Number: AI-102
Passing Score: 800
Time Limit: 120
File Version: 41.0

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
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:

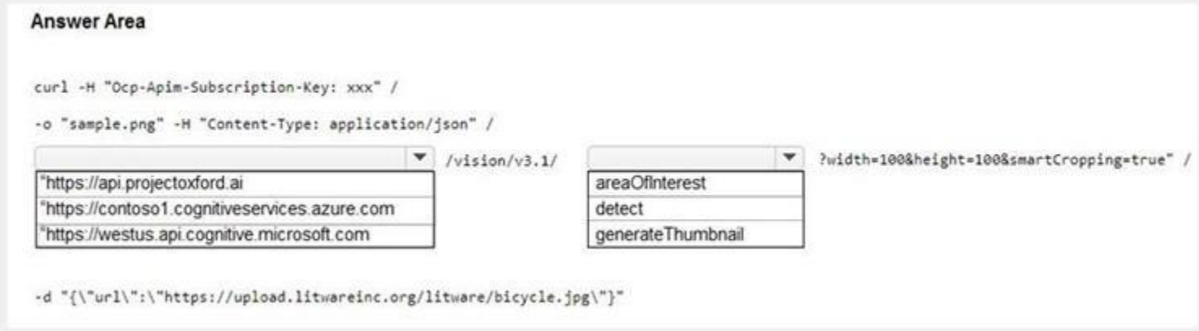

Answer Area:

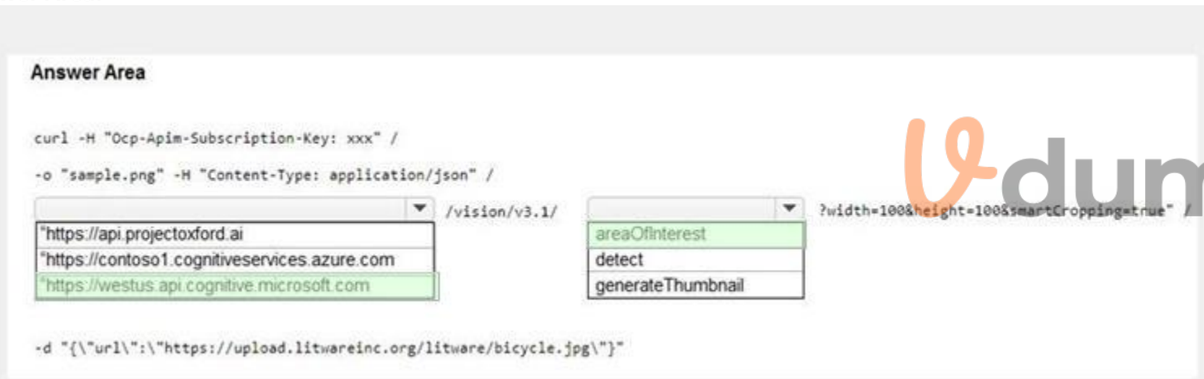

Section:
Explanation:
Reference:
https://westus.dev.cognitive.microsoft.com/docs/services/computer-vision-v3-2/operations/56f91f2e778daf14a499f21b
https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-generating-thumbnails#examples

QUESTION 3
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:

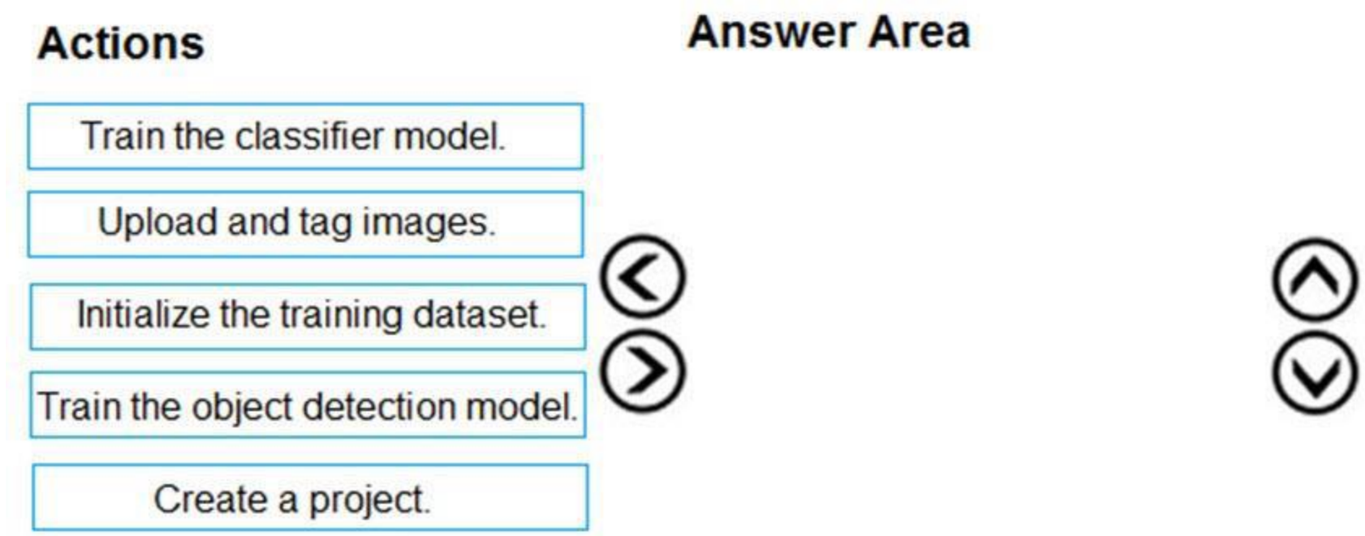

Correct Answer:

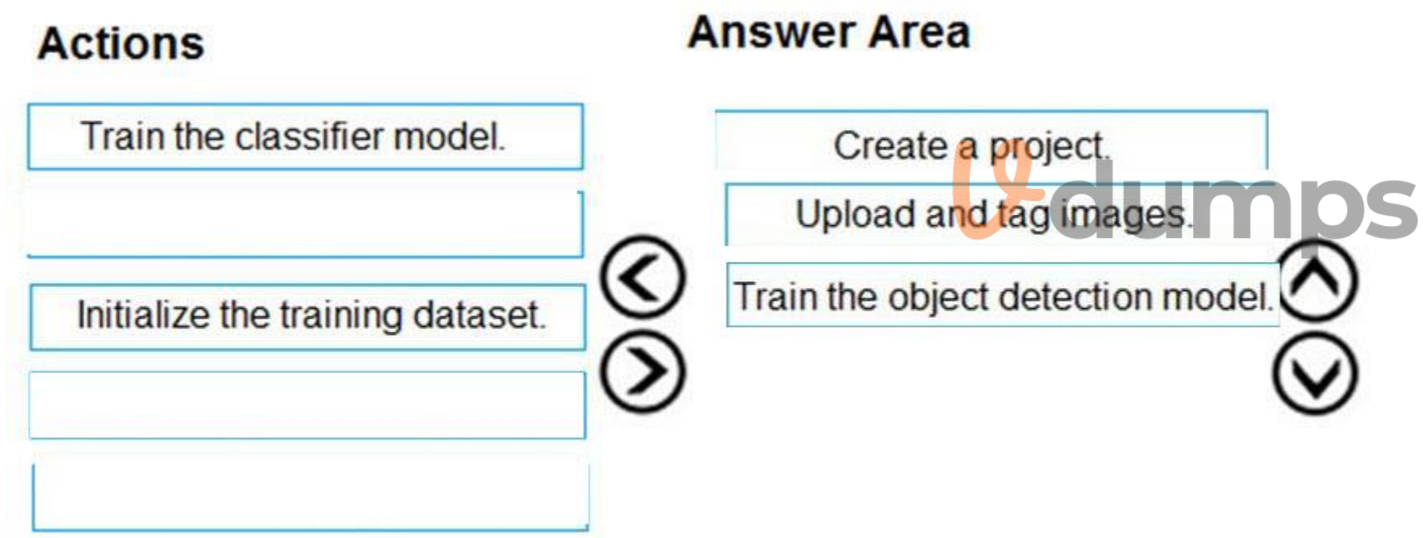

Section:
Explanation:
Step 1: Create a project
Create a new project.
Step 2: Upload and tag the images
Choose training images. Then upload and tag the images.
Step 3: Train the classifier model.
Train the classifier
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/getting-started-build-a-classifier

QUESTION 4
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:

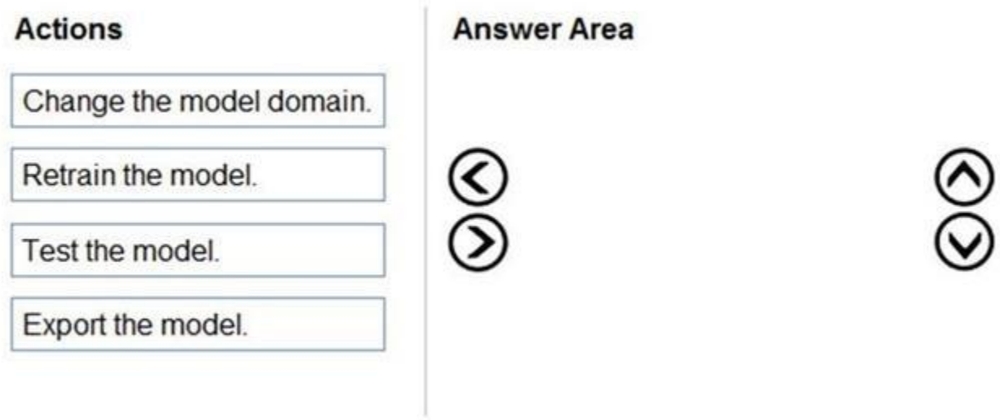

Correct Answer:

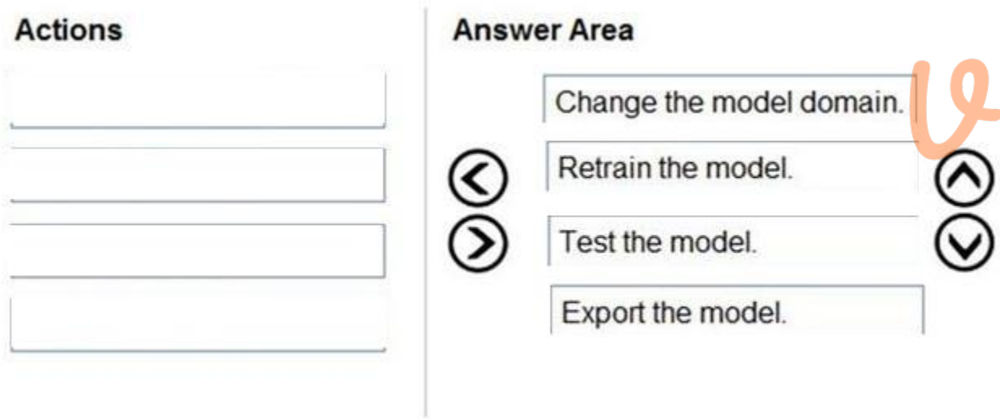

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/export-your-model

QUESTION 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

Statements	Yes	No
The code will add a face image to a person object in a person group.	○	○
The code will work for a group of 10,000 people.	○	○
`AddFace` can be called multiple times to add multiple face images to a person object.	○	○

Answer Area:

Answer Area

Statements	Yes	No
The code will add a face image to a person object in a person group.	○	○
The code will work for a group of 10,000 people.	○	○
`AddFace` can be called multiple times to add multiple face images to a person object.	○	○

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:

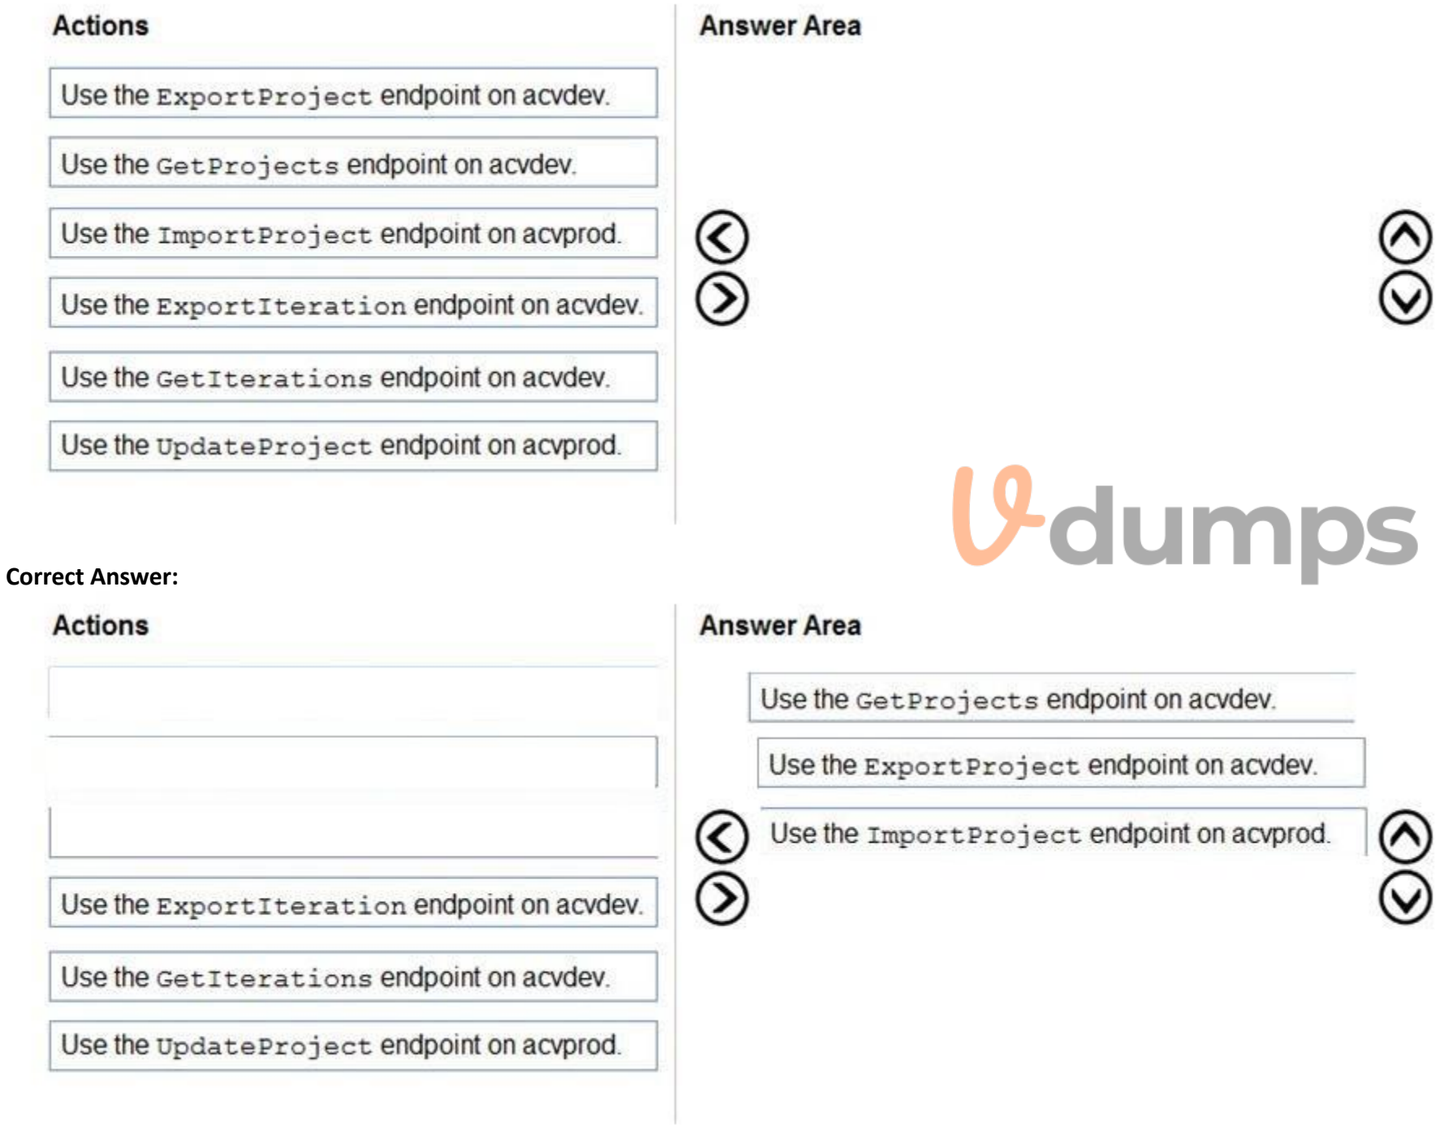

Correct Answer:

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

QUESTION 7
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 8
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 9
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 10
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 11
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 12
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

Statements	Yes	No
The code will return the name of each detected brand with a confidence equal to or higher than 75 percent.	○	○
The code will return coordinates for the top-left corner of the rectangle that contains the brand logo of the displayed brands.	○	○
The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.	○	○

Answer Area:

Answer Area

Statements	Yes	No
The code will return the name of each detected brand with a confidence equal to or higher than 75 percent.	○	○
The code will return coordinates for the top-left corner of the rectangle that contains the brand logo of the displayed brands.	○	○
The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.	○	○

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 13
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))
                  File
                  Stream
                  Uri
                  Url

