographyV2
The prebuilt geographyV2 entity detects places. Because this entity is already trained, you do not need to add example utterances containing GeographyV2 to the application intents.
Box 2: Email
Email prebuilt entity for a LUIS app: Email extraction includes the entire email address from an utterance. Because this entity is already trained, you do not need to add example utterances containing email to the application intents.
Box 3: Machine learned
The machine-learning entity is the preferred entity for building LUIS applications.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-geographyv2
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-email
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/reference-entity-machine-learned-entity

QUESTION 16
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 build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.
A conversational expert provides you with the following list of phrases to use for training.
Find contacts in London.
Who do I know in Seattle?
Search for contacts in Ukraine.
You need to implement the phrase list in Language Understanding.
Solution: You create a new intent for location.
Does this meet the goal?

A. Yes
B. No

Correct Answer: A
Section:

Explanation:
An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance. Define a set of intents that corresponds to actions users want to take in your application.
Reference:
https://docs.microsoft.com/en-us/azure/coqnitive-services/luis/luis-concept-intent

QUESTION 17
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 build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.
A conversational expert provides you with the following list of phrases to use for training.
Find contacts in London.
Who do I know in Seattle?
Search for contacts in Ukraine.
You need to implement the phrase list in Language Understanding.
Solution: You create a new entity for the domain.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
Instead use a new intent for location.
Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance. Define a set of intents that corresponds to actions users want to take in your application.
Reference:
https://docs.microsoft.com/en-us/azure/coqnitive-services/luis/luis-concept-intent

QUESTION 18
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 build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.
A conversational expert provides you with the following list of phrases to use for training.
Find contacts in London.
Who do I know in Seattle?
Search for contacts in Ukraine.
You need to implement the phrase list in Language Understanding.
Solution: You create a new pattern in the FindContact intent.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance. Define a set of intents that corresponds to actions users want to take in your application.
Reference:
https://docs.microsoft.com/en-us/azure/coqnitive-services/luis/luis-concept-intent

QUESTION 19
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an application to identify species of flowers by training a Custom Vision model.
You receive images of new flower species.
You need to add the new images to the classifier.
Solution: You add the new images, and then use the Smart Labeler tool.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
The model need to be extended and retrained.
Note: Smart Labeler to generate suggested tags for images. This lets you label a large number of images more quickly when training a Custom Vision model.

QUESTION 20
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an application to identify species of flowers by training a Custom Vision model.
You receive images of new flower species.
You need to add the new images to the classifier.
Solution: You add the new images and labels to the existing model. You retrain the model, and then publish the model.
Does this meet the goal?

A. Yes
B. No

Correct Answer: A
Section:
Explanation:
The model needs to be extended and retrained.

QUESTION 21
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an application to identify species of flowers by training a Custom Vision model.
You receive images of new flower species.
You need to add the new images to the classifier.
Solution: You create a new model, and then upload the new images and labels.

Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:
Explanation:
E xp l an ati o n/Ref e re n ce:
The model needs to be extended and retrained.

02 - Implement Natural Language Processing Solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Overview
A company named Wide World Importers is developing an e-commerce platform.
You are working with a solutions architect to design and implement the features of the e-commerce platform. The platform will use microservices and a serverless environment built on Azure.
Wide World Importers has a customer base that includes English, Spanish, and Portuguese speakers.
Existing Environment
Applications
Wide World Importers has an App Service plan that contains the web apps shown in the following table.

Name	Description
Product Management	An app used by employees to create and manage products. The app and the expected inputs from the employees are in English.
Inventory Tracking	An app used by employees to manage inventory when dispatching orders, receiving refunds, and receiving consignments from suppliers.

Azure Resources
You have the following resources:
▪ An Azure Active Directory (Azure AD) tenant
An Azure Active Directory (Azure AD) tenant
- The tenant supports internal authentication.
- All employees belong to a group named AllUsers.
- Senior managers belong to a group named LeadershipTeam.
An Azure Functions resource
- A function app posts to Azure Event Grid when stock levels of a product change between OK, Low Stock, and Out of Stock. The function app uses the Azure Cosmos DB change feed. An Azure Cosmos DB account
- The account uses the Core (SQL) API.
- The account stores data for the Product Management app and the Inventory Tracking app.
▪ An Azure Storage account
An Azure Storage account
- The account contains blob containers for assets related to products.
- The assets include images, videos, and PDFs.
An Azure Cognitive Services resource named wwics
A Video Indexer resource named wwivi

Requirements
Business Goals
Wide World Importers wants to leverage AI technologies to differentiate itself from its competitors.
Planned Changes
Wide World Importers plans to start the following projects:
A product creation project: Help employees create accessible and multilingual product entries, while expediting product entry creation. A smart e-commerce project: Implement an Azure Cognitive Search solution to display products for customers to browse. A shopping on-the-go project: Build a chatbot that can be integrated into smart speakers to support customers.
Business Requirements
Wide World Importers identifies the following business requirements for all the projects:
Provide a multilingual customer experience that supports English, Spanish, and Portuguese.
Whenever possible, scale based on transaction volumes to ensure consistent performance.
Minimize costs.
Governance and Security Requirements
Wide World Importers identifies the following governance and security requirements:
Data storage and processing must occur in datacenters located in the United States.
Azure Cognitive Services must be inaccessible directly from the internet.
Accessibility Requirements
Wide World Importers identifies the following accessibility requirements:
All images must have relevant alt text.
All videos must have transcripts that are associated to the video and included in product descriptions. Product descriptions, transcripts, and alt text must be available in English, Spanish, and Portuguese.
Product Creation Requirements
Wide World Importers identifies the following requirements for improving the Product Management app:
Minimize how long it takes for employees to create products and add assets.
Remove the need for manual translations.
Smart E-Commerce Requirements
Wide World Importers identifies the following requirements for the smart e-commerce project:
Ensure that the Cognitive Search solution meets a Service Level Agreement (SLA) of 99.9% availability for searches and index writes. Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products. Support autocompletion and autosuggestion based on all product name variants.
Store all raw insight data that was generated, so the data can be processed later.
Update the stock level field in the product index immediately upon changes.
Update the product index hourly.
Shopping On-the-Go Requirements
Wide World Importers identifies the following requirements for the shopping on-the-go chatbot:
Answer common questions.
Support interactions in English, Spanish, and Portuguese.
Replace an existing FAQ process so that all Q&A is managed from a central location.
Provide all employees with the ability to edit Q&As. Only senior managers must be able to publish updates. Support purchases by providing information about relevant products to customers. Product displays must include images and warnings when stock levels are low or out of stock.
Product JSON Sample
You have the following JSON sample for a product.

```
{
    "sku": "b1",
    "name": {
        "en":"Bicycle",
        "es": "Bicicleta",
        "pt": "Bicicleta"
    },
    "stocklevel": "Out of Stock",
    "description": {
        "en":"Bicycle",
        "es": "Bicicleta",
        "pt": "Bicicleta"
    },
    "image":
{"uri": "https://upload.worldwideimporters.org/bicycle.jpg",
           "alttext": {
                "en":"Bicycle",
                "es": "Bicicleta",
                "pt": "Bicicleta"
        }
       },
    "createdUtc": "2020-02-14T06:08:39Z",
    "language": "en"

}
```


QUESTION 1
HOTSPOT
You are planning the product creation project.
You need to build the REST endpoint to create the multilingual product descriptions.
How should you complete the URI? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

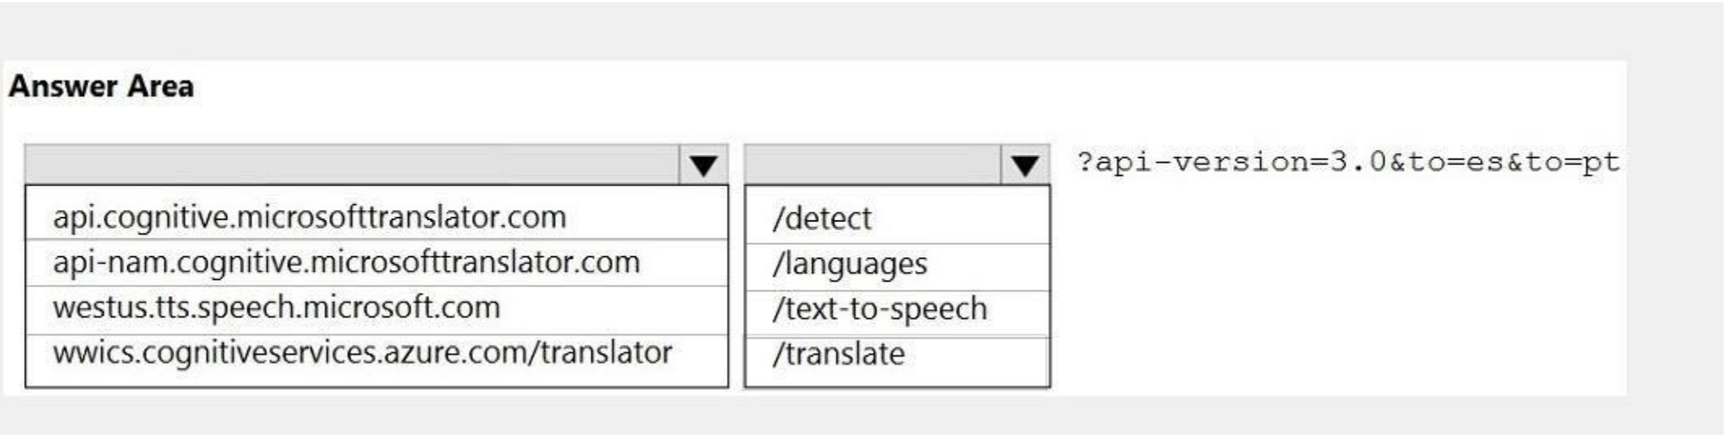

Answer Area:

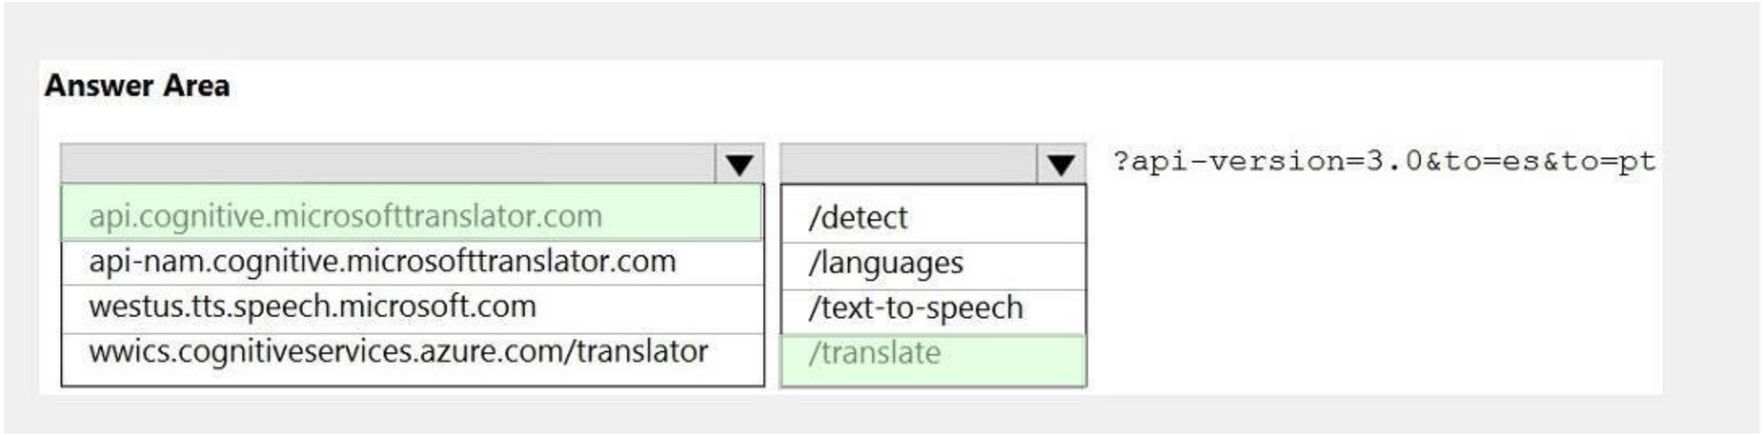

Section:
Explanation:
Box 1: api.cognitive.microsofttranslator.com
Translator 3.0: Translate. Send a POST request to:
https://api.cognitive.microsofttranslator.com/translate?api-version=3.0
Box 2: /translate
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/translator/reference/v3-0-translate

01 - Implement Knowledge Mining Solutions

QUESTION 1
HOTSPOT
You are developing a text processing solution.
You develop the following method.

```
static void GetKeyPhrases(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.ExtractKeyPhrases(text);
    Console.WriteLine("Key phrases:");

    foreach (string keyphrase in response.Value)
    {
        Console.WriteLine($"\t{keyphrase}");
    }
}
```

You call the method by using the following code.
GetKeyPhrases(textAnalyticsClient, "the cat sat on the mat");
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
The call will output key phrases from the input string to the console.	○	○
The output will contain the following words: the, cat, sat, on, and mat.	○	○
The output will contain the confidence level for key phrases.	○	○

Answer Area:

Answer Area

Statements	Yes	No
The call will output key phrases from the input string to the console.	○	○
The output will contain the following words: the, cat, sat, on, and mat.	○	○
The output will contain the confidence level for key phrases.	○	○

Section:
Explanation:
Box 1: Yes
The Key Phrase Extraction API evaluates unstructured text, and for each JSON document, returns a list of key phrases.
Box 2: No
'the' is not a key phrase.
This capability is useful if you need to quickly identify the main points in a collection of documents. For example, given input text "The food was delicious and there were wonderful staff", the service returns the main talking points: "food" and "wonderful staff".
Box 3: No
Key phrase extraction does not have confidence levels.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-keyword-extraction

QUESTION 2
HOTSPOT
You are creating an enrichment pipeline that will use Azure Cognitive Search. The knowledge store contains unstructured JSON data and scanned PDF documents that contain text.
Which projection type should you use for each data type? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

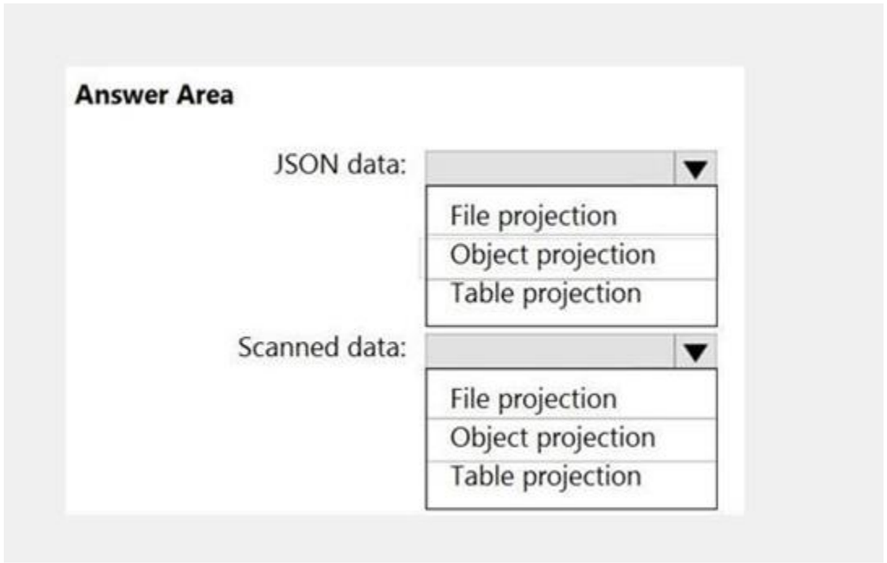

Answer Area:

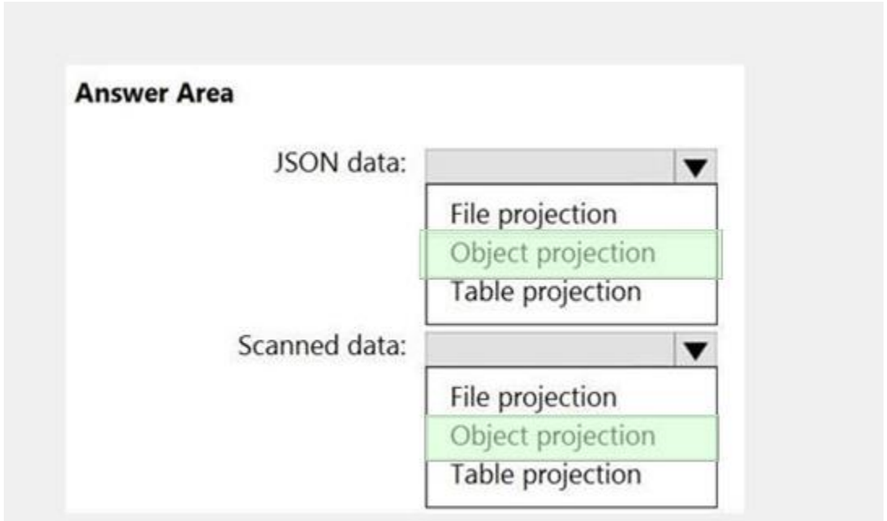

Section:
Explanation:
Box 1: Object projection
Object projections are JSON representations of the enrichment tree that can be sourced from any node.
Box 2: File projection
File projections are similar to object projections and only act on the normalized_images collection.
Reference:
https://docs.microsoft.com/en-us/azure/search/knowledge-store-projection-overview

QUESTION 3
HOTSPOT
You are building an Azure Cognitive Search custom skill.