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


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))
        {
          await faceClient.PersonGroupPerson.
(personGroupId, personId, t);
        }
});
```

File
Stream
Uri
Url

AddFaceFromStreamAsync
AddFaceFromUrlAsync
CreateAsync
GetAsync

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 14
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:

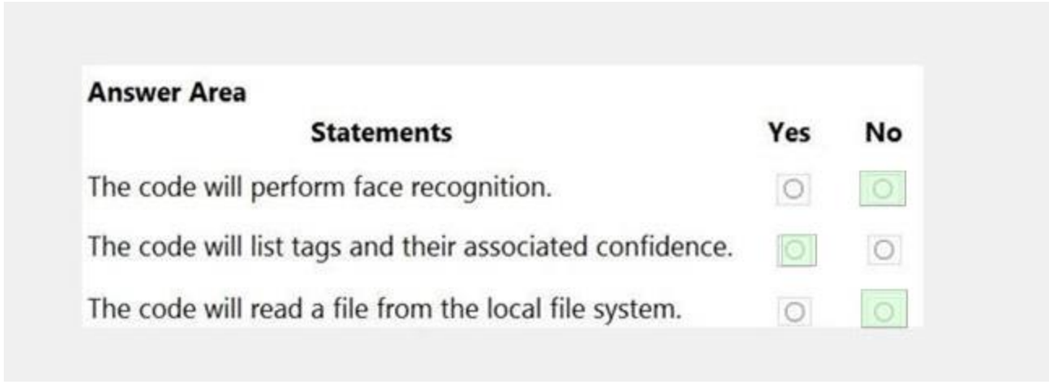

Section:
Explanation:
Box 1: No
Box 2: Yes
The ComputerVision.analyzeImageInStreamAsync operation extracts a rich set of visual features based on the image content.
Box 3: No
Images will be read from a stream.
Reference:
https://docs.microsoft.com/en-us/java/api/com.microsoft.azure.cognitiveservices.vision.computervision.computervision.analyzeimageinstreamasync

QUESTION 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

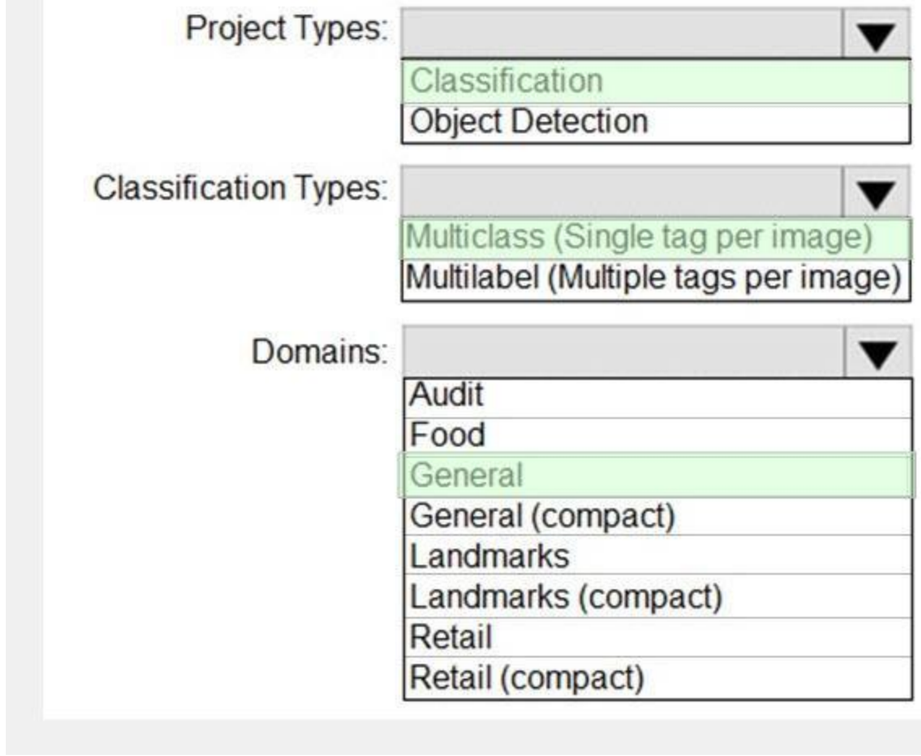

Section:
Explanation:
Box 1: Classification
Incorrect Answers:
An object detection project is for detecting which objects, if any, from a set of candidates are present in an image.
Box 2: Multiclass
A multiclass classification project is for classifying images into a set of tags, or target labels. An 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:

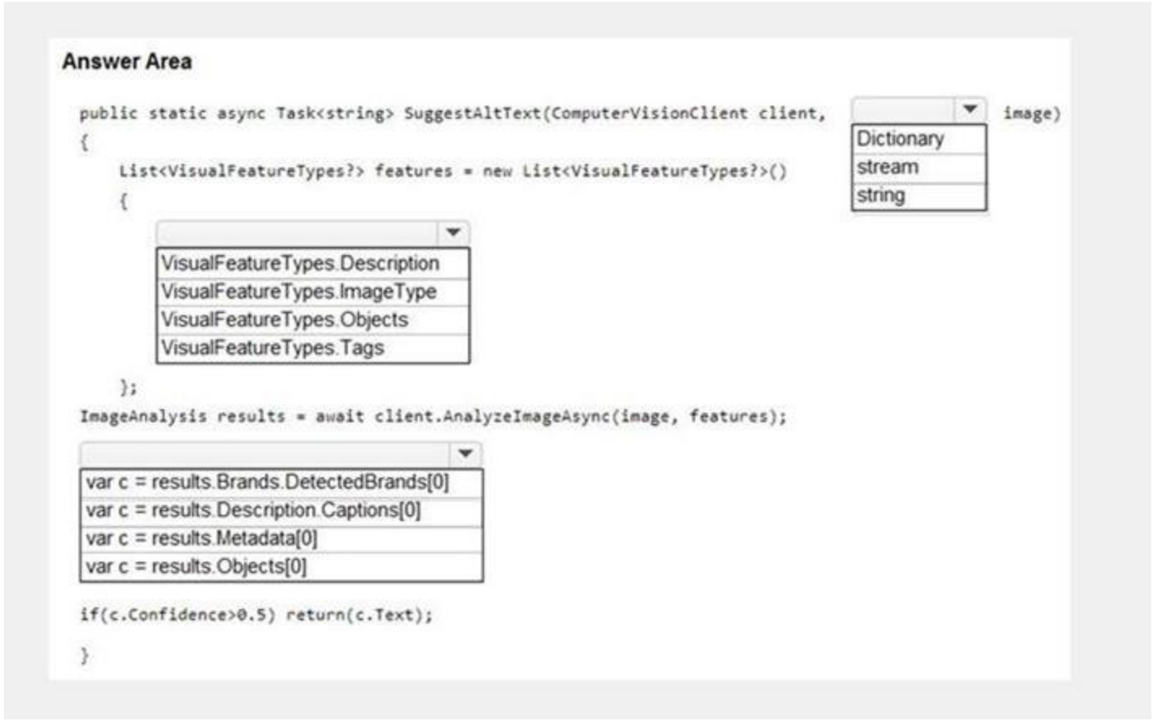

Answer Area:

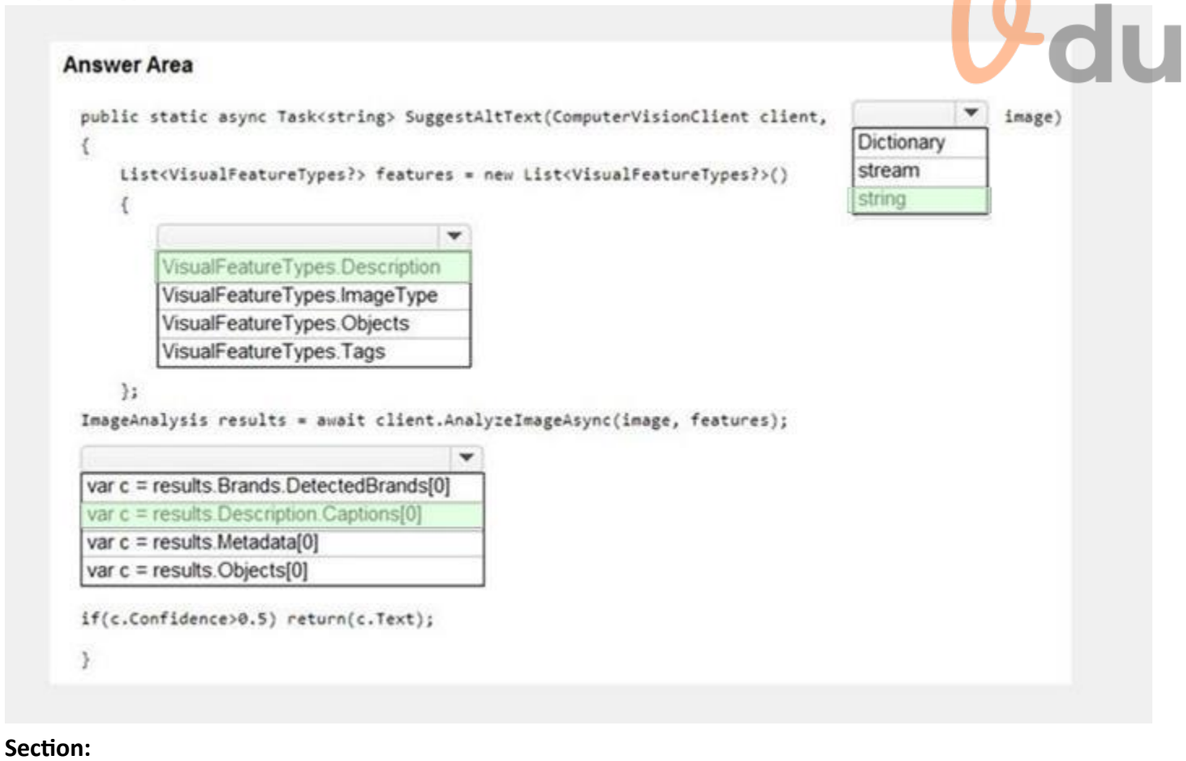

Section:
Explanation:
Reference:

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
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

```
{
    "@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"}
    ]
}
```

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

QUESTION 2

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 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
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

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:

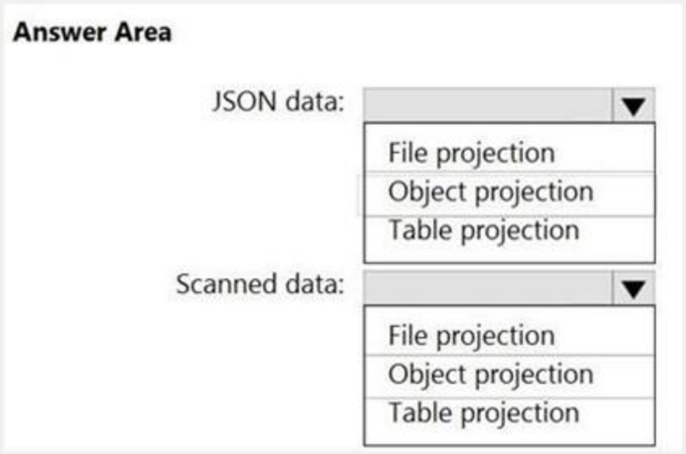

Answer Area:

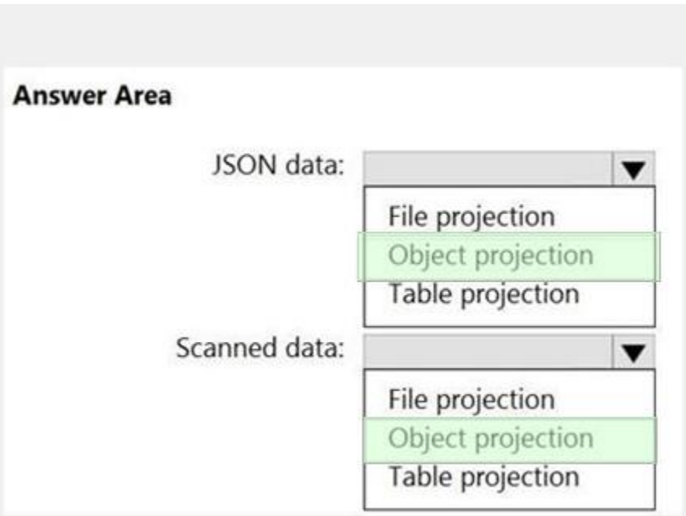

Section:
Explanation:
Box 1: Object projection
Object projections are JSON representations of the enrichment tree that can be sourced from any node.
Box 2: File projection
File projections are similar to object projections and only act on the normalized_images collection.
Reference:
https://docs.microsoft.com/en-us/azure/search/knowledge-store-projection-overvie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

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

Statements	Yes	No
CompanyDescription is available for indexing.	○	○
The definition calls a web API as part of the enrichment process.	○	○
The enrichment step is called only for the first organization under "/document/organizations".	○	○

Answer Area:

Answer Area

Statements	Yes	No
CompanyDescription is available for indexing.	○	○
The definition calls a web API as part of the enrichment process.	○	○
The enrichment step is called only for the first organization under "/document/organizations".	○	○

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 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 6
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

QUESTION 7
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 8
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 9
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:

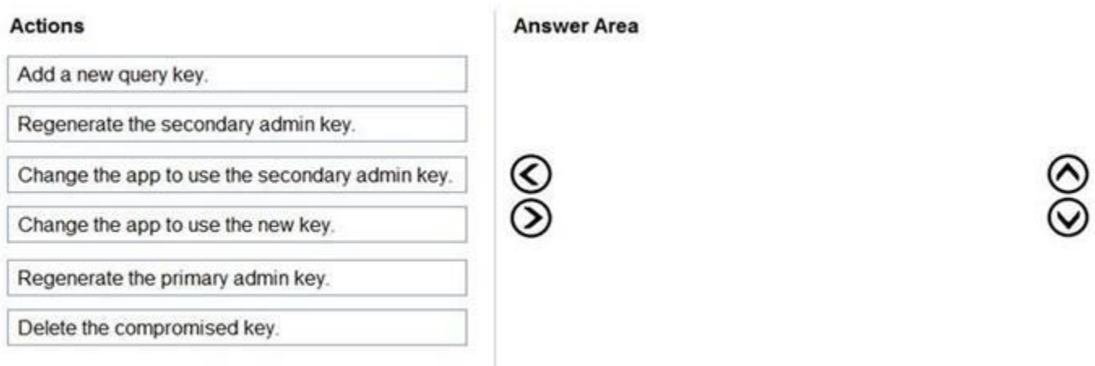

Correct Answer:

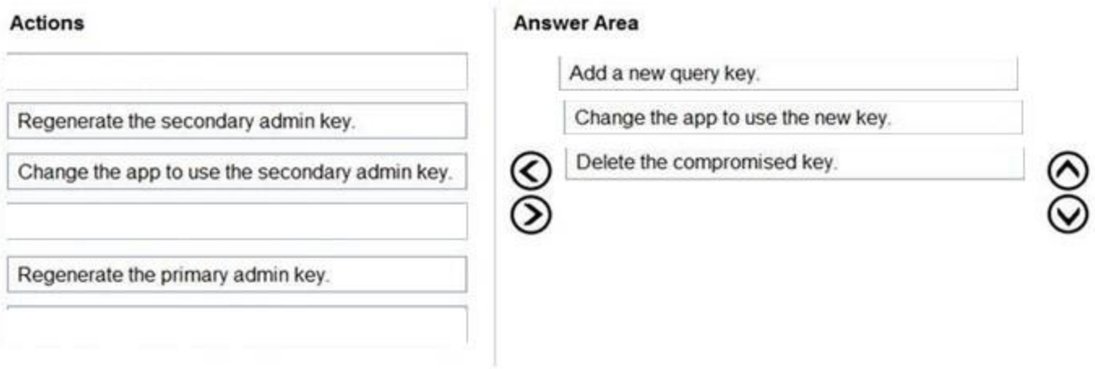

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/search/search-security-api-keys

QUESTION 10
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

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

01 - Implement Conversational AI Solutions

QUESTION 1
HOTSPOT
You are building a chatbot by using the Microsoft Bot Framework SDK.
You use an object named UserProfile to store user profile information and an object named ConversationData to store information related to a conversation.
You create the following state accessors to store both objects in state.
var userStateAccessors = _userState.CreateProperty<UserProfile>(nameof(UserProfile));
var conversationStateAccessors = _conversationState.CreateProperty<ConversationData>(nameof(ConversationData));
The state storage mechanism is set to Memory Storage.
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 create and maintain the `UserProfile` object in the underlying storage layer.	○	○
The code will create and maintain the `ConversationData` object in the underlying storage layer.	○	○
The `UserProfile` and `ConversationData` objects will persist when the Bot Framework runtime terminates.	○	○

Answer Area:

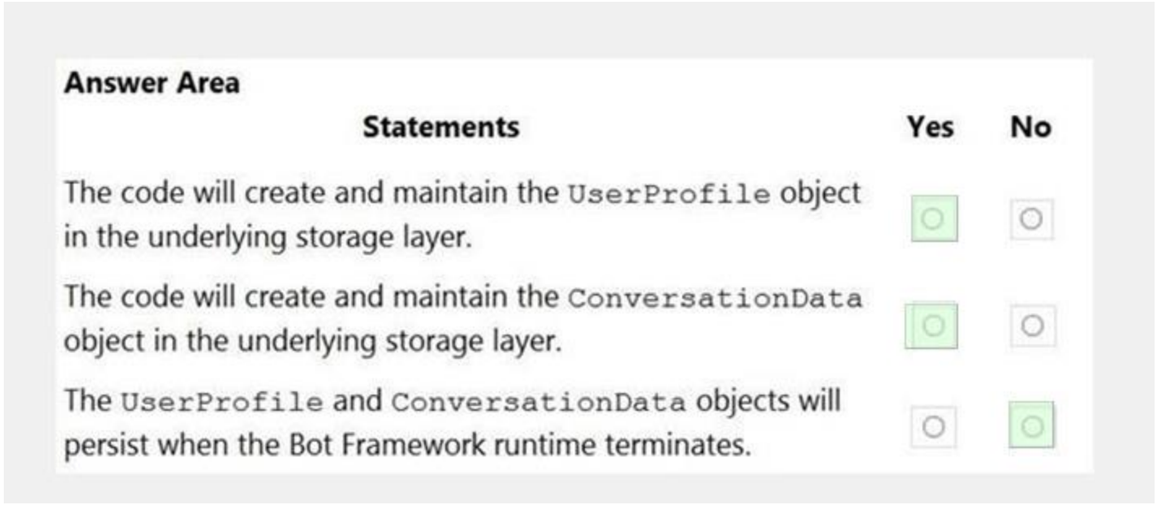

Section:
Explanation:
Box 1: Yes
You create property accessors using the CreateProperty method that provides a handle to the BotState object. Each state property accessor allows you to get or set the value of the associated state property.
Box 2: Yes
Box 3: No
Before you exit the turn handler, you use the state management objects' SaveChangesAsync() method to write all state changes back to storage.
Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-v4-state

QUESTION 2
HOTSPOT
You are building a chatbot that will provide information to users as shown in the following exhibit.

Passengers
Sarah Hum
Jeremy Goldberg
Evan Litvak

2 Stops
Tue, May 30, 2017 10:25 PM

San Francisco
Amsterdam

San Francisco
Amsterdam

SFO
AMS

SFO
AMS

Non-Stop
Fri, Jun 2, 2017 11:55 PM

San Francisco
Amsterdam

San Francisco
Amsterdam

SFO
AMS

SFO
AMS

Total **$4,032.54**

Use the drop-down menus to 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:

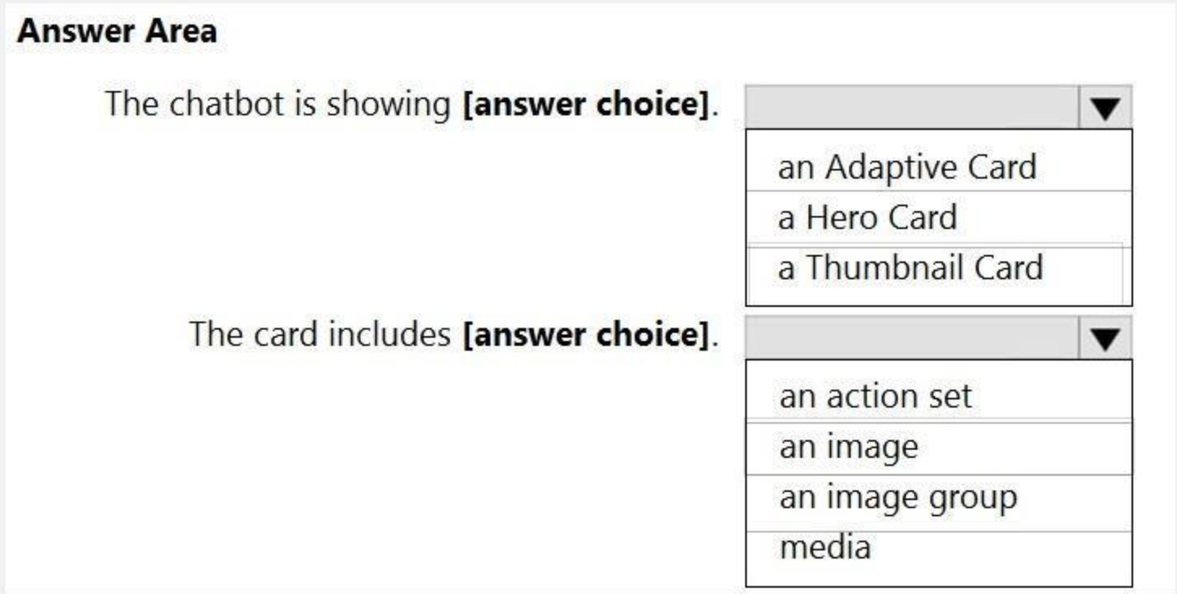

Answer Area:

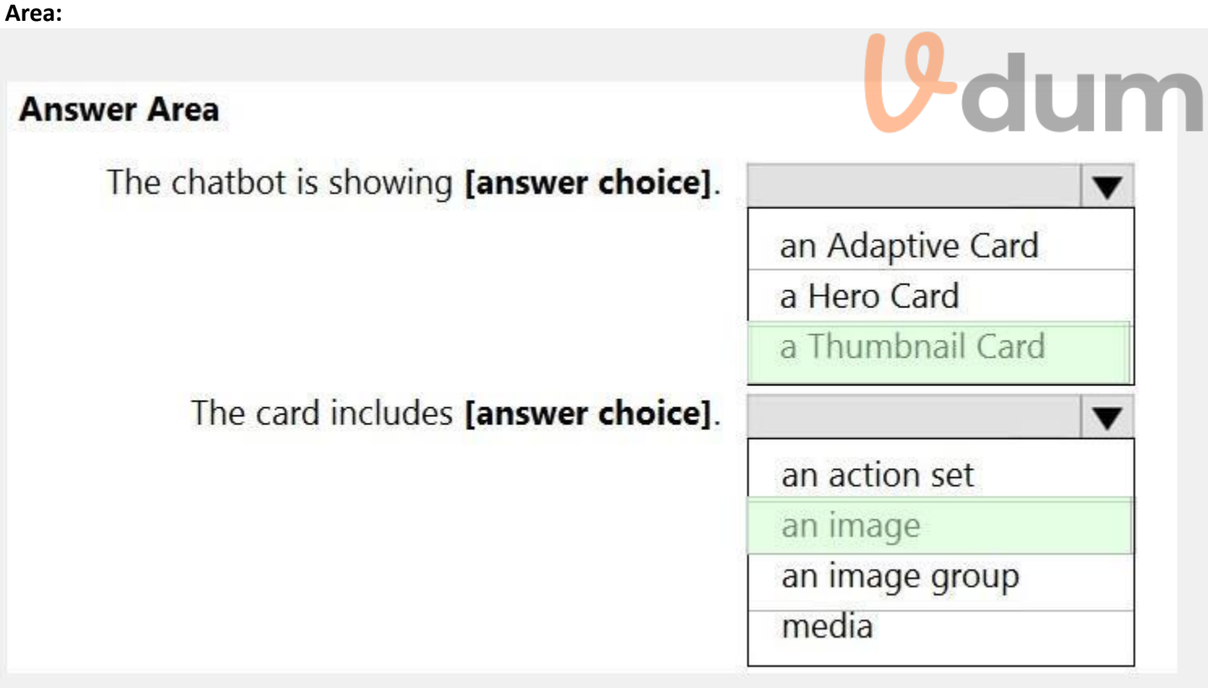

Section:
Explanation:
Box 1: A Thumbnail card
A Thumbnail card typically contains a single thumbnail image, some short text, and one or more buttons.
Incorrect Answers:

an Adaptive card is highly customizable card that can contain any combination of text, speech, images, buttons, and input fields.
a Hero card typically contains a single large image, one or more buttons, and a small amount of text.
Box 2: an image
Reference:
https://docs.microsoft.com/en-us/microsoftteams/platform/task-modules-and-cards/cards/cards-reference

QUESTION 3
HOTSPOT
You are building a bot and that will use Language Understanding.
You have a LUDown file that contains the following content.