You have the following custom skill schema definition.

```
{
    "@odata.type": "#Microsoft.Skills.Custom.WebApiSkill",
    "description": "My custom skill description",
    "uri": "https://contoso-webskill.azurewebsites.net/api/process",
    "context": "/document/organizations/*",
    "inputs": [
      {
        "name": "companyName",
        "source": "/document/organizations/*"
      }
    ],
    "outputs": [
      {
        "name": "companyDescription",
      }
    ]
}
```

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

Hot Area:

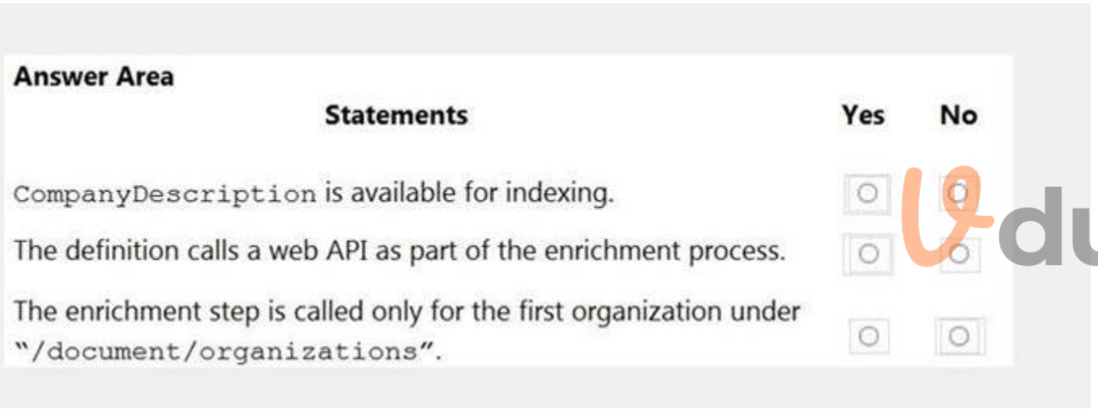

Answer Area:

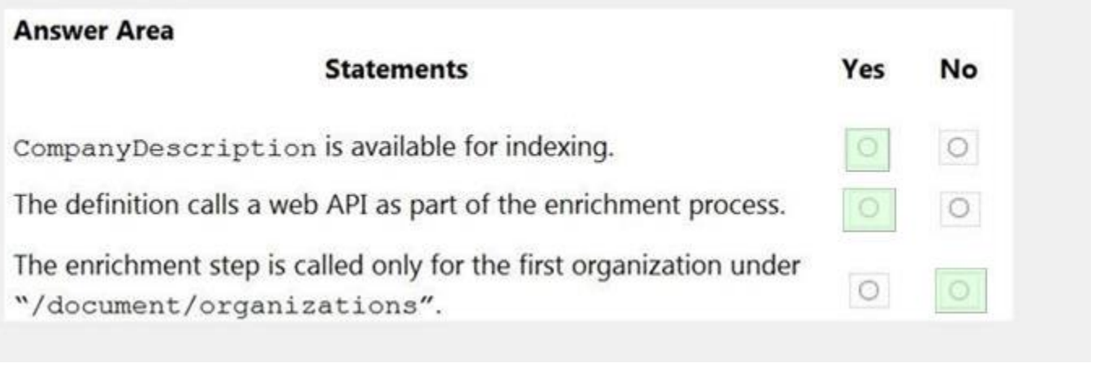

Section:
Explanation:
Box 1: Yes
Once you have defined a skillset, you must map the output fields of any skill that directly contributes values to a given field in your search index.
Box 2: Yes

The definition is a custom skill that calls a web API as part of the enrichment process.
Box 3: No
For each organization identified by entity recognition, this skill calls a web API to find the description of that organization.
Reference:
https://docs.microsoft.com/en-us/azure/search/cognitive-search-output-field-mapping

QUESTION 4
You deploy a web app that is used as a management portal for indexing in Azure Cognitive Search. The app is configured to use the primary admin key. During a security review, you discover unauthorized changes to the search index. You suspect that the primary access key is compromised. You need to prevent unauthorized access to the index management endpoint. The solution must minimize downtime. What should you do next?

A. Regenerate the primary admin key, change the app to use the secondary admin key, and then regenerate the secondary admin key.
B. Change the app to use a query key, and then regenerate the primary admin key and the secondary admin key.
C. Regenerate the secondary admin key, change the app to use the secondary admin key, and then regenerate the primary key.
D. Add a new query key, change the app to use the new query key, and then delete all the unused query keys.

Correct Answer: A
Section:
Explanation:
Regenerate admin keys.
Two admin keys are created for each service so that you can rotate a primary key, using the secondary key for business continuity. 1. In the Settings >Keys page, copy the secondary key.
2. For all applications, update the API key settings to use the secondary key.
3. Regenerate the primary key.
4. Update all applications to use the new primary key.
Note: Two admin api-keys, referred to as primary and secondary keys in the portal, are automatically generated when the service is created and can be individually regenerated on demand. Having two keys allows you to roll over one key while using the second key for continued access to the service.
Reference:
https://docs.microsoft.com/en-us/azure/search/search-security-api-keys#regenerate-admin-keys

QUESTION 5
You have the following data sources:
- Finance: On-premises Microsoft SQL Server database
- Sales: Azure Cosmos DB using the Core (SQL) API
- Logs: Azure Table storage
- HR: Azure SQL database
You need to ensure that you can search all the data by using the Azure Cognitive Search REST API.
What should you do?

A. Configure multiple read replicas for the data in Sales.
B. Mirror Finance to an Azure SQL database.
C. Migrate the data in Sales to the MongoDB API.
D. Ingest the data in Logs into Azure Sentinel.

Correct Answer: B
Section:
Explanation:
On-premises Microsoft SQL Server database cannot be used as an index data source.
Note: Indexer in Azure Cognitive Search: : Automate aspects of an indexing operation by configuring a data source and an indexer that you can schedule or run on demand. This feature is supported for a limited number of data source types on Azure.
Indexers crawl data stores on Azure.

Azure Blob Storage
Azure Data Lake Storage Gen2 (in preview)
Azure Table Storage
Azure Cosmos DB
Azure SQL Database
SQL Managed Instance
SQL Server on Azure Virtual Machines
Reference:
https://docs.microsoft.com/en-us/azure/search/search-indexer-overview#supported-data-sources

QUESTION 6
You are developing a solution to generate a word cloud based on the reviews of a company's products. Which Text Analytics REST API end point should you use?

A. keyPhrases
B. sentiment
C. languages
D. entities/recognition/general

Correct Answer: A
Section:
Explanation:
Reference: https://docs.microsoft.com/en-us/azure/coqnitive-services/text-analytics/overview

QUESTION 7
DRAG DROP
You have a web app that uses Azure Cognitive Search.
When reviewing billing for the app, you discover much higher than expected charges. You suspect that the query key is compromised.
You need to prevent unauthorized access to the search endpoint and ensure that users only have read only access to the documents collection. The solution must minimize app downtime.
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:

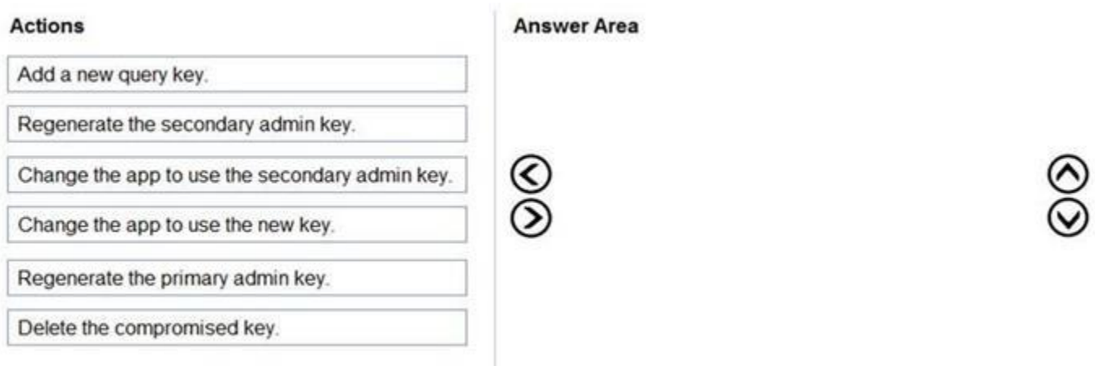

Correct Answer:

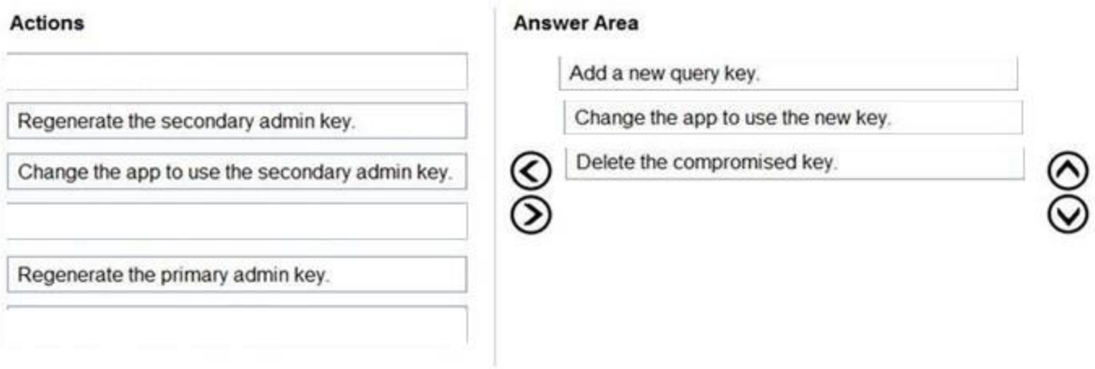

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/search/search-security-api-keys

QUESTION 8
You are developing an application that will use Azure Cognitive Search for internal documents.
You need to implement document-level filtering for Azure Cognitive Search.
Which three actions should you include in the solution? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Send Azure AD access tokens with the search request.
B. Retrieve all the groups.
C. Retrieve the group memberships of the user.
D. Add allowed groups to each index entry.
E. Create one index per group.
F. Supply the groups as a filter for the search requests.

Correct Answer: C, D
Section:
Explanation:
Your documents must include a field specifying which groups have access. This information becomes the filter criteria against which documents are selected or rejected from the result set returned to the issuer.
D: A query request targets the documents collection of a single index on a search service.
CF: In order to trim documents based on group_ids access, you should issue a search query with a group_ids/any(g:search.in(g, 'group_id1, group_id2,...')) filter, where 'group_id1, group_id2,...' are the groups to which the search request issuer belongs.
Reference:
https://docs.microsoft.com/en-us/azure/search/search-security-trimming-for-azure-search

QUESTION 9
You have an existing Azure Cognitive Search service.
You have an Azure Blob storage account that contains millions of scanned documents stored as images and PDFs. You need to make the scanned documents available to search as quickly as possible.
What should you do?

A. Split the data into multiple blob containers. Create a Cognitive Search service for each container. Within each indexer definition, schedule the same runtime execution pattern.
B. Split the data into multiple blob containers. Create an indexer for each container. Increase the search units. Within each indexer definition, schedule a sequential execution pattern.
C. Create a Cognitive Search service for each type of document.

D. Split the data into multiple virtual folders. Create an indexer for each folder. Increase the search units. Within each indexer definition, schedule the same runtime execution pattern.

Correct Answer: D
Section:
Explanation:
Incorrect Answers:
A: Need more search units to process the data in parallel.
B: Run them in parallel, not sequentially.
C: Need a blob indexer.
Note: A blob indexer is used for ingesting content from Azure Blob storage into a Cognitive Search index.
Index large datasets
Indexing blobs can be a time-consuming process. In cases where you have millions of blobs to index, you can speed up indexing by partitioning your data and using multiple indexers to process the data in parallel. Here's how you can set this up:
Partition your data into multiple blob containers or virtual folders
Set up several data sources, one per container or folder.
Create a corresponding indexer for each data source. All of the indexers should point to the same target search index. One search unit in your service can run one indexer at any given time. Creating multiple indexers as described above is only useful if they actually run in parallel.
Reference:
https://docs.microsoft.com/en-us/azure/search/search-howto-indexing-azure-blob-storage

QUESTION 10
You need to implement a table projection to generate a physical expression of an Azure Cognitive Search index. Which three properties should you specify in the skillset definition JSON configuration table node? Each correct answer presents part of the solution. (Choose three.) NOTE: Each correct selection is worth one point.

A. tableName
B. generatedKeyName
C. datasource
D. dataSourceConnection
E. source

Correct Answer: A, B, E
Section:
Explanation:
Defining a table projection.
Each table requires three properties:
tableName: The name of the table in Azure Storage.
generatedKeyName: The column name for the key that uniquely identifies this row.
source: The node from the enrichment tree you are sourcing your enrichments from. This node is usually the output of a shaper, but could be the output of any of the skills.
Reference:
https://docs.microsoft.com/en-us/azure/search/knowledge-store-projection-overview

02 - Implement Knowledge Mining Solutions
Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Overview
A company named Wide World Importers is developing an e-commerce platform.
You are working with a solutions architect to design and implement the features of the e-commerce platform. The platform will use microservices and a serverless environment built on Azure.
Wide World Importers has a customer base that includes English, Spanish, and Portuguese speakers.
Existing Environment
Applications
Wide World Importers has an App Service plan that contains the web apps shown in the following table.

Name	Description
Product Management	An app used by employees to create and manage products. The app and the expected inputs from the employees are in English.
Inventory Tracking	An app used by employees to manage inventory when dispatching orders, receiving refunds, and receiving consignments from suppliers.

Azure Resources
You have the following resources:
▪ An Azure Active Directory (Azure AD) tenant
An Azure Active Directory (Azure AD) tenant
- The tenant supports internal authentication.
- All employees belong to a group named AllUsers.
- Senior managers belong to a group named LeadershipTeam.
An Azure Functions resource
- A function app posts to Azure Event Grid when stock levels of a product change between OK, Low Stock, and Out of Stock. The function app uses the Azure Cosmos DB change feed. An Azure Cosmos DB account
- The account uses the Core (SQL) API.
- The account stores data for the Product Management app and the Inventory Tracking app.
▪ An Azure Storage account
An Azure Storage account
- The account contains blob containers for assets related to products.
- The assets include images, videos, and PDFs.
An Azure Cognitive Services resource named wwics
A Video Indexer resource named wwivi
Requirements
Business Goals
Wide World Importers wants to leverage AI technologies to differentiate itself from its competitors.
Planned Changes
Wide World Importers plans to start the following projects:
A product creation project: Help employees create accessible and multilingual product entries, while expediting product entry creation. A smart e-commerce project: Implement an Azure Cognitive Search solution to display products for customers to browse. A shopping on-the-go project: Build a chatbot that can be integrated into smart speakers to support customers.
Business Requirements
Wide World Importers identifies the following business requirements for all the projects:
Provide a multilingual customer experience that supports English, Spanish, and Portuguese.
Whenever possible, scale based on transaction volumes to ensure consistent performance.
Minimize costs.
Governance and Security Requirements
Wide World Importers identifies the following governance and security requirements:
Data storage and processing must occur in datacenters located in the United States.
Azure Cognitive Services must be inaccessible directly from the internet.
Accessibility Requirements
Wide World Importers identifies the following accessibility requirements:
All images must have relevant alt text.
All videos must have transcripts that are associated to the video and included in product descriptions. Product descriptions, transcripts, and alt text must be available in English, Spanish, and Portuguese.
Product Creation Requirements

Wide World Importers identifies the following requirements for improving the Product Management app:
Minimize how long it takes for employees to create products and add assets.
Remove the need for manual translations.
Smart E-Commerce Requirements
Wide World Importers identifies the following requirements for the smart e-commerce project:
Ensure that the Cognitive Search solution meets a Service Level Agreement (SLA) of 99.9% availability for searches and index writes. Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products. Support autocompletion and autosuggestion based on all product name variants.
Store all raw insight data that was generated, so the data can be processed later.
Update the stock level field in the product index immediately upon changes.
Update the product index hourly.
Shopping On-the-Go Requirements
Wide World Importers identifies the following requirements for the shopping on-the-go chatbot:
Answer common questions.
Support interactions in English, Spanish, and Portuguese.
Replace an existing FAQ process so that all Q&A is managed from a central location.
Provide all employees with the ability to edit Q&As. Only senior managers must be able to publish updates. Support purchases by providing information about relevant products to customers. Product displays must include images and warnings when stock levels are low or out of stock.
Product JSON Sample
You have the following JSON sample for a product.

```
{
    "sku": "b1",
    "name": {
        "en":"Bicycle",
        "es": "Bicicleta",
        "pt": "Bicicleta"
    },
    "stocklevel": "Out of Stock",
    "description": {
        "en":"Bicycle",
        "es": "Bicicleta",
        "pt": "Bicicleta"
    },
    "image":
{"uri": "https://upload.worldwideimporters.org/bicycle.jpg",
            "alttext": {
                "en":"Bicycle",
                "es": "Bicicleta",
                "pt": "Bicicleta"
        }
       },
    "createdUtc": "2020-02-14T06:08:39Z",
    "language": "en"

}
```

QUESTION 1
You are developing the smart e-commerce project.
You need to implement autocompletion as part of the Cognitive Search solution.
Which three actions should you perform? Each correct answer presents part of the solution. (Choose three.) NOTE: Each correct selection is worth one point.

A. Make API queries to the autocomplete endpoint and include suqqesterName in the body.
B. Add a suggesterthat has the three product name fields as source fields.
C. Make API queries to the search endpoint and include the product name fields in the searchFieids query parameter.
D. Add a suggester for each of the three product name fields.

E. Set the searchAnalyzer property for the three product name variants.
F. Set the analyzer property for the three product name variants.

Correct Answer: A, B, F
Section:
Explanation:
Scenario: Support autocompletion and autosuggestion based on all product name variants.
A: Call a suggester-enabled query, in the form of a Suggestion request or Autocomplete request, using an API. API usage is illustrated in the following call to the Autocomplete REST API.
POST /indexes/myxboxgames/docs/autocomplete?search&api-version=2020-06-30
{
"search": "minecraf",
"suggesterName": "sg"
}
B: In Azure Cognitive Search, typeahead or "search-as-you-type" is enabled through a suggester. A suggester provides a list of fields that undergo additional tokenization, generating prefix sequences to support matches on partial terms. For example, a suggester that includes a City field with a value for "Seattle" will have prefix combinations of "sea", "seat", "seatt", and "seattl" to support typeahead.
F. Use the default standard Lucene analyzer ("analyzer": null) or a language analyzer (for example, "analyzer": "en.Microsoft") on the field.
Reference:
https://docs.microsoft.com/en-us/azure/search/index-add-suggesters

Exam A

QUESTION 1
You have a SQL query that combines customer data and order dat a. The query includes calculated columns. You need to create a database object that would allow other users to rerun the same SOL query. What should you create?

A. an Index
B. a view
C. a scalar function
D. a table

Correct Answer: B
Section:

QUESTION 2
What ate two uses of data visualization? Each correct answer presents a complete solution. NOTE:
Each correct selection is worth one point.

A. Communicate the significance of data.
B. Represent trends and patterns over time.
C. Implement machine learning to predict future values.
D. Enforce business logic across reports.

Correct Answer: A, B
Section:

QUESTION 3
What should you use to build a Microsoft Power Bi paginated report?

A. Power BI Report Builder
B. Charciculator
C. Power BI Desktop
D. the Power BI service

Correct Answer: A
Section:

QUESTION 4
Which scenario is an example of a streaming workload?

A. sending transactions daily from point of sale (POS) devices
B. sending cloud infrastructure metadata every 30 minutes
C. sending transactions that are older than a month to an archive
D. sending telemetry data from edge devices

Correct Answer: D
Section:

QUESTION 5
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:

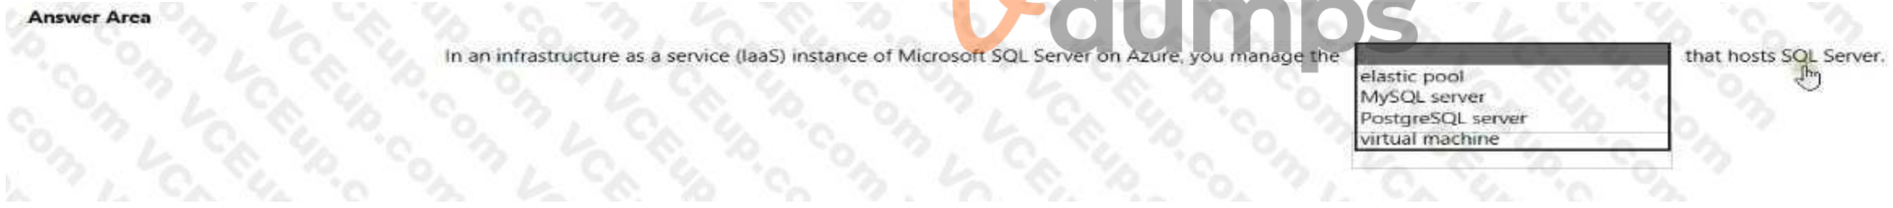

Answer Area:

Section:
Explanation:

QUESTION 6
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A relational database is appropriate for scenarios that involve a high volume of

changes to relationships between entities.
geographically distributed writes.
transactional writes.
writes that have varying data structures.