```
## Confirm
- confirm
- ok
- yes

## ExtractName
- call me steve !
- i am anna
- (i'm|i am) {@PersonName.Any}[.]
- my name is {@PersonName.Any}[.]

## Logout
- forget me
- log out

## SelectItem
- choose last
- choose the {@DirectionalReference=bottom left}
- choose {@DirectionalReference=top right}
- i like {@DirectionalReference=left} one

## SelectNone
- none

@ ml DirectionalReference
@ prebuilt personName
```

Use the drop-down menus to 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:

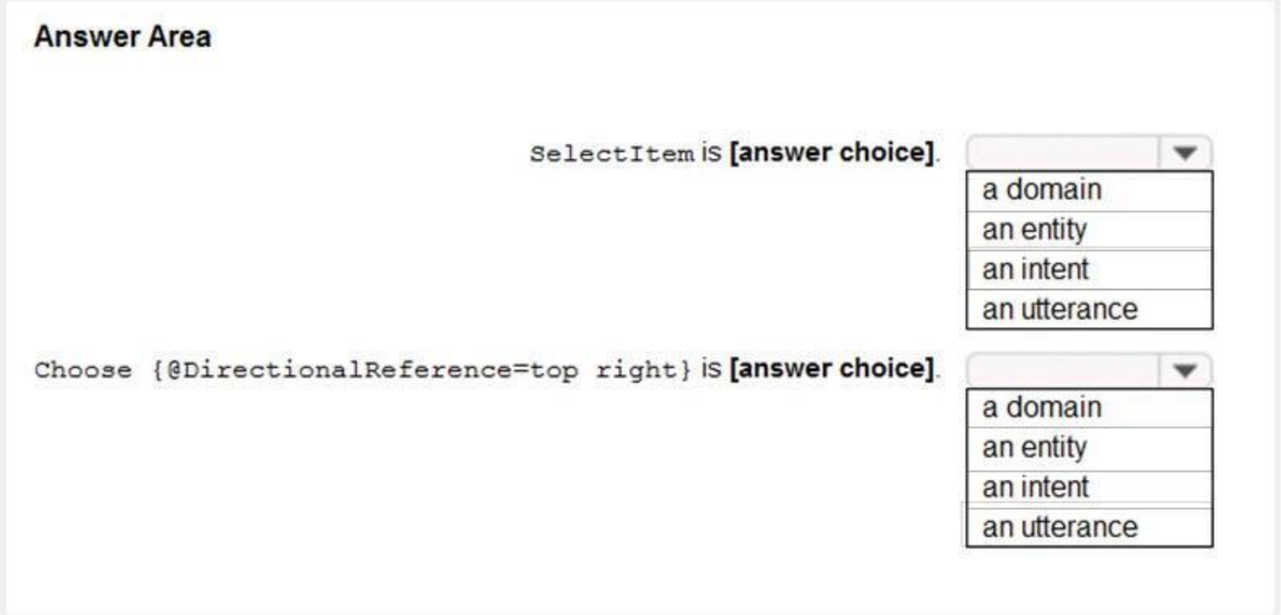

Answer Area:

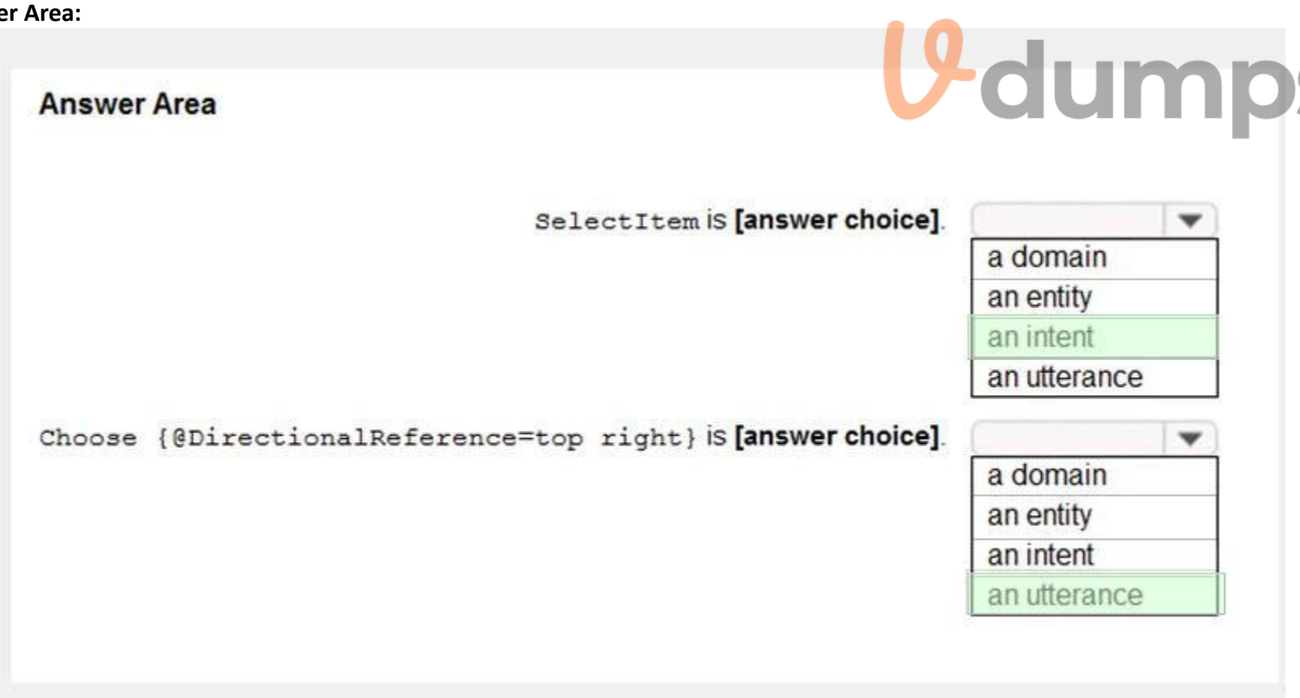

Section:
Explanation:
Reference:
https://github.com/solliancenet/tech-immersion-data-ai/blob/master/ai-exp1/README.md

QUESTION 4
HOTSPOT
You are designing a conversation flow to be used in a chatbot.
You need to test the conversation flow by using the Microsoft Bot Framework Emulator.
How should you complete the .chat file? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

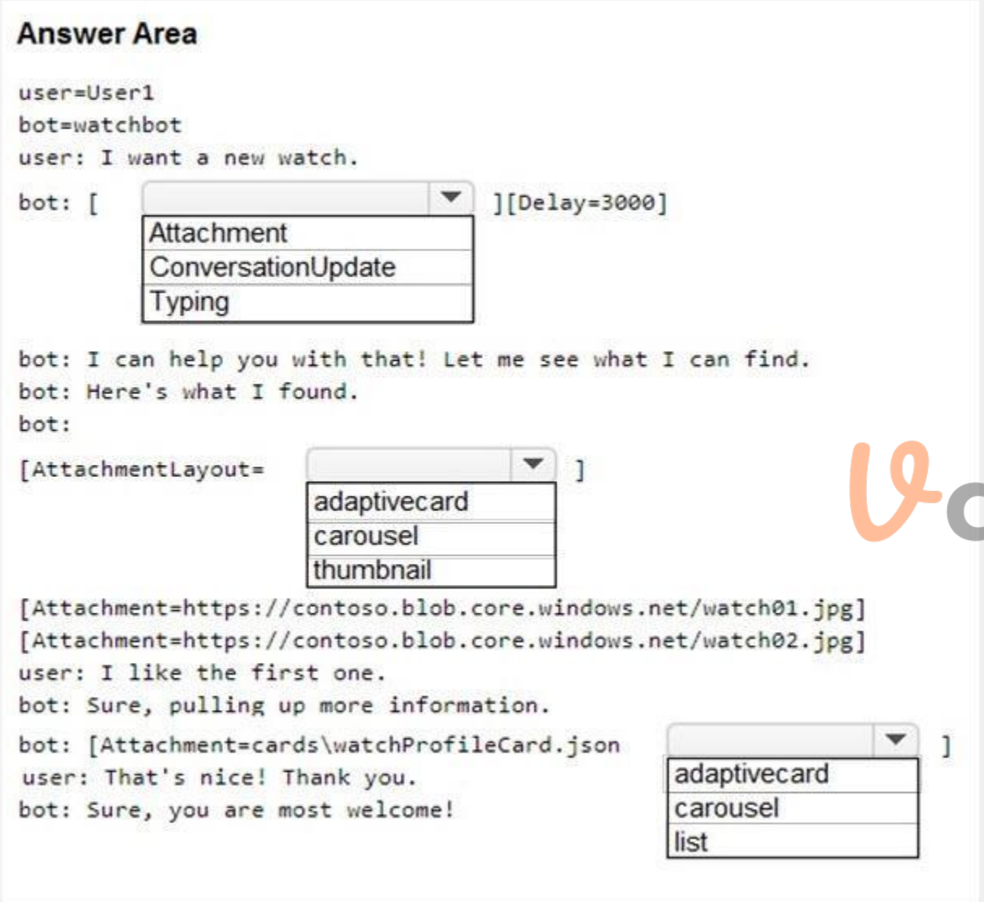

Answer Area:

Answer Area

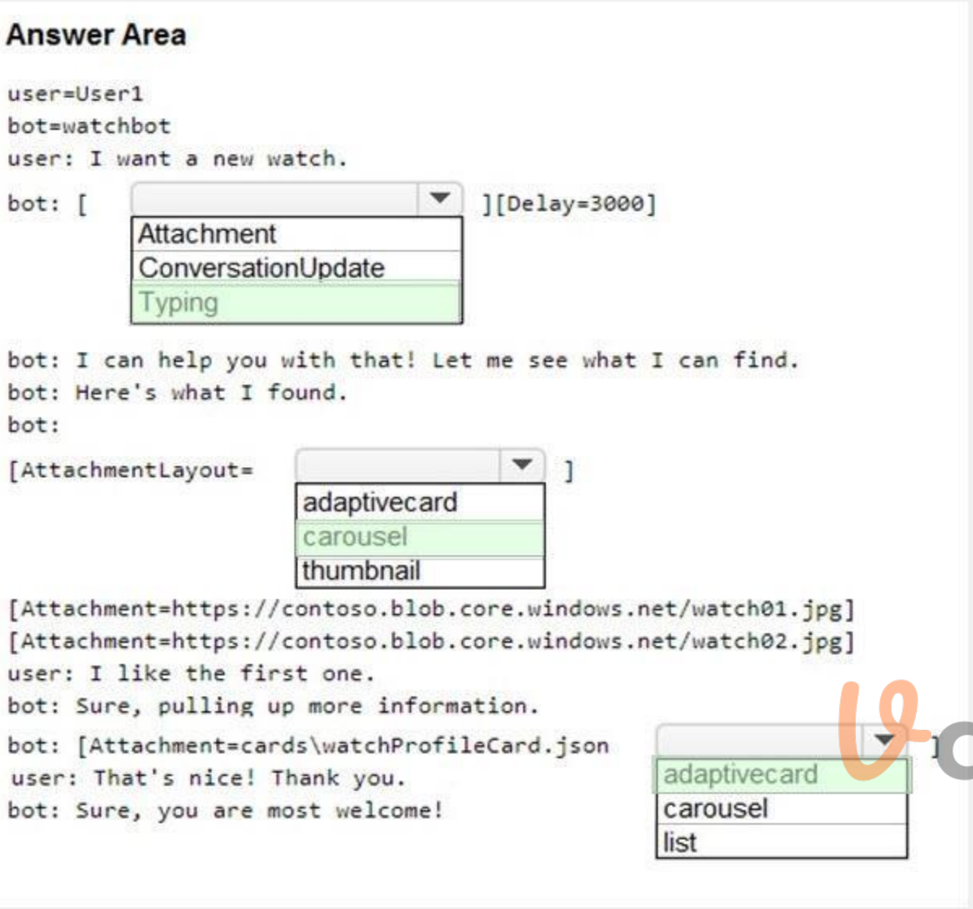

Section:
Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-add-media-attachments?view=azure-bot-service-4.0&tabs=csharp

QUESTION 5
You are building a multilingual chatbot
You need to send a different answer for positive and negative messages.
Which two Text Analytics APIs should you use? Each correct answer presents part of the solution. (Choose two.) NOTE: Each correct selection is worth one point.

A. Linked entities from a well-known knowledge base
B. Sentiment Analysis
C. Key Phrases
D. Detect Language
E. Named Entity Recognition

Correct Answer: B, D

Section:
Explanation:
B: The Text Analytics API's Sentiment Analysis feature provides two ways for detecting positive and negative sentiment. If you send a Sentiment Analysis request, the API will return sentiment labels (such as "negative", "neutral" and "positive") and confidence scores at the sentence and document-level.
D: The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.
This capability is useful for content stores that collect arbitrary text, where language is unknown.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-sentiment-analysis?tabs=version-3-1
https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-language-detection

QUESTION 6
You are building a chatbot by using the Microsoft Bot Framework Composer as shown in the exhibit. (Click the Exhibit tab.)

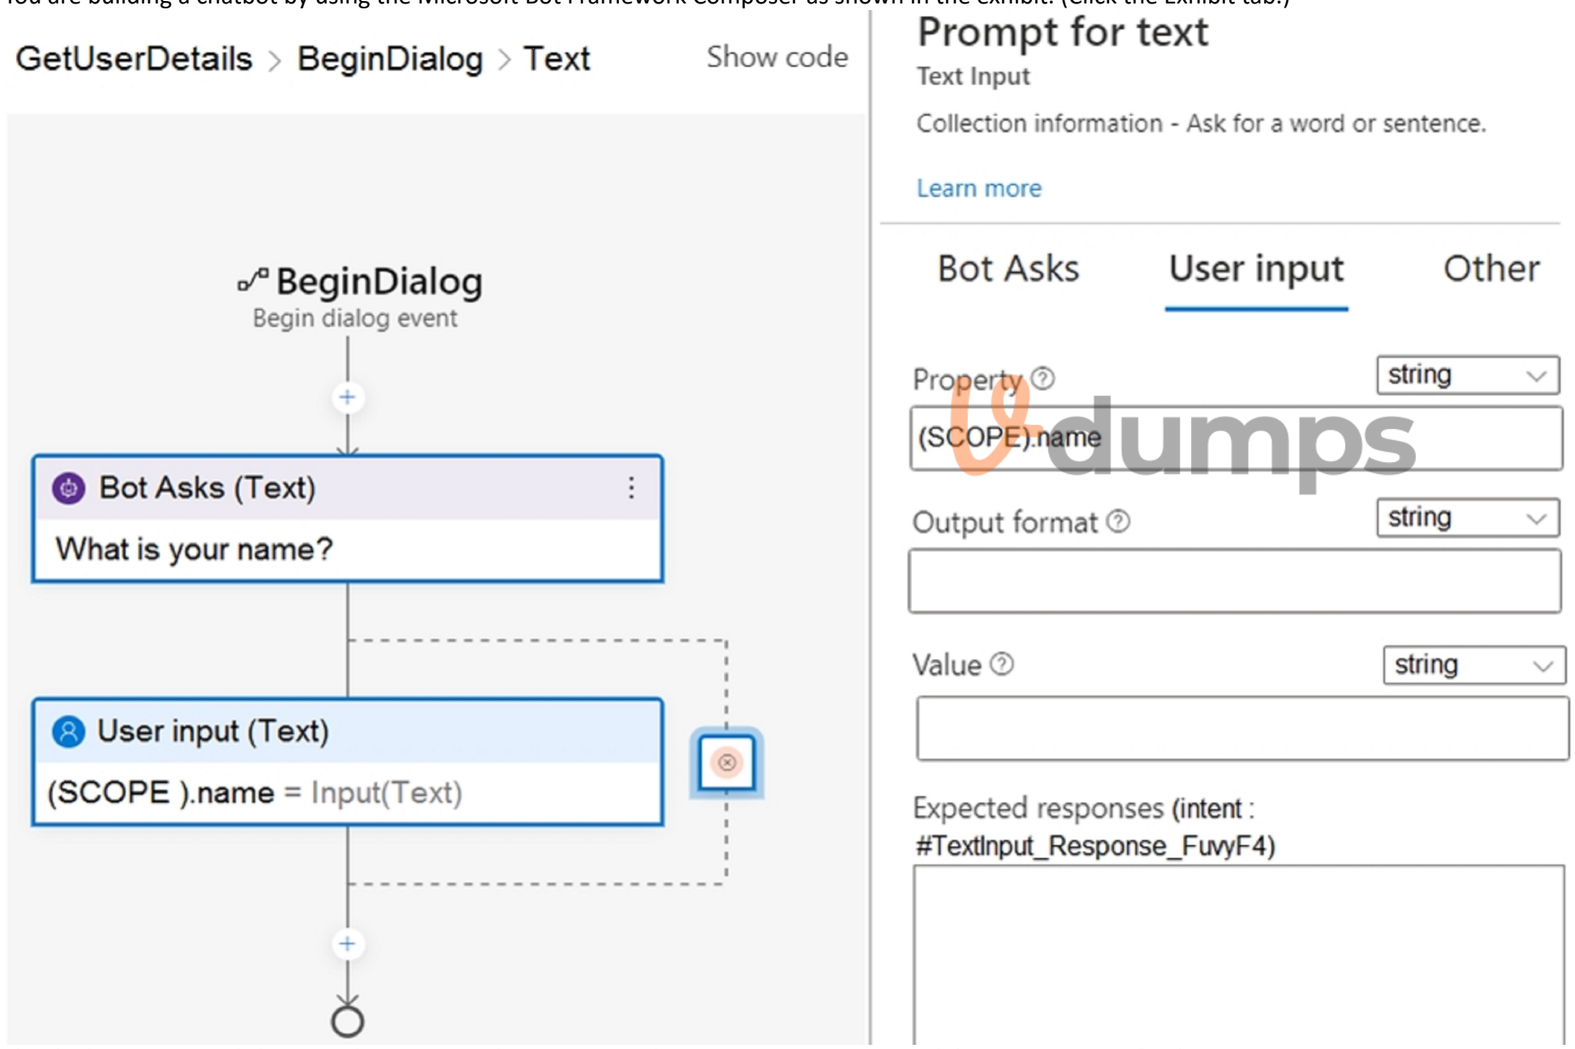

The chatbot contains a dialog named GetUserDetails. GetUserDetails contains a TextInput control that prompts users for their name.
The user input will be stored in a property named name.
You need to ensure that you can dispose of the property when the last active dialog ends.
Which scope should you assign to name?

A. dialog

B. user
C. turn
D. conversation

Correct Answer: A
Section:
Explanation:
The dialog scope associates properties with the active dialog. Properties in the dialog scope are retained until the dialog ends.
Incorrect Answers:
A: The conversation scope associates properties with the current conversation. Properties in the conversation scope have a lifetime of the conversation itself. These properties are in scope while the bot is processing an activity associated with the conversation (for example, multiple users together in a Microsoft Teams channel).
B: The user scope associates properties with the current user. Properties in the user scope do not expire. These properties are in scope while the bot is processing an activity associated with the user.
C: The turn scope associates properties with the current turn. Properties in the turn scope expire at the end of the turn.
Reference:
https://docs.microsoft.com/en-us/composer/concept-memory?tabs=v2x

QUESTION 7
DRAG DROP
You have a chatbot that uses a QnA Maker application.
You enable active learning for the knowledge base used by the QnA Maker application.
You need to integrate user input into 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:

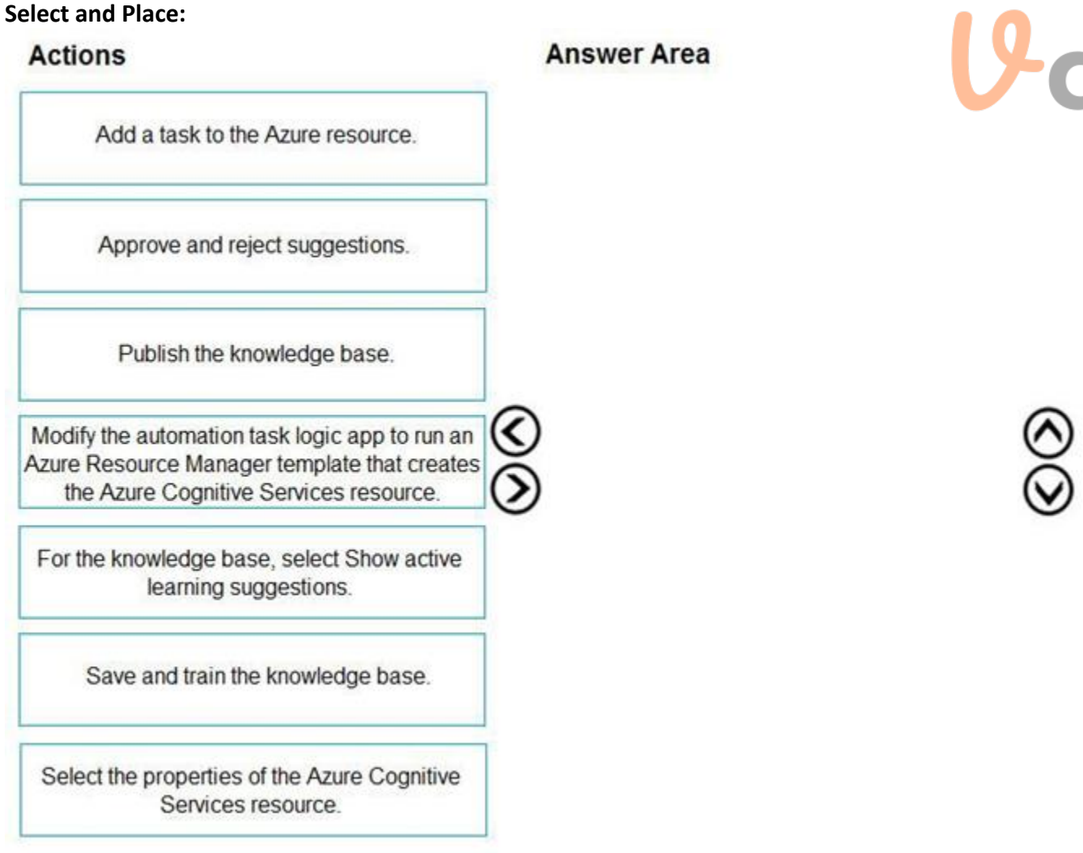

Correct Answer:

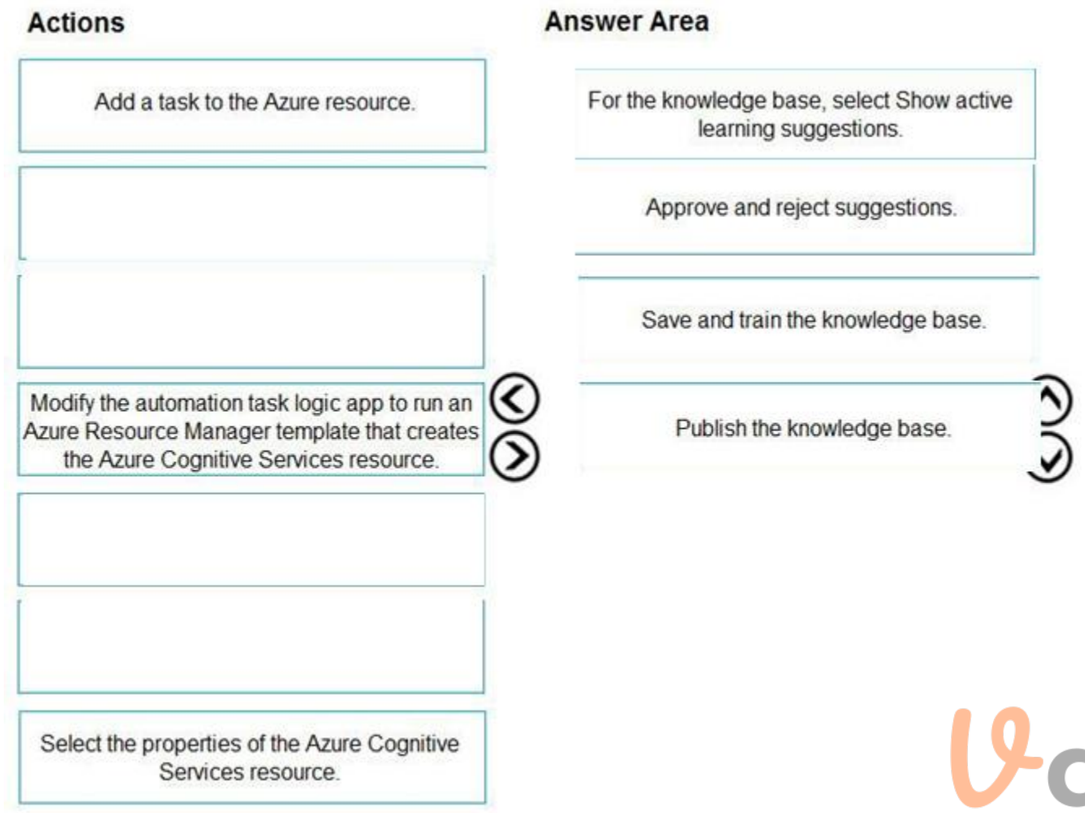

Section:
Explanation:
Step 1: For the knowledge base, select Show active learning suggestions.
In order to see the suggested questions, on the Edit knowledge base page, select View Options, then select Show active learning suggestions.
Step 2: Approve and reject suggestions.
Each QnA pair suggests the new question alternatives with a check mark, ?, to accept the question or an x to reject the suggestions. Select the check mark to add the question.
Step 3: Save and train the knowledge base.
Select Save and Train to save the changes to the knowledge base.
Step 4: Publish the knowledge base.
Select Publish to allow the changes to be available from the GenerateAnswer API.
When 5 or more similar queries are clustered, every 30 minutes, QnA Maker suggests the alternate questions for you to accept or reject.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/improve-knowledge-base

QUESTION 8
You need to enable speech capabilities for a chatbot.
Which three actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Enable WebSockets for the chatbot app.
B. Create a Speech service.
C. Register a Direct Line Speech channel.

D. Register a Cortana channel.
E. Enable CORS for the chatbot app.
F. Create a Language Understanding service.

Correct Answer: A, B, C
Section:
Explanation:
You can use the Speech service to voice-enable a chat bot.
The Direct Line Speech channel uses the text-to-speech service, which has neural and standard voices.
You'll need to make a small configuration change so that your bot can communicate with the Direct Line Speech channel using web sockets.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/tutorial-voice-enable-your-bot-speech-sdk

QUESTION 9
You are building a bot on a local computer by using the Microsoft Bot Framework. The bot will use an existing Language Understanding model. You need to translate the Language Understanding model locally by using the Bot Framework CU.
What should you do first?

A. From the Language Understanding portal, clone the model.
B. Export the model as an .lu file.
C. Create a new Speech service.
D. Create a new Language Understanding service.

Correct Answer: B
Section:
Explanation:
You might want to manage the translation and localization for the language understanding content for your bot independently.
Translate command in the @microsoft/bf-lu library takes advantage of the Microsoft text translation API to automatically machine translate .lu files to one or more than 60+ languages supported by the Microsoft text translation cognitive service.
What is translated?
An .lu file and optionally translate
Comments in the lu file
LU reference link texts
List of .lu files under a specific path.
Reference:
https://github.com/microsoft/botframework-cli/blob/main/packages/luis/docs/translate-command.md

QUESTION 10
You build a conversational bot named bot1.
You need to configure the bot to use a QnA Maker application.
From the Azure Portal, where can you find the information required by bot1 to connect to the QnA Maker application?

A. Access control (IAM)
B. Properties
C. Keys and Endpoint
D. Identity

Correct Answer: C

Section:
Explanation:
Obtain values to connect your bot to the knowledge base
1. In the QnA Maker site, select your knowledge base.
2. With your knowledge base open, select the SETTINGS tab. Record the value shown for service name. This value is useful for finding your knowledge base of interest when using the QnA Maker portal interface. It's not used to connect your bot app to this knowledge base.
3. Scroll down to find Deployment details and record the following values from the Postman sample HTTP request:
4. POST /knowledgebases/<knowledge-base-id>/generateAnswer
5. Host: <your-host-url>
6. Authorization: EndpointKey <your-endpoint-key>
Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-qna

QUESTION 11
You build a bot by using the Microsoft Bot Framework SDK and the Azure Bot Service.
You plan to deploy the bot to Azure.
You register the bot by using the Bot Channels Registration service.
Which two values are required to complete the deployment? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. botId
B. tenantId
C. appId
D. object Id
E. appSecret

Correct Answer: C, E
Section:
Explanation:
Reference:
https://github.com/MicrosoftDocs/bot-docs/blob/li ve/arti cles/bot-servi ce-qui ckstart-reqi strati on.md

QUESTION 12
DRAG DROP
You plan to build a chatbot to support task tracking.
You create a Language Understanding service named lu1.
You need to build a 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. (Choose four.)

Select and Place:

Actions	Answer Area
Train the application.	
Publish the application.	
Add a new application.	
Add example utterances.	
Add the prebuilt domain ToDo.	

Correct Answer:

Actions	Answer Area
	Add a new application.
	Add example utterances.
	Train the application.
	Publish the application.
Add the prebuilt domain ToDo.	

Section:

Explanation:

Step 1: Add a new application

Create a new app

1. Sign in to the LUIS portal with the URL of https://www.luis.ai.
2. Select Create new app.
3. Etc.

Step 2: Add example utterances.

In order to classify an utterance, the intent needs examples of user utterances that should be classified with this intent.

Step 3: Train the application

Step 4: Publish the application

In order to receive a LUIS prediction in a chat bot or other client application, you need to publish the app to the prediction endpoint.

Reference:

https://docs.microsoft.com/en-us/azure/cognitive-services/luis/tutorial-intents-only

QUESTION 13

DRAG DROP

You are using a Language Understanding service to handle natural language input from the users of a web-based customer agent.

The users report that the agent frequently responds with the following generic response: "Sorry, I don't understand that." You need to improve the ability of the agent to respond to requests.

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 prebuilt domain models as required.
- Validate the utterances logged for review and modify the model.
- Migrate authoring to an Azure resource authoring key.
- Enable active learning.
- Enable log collection by using Log Analytics.
- Train and republish the Language Understanding model.

Answer Area

Correct Answer:

Actions

- Validate the utterances logged for review and modify the model.
- Migrate authoring to an Azure resource authoring key.
- Enable log collection by using Log Analytics.

Answer Area

- Add prebuilt domain models as required.
- Enable active learning.
- Train and republish the Language Understanding model.

Section:

Explanation:

Step 1: Add prebuilt domain models as required.

Prebuilt models provide domains, intents, utterances, and entities. You can start your app with a prebuilt model or add a relevant model to your app later. Note: Language Understanding (LUIS) provides prebuilt domains,

which are pre-trained models of intents and entities that work together for domains or common categories of client applications.
The prebuilt domains are trained and ready to add to your LUIS app. The intents and entities of a prebuilt domain are fully customizable once you've added them to your app.
Step 2: Enable active learning
To enable active learning, you must log user queries. This is accomplished by calling the endpoint query with the log=true querystring parameter and value.
Step 3: Train and republish the Language Understanding model
The process of reviewing endpoint utterances for correct predictions is called Active learning. Active learning captures endpoint queries and selects user's endpoint utterances that it is unsure of. You review these utterances to select the intent and mark entities for these real-world utterances. Accept these changes into your example utterances then train and publish. LUIS then identifies utterances more accurately.
Incorrect Answers:
Enable log collection by using Log Analytics
Application authors can choose to enable logging on the utterances that are sent to a published application. This is not done through Log Analytics.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-how-to-review-endpoint-utterances#log-user-queries-to-enable-active-learning
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-prebuilt-model

QUESTION 14
HOTSPOT
You are building a chatbot by using the Microsoft Bot Framework Composer.
You have the dialog design shown in the following exhibit.

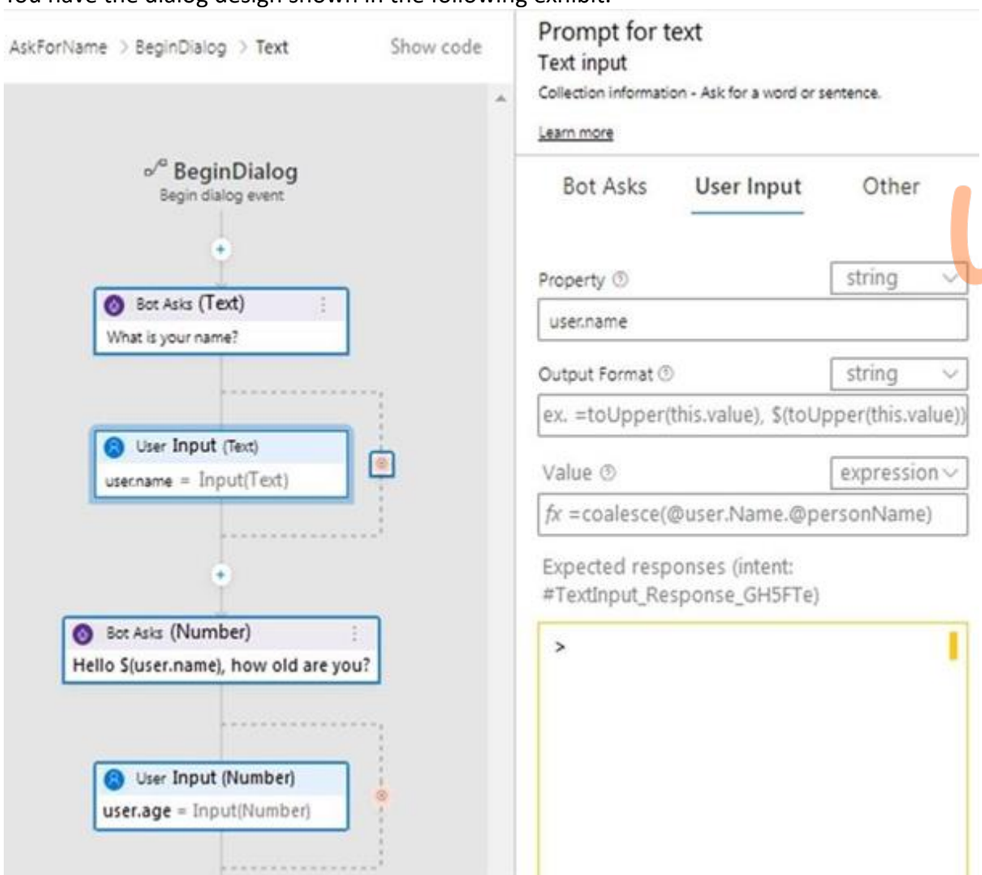

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

Section:
Explanation:
Box 1: No
User.name is a property.
Box 2: Yes
Box 3: Yes
The coalesce() function evaluates a list of expressions and returns the first non-null (or non-empty for string) expression.
Reference:
https://docs.microsoft.com/en-us/composer/concept-language-generation
https://docs.microsoft.com/en-us/azure/data-explorer/kusto/query/coalescefunction

QUESTION 15
HOTSPOT
You are building a chatbot for a Microsoft Teams channel by using the Microsoft Bot Framework SDK. The chatbot will use the following code.

```
protected override async Task OnMembersAddedAsync(IList<ChannelAccount>
membersAdded, ITurnContext<IConversationUpdateActivity> turnContext,
CancellationToken cancellationToken)
{
    foreach (var member in membersAdded)
        if (member.Id != turnContext.Activity.Recipient.Id)
            await turnContext.SendActivityAsync($"Hi there - {member.Name}.
{WelcomeMessage}", cancellationToken: cancellationToken);
}
```

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
`OnMembersAddedAsync` will be triggered when a user joins the conversation.	○	○
When a new user joins the conversation, the existing users in the conversation will see the chatbot greeting.	○	○
`OnMembersAddedAsync` will be initialized when a user sends a message.	○	○

Answer Area:

Answer Area

Statements	Yes	No
`OnMembersAddedAsync` will be triggered when a user joins the conversation.	○	○
When a new user joins the conversation, the existing users in the conversation will see the chatbot greeting.	○	○
`OnMembersAddedAsync` will be initialized when a user sends a message.	○	○

Section:
Explanation:
Box 1: Yes
The ActivityHandler.OnMembersAddedAsync method overrides this in a derived class to provide logic for when members other than the bot join the conversation, such as your bot's welcome logic.
Box 2: Yes
membersAdded is a list of all the members added to the conversation, as described by the conversation update activity. Box 3: No
Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.bot.builder.activityhandler.onmembersaddedasync?view=botbuilder-dotnet-stable

QUESTION 16
HOTSPOT
You are reviewing the design of a chatbot. The chatbot includes a language generation file that contains the following fragment.
Greet(user)
- ${Greeting()}, ${user.name}

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

Hot Area:

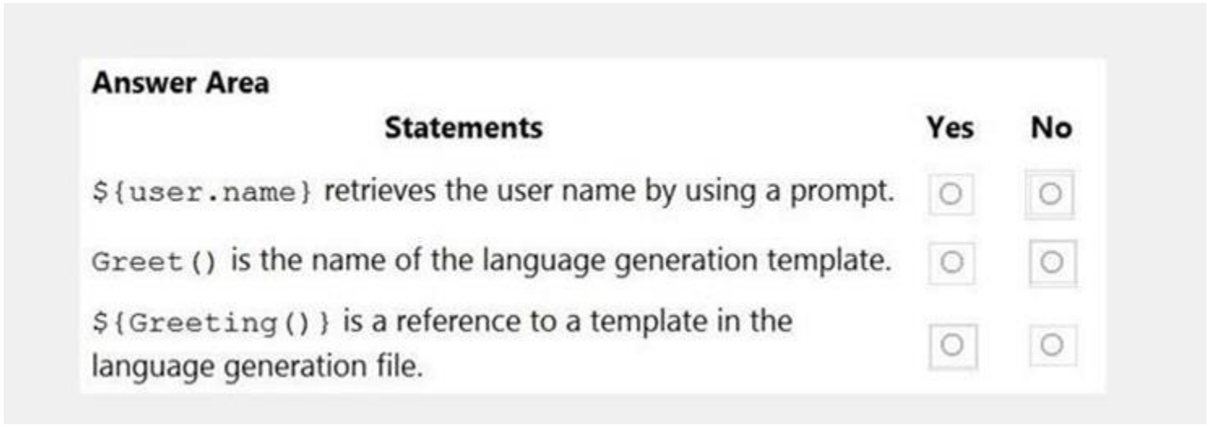

Answer Area:

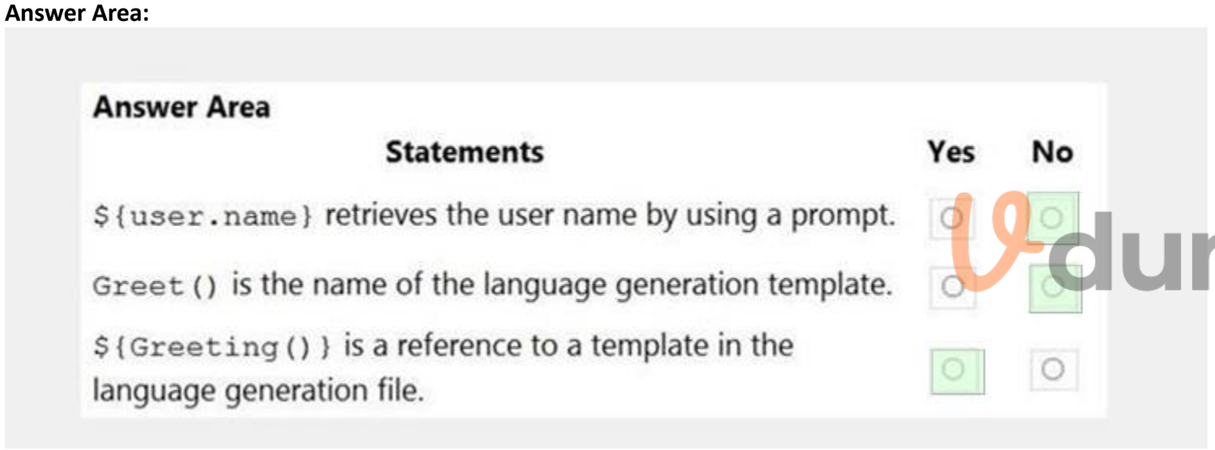

Section:
Explanation:
Box 1: No
Example: Greet a user whose name is stored in `user.name`
- ${ welcomeUser(user.name) }
Example: Greet a user whose name you don't know:
- ${ welcomeUser() }
Box 2: No
Greet(User) is a Send a response action.
Box 3: Yes
Reference:
https://docs.microsoft.com/en-us/composer/how-to-ask-for-user-input

02 - Implement Conversational AI 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
- 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 developing the shopping on-the-go project.
You need to build the Adaptive Card for the chatbot.
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:

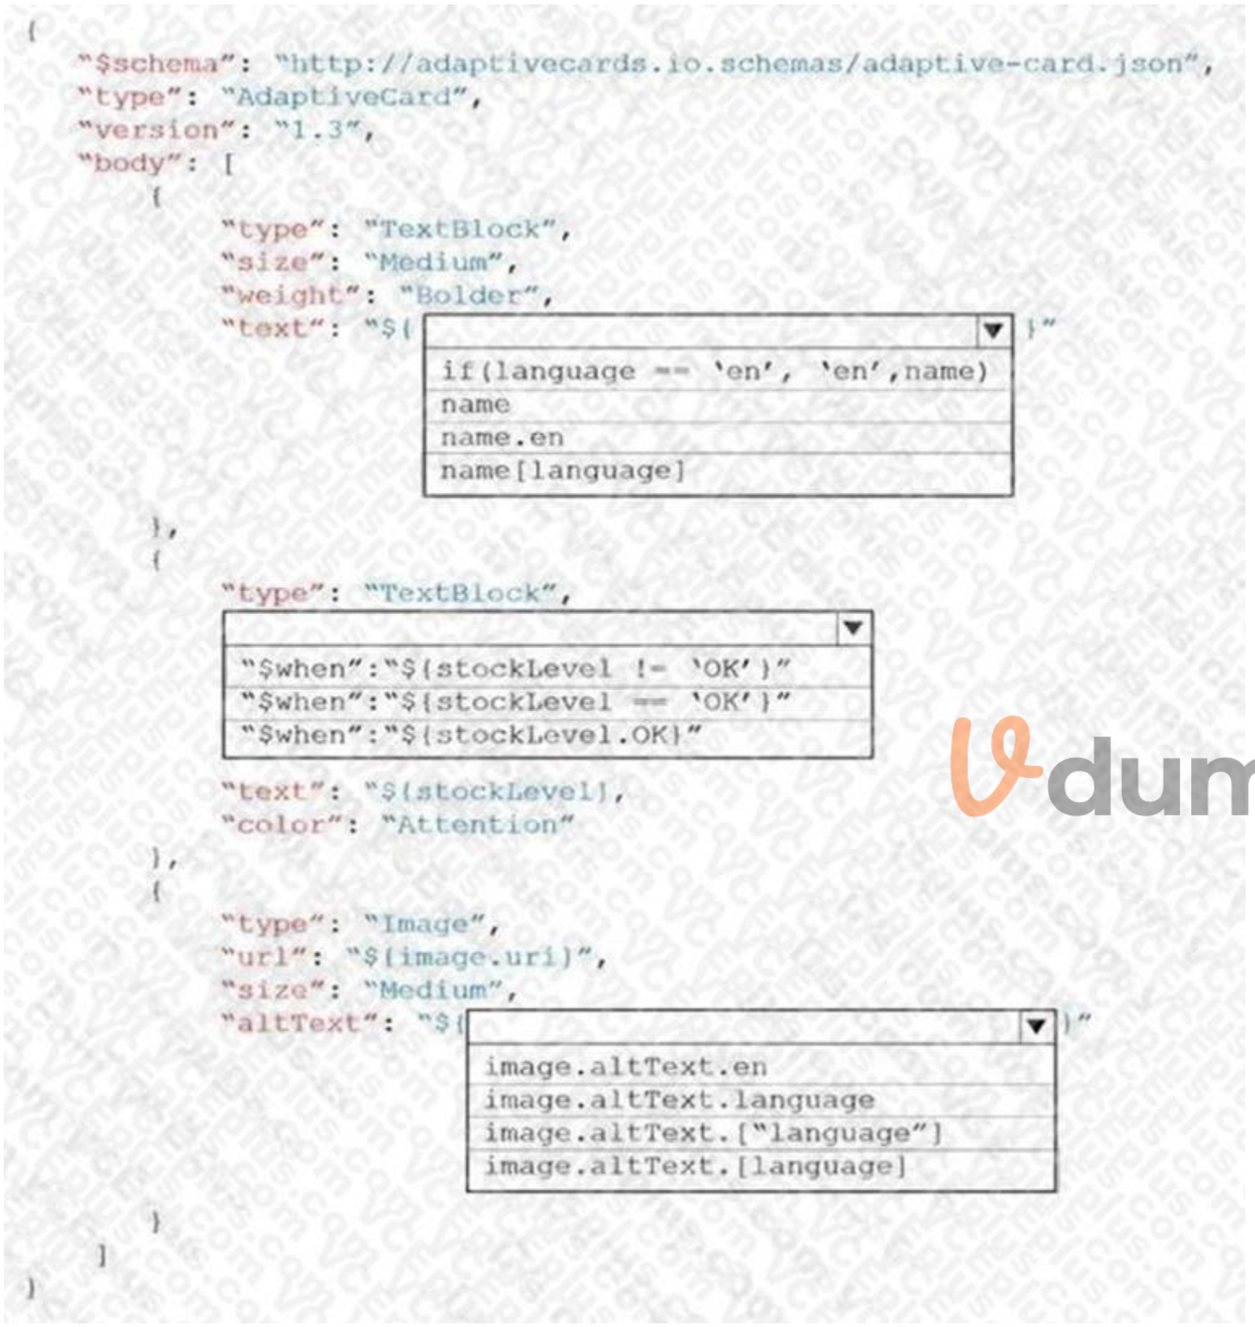

Answer Area:

Answer Area

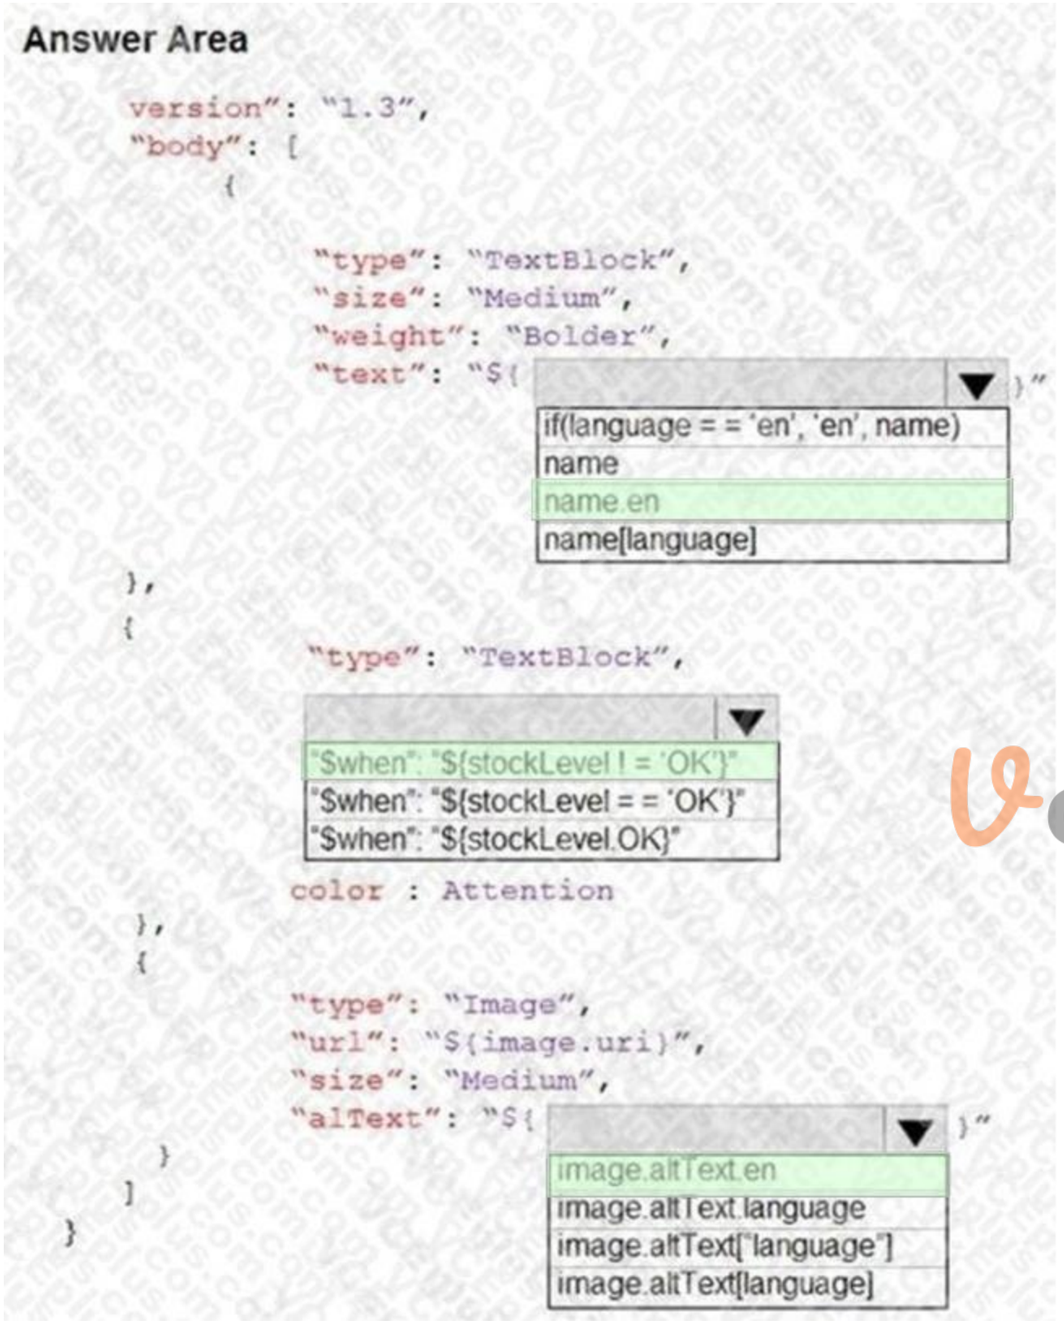

Section:
Explanation:
Box 1: name [language]
Chatbot must support interactions in English, Spanish, and Portuguese.
Box 2: "$when:${stockLevel != 'OK'}"
Product displays must include images and warnings when stock levels are low or out of stock.
Box 3: image.altText[language]

QUESTION 2
HOTSPOT
You are developing the shopping on-the-go project.
You are configuring access to the QnA Maker resources.
Which role should you assign to AllUsers and LeadershipTeam? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

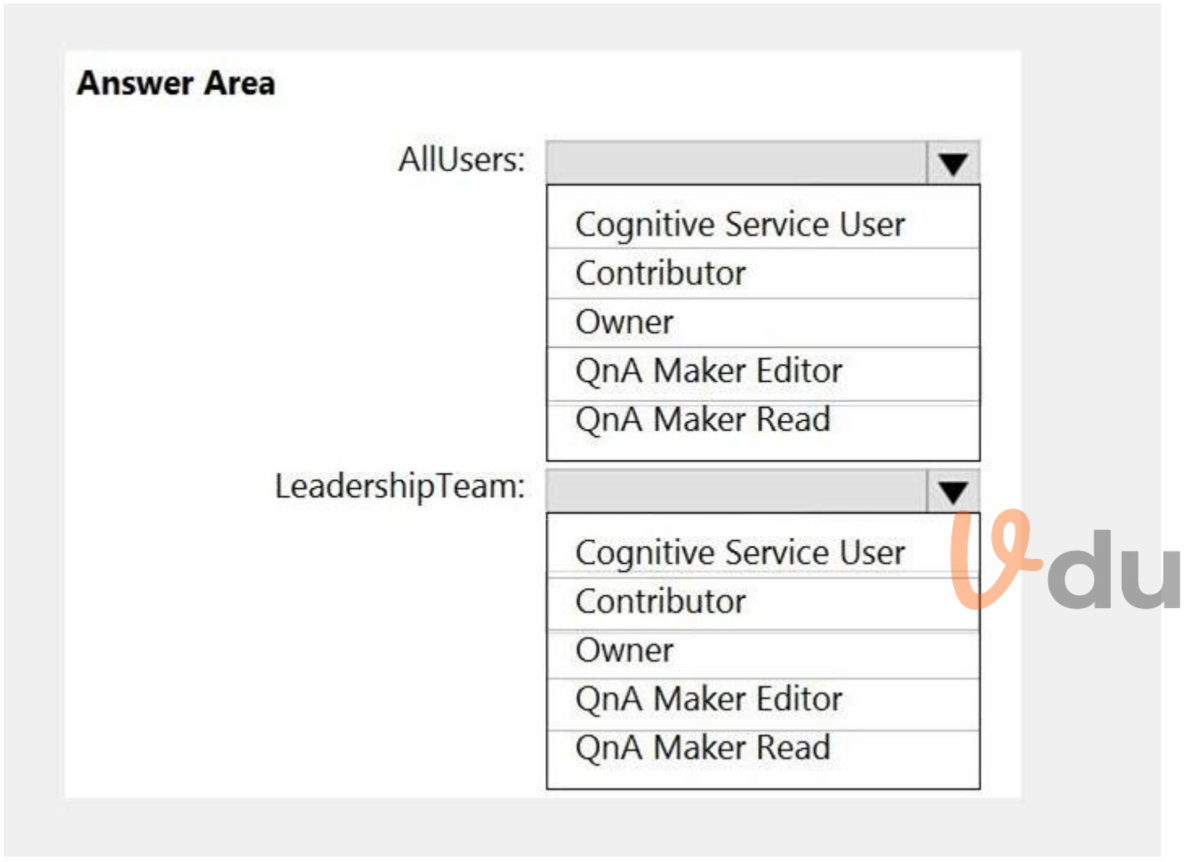

Answer Area:

Answer Area

Section:

Explanation:

Box 1: QnA Maker Editor

Scenario: Provide all employees with the ability to edit Q&As.

The QnA Maker Editor (read/write) has the following permissions:

Create KB API

Update KB API

Replace KB API

Replace Alterations

"Train API" [in new service model v5]

Box 2: Contributor

Scenario: Only senior managers must be able to publish updates.

Contributor permission: All except ability to add new members to roles

Reference:

https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/reference-role-based-access-control

QUESTION 3

HOTSPOT

You are developing the shopping on-the-go project.

You need to build the Adaptive Card for the chatbot.

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:

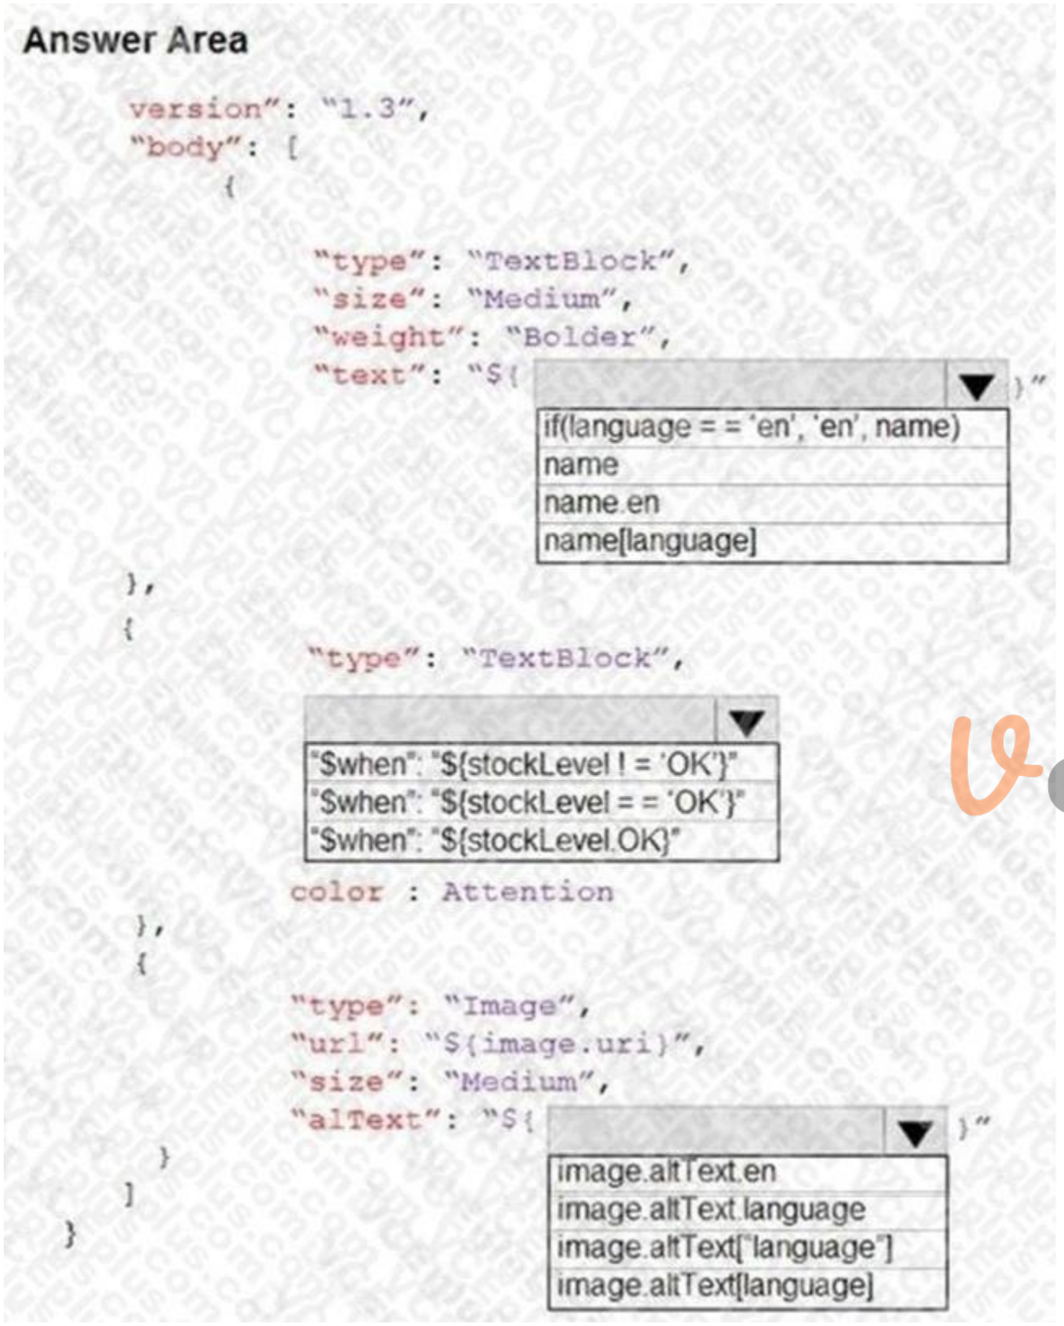

Answer Area:

Answer Area

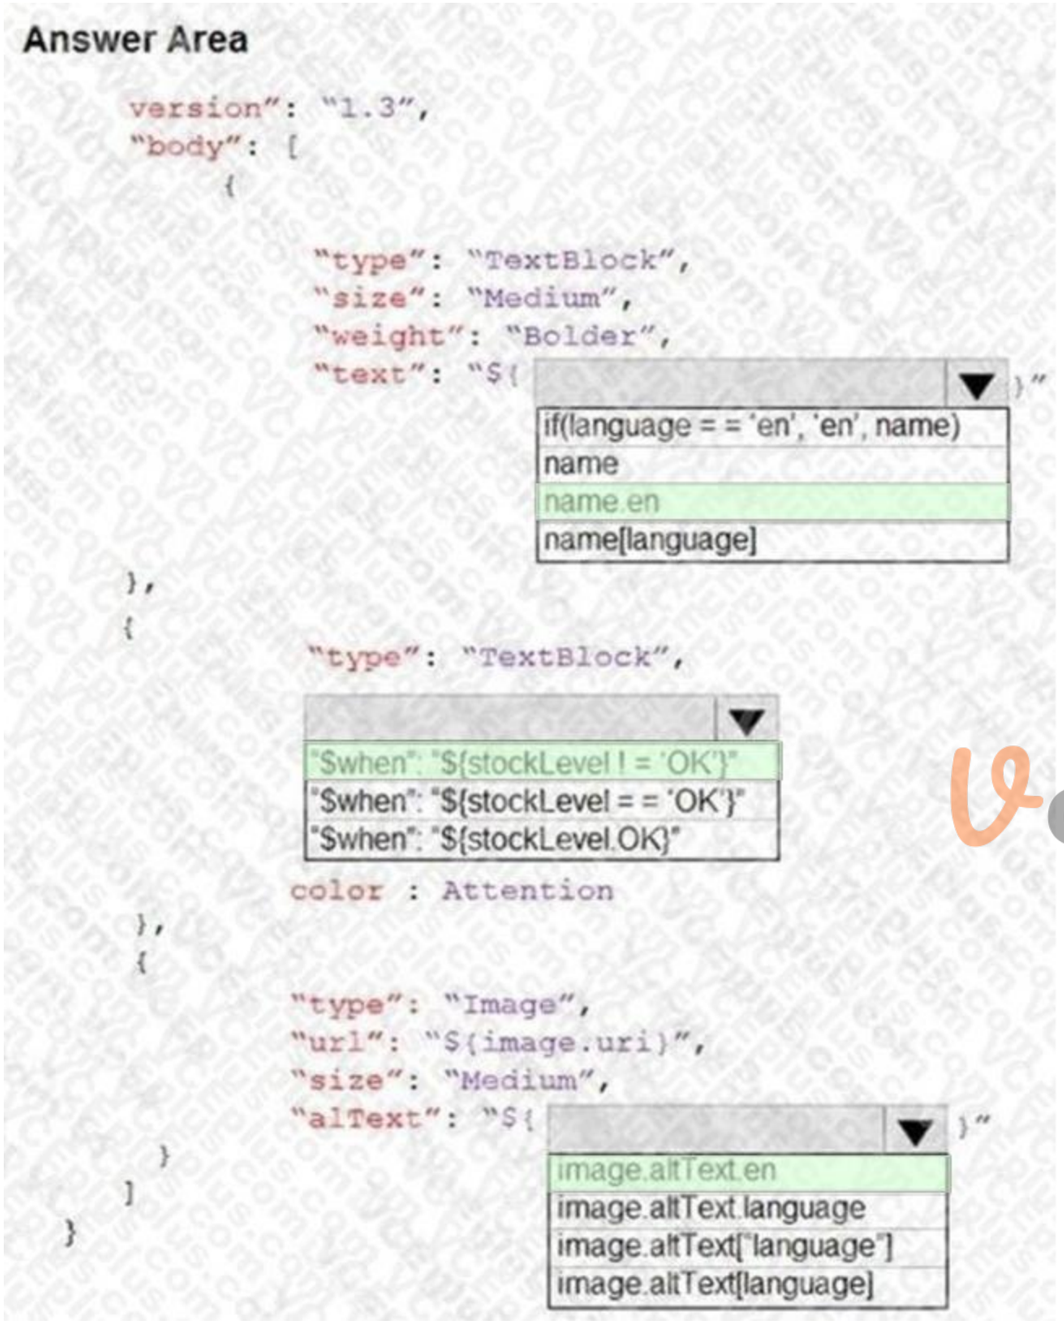

Section:
Explanation:
Box 1: name.en
Box 2: "$when": "${stockLevel != 'OK'}"
Product displays must include images and warnings when stock levels are low or out of stock.
Box 3:image.altText.en

Exam H

QUESTION 1
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 2
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:

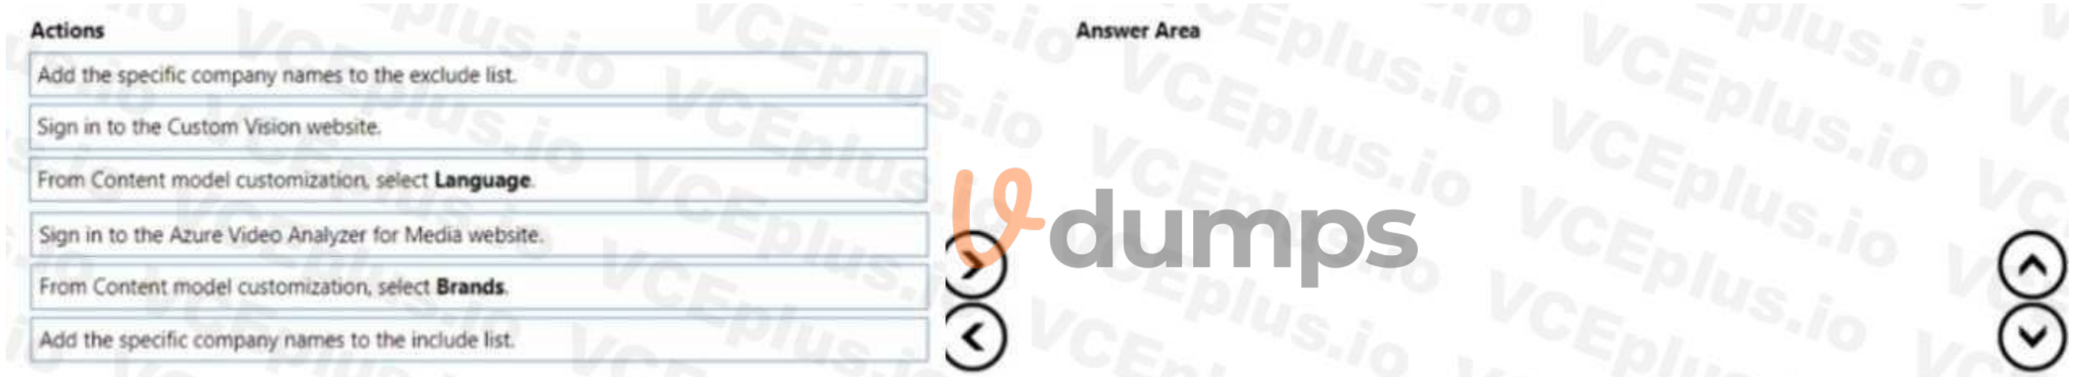

Correct Answer:

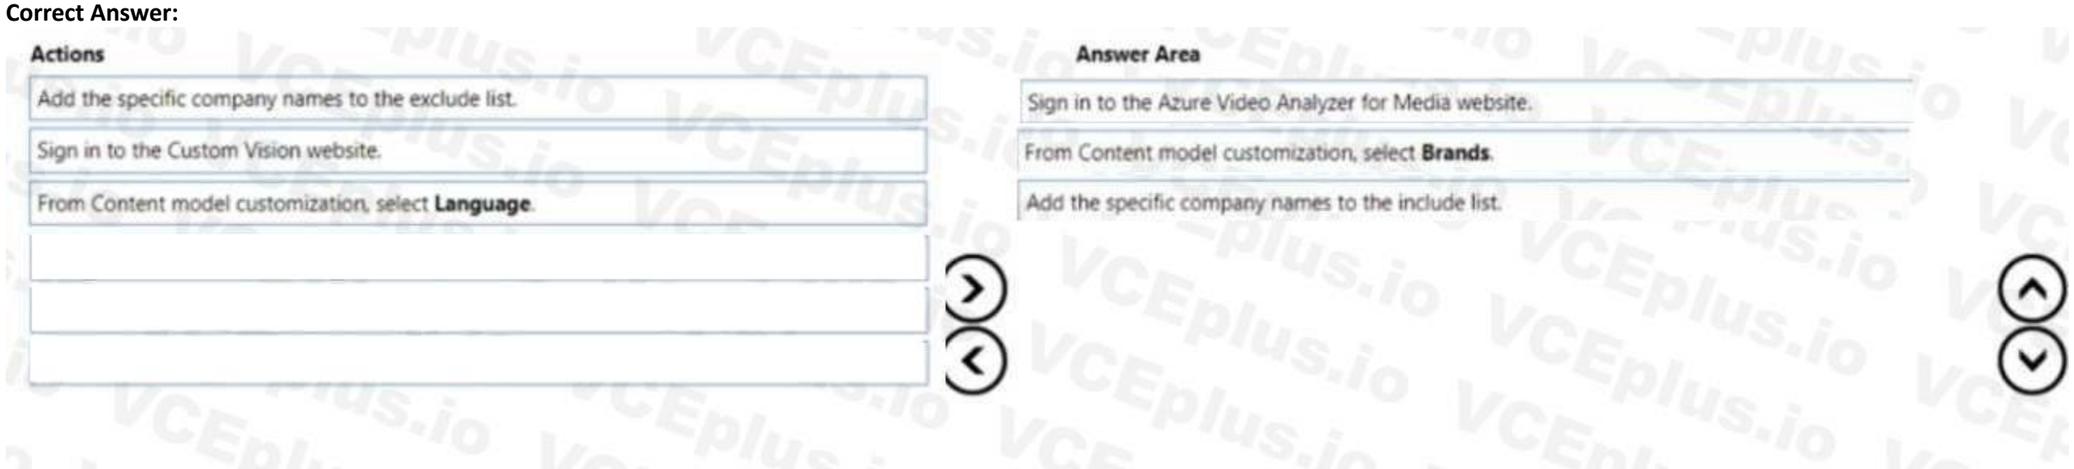

Section:
Explanation:

QUESTION 3
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 4
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

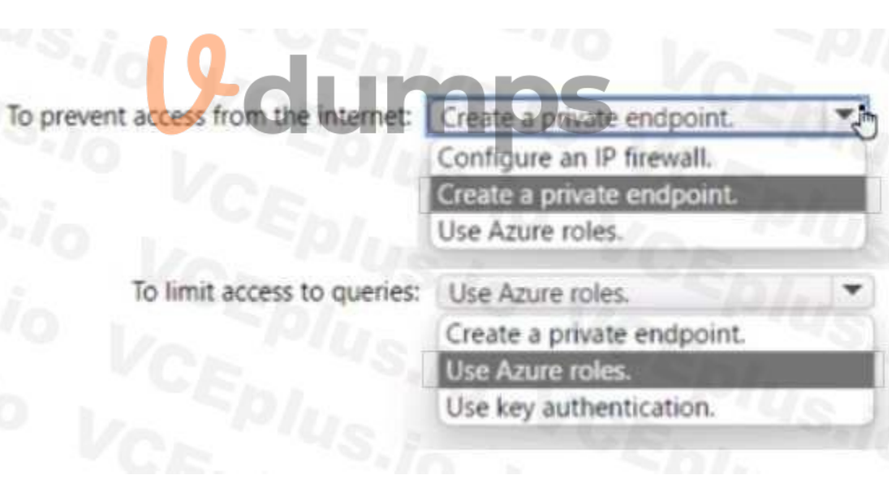