Answer Area:

Answer Area

A relational database is appropriate for scenarios that involve a high volume of

changes to relationships between entities.
geographically distributed writes.
transactional writes.
writes that have varying data structures.

Section:
Explanation:

QUESTION 7
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A JSON document is an example of

graph data.
relational data.
semi-structured data.
unstructured data.

Answer Area:

Answer Area

A JSON document is an example of

graph data.
relational data.
semi-structured data.
unstructured data.

Section:
Explanation:

QUESTION 8
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

Relational data is stored in

a file system as unstructured data.
a hierarchal folder structure.
a tabular form of rows and columns.
comma-separated value (CSV) files.

Answer Area:

Answer Area

Relational data is stored in

- a file system as unstructured data.
- a hierarchal folder structure.
- a tabular form of rows and columns.
- comma-separated value (CSV) files.

Section:
Explanation:

QUESTION 9
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a

- graph.
- table.
- partition key.
- document.

Answer Area:

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a

- graph.
- table.
- partition key.
- document.

Section:
Explanation:

QUESTION 10
DRAG DROP
Match the Azure Cosmos DB APIs to the appropriate data structures.
To answer, drag the appropriate API from the column on the left to its data structure on the right. Each API may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Select and Place:

APIs

- Cassandra API
- Gremlin API
- MongoDB API
- Table API

Answer Area

API	Graph data
API	JSON documents
API	Key/value data

Correct Answer:

Section:
Explanation:

QUESTION 11
HOTSPOT
Select the answer that correctly completes the sentence.

Hot Area:

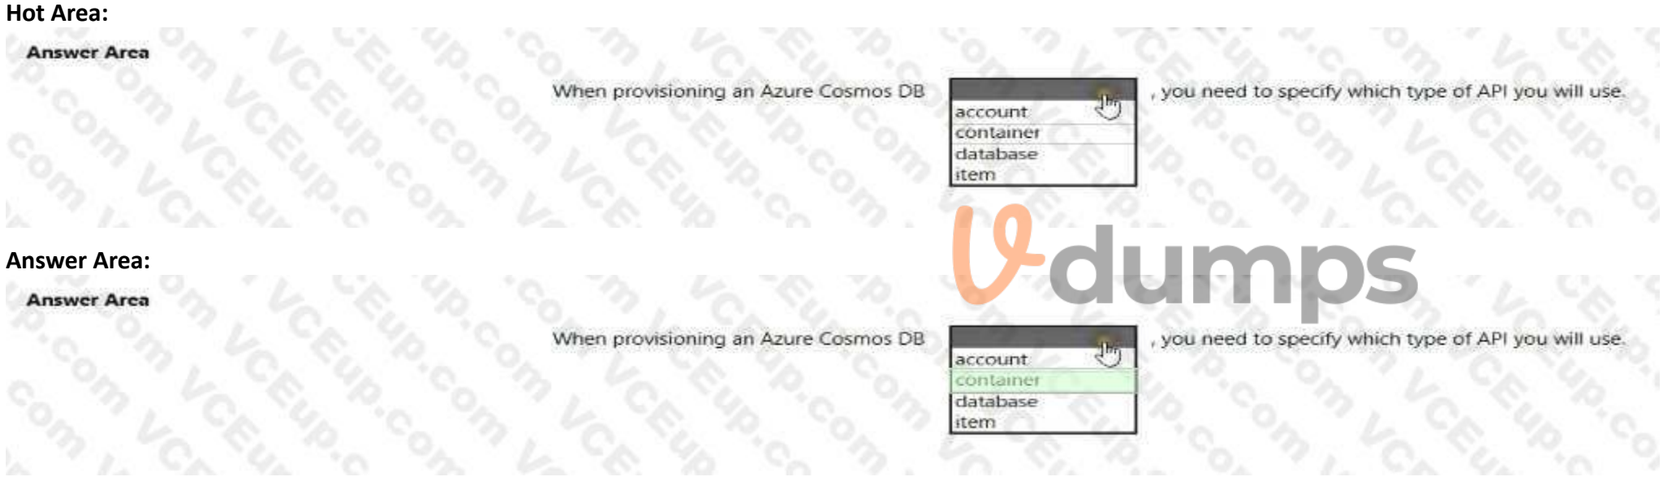

Answer Area:

Section:
Explanation:

QUESTION 12
Which statement is an example of Data Manipulation Language (DML)?

A. Revoke
B. UPDATE
C. DROP
D. CREATE

Correct Answer: B
Section:

QUESTION 13
Your company needs to implement a relational database in Azure. The solution must minimize ongoing maintenance. Which Azure service should you use?

A. SQL Server on Azure Virtual Machines
B. Azure SOL Database
C. Azure HDInsight
D. Azure Cosmos DB

Correct Answer: B
Section:

QUESTION 14
What are two benefits of platform as a service (PaaS) relational database offerings in Azure, such as Azure SQL Database? Each correct answer presents a complete solution. NOTE: Each correct selection Is worth one point.

A. reduced administrative effort for managing the server infrastructure
B. complete control over backup and restore processes
C. in-database machine learning services S3
D. access to the latest features

Correct Answer: A, D
Section:

QUESTION 15
You have data saved in the following format.

```
FirstName,LastName,Age,LeisureHobby,SportsHobby
John,Smith,23,Reading,Basketball
Ben,Smith,21,Guitar,Curling
```

Which format was used?

A. CSV
B. JSON
C. HTML
D. YAML

Correct Answer: A
Section:

QUESTION 16
What is a primary characteristic of a relational database?

A. data is queried and manipulated by using a variant of the SOL language
B. a lack of dependencies between tables
C. a flexible data structure
D. a large amount of duplicate data

Correct Answer: C
Section:

QUESTION 17

What is the primary purpose of a data warehouse?

A. to provide storage for transactional line-of-business (LOB) applications
B. to provide transformation services between source and target data stores
C. to provide read only storage of relational and non relational historical data
D. to provide answers to complex queries that rely on data from multiple sources

Correct Answer: C
Section:

QUESTION 18
You need to develop a solution to provide data to executives. The solution must provide an interactive graphical interface, depict various key performance indicators, and support data exploration by using drill down. What should you use in Microsoft Power Bl?

A. a report
B. Microsoft Power Apps
C. a view
D. a dataflow

Correct Answer: C
Section:

QUESTION 19
Your company has a repotting solution that has paginated reports. The reports query a dimensional model in a data warehouse. Which type of processing does the reporting solution use?

A. Online Transaction Processing (OLTP)
B. Online Analytical Processing (OLAP)
C. batch processing
D. stream processing

Correct Answer: B
Section:

QUESTION 20
You manage an application that stores data in a shared folder on a Windows server. You need to move the shared folder to Azure Storage. Which type of Azure Storage should you use?

A. table
B. queue
C. file
D. blob

Correct Answer: C
Section:

QUESTION 21
You need to recommend a non-relational data store that is optimized for storing and retrieving text files, videos, audio streams, and virtual disk images. The data store must store data, some metadata, and a unique ID for each file. Which type of data store should you recommend?

A. columnar
B. key/value
C. document
D. object

Correct Answer: D
Section:

QUESTION 22
You need to store event log data that is semi-structured and received as the logs occur. What should you use?

A. Azure Table storage
B. Azure Queue storage
C. Azure Files

Correct Answer: A
Section:

QUESTION 23
What should you use to automatically delete blobs from Azure Blob Storage?

A. the change feed
B. a lifecycle management policy
C. soft delete
D. archive storage

Correct Answer: D
Section:

QUESTION 24
What is a characteristic of a non-relational database?

A. full support for Transact-SGL
B. a fixed schema
C. self describing entities

Correct Answer: C
Section:

QUESTION 25
You are building a retail kiosk system that will use a custom neural voice. You acquire audio samples and consent from the voice talent. You need to create a voice talent profile. What should you upload to the profile?

A. a five-minute wav or mp3 file of the voce talent describing the kiosk system
B. a five-minute wav or mp3 file of the voce talent describing the kiosk system
C. a five-minute .flac audio file and the associated transcript as a w file
D. a .wav or mp3 file of the voice talent consenting to the creation of a synthetic version of their voice

E. a.zip file that contains 10-second .wav files and the associated transcripts as .txt files

Correct Answer: D
Section:

QUESTION 26
HOTSPOT
You are building an app that will enable users to upload images. The solution must meet the following requirements:
• Automatically suggest alt text for the images.
• Detect inappropriate images and block them.
• Minimize development effort.
You need to recommend a computer vision endpoint for each requirement.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

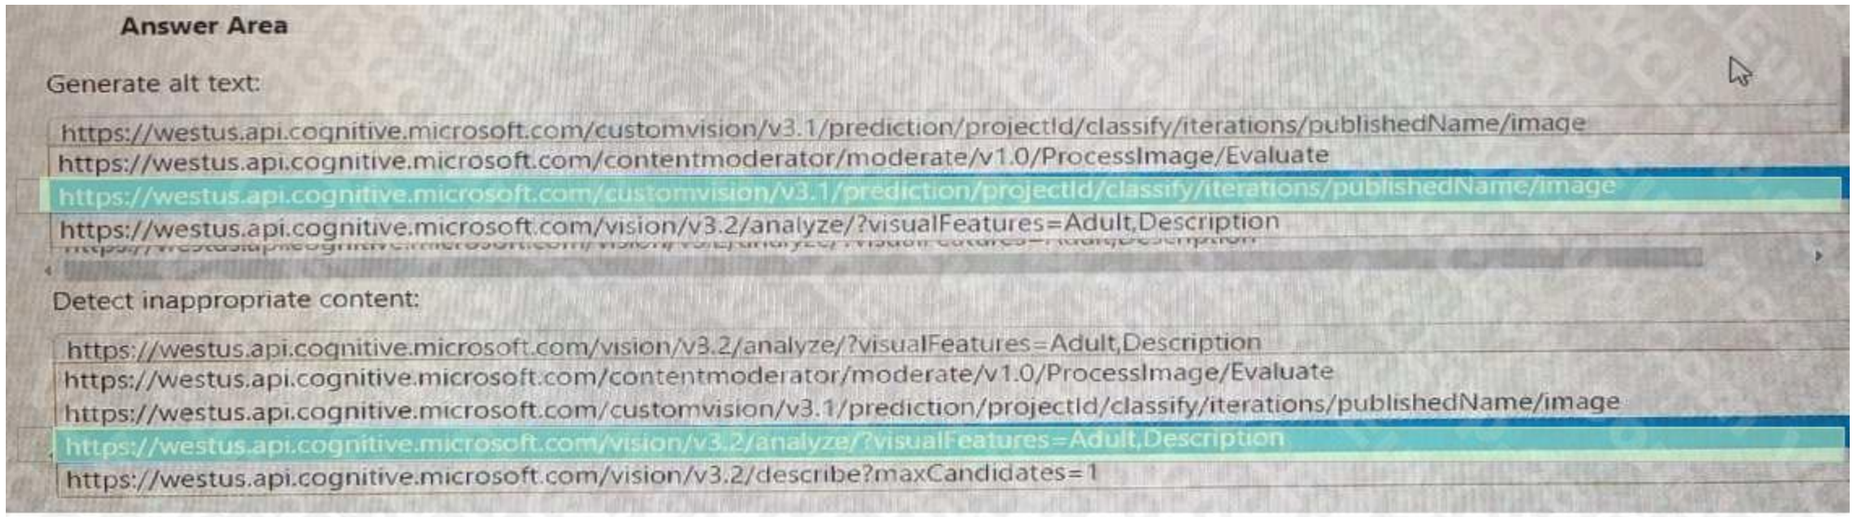

Section:
Explanation:

QUESTION 27
HOTSPOT
You are building content for a video training solution.
You need to create narration to accompany the video content. The solution must use Custom Neural Voice. What should you use to create a custom neural voice, and which service should you use to generate the narration?
To answer, select the appropriate options in the answer area. NOTE: Each correct answer is worth one point.

Hot Area:

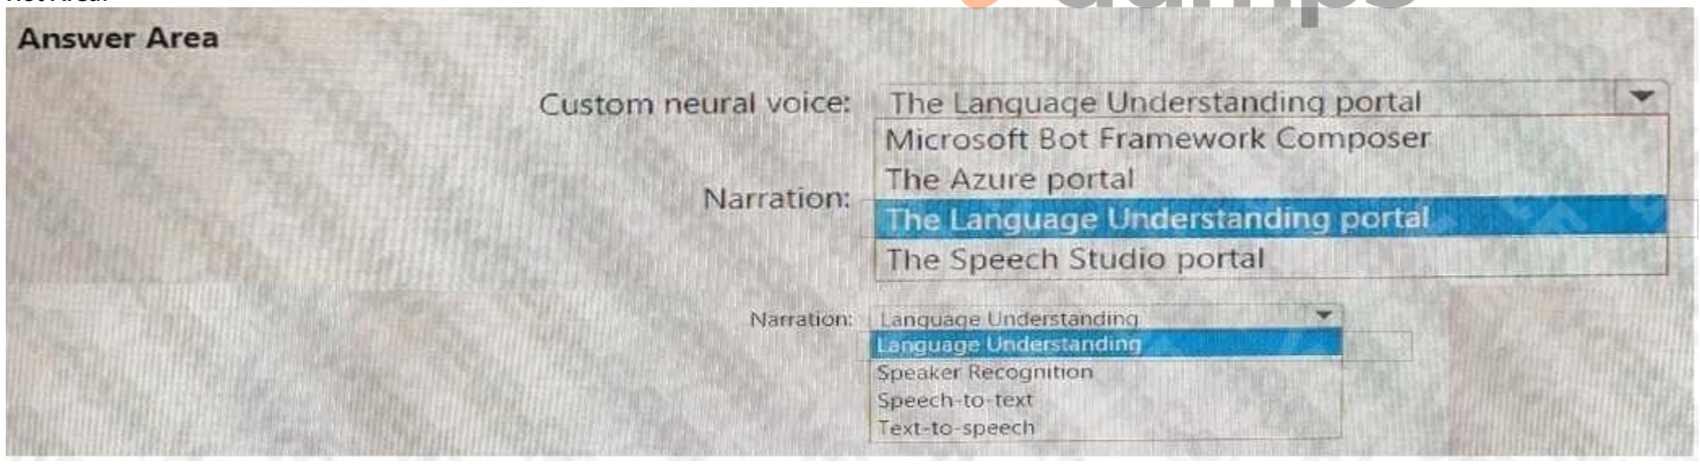

Answer Area:

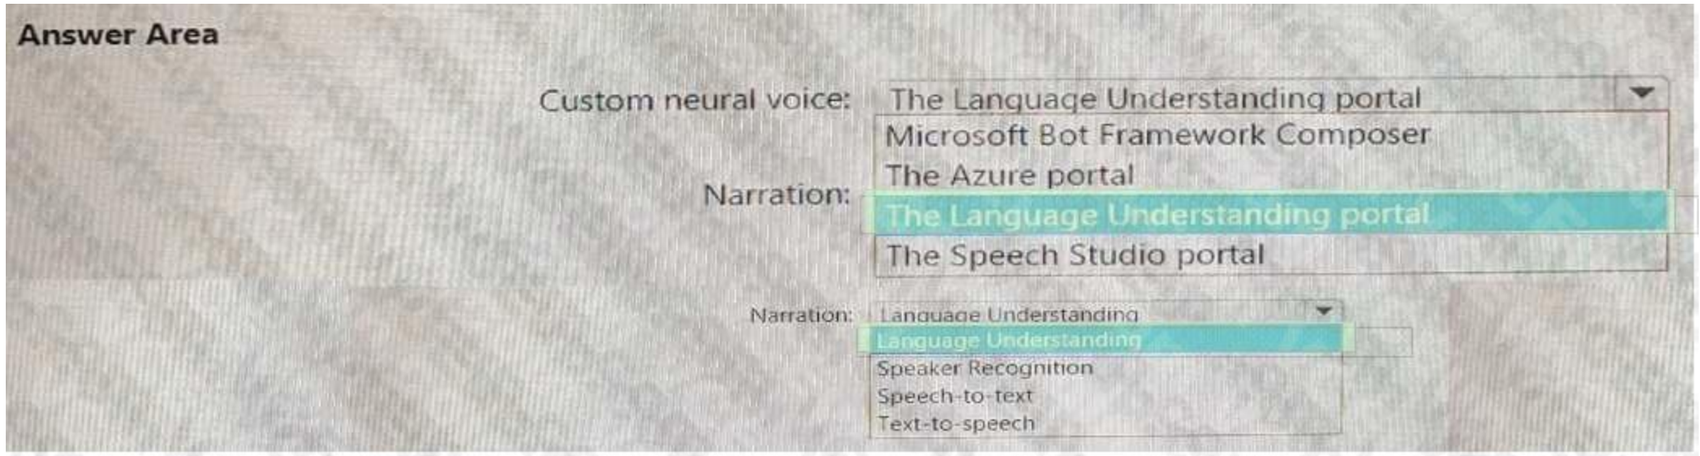

Section:
Explanation:

QUESTION 28
You need to create a visualization of running sales totals per quarter as shown in the following exhibit.

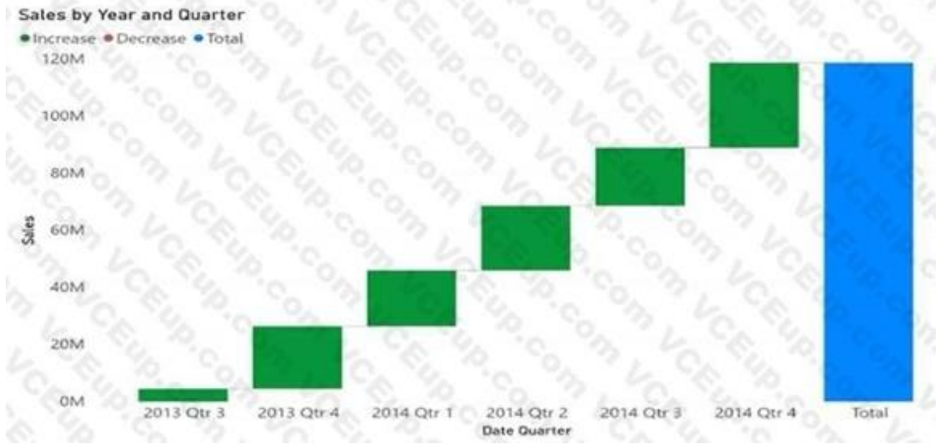

What should you create in Cower Bl Desktop;1

A. a waterfall chart
B. a ribbon chart
C. a bar chart
D. a decomposition tree

Correct Answer: C
Section:

QUESTION 29
Which Azure Storage service implements the key/value model?