Answer Area:

Answer Area

Section:
Explanation:

QUESTION 5
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 6
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 7
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 8
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:

Correct Answer:

Section:
Explanation:

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

Hot Area:

Answer Area:

Answer Area

When provisioning an Azure Cosmos DB [account / container / database / item], you need to specify which type of API you will use.

Section:
Explanation:

QUESTION 10
Which statement is an example of Data Manipulation Language (DML)?

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

Correct Answer: B
Section:

QUESTION 11
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 12
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 13
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 14
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 15
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 16
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 17

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 18
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 19
You need to create a visualization of running sales totals per quarter as shown in the following exhibit.

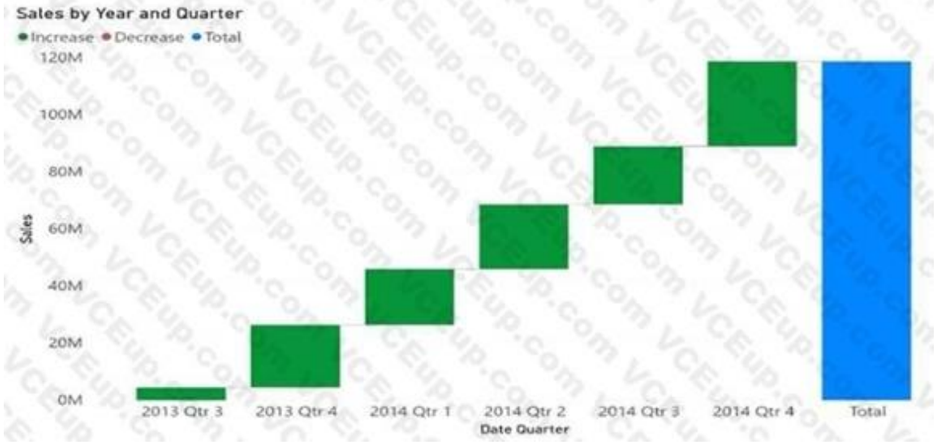

What should you create in Cower Bl Desktop;1

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

Correct Answer: C
Section:

QUESTION 20

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 21
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 22
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 23
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 24
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 25
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 26
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 27
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 28
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:

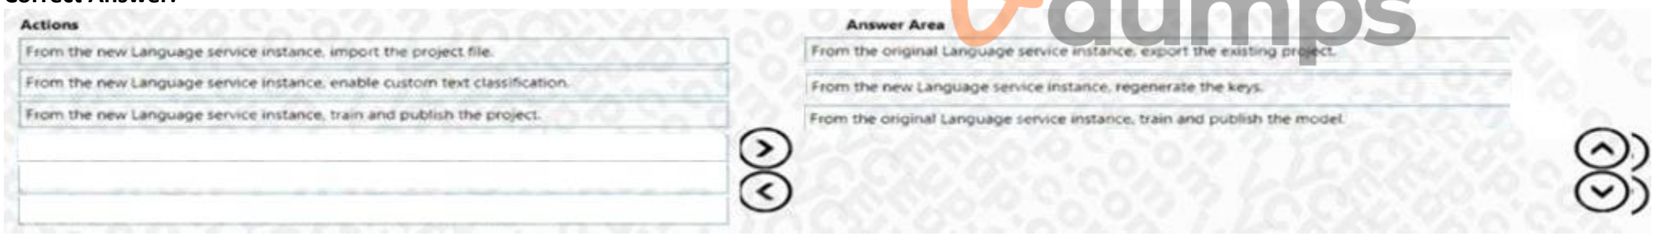

Section:
Explanation:

QUESTION 29
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:

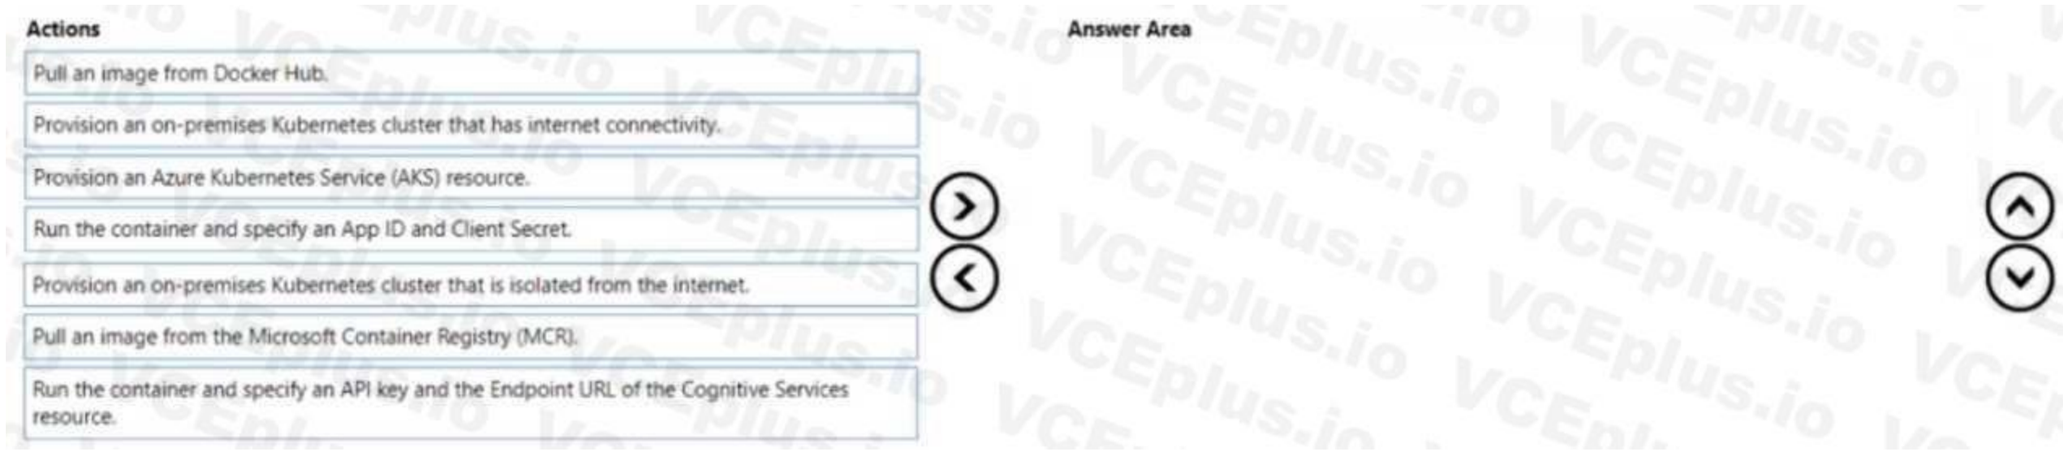

Correct Answer:

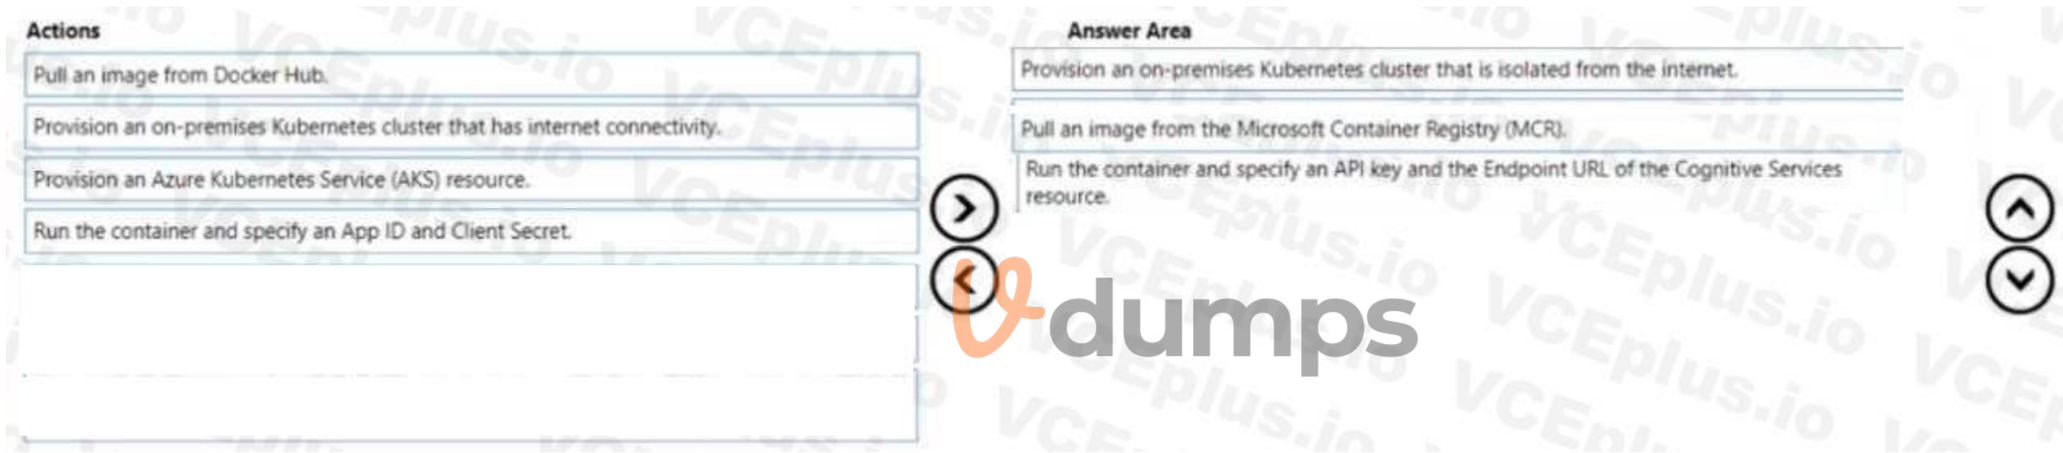

Section:

Explanation:

Provision an on-premises Kubernetes...
Pull an image from the Microsoft Container Registry...
Run the container and specify an API key...

QUESTION 30

You have a factory that produces food products.

You need to build a monitoring solution for staff compliance with personal protective equipment (PPE) requirements. The solution must meet the following requirements:

- identify staff who have removed masks or safety glasses.
- Perform a compliance check every 15 minutes.
- Minimize development effort.
- Minimize costs.

Which service should you use?

A. Face
B. Computer Vision
C. Azure Video Analyzer for Media (formerly Video indexer)

Correct Answer: A
Section:

QUESTION 31
HOTSPOT
You have an Azure subscription that has the following configurations:
• Subscription ID: 8d3591aa-96b8-4737-ad09-OOf9b1ed35ad
• Tenant ID: 3edfe572-cbS4-3ced-ae12-c5c177f39a12
You plan to create a resource that will perform sentiment analysis and optical character recognition (OCR). You need to use an HTTP request to create the resource in the subscription. The solution must use a single key and endpoint. How should you complete the request? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area:

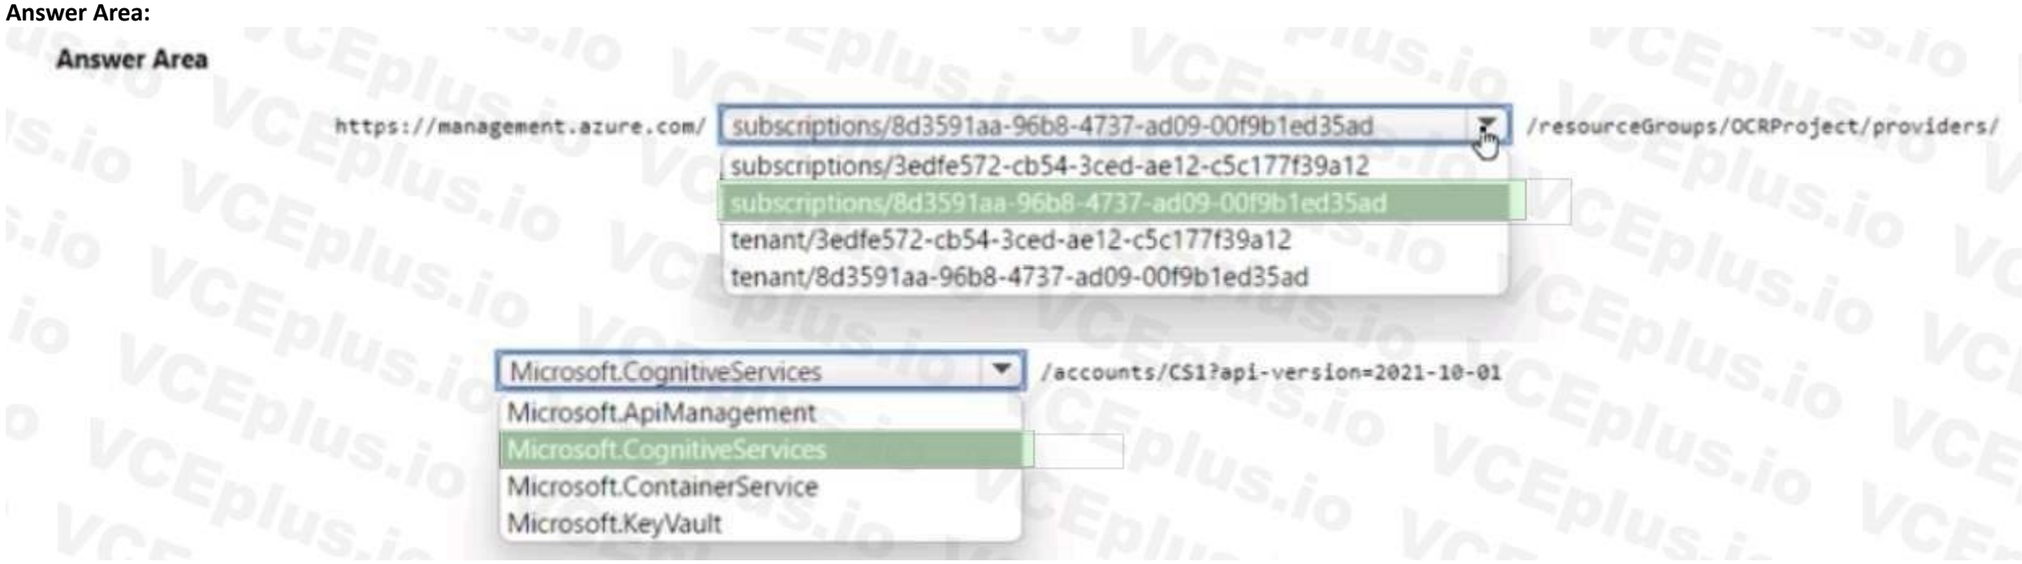

Section:
Explanation:

QUESTION 32
You have a text-based chatbot.
You need to enable content moderation by using the Text Moderation API of Content Moderator.
Which two service responses should you use? Each correct answer presents part of the solution NOTE: Each correct selection is worth one point.

A. the adult classification score
B. optical character recognition (OCR)
C. personal data
D. text classification
E. the racy classification score

Correct Answer: C, D
Section:

QUESTION 33
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 34
You are developing a system that will monitor temperature data from a data stream. The system must generate an alert in response to atypical values. The solution must minimize development effort.
What should you include in the solution?

A. Univariate Anomaly Detection
B. Azure Stream Analytics
C. metric alerts in Azure Monitor
D. Multivariate Anomaly Detection

Correct Answer: D
Section:

QUESTION 35
HOTSPOT
You are building a language learning solution.
You need to recommend which Azure services can be used to perform the following tasks:
* Analyze lesson plans submitted by teachers and extract key fields, such as lesson times and required texts.