A. Azure Files
B. Azure Blob

C. Azure Table
D. Azure Queue

Correct Answer: C
Section:

QUESTION 30
You have an app named App1 that uses an Azure Cognitive Services model to identify anomalies in a time series data stream. You need to run App1 in a location that has limited connectivity. The solution must minimize costs.
What should you use to host the model?

A. Azure Kubernetes Services (AKS)
B. a Kubernetes cluster hosted in an Azure Stack Hub integrated system
C. Azure Container instances
D. the Docker Engine

Correct Answer: B
Section:

QUESTION 31
You have an Azure loT hub that receives series data from machinery. You need to build an app that will perform the following actions:
• Perform anomaly detection across multiple correlated sensors
• Identify the root cause of process stops.
• Send incident alerts
The solution must minimize development time. Which Azure service should you use?

A. Azure Metrics Advisor
B. Form Recognizer
C. Azure Machine teaming
D. Anomaly Detector

Correct Answer: D
Section:

QUESTION 32
You build a language model by using Conversational Language Understanding. The language model is used to search for information on a contact list by using an intent named Findcontact. A conversational expert provides you with the following list of phrases to use for training
• Find contacts in London.
• Who do I know in Seattle?
• Search for contacts m Ukraine.
You need to implement the phrase list in Conversational Language Understanding.
Solution: You create a new utterance for each phrase in the FindContact intent.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:

QUESTION 33
You develop a Conversational Language Undemanding model by using Language Studio
During testing, users receive incorrect responses to requests that do NOT relate to the capabilities of the model. You need to ensure that the model identifies spurious requests.
What should you do?

A. Enable active learning.
B. Add examples to the custom intents.
C. Add examples to the None intent
D. Add entities.

Correct Answer: A
Section:

QUESTION 34
You have an Azure Cognitive Search instance that indexes purchase orders by using Form Recognizer
You need to analyze the extracted information by using Microsoft Power Bl. The solution must minimize development effort.
What should you add to the indexer?

A. a table projection
B. a projection group
C. an object projection
D. a file projection

Correct Answer: C
Section:

QUESTION 35
You develop a custom question answering project in Azure Cognitive Service for Language. The project will be used by a chatbot. You need to configure the project to engage in multi-turn conversations. What should you do?

A. Add follow-up prompts.
B. Enable active learning.
C. Add alternate questions.
D. Enable chit-chat.

Correct Answer: A
Section:

QUESTION 36
You need to measure the public perception of your brand on social media by using natural language processing. Which Azure service should you use?

A. Content Moderator
B. Form Recognizer
C. Computer Vision
D. Language service

Correct Answer: D
Section:

QUESTION 37
DRAG DROP
You plan to build a chatbot to support task tracking.
You create a Conversational Language Understanding service named Iu1.
You need to build a Conversational Language Understanding model to Integrate into the chatbot. The solution must minimize development time to build the model. 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.

Select and Place:

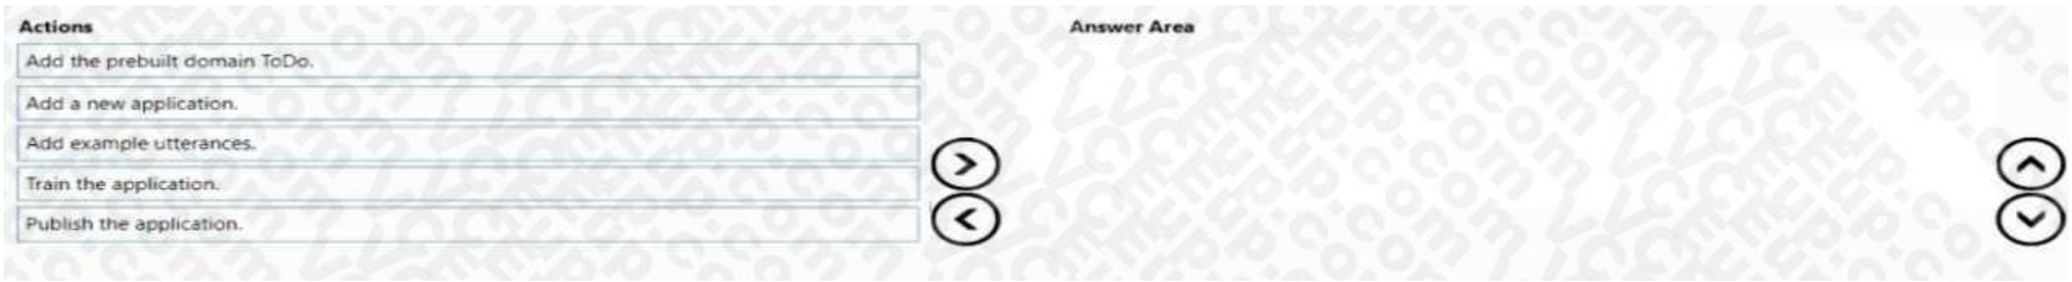

Correct Answer:

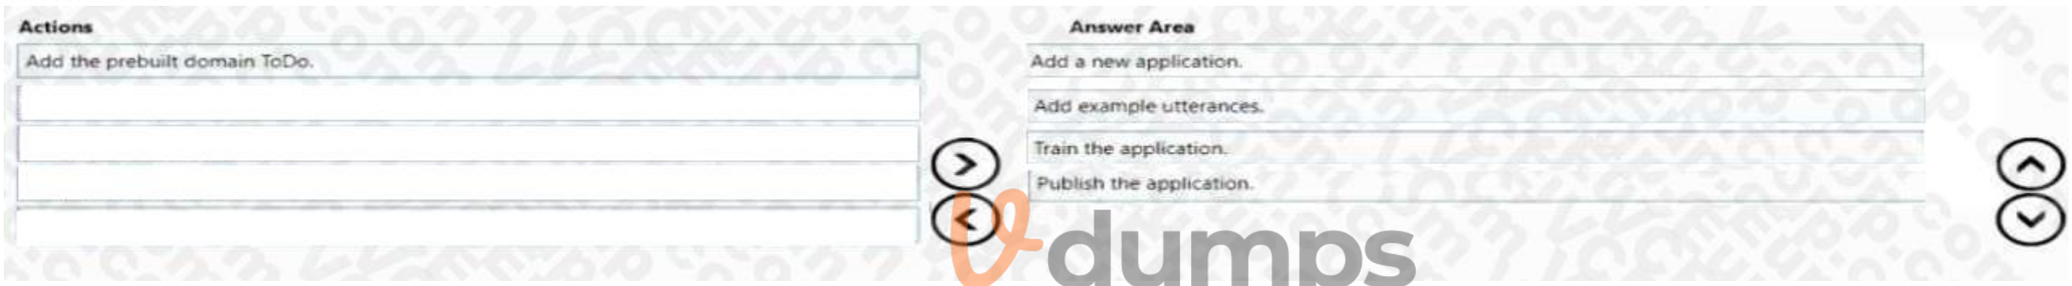

Section:
Explanation:

QUESTION 38
DRAG DROP
You build a bot by using the Microsoft Bot Framework SDK.
You need to test the bot interactively on a local machine.
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 cor'ect 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:

Section:
Explanation:

QUESTION 39
DRAG DROP
You have a Custom Vision service project that performs object detection. The project uses the General domain for classification and contains a trained model. You need to export the model for use on a network that is disconnected from the internet.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list o' actions to the answer area and arrange them in the correct order.

Select and Place:

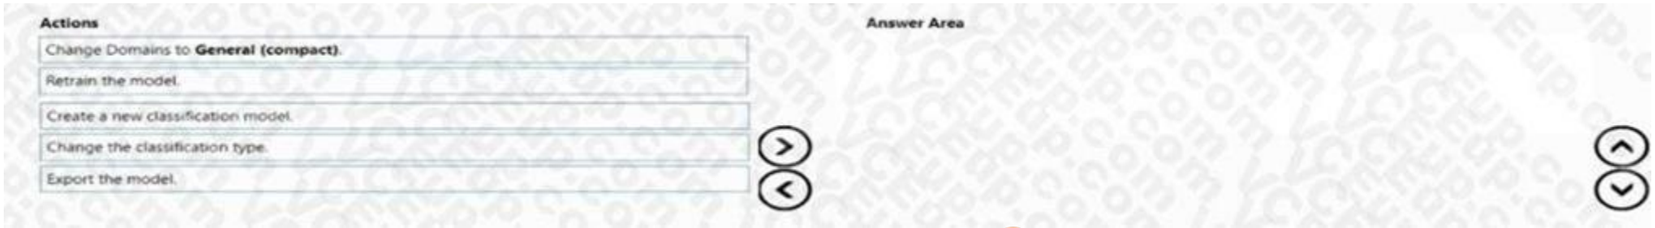

Correct Answer:

Section:
Explanation:

QUESTION 40
DRAG DROP
You have a question answering project in Azure Cognitive Service for Language.
You need to move the project to a Language service instance in a different Azure region.
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:

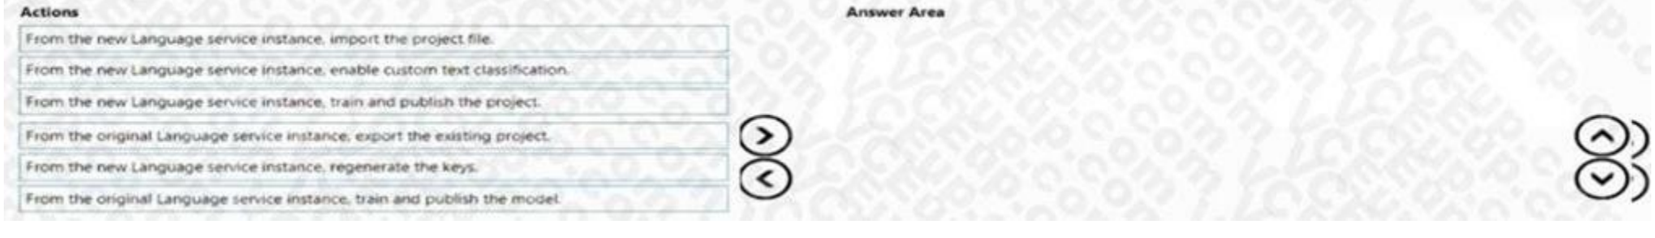

Correct Answer:

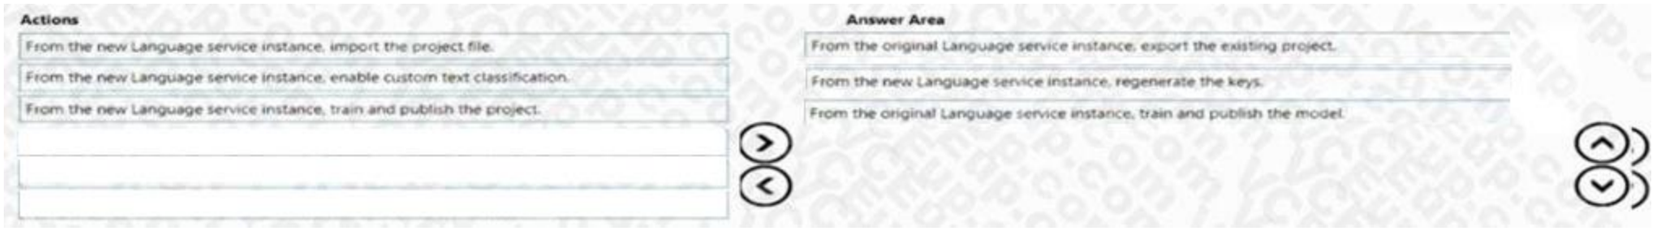

Section:
Explanation:

QUESTION 41
DRAG DROP
You are building an app that will scan confidential documents and use the Language service to analyze the contents. You provision an Azure Cognitive Services resource.
You need to ensure that the app can make requests to the Language service endpoint. The solution must ensure that confidential documents remain on-premises. 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:

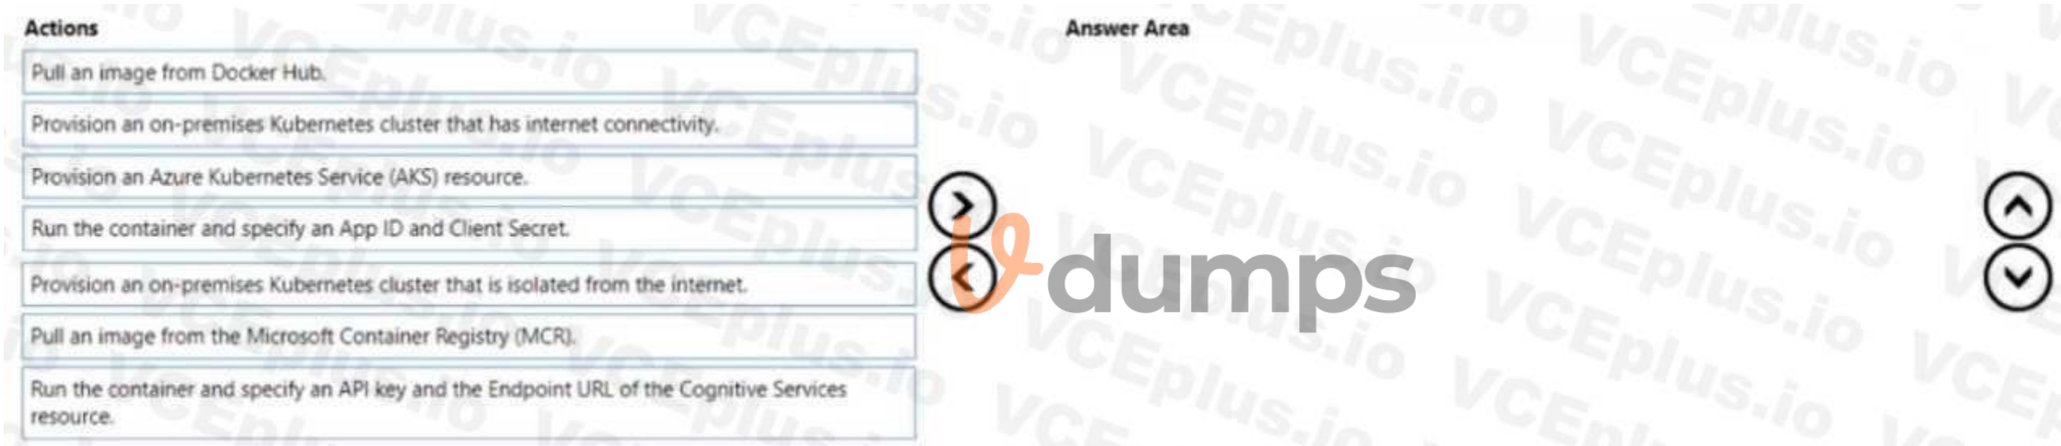

Correct Answer:

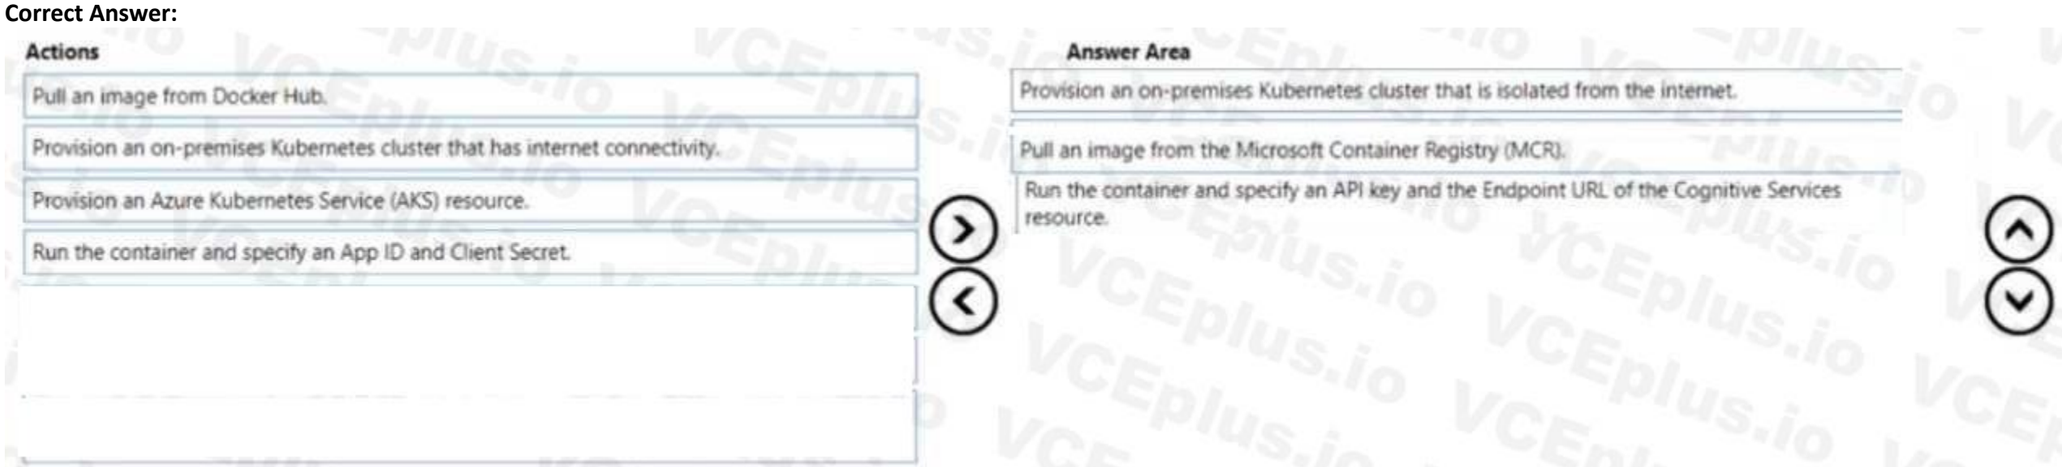

Section:
Explanation:
Provision an on-premises Kubernetes...
Pull an image from the Microsoft Container Registry...

Run the container and specify an API key...

QUESTION 42
You have an Azure subscription that contains an Azure Cognitive Service for Language resource. You need to identify the URL of the REST interface for the Language service. Which blade should you use in the Azure portal?

A. Identity
B. Keys and Endpoint
C. Properties
D. Networking

Correct Answer: B
Section:

QUESTION 43
You are building a chatbot by using Microsoft Bot Framework Composer.
You need to configure the chatbot to present a list of available options. The solution must ensure that an image is provided for each option. Which two features should you use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. an Azure function
B. an adaptive card
C. an entity
D. a dialog
E. an utterance

Correct Answer: B, D
Section:

QUESTION 44
You are building a chatbot.
You need to configure the chatbot to query a knowledge base.
Which dialog class should you use?