* Analyze learning content and provide students with pictures that represent commonly used words or phrases in the text
The solution must minimize development effort.
Which Azure service should you recommend for each task? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

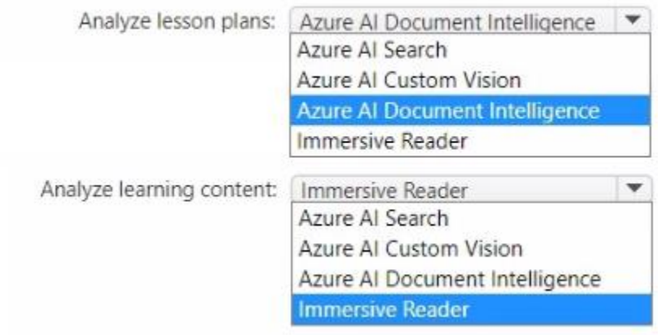

Answer Area:

Answer Area

Section:
Explanation:

QUESTION 36
You have a product knowledgebase that contains multiple PDF documents.
You need to build a chatbot that will provide responses based on data in the knowledgebase. The solution must minimize development effort and costs.
What should you include in the solution?

A. Azure Al Language conversational language understanding (CLU)
B. Azure Al language detection
C. Azure Al Language custom question answering
D. Azure OpenAI

Correct Answer: C
Section:

QUESTION 37
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:

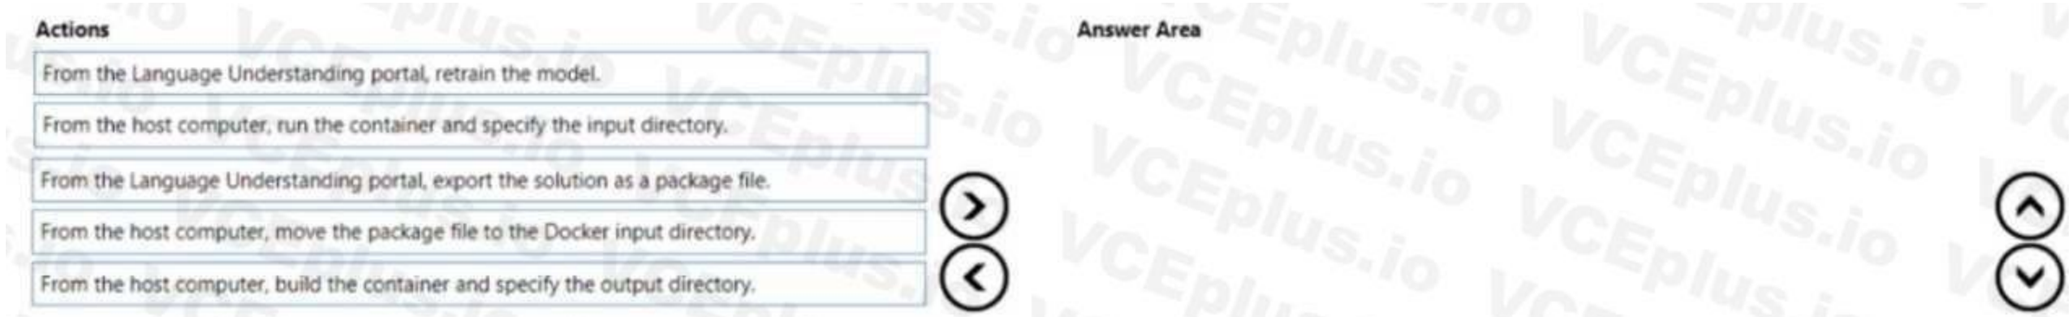

Correct Answer:

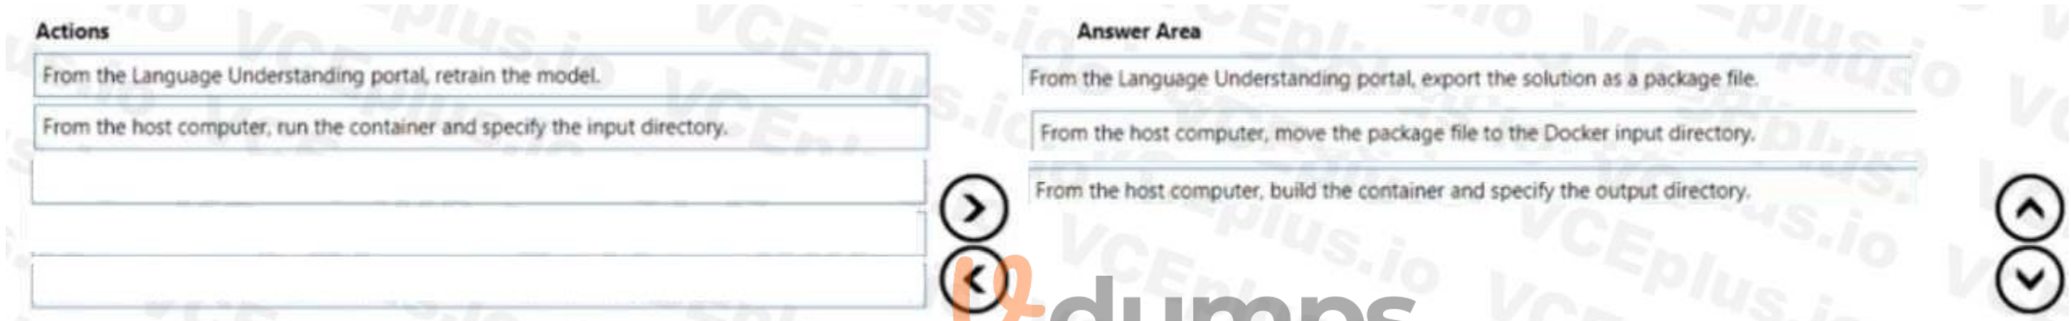

Section:
Explanation:

QUESTION 38
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 39
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 40
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 41
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 42
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 43
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 44
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 45
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 46
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 47
You have an Azure subscription that contains a Language service resource named ta1 and a virtual network named vnet1. You need to ensure that only resources in vnet1 can access ta1. What should you configure?

A. a network security group (NSG) for vnet1
B. Azure Firewall for vnet1
C. the virtual network settings for ta 1
D. a Language service container for ta1

Correct Answer: B
Section:

QUESTION 48
HOTSPOT
You have a collection of press releases stored as PDF files.
You need to extract text from the files and perform sentiment analysis.
Which service should you use for each task? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

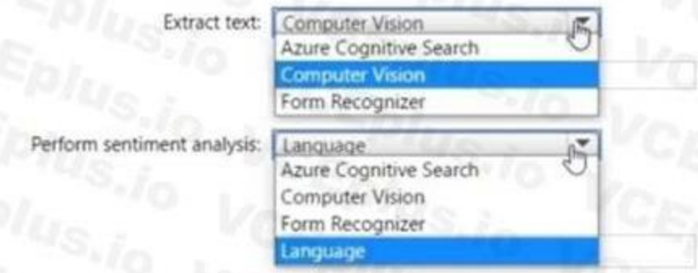

Answer Area:

Answer Area

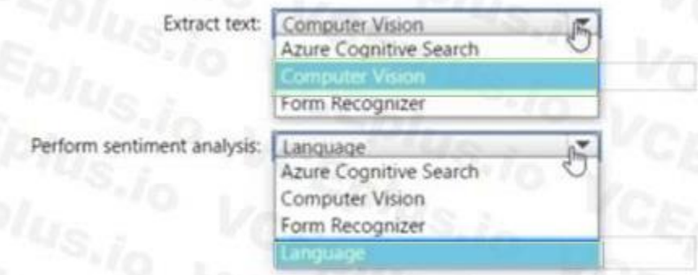

Section:
Explanation:

QUESTION 49
HOTSPOT
You are building a chatbot.
You need to use the Content Moderator service to identify messages that contain sexually explicit language. Which section in the response from the service will contain the category score, and which category will be assigned to the message? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

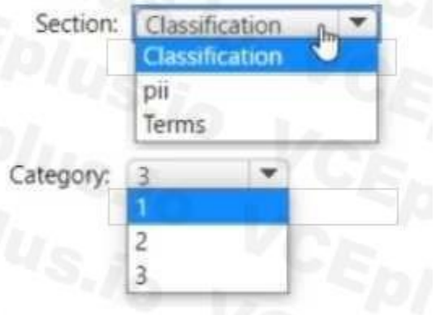

Answer Area:

Section:
Explanation:

QUESTION 50
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:

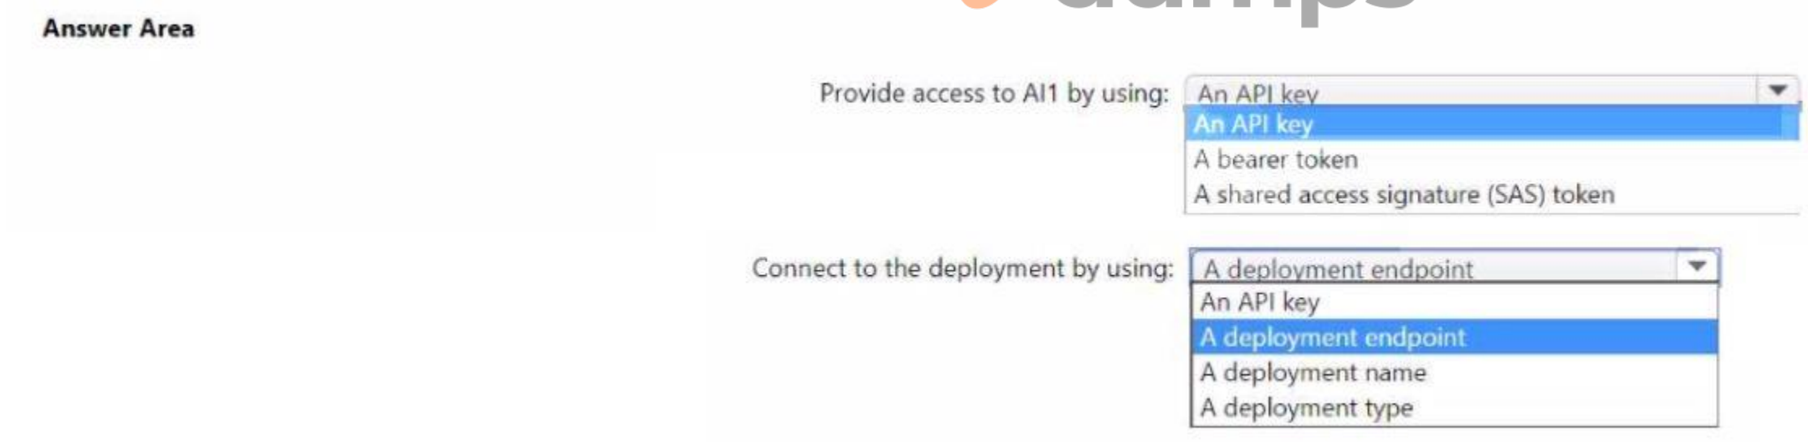

Answer Area:

Answer Area

Section:
Explanation:

QUESTION 51
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:

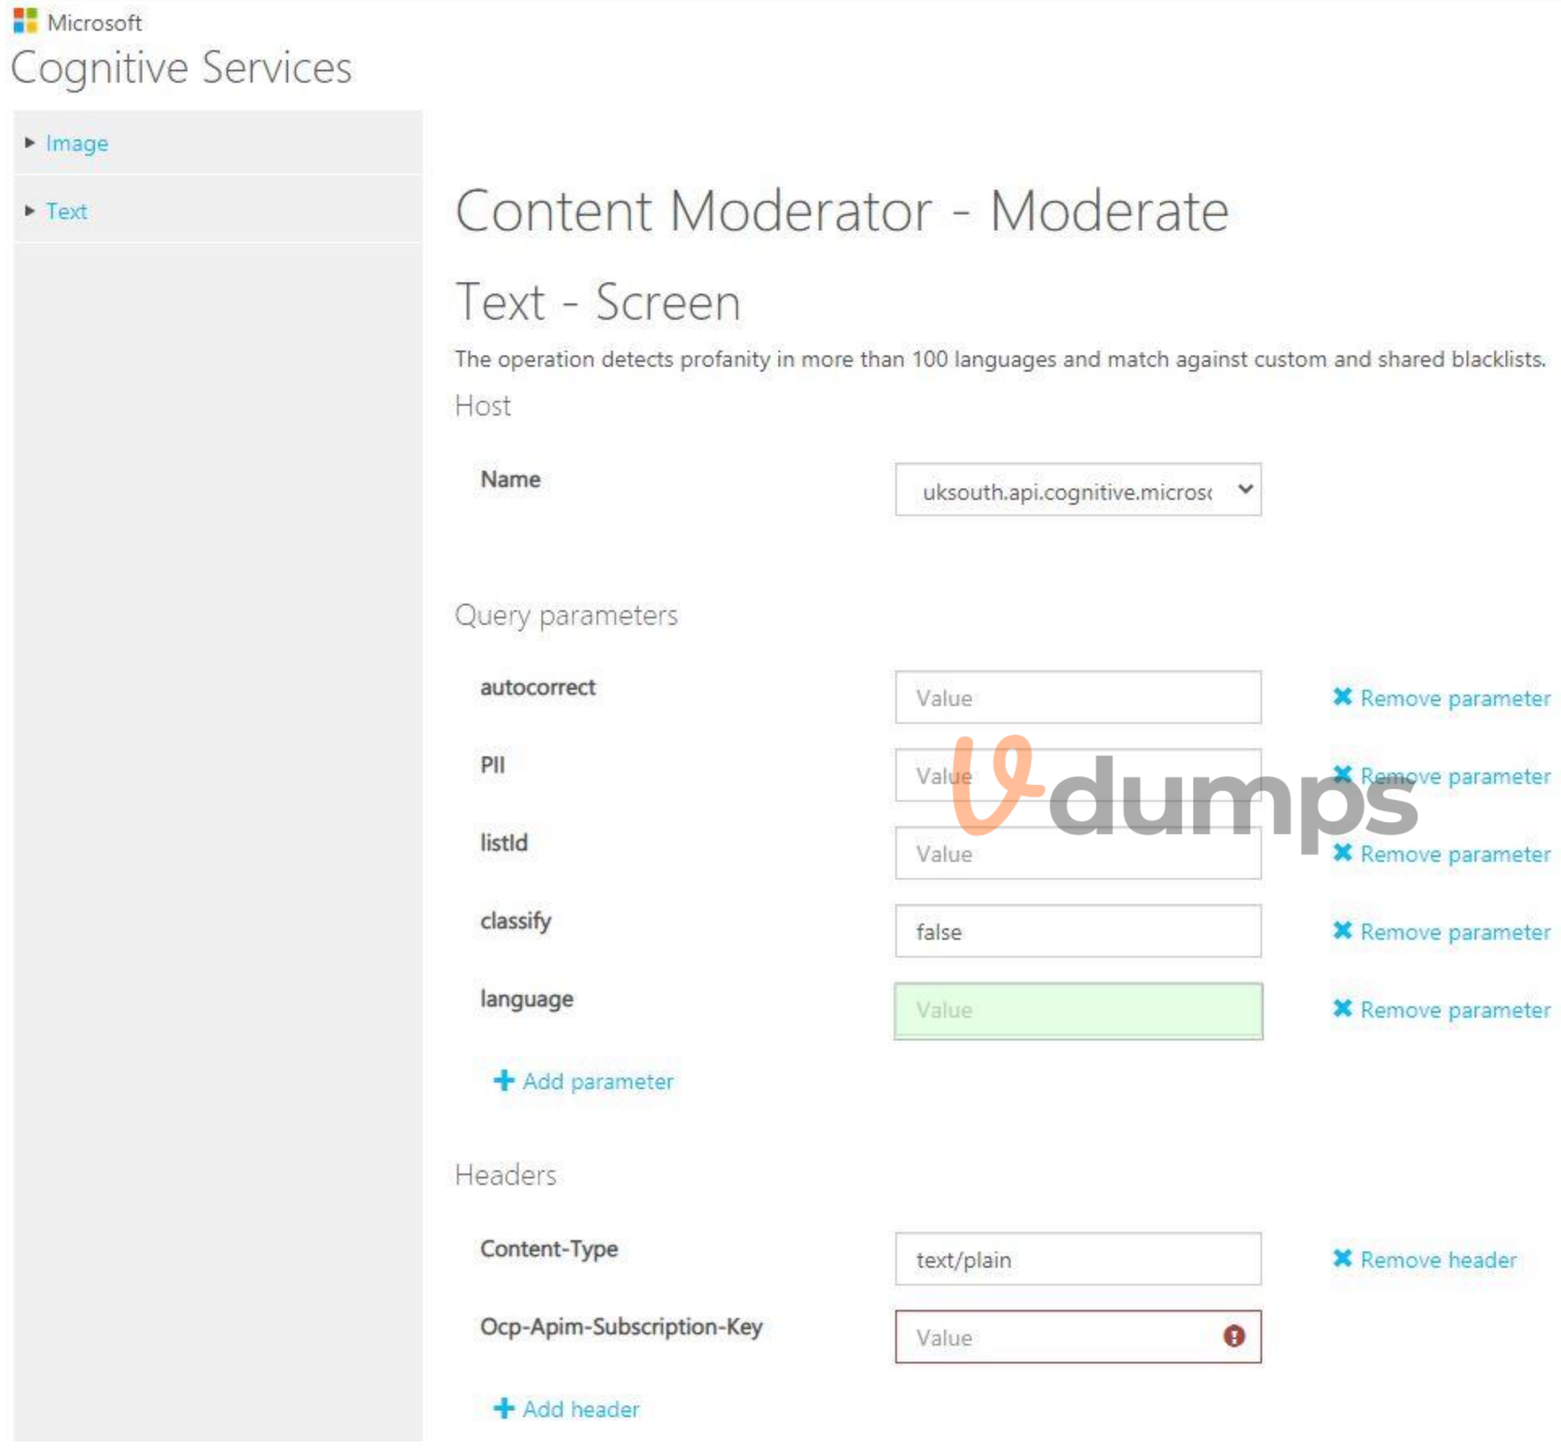

Section:
Explanation:

QUESTION 52

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 53
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" | "count" | "maxValue" | "size"

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" : 600

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

1
60
100
600

Section:
Explanation:

QUESTION 54
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 55
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 56
DRAG DROP
You have a monitoring solution that uses the Azure AI Anomaly Detector service.
You provision a server named Server! that has intermittent internet access.
You need to deploy the Azure AI Anomaly Detector to Server 1.
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:

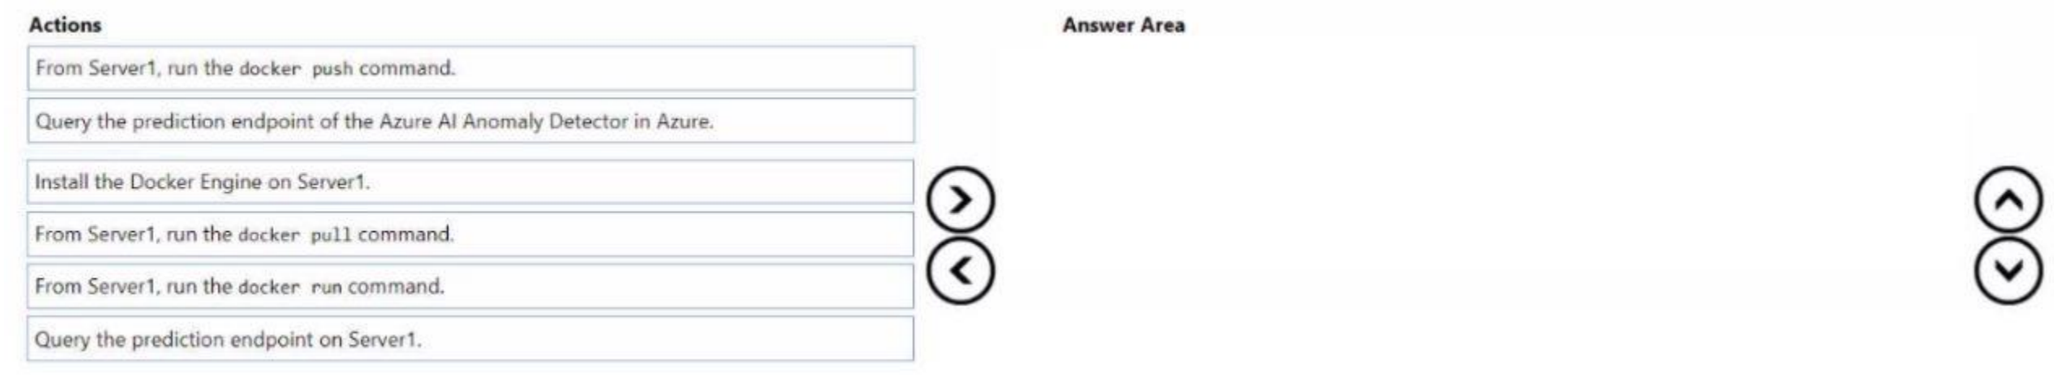

Correct Answer:

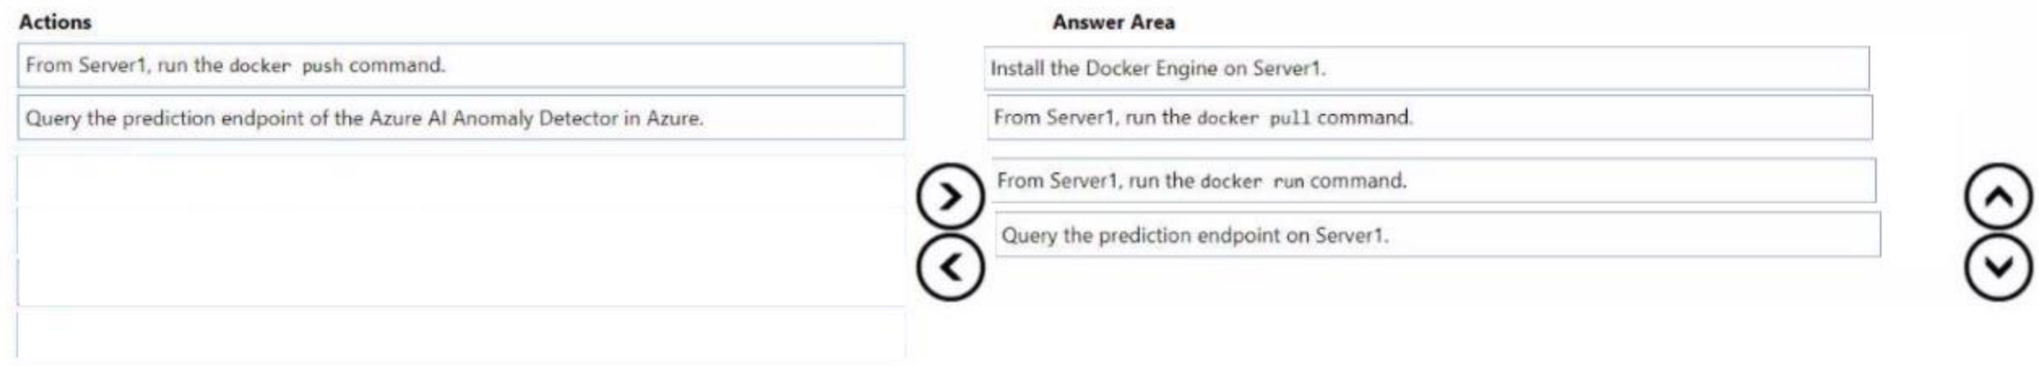

Section:
Explanation:
Install the Docker Engine on Server1.
From Server1, rund the docker pull command.
From Server1, run the docker run command.
Query the prediction endpoint on Server1.

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

QUESTION 57
You have an Azure Cognitive Search solution and a collection of handwritten letters stored as JPEG files.
You plan to index the collection. The solution must ensure that queries can be performed on the contents of the letters.
You need to create an indexer that has a skillset.
Which skill should you include?

A. key phrase extraction
B. optical character recognition (OCR)
C. document extraction
D. image analysis

Correct Answer: B
Section:

QUESTION 58
HOTSPOT
You plan to provision Azure Cognitive Services resources by using the following method.

```
{
  CognitiveServicesAccount parameters =
    new CognitiveServicesAccount(null, null, kind, location, name,
      new CognitiveServicesAccountProperties(), new Sku(tier));
  result = client.Accounts.Create(resource_group_name, tier, parameters);
}
```

You need to create a Standard tier resource that will convert scanned receipts into text.
How should you call the method? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

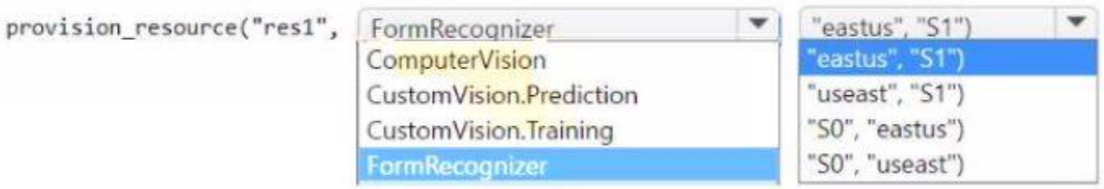

Answer Area:

Answer Area

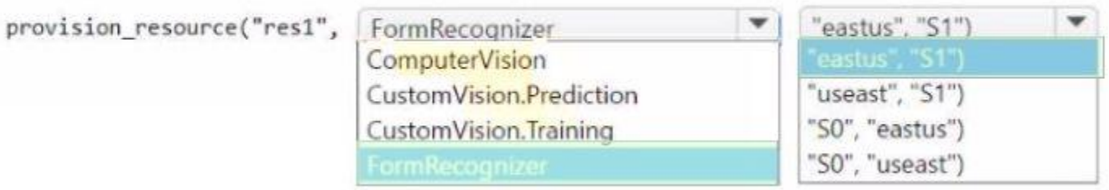

Section:
Explanation:

QUESTION 59
HOTSPOT
You have an Azure subscription that contains an Azure AI Document Intelligence resource named DM.
You build an app named App1 that analyzes PDF files for handwritten content by using DM.
You need to ensure that App1 will recognize the handwritten content.
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