A. AdaptiveDialog
B. QnAMakerDialog
C. ComponentDialog
D. SkillDialog

Correct Answer: B
Section:

QUESTION 45
DRAG DROP
You need to analyze video content to identify any mentions of specific company names.
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:

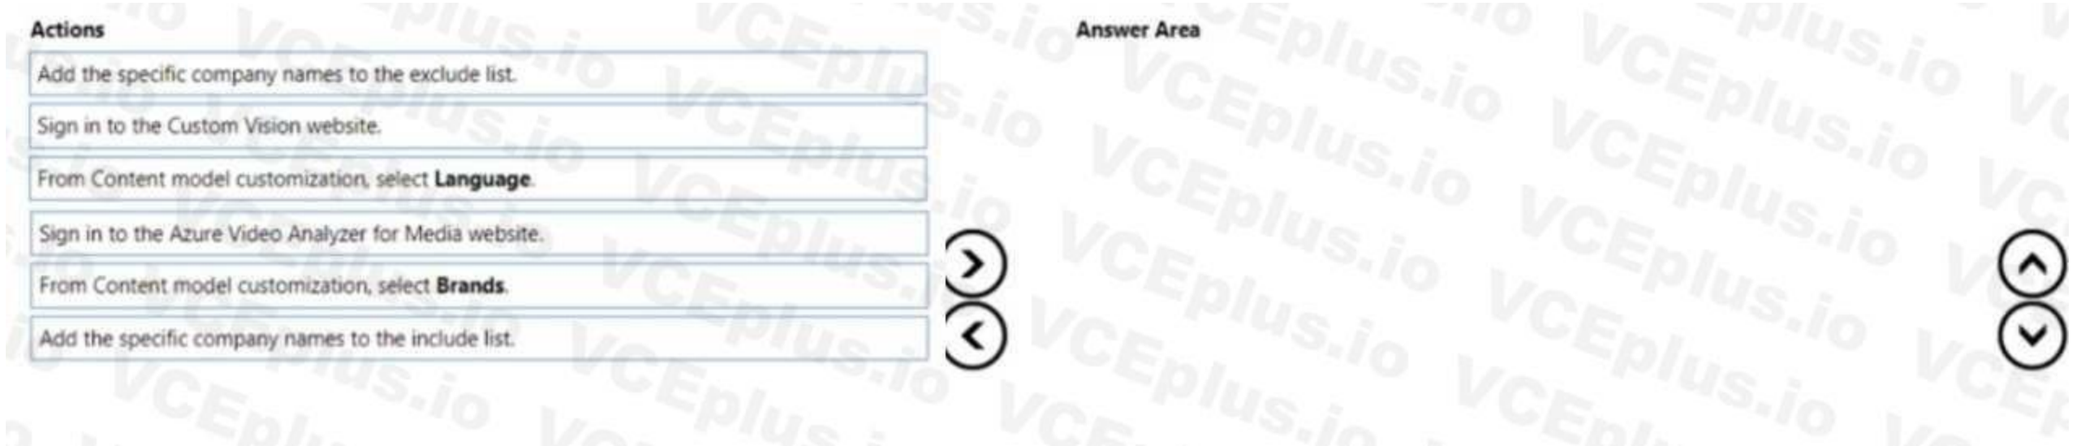

Correct Answer:

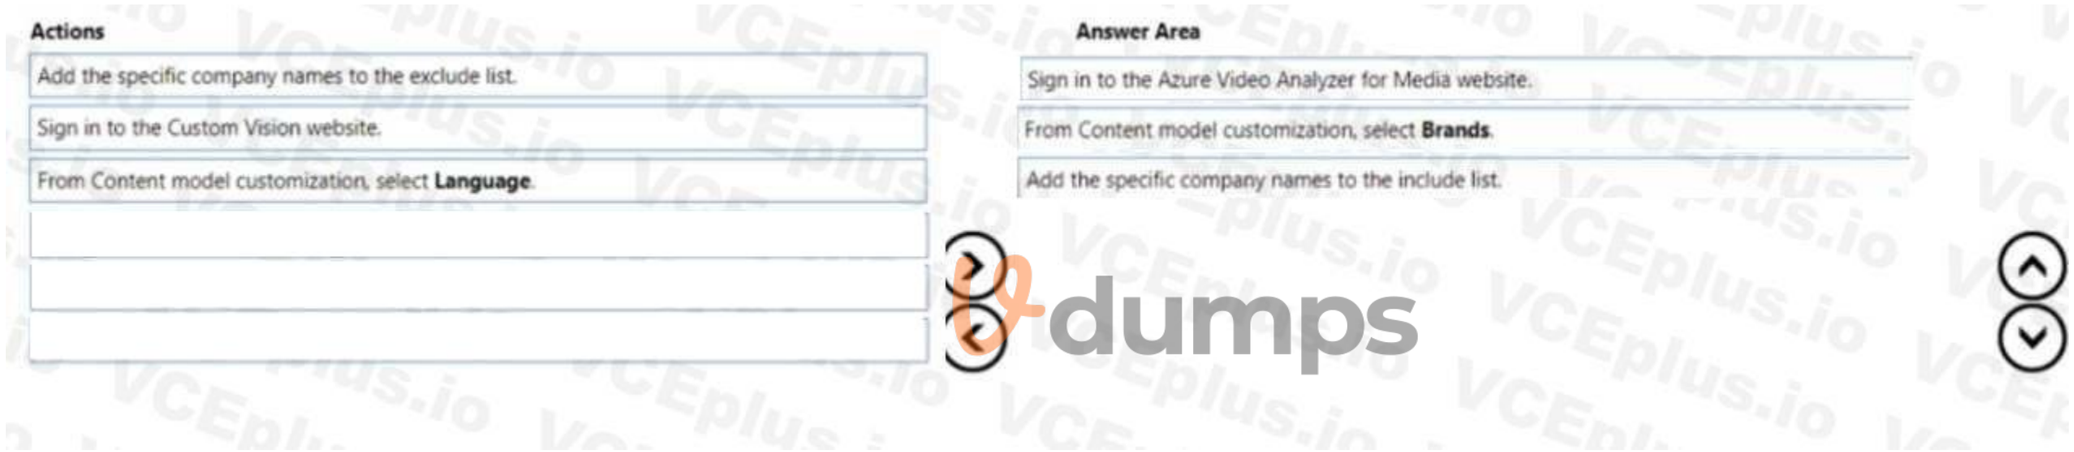

Section:
Explanation:

QUESTION 46
You plan to build an app that will generate a list of tags for uploaded images. The app must meet the following requirements:
• Generate tags in a users preferred language.
• Support English, French, and Spanish.
• Minimize development effort
You need to build a function that will generate the tags for the app. Which Azure service endpoint should you use?

A. Custom Vision image classification
B. Content Moderator Image Moderation
C. Custom Translator
D. Computer Vision image Analysis

Correct Answer: D
Section:

QUESTION 47
HOTSPOT

You have an Azure Cognitive Search resource named Search 1 that is used by multiple apps You need to secure Search 1. The solution must meet the following requirements:
• Prevent access to Search1 from the internet.
• Limit the access of each app to specific queries
What should you do? To answer, select the appropriate options in the answer area NOTE Each correct answer is worth one point.

Hot Area:

Answer Area

Answer Area:

Answer Area

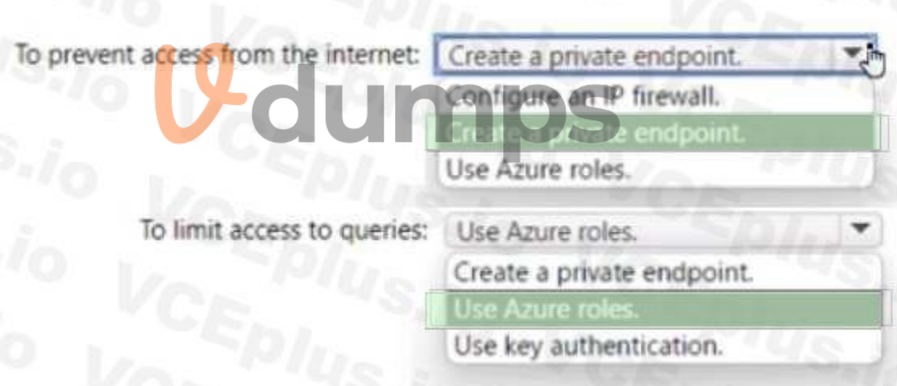

Section:
Explanation:

QUESTION 48
You need to measure the public perception of your brand on social media by using natural language processing. Which Azure service should you use?

A. Content Moderator
B. Form Recognizer
C. Computer Vision
D. Language service

Correct Answer: D
Section:

QUESTION 49
You are examining the Language service output of an application.
The text analyzed is: Our tour guide took us up the Space Needle during our trip to Seattle last week. The response contains the data shown in the following table.

Text	Category	ConfidenceScore
Tour guide	PersonType	0.45
Space Needle	Location	0.38
Trip	Event	0.78
Seattle	Location	0.78
Last week	DateTime	0.80

Which Language service API is used 10 analyze the Text?

A. Entity Linking
B. Named Entity Recognition
C. Key Phrase Extraction
D. Sentiment Analysis

Correct Answer: B
Section:

QUESTION 50
You are building a Conversational Language Understanding model.
You need to ensure that the model will support the following sample utterances:
- Set all the lights to on.
- Turn off the lights in the living room.
- What is the current thermostat temperature?
- Lower the temperature of the thermostat by five degrees.
Which three elements should you add to the model?
Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. a location Intent
B. a change setting entity
C. a device intent
D. a change setting intent
E. a query setting intent
F. a device entity

Correct Answer: B, C, F
Section:

QUESTION 51
HOTSPOT
You have a bot that was built by using the Microsoft Bot Framework composer as shown in the following exhibit.

Use the drop-down menus 10 select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

Section:
Explanation:

QUESTION 52
HOTSPOT
You are building a chatbot by using the Microsoft Bot Framework Composer. You have the dialog design shown in the following exhibit.

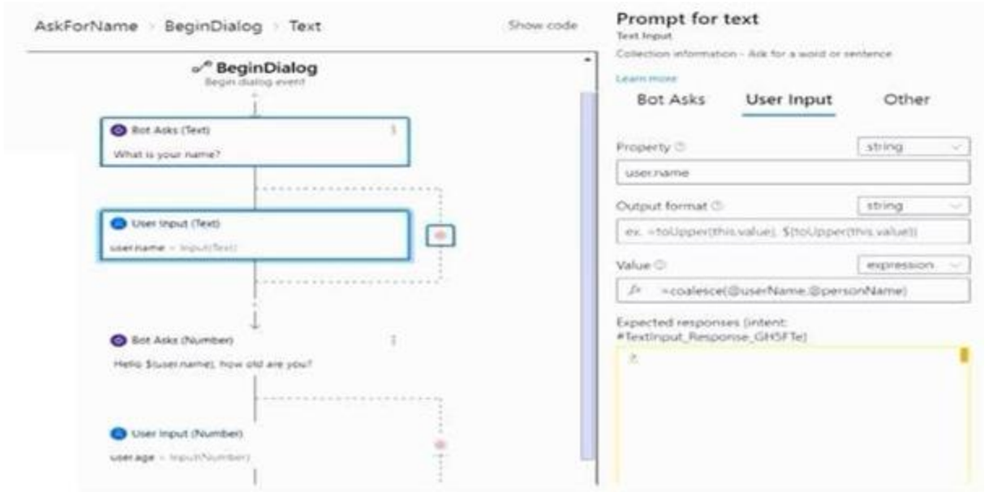

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection Is worth one point.

Hot Area:

Answer Area:

Section:
Explanation:

QUESTION 53
You have a mobile app that manages printed forms.
You need the app to send images of the forms directly to Forms Recognizer to extract relevant information. For compliance reasons, the image files must not be stored in the cloud. In which format should you send the images to the Form Recognizer API endpoint?

A. raw image binary
B. form URL encoded
C. JSON

Correct Answer: A
Section:

QUESTION 54
DRAG DROP
You have a factory that produces cardboard packaging for food products. The factory has intermittent internet connectivity. The packages are required to include four samples of each product.
You need to build a Custom Vision model that will identify defects in packaging and provide the location of the defects to an operator. The model must ensure that each package contains the four products. Which project type and domain should you use? To answer, drag the appropriate options to the correct targets. Each option 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:

Section:
Explanation:

QUESTION 55
DRAG DROP
You are building a customer support chatbot.
You need to configure the bot to identify the following:
• Code names for internal product development
• Messages that include credit card numbers
The solution must minimize development effort.
Which Azure Cognitive Service for Language feature should you use for each requirement? 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.

Select and Place:

Correct Answer:

Section:
Explanation:

QUESTION 56
You are building a chatbot.
You need to configure the bot to guide users through a product setup process.
Which type of dialog should you use?

A. component
B. waterfall
C. adaptive
D. action

Correct Answer: B
Section:

QUESTION 57
You have a Language service resource that performs the following:
• Sentiment analysis
• Named Entity Recognition (NER)
• Personally Identifiable Information (Pll) identification
You need to prevent the resource from persisting input data once the data is analyzed. Which query parameter in the Language service API should you configure?

A. loggingOptOut

B. piiCategories
C. showStats
D. Model-version

Correct Answer: A
Section:

QUESTION 58
You create a bot by using the Microsoft Bot Framework SDK.
You need to configure the bot to respond to events by using custom text responses.
What should you use?

A. an adaptive card
B. an activity handler
C. a dialog
D. a skill

Correct Answer: B
Section:

QUESTION 59
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 have a chatbot that uses question answering in Azure Cognitive Service for Language.
Users report that the responses of the chatbot lack formality when answering random questions that are outside the scope of the knowledge base.
You need to ensure that the chatbot provides formal responses to these spurious questions.
Solution: From Language Studio, you modify the question and answer pairs for the custom intents, and then retrain and republish the model.
Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:

QUESTION 60
DRAG DROP
You have a Docker host named Host1 that contains a container base image.
You have an Azure subscription that contains a custom speech-to-text model named model1.
You need to run model1 on Host1.
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:

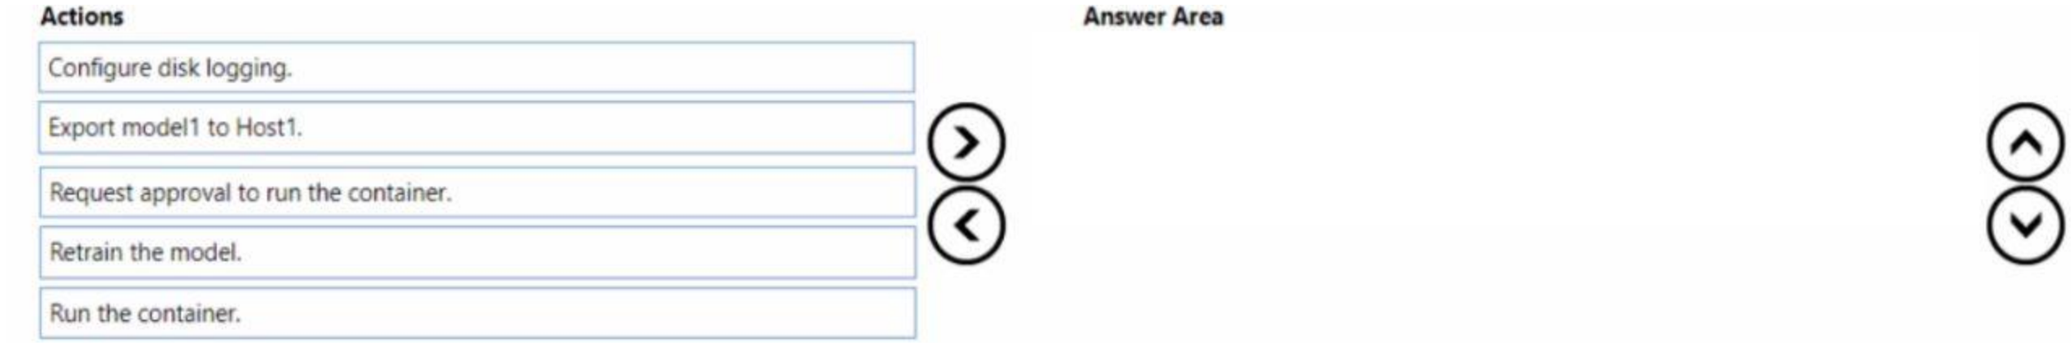

Correct Answer:

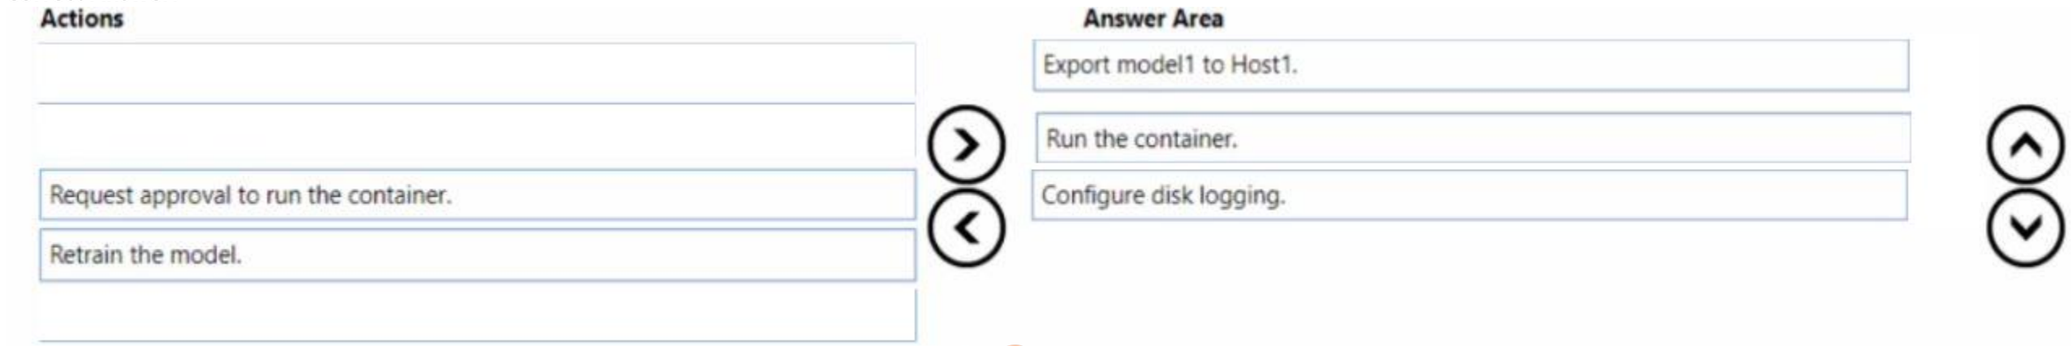

Section:
Explanation:
Export model1 to Host1
Run the container
Configure disk logging

QUESTION 61
DRAG DROP
You are building a transcription service for technical podcasts.
Testing reveals that the service fails to transcribe technical terms accurately.
You need to improve the accuracy of the service.
Which five 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:

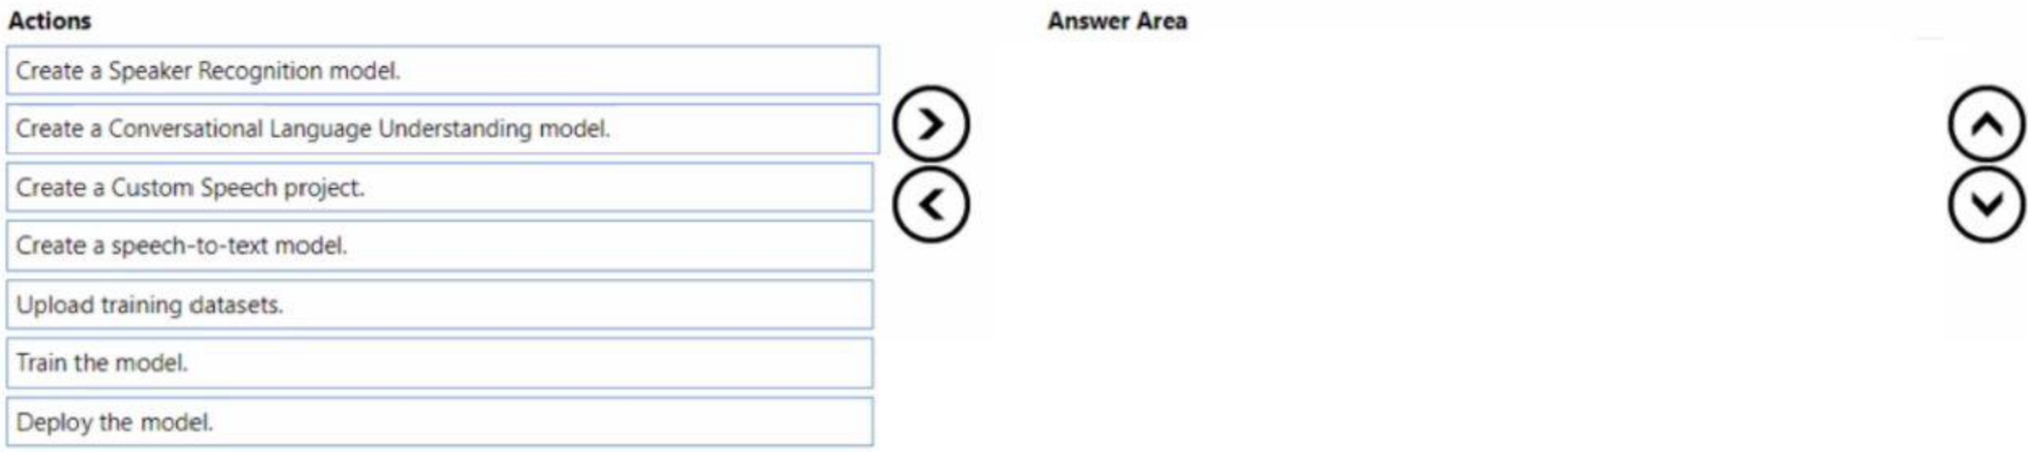

Correct Answer:

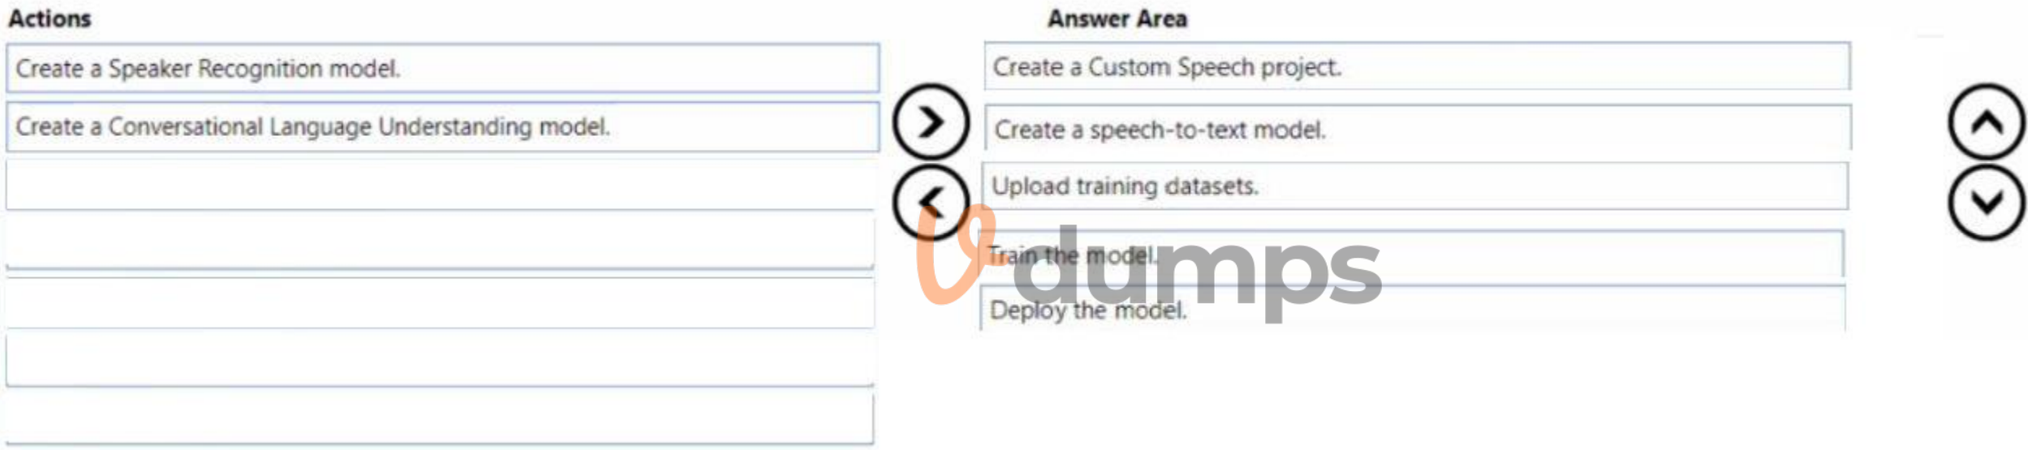

Section:
Explanation:
Create a Custom Speech project.
Create a speech-to-text model.
Upload training datasets.
Train the model.
Deploy the model.

QUESTION 62
You have a file share that contains 5,000 images of scanned invoices.
You need to analyze the images. The solution must extract the following data:
* Invoice items
* Sales amounts
* Customer details
What should you use?

A. Custom Vision

B. Computer Vision
C. Immersive Reader
D. Form Recognizer

Correct Answer: C
Section:

QUESTION 63
You have an Azure Cognitive Search solution and a collection of blog posts that include a category field. You need to index the posts. The solution must meet the following requirements:
* Include the category field in the search results.
* Ensure that users can search for words in the category field.
* Ensure that users can perform drill down filtering based on category.
Which index attributes should you configure for the category field?

A. searchable, facetable, and retrievable
B. retrievable, filterable, and sortable
C. retrievable, facetable, and key
D. searchable, sortable, and retrievable

Correct Answer: B
Section:

QUESTION 64
You are building a social media extension that will convert text to speech. The solution must meet the following requirements:
- Support messages of up to 400 characters.
- Provide users with multiple voice options.
- Minimize costs.
You create an Azure Cognitive Services resource.
Which Speech API endpoint provides users with the available voice options?

A. https://uksouth.customvoice.apispeech.microsoft.eom/api/texttospeech/v3.0/longaudiosynthesis/voices
B. https://uksouth.tts.speech.microsoft.conT/cognitiveservices/voices/list
C. https://uksouth.voice.speech.microsoft.com/cognitiveservices/v1?deploymentId={deploymentId>
D. https://uksouth.api.cognitive.microsoft.eom/speechtotext/v3.0/models/base

Correct Answer: C
Section:

QUESTION 65
You are building a chatbot that will use question answering in Azure Cognitive Service for Language. You have a PDF named Docl.pdf that contains a product catalogue and a price list
You upload Docl.pdf and train the model.
During testing, users report that the chatbot responds correctly to the following QWhat is the price of < product>? The chatbot fails to respond to the following QHow much does <product* cost?
You need to ensure that the chatbot responds correctly to both questions.
Solution: From Language Studio, you create an entity for cost, and then retrain and republish the model. Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section:

QUESTION 66
What is used to define a query in a stream processing jobs in Azure Stream Analytics?

A. SQL
B. XML
C. YAML
D. KOL

Correct Answer: A
Section:

QUESTION 67
Which property of a transactional workload guarantees that each transaction is treated as a single unit that either succeeds completely or tails completely?

A. isolation
B. atomicity
C. consistency
D. durability

Correct Answer: B
Section:

QUESTION 68
Which database transaction property ensures that transactional changes to a database are preserved during unexpected operating system restarts?

A. durability
B. atomicity
C. consistency
D. isolation

Correct Answer: A
Section:

QUESTION 69
Which database transaction property ensures that individual transactions are executed only once and either succeed in their entirety or roll back?

A. consistency
B. isolation
C. atomicity
D. durability

Correct Answer: A
Section:

QUESTION 70
You have an Azure subscription that contains a multi-service Azure Cognitive Services Translator resource named Translator1.
You are building an app that will translate text and documents by using Translator1.
You need to create the REST API request for the app.
Which headers should you include in the request?

A. the subscription key and the client trace ID
B. the subscription key, the subscription region, and the content type
C. the resource ID and the content language
D. the access control request, the content type, and the content length

Correct Answer: B
Section:

QUESTION 71
You are building an app that will use the Azure Video Indexer service.
You plan to train a language model to recognize industry-specific terms.
You need to upload a file that contains the industry-specific terms.
Which file format should you use?

A. PDF
B. XML
C. TXT
D. XLS

Correct Answer: C
Section:

QUESTION 72
HOTSPOT
You are building an Azure web app named App1 that will translate text from English to Spanish.
You need to use the Text Translation REST API to perform the translation. The solution must ensure that you have data sovereignty in the United States.
How should you complete the URI? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

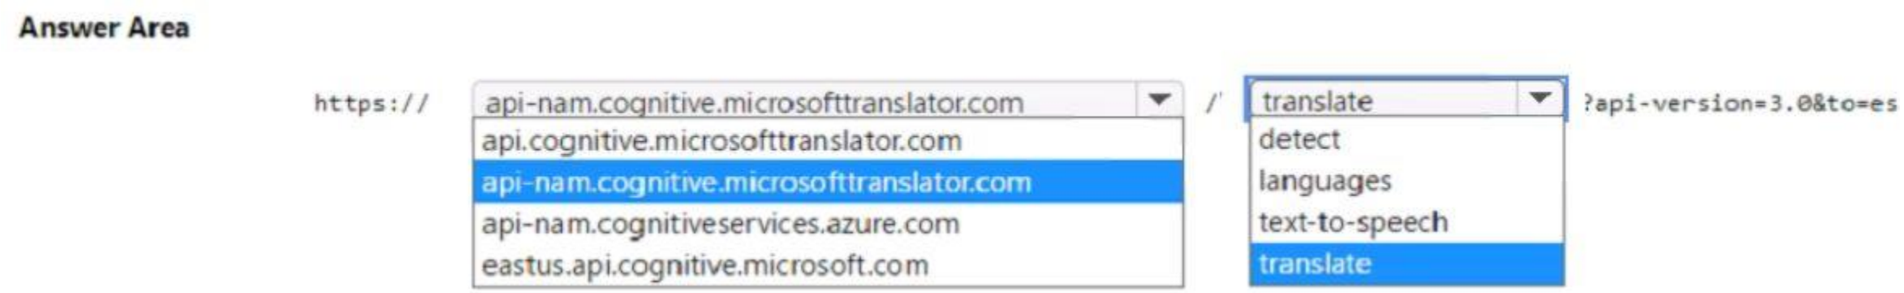

Answer Area:

Answer Area

Section:
Explanation:

QUESTION 73
HOTSPOT
You build a chatbot by using Azure OpenAI Studio.
You need to ensure that the responses are more deterministic and less creative.
Which two parameters should you configure? To answer, select the appropriate parameters in the answer area.
NOTE: Each correct answer is worth one point.

Hot Area:

Answer Area

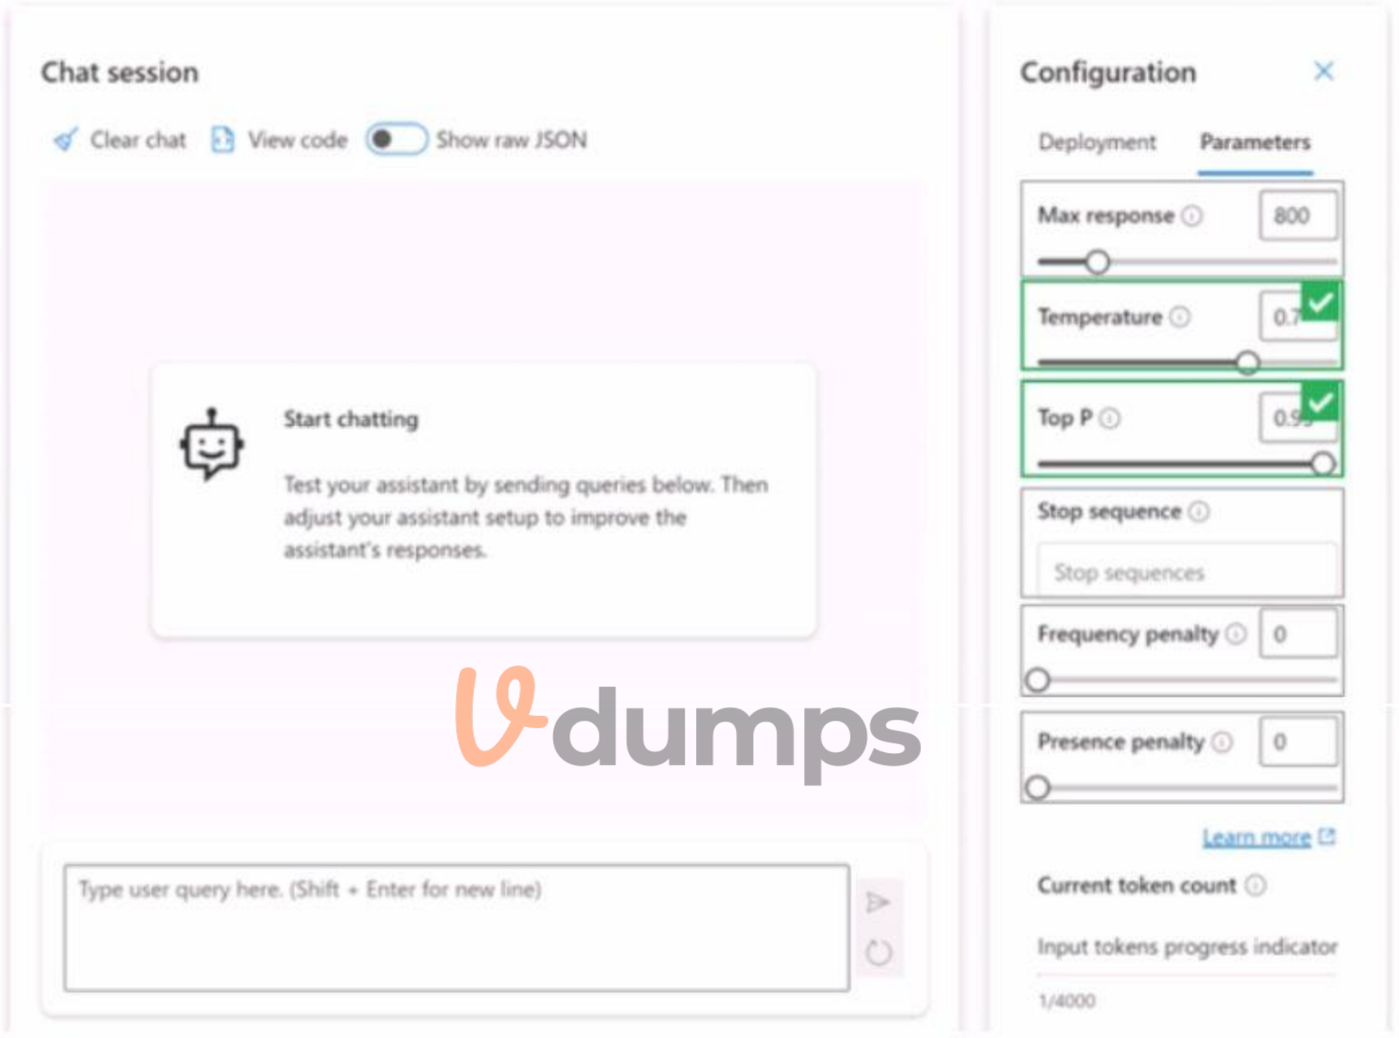

Answer Area:

Answer Area

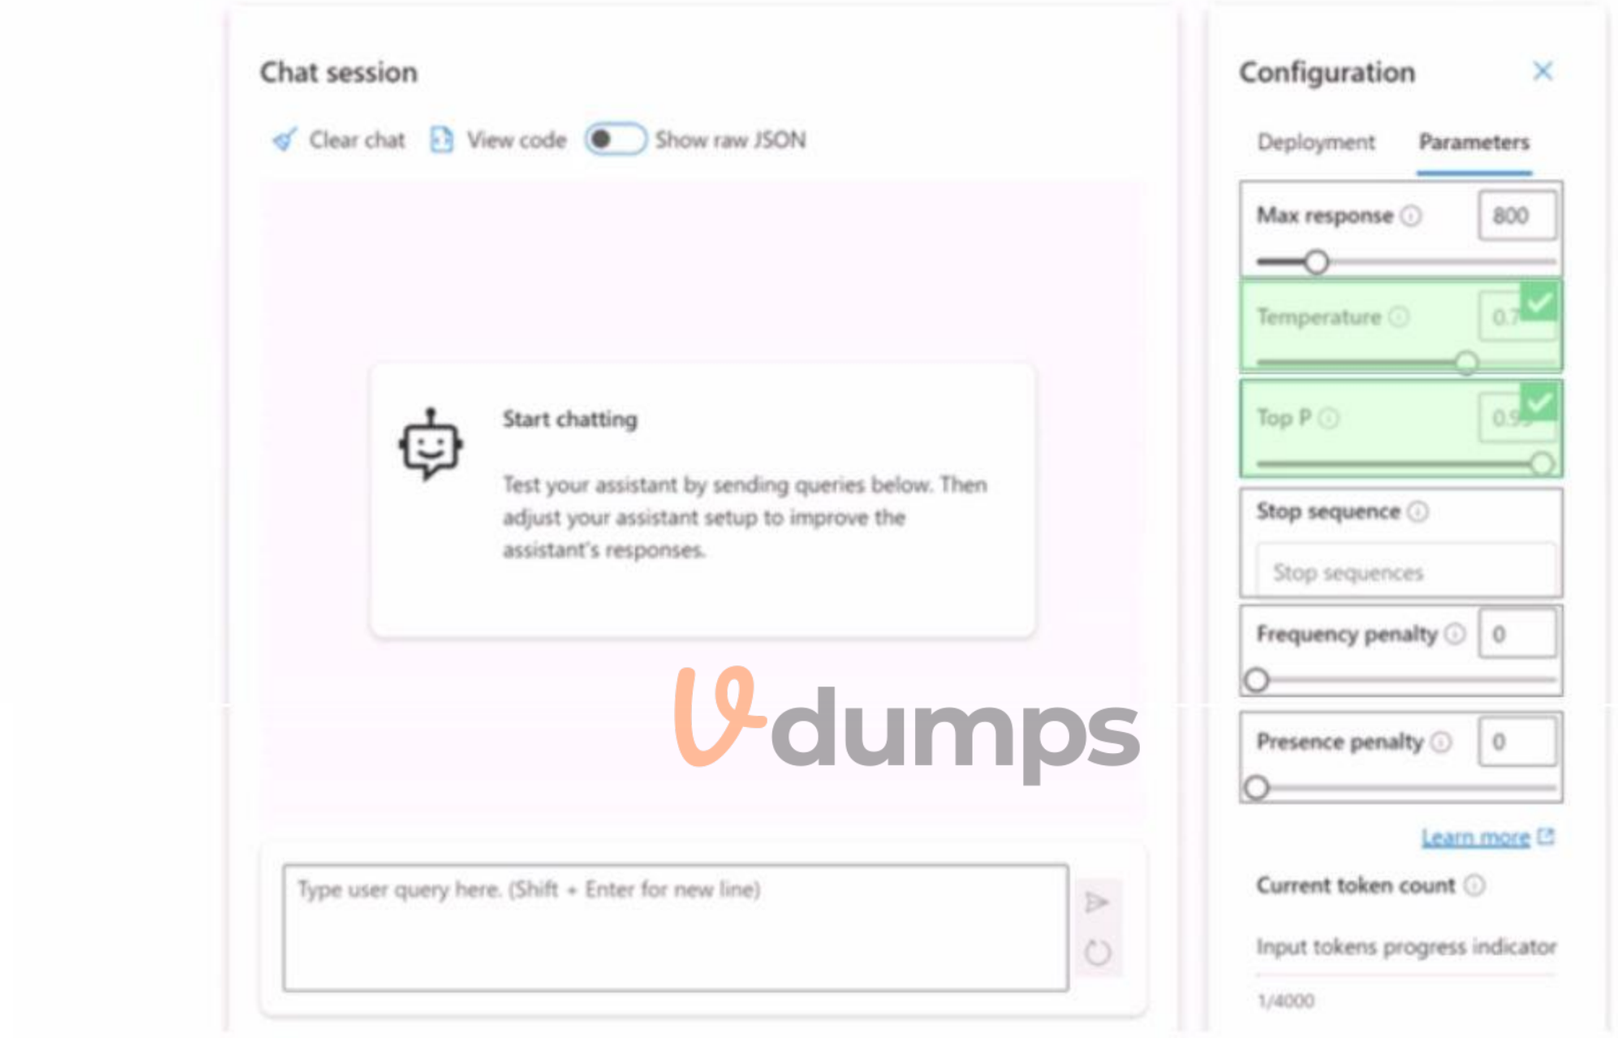

Section:
Explanation:

QUESTION 74
You build a chatbot that uses the Azure OpenAI GPT 3.5 model.
You need to improve the quality of the responses from the chatbot. The solution must minimize development effort.
What are two ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct answer is worth one point.

A. Fine-tune the model.
B. Provide grounding content.
C. Add sample request/response pairs.
D. Retrain the language model by using your own data.

E. Train a custom large language model (LLM).

Correct Answer: B, C
Section:

QUESTION 75
HOTSPOT
You have an Azure subscription that contains an Azure OpenAI resource named AH.
You build a chatbot that will use AI1 to provide generative answers to specific questions.
You need to ensure that the responses are more creative and less deterministic.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

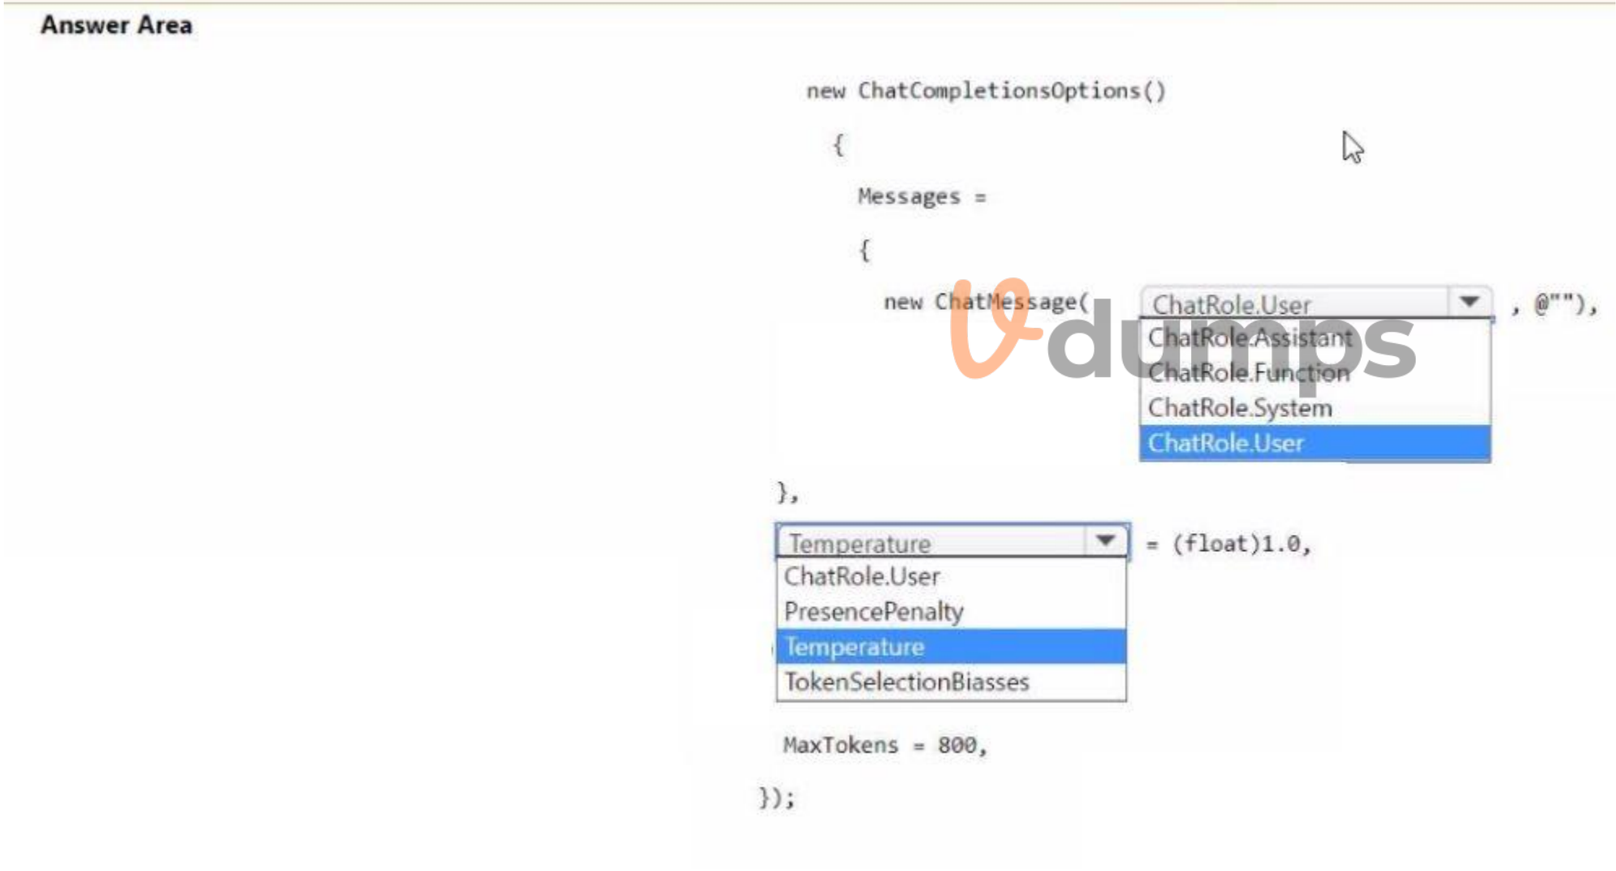

Answer Area:

Answer Area

```
new ChatCompletionsOptions()
  {
    Messages =
    {
      new ChatMessage( [ChatRole.User] , @""),
    },
    [Temperature] = (float)1.0,
    MaxTokens = 800,
});
```

ChatRole.User
- ChatRole.Assistant
- ChatRole.Function
- ChatRole.System
- ChatRole.User

Temperature
- ChatRole.User
- PresencePenalty
- Temperature
- TokenSelectionBiasses

Section:
Explanation:

QUESTION 76
DRAG DROP
You have an Azure subscription that contains an Azure OpenA1 resource named AH.
You plan to build an app named App1 that will write press releases by using AM.
You need to deploy an Azure OpenA1 model for Appl. The solution must minimize development effort.
Which three actions should you perform in sequence in Azure OpenAI Studio? 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:

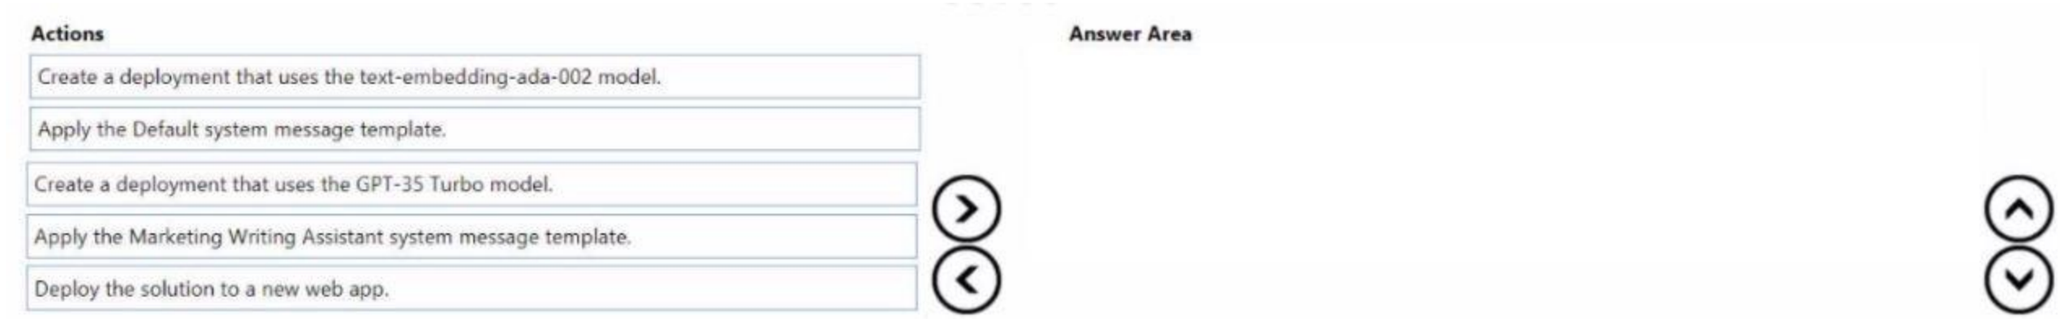

Correct Answer:

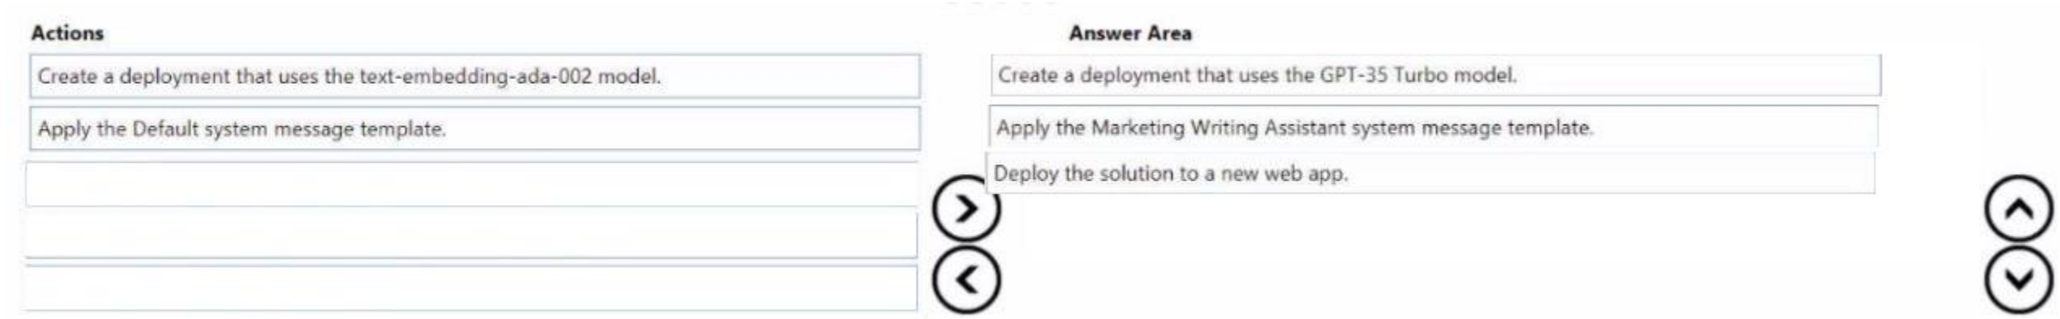

Section:
Explanation:
Create a deployment that uses the GPT-35 Turbo model.
Apply the Marketing Writing Assistant system message template.
Deploy the solution to a new web app.

QUESTION 77
You are building a chatbot for a travel agent. The chatbot will use the Azure OpenAI GPT 3.5 model and will be used to make travel reservations.
You need to maximize the accuracy of the responses from the chatbot.
What should you do?

A. Configure the model to include data from the travel agent's database.
B. Set the Top P parameter for the model to 0.
C. Set the Temperature parameter for the model to 0.
D. Modify the system message used by the model to specify that the answers must be accurate.

Correct Answer: A
Section:

QUESTION 78
HOTSPOT
You are developing a text processing solution.
You have the function shown below.

```
static void GetKeyWords(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.RecognizeEntities (text);
    Console.WriteLine("Key words:");

    foreach (CategorizedEntity entity in response.Value)
    {
        Console.WriteLine($"\t{entity.Text}");
    }
}
```

For the second argument, you call the function and specify the following string.
Our tour of Paris included a visit to the Eiffel Tower.
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
The output will include the following words: our and included.	○	○
The output will include the following words: Paris, Eiffel, and Tower.	○	○
The function will output all the key phrases from the input string to the console.	○	○

Answer Area:

Answer Area

Statements	Yes	No
The output will include the following words: our and included.	○	○
The output will include the following words: Paris, Eiffel, and Tower.	○	○
The function will output all the key phrases from the input string to the console.	○	○

Section:
Explanation:

QUESTION 79
HOTSPOT
A1 You have an Azure OpenA1 resource named AH that hosts three deployments of the GPT 3.5 model. Each deployment is optimized for a unique workload.
You plan to deploy three apps. Each app will access AM by using the REST API and will use the deployment that was optimized for the apps intended workload.
You need to provide each app with access to AH and the appropriate deployment. The solution must ensure that only the apps can access AM.
What should you use to provide access to AM. and what should each app use to connect to its appropriate deployment? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Provide access to AI1 by using: An API key
- An API key
- A bearer token
- A shared access signature (SAS) token

Connect to the deployment by using: A deployment endpoint
- An API key
- A deployment endpoint
- A deployment name
- A deployment type

Answer Area:

Answer Area

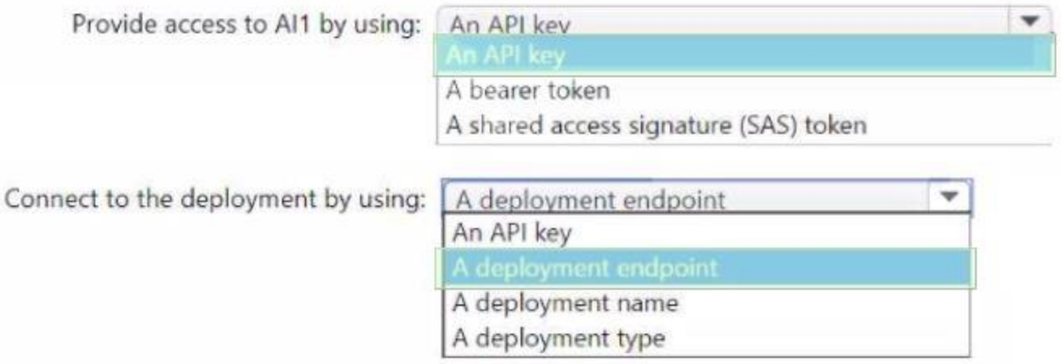

Section:
Explanation:

QUESTION 80
HOTSPOT
You are building a chatbot.
You need to use the Content Moderator API to identify aggressive and sexually explicit language.
Which three settings should you configure? To answer, select the appropriate settings in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Microsoft
Cognitive Services

- Image
- Text

Content Moderator - Moderate

Text - Screen

The operation detects profanity in more than 100 languages and match against custom and shared blacklists.