```
Uri fileUri = new Uri("<fileUri>");
AnalyzeDocumentOperation operation = await client.AnalyzeDocumentFromUriAsync(WaitUntil.Completed, "prebuilt-document ", fileUri);
AnalyzeResult result = operation.Value;
foreach (DocumentStyle style in result.Styles)
{
 bool isHandwritten = style.IsHandwritten.HasValue && style.IsHandwritten == true;
  if (isHandwritten && style.Confidence > 0.75 )
 {

Console.WriteLine($"Handwritten content found:");
foreach (DocumentSpan span in style.Spans)
{
  Console.WriteLine($"  Content: {result.Content.Substring(span.Index, span.Length)}");
    }
  }
}
```

"prebuilt-document ", "prebuilt-contract", "prebuilt-read"

0.1, 0.75, 1.0

Answer Area:

Answer Area

```
Uri fileUri = new Uri("<fileUri>");
AnalyzeDocumentOperation operation = await client.AnalyzeDocumentFromUriAsync(WaitUntil.Completed, "prebuilt-document ", fileUri);
AnalyzeResult result = operation.Value;
foreach (DocumentStyle style in result.Styles)
{
 bool isHandwritten = style.IsHandwritten.HasValue && style.IsHandwritten == true;
  if (isHandwritten && style.Confidence > 0.75 )
 {
Console.WriteLine($"Handwritten content found:");
foreach (DocumentSpan span in style.Spans)
{
  Console.WriteLine($"  Content: {result.Content.Substring(span.Index, span.Length)}");
    }
  }
}
```

"prebuilt-document ",
"prebuilt-contract",
"prebuilt-read",

0.1
0.75
1.0

Section:
Explanation:

QUESTION 60
HOTSPOT
You have an Azure subscription.
You need to deploy an Azure AI Document Intelligence resource.
How should you complete the Azure Resource Manager (ARM) template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

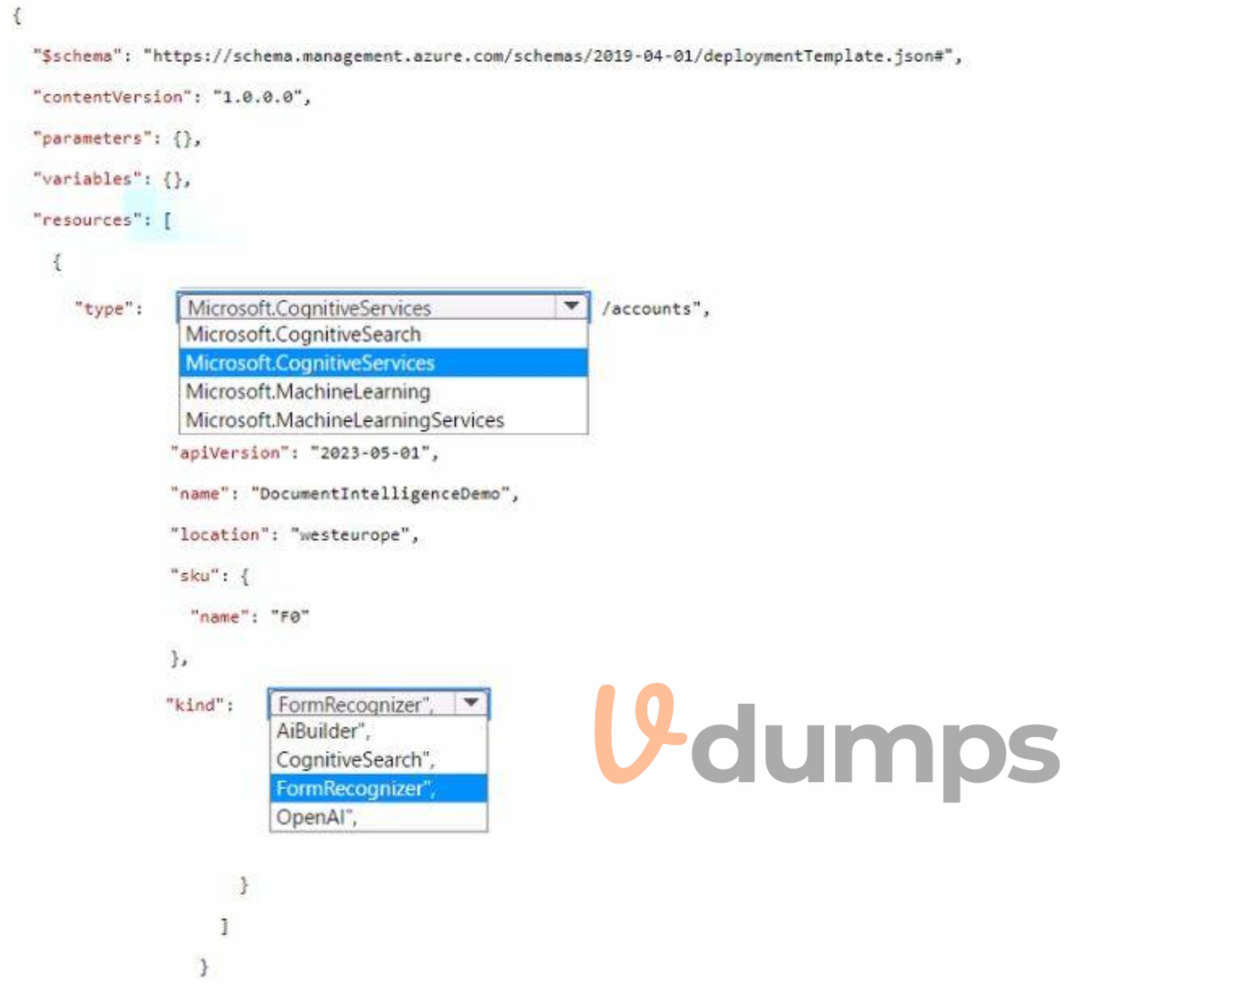

Answer Area:

Answer Area

Section:
Explanation:

QUESTION 61
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it as a result, these questions will not appear in the review screen.
You are building a chatbot that will use question answering in Azure Cognitive Service for Language.
You upload Doc1.pdf and train that contains a product catalogue and a price list.
During testing, users report that the chatbot responds correctly to the following question: What is the price of ?
The chatbot fails to respond to the following question: How much does cost?
You need to ensure that the chatbot responds correctly to both questions.
Solution: from Language Studio, you create an entity for price, and then retrain and republish the model.
Does this meet the goal?

A. Yes

B. No

Correct Answer: B
Section:

QUESTION 62
You need to ensure that the chatbot can classify user input into separate categories. The categones must be dynamic and defined at the time of inference.
Which service should you use to classify the input?

A. Azure OpenAI text summarization
B. Azure OpenAI text classification
C. Azure Al Language custom named entity recognition (NER)
D. Azure Al Language custom text classification

Correct Answer: D
Section