Host

Name uksouth.api.cognitive.micros

Query parameters

autocorrect	Value	Remove parameter
PII	Value	Remove parameter
listId	Value	Remove parameter
classify	false	Remove parameter
language	Value	Remove parameter

Add parameter

Headers

Content-Type	text/plain	Remove header
Ocp-Apim-Subscription-Key	Value	

Add header

Answer Area:

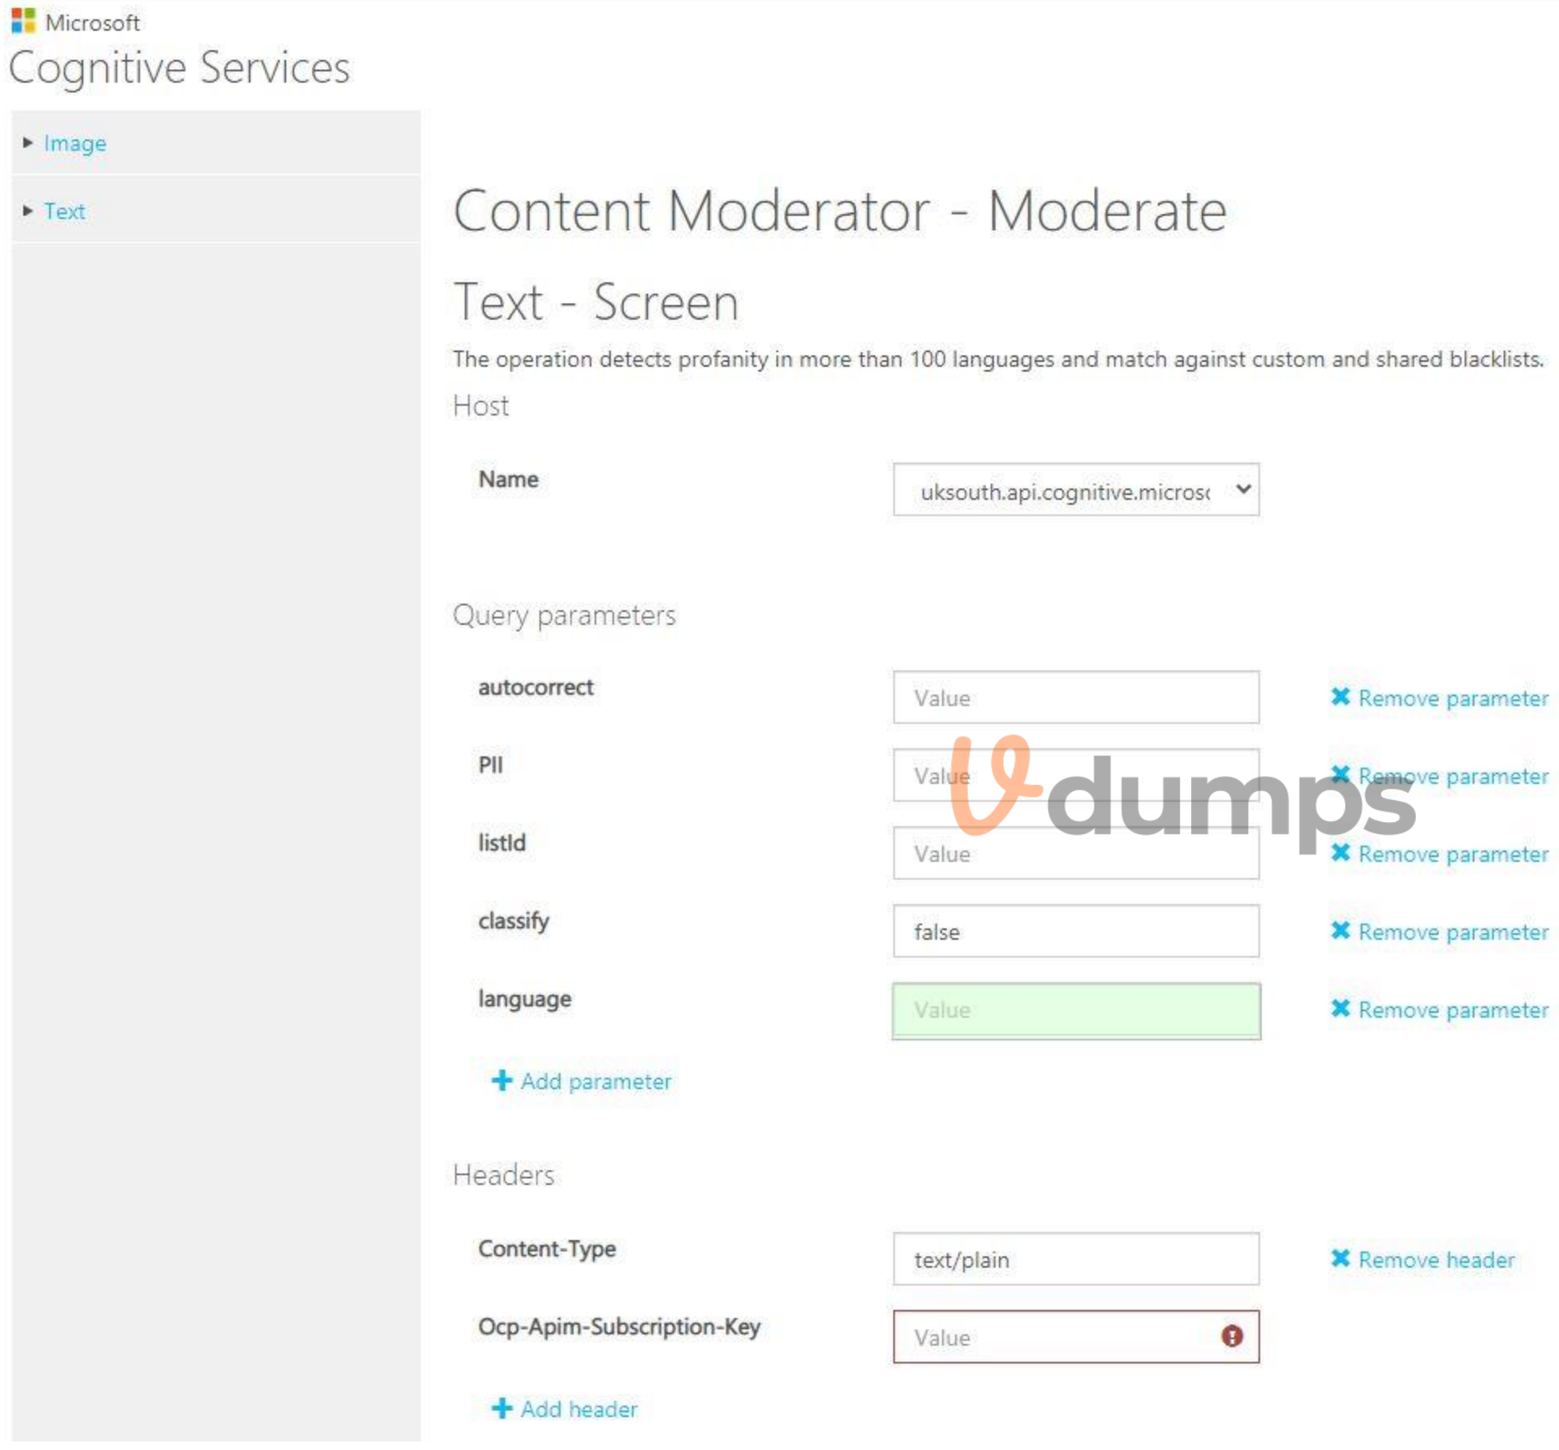

Section:
Explanation:

QUESTION 81

HOTSPOT
You have an app named App1 that uses Azure Al Document Intelligence to analyze medical You have an app named App1 that uses Azure Al Document Intelligence for patients. You send a request to App1 and receive the following response.

```
{
  "status": "succeeded",
  "createdDateTime": "2023-09-14T21:01:02Z",
  "lastUpdatedDateTime": "2023-09-14T21:01:03Z",
  "analyzeResult": {
    "apiVersion": "2023-07-31",
    "modelId": "prebuilt-healthInsuranceCard.us",
    "stringIndexType": "utf16CodeUnit",
    "content": "Blood Pressure 118/72",
    "pages": [
      {
      ...
      "words": [
        {
          "content": "Blood",
          "polygon": [ ... ],
          "confidence": 0.766,
          "span": { ... }
        },
        {
          "content": "Pressure",
          "polygon": [ ... ],
          "confidence": 0.716,
          "span": { ...  }
        },
        {
          "content": "118/72",
          "polygon": [ ... ],
          "confidence": 0.761,
          "span": { ... }
        }
      ],
      ...
  "documents": [
      {
        "docType": "healthInsuranceCard.us",
        "boundingRegions": [ ... ]
```

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

Hot Area:

Answer Area

Statements	Yes	No
The chosen model is suitable for the intended use case.	○	○
The text content was recognized with greater than 70 percent confidence.	○	○
The form elements were recognized with greater than 70 percent confidence.	○	○

Answer Area:

Answer Area

Statements	Yes	No
The chosen model is suitable for the intended use case.	○	○
The text content was recognized with greater than 70 percent confidence.	○	○
The form elements were recognized with greater than 70 percent confidence.	○	○

Section:
Explanation:

QUESTION 82
HOTSPOT
You plan to deploy an Azure OpenAI resource by using an Azure Resource Manager (ARM) template.
You need to ensure that the resource can respond to 600 requests per minute.
How should you complete the template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity" : 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
...
}
```

"capacity" | "capacity" | "count" | "maxValue" | "size"

600 | 1 | 60 | 100 | 600

Answer Area:

Answer Area

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity" : 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
...
}
```

"capacity" dropdown: "capacity", "count", "maxValue", "size"

600 dropdown: 1, 60, 100, 600

Section:
Explanation:

QUESTION 83
You have an app that analyzes images by using the Computer Vision API.
You need to configure the app to provide an output for users who are vision impaired. The solution must provide the output in complete sentences. Which API call should you perform?

A. readInStreamAsync
B. describeImageInStreamAsync
C. toglfiageInStreaaAsync

D. analyreImageByDomainInStreamAsync

Correct Answer: D
Section:

QUESTION 84
You are designing a conversational interface for an app that will be used to make vacation requests. The interface must gather the following data:
• The start date of a vacation
• The end date of a vacation
• The amount of required paid time off
The solution must minimize dialog complexity. Which type of dialog should you use?

A. Skill
B. waterfall
C. adaptive
D. component

Correct Answer: D
Section:

QUESTION 85
You are building an Al solution that will use Sentiment Analysis results from surveys to calculate bonuses for customer service staff. You need to ensure that the solution meets the Microsoft responsible Al principles. What should you do?

A. Add a human review and approval step before making decisions that affect the staffs financial situation
B. Include the Sentiment Analysis results when surveys return a low confidence score.
C. Use all the surveys, including surveys by customers who requested that their account be deleted and their data be removed.
D. Publish the raw survey data to a central location and provide the staff with access to the location.

Correct Answer: A
Section:

QUESTION 86
DRAG DROP
You have a Language Understanding solution that runs in a Docker container.
You download the Language Understanding container image from the Microsoft Container Registry (MCR). You need to deploy the container image to a host computer.
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:

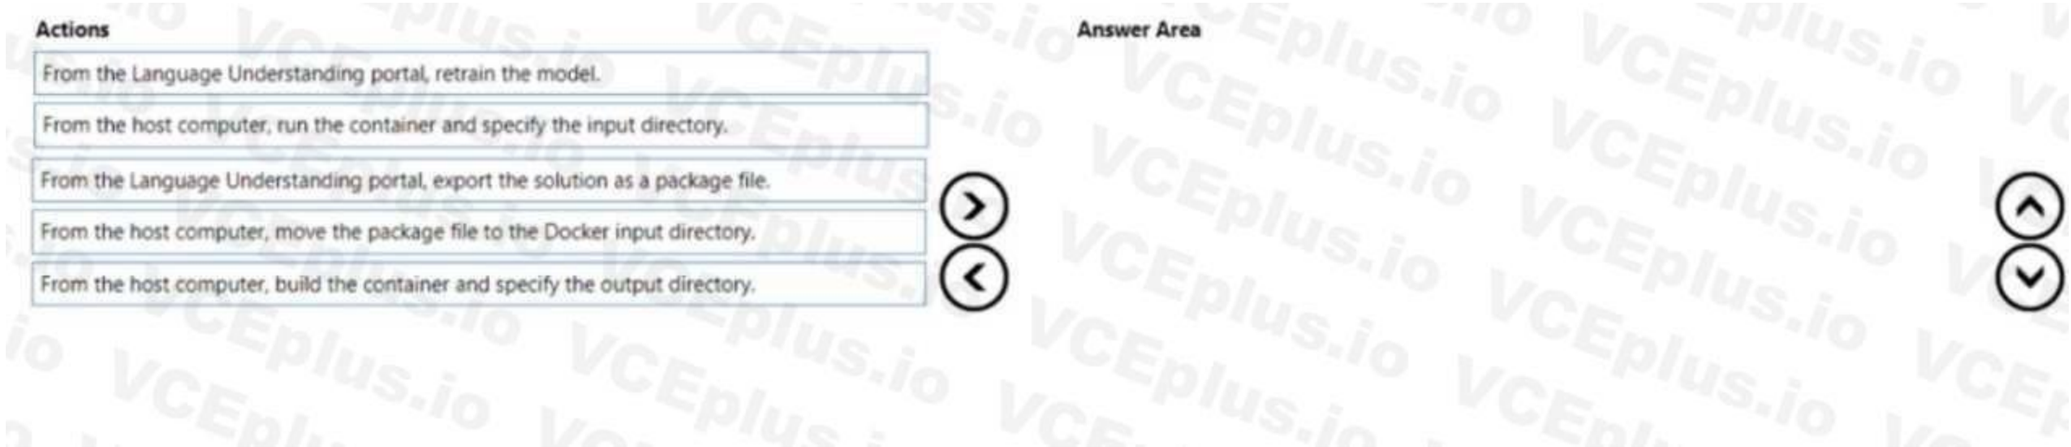

Correct Answer:

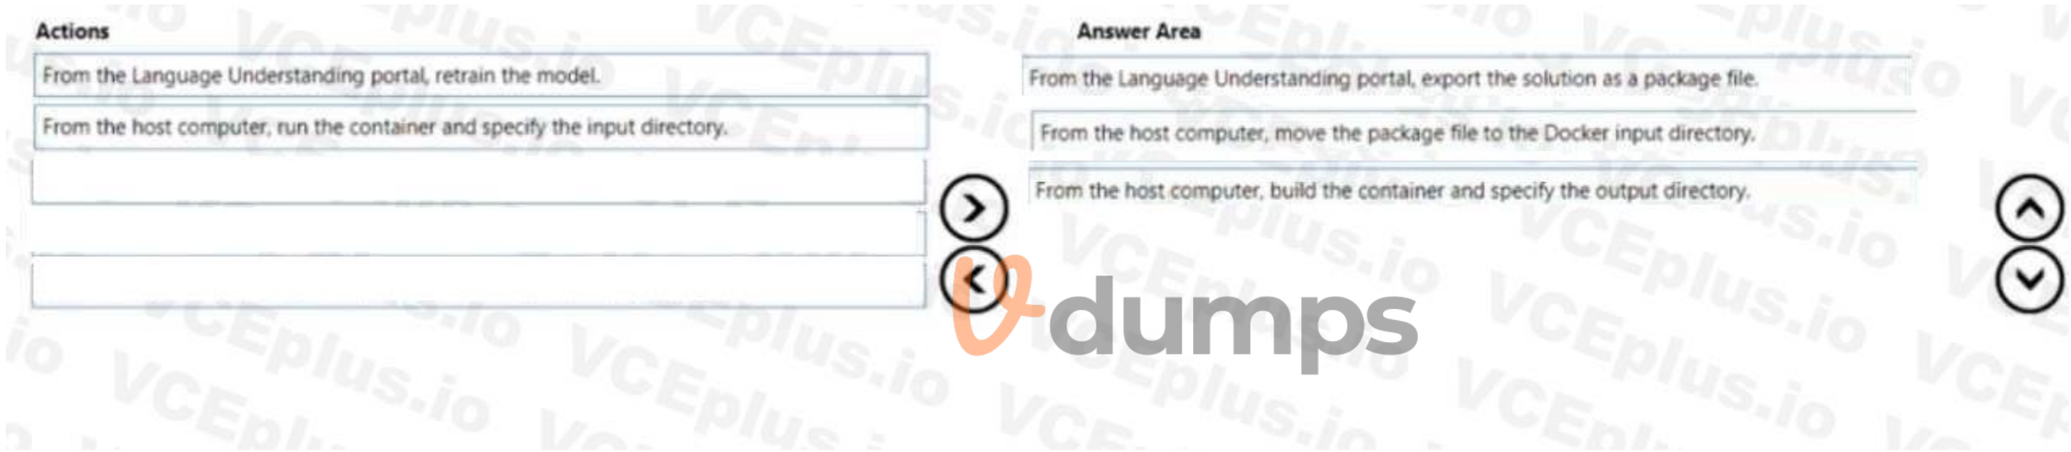

Section:
Explanation:

QUESTION 87
You have an Azure subscription. The subscription contains an Azure OpenAI resource that hosts a GPT-4 model named Modell and an app named Appl. App1 uses Model!
You need to ensure that App1 will NOT return answers that include hate speech.
What should you configure for Model1?

A. the Frequency penalty parameter
B. abuse monitoring
C. a content filter
D. the Temperature parameter

Correct Answer: B
Section:

QUESTION 88
You have an Azure subscription. The subscription contains an Azure OpenAI resource that hosts a GPT-3.5 Turbo model named Model1.
You configure Model1 to use the following system message: 'You are an AI assistant that helps people solve mathematical puzzles. Explain your answers as if the request is by a 4-year-old.'
Which type of prompt engineering technique is this an example of?

A. few-shot learning
B. affordance
C. chain of thought
D. priming

Correct Answer: D
Section:

QUESTION 89
DRAG DROP
You have an Azure subscription that contains an Azure Al Search resource named AS1.
You implement a custom skill in AS1 that performs language and sentiment analysis of documents.
You are evaluating the use of AS1 as part of an enrichment pipeline.
In which order will AS1 index the documents? To answer, move all indexing stages from the list of stages to the answer area and arrange them in the correct order.

Select and Place:

Correct Answer:

Stages

Answer Area

- document cracking
- skillset execution
- field mappings
- output field mappings
- push to index

Section:
Explanation:

QUESTION 90
You have an Azure subscription that contains an Azure OpenAI resource named AI1 and a user named User1.
You need to ensure that User1 can perform the following actions in Azure OpenAI Studio:
. Identify resource endpoints.
. View models that are available for deployment.
. Generate text and images by using the deplayed models.
The solution must follow the principle of least privilege.
Which role should you assign to User1?

A. Cognitive Services Contributor
B. Cognitive Services OpenAI Contributor
C. Contributor
D. Cognitive Services OpenAI User

Correct Answer: D
Section:

QUESTION 91
You are building an app that will use the Azure AI Speech service.
You need to ensure that the app can authenticate to the service by using a Microsoft Entra ID token.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Request an X.509 certificate.
B. Create a private endpoint.
C. Create a Conditional Access policy.
D. Configure a custom subdomain.
E. Enable a virtual network service endpoint.

Correct Answer: B, D
Section:

QUESTION 92
You are developing a text processing solution.
You have the following function.
static void GotKeyWords(TextAnalyticsClient textAnalyticsClient, string text)
var response = textAnalyticsClient.RecognizeEntities (text);
Console.WriteLine('Key words:');
foreach (CategorizedEntity entity in response.Value)
Console.WriteLine($'\t{entity.Text}');
You call the function and use the following string as the second argument.
Our tour of London included a visit to Buckinghan Palace
What will the function return?

A. Tour and visit only
B. London and Tour only
C. Our tour of London included a visit to Buckingham Palace
D. London and Buckinghan Palace only

Correct Answer: D
Section:
Explanation:
The function GotKeyWords uses the RecognizeEntities method from the TextAnalyticsClient class in Azure's Text Analytics service to identify named entities in the given text. The method extracts entities such as locations, organizations, dates, and other significant items.
Given the string 'Our tour of London included a visit to Buckinghan Palace', the RecognizeEntities method will identify named entities within the text.
The named entities in this text are:
London (a location)
Buckingham Palace (a location, but note the typo 'Buckinghan Palace' instead of 'Buckingham Palace')
However, the typo in 'Buckinghan Palace' might prevent the service from recognizing it correctly. Assuming the typo is fixed, the recognized entities would be 'London' and 'Buckingham Palace'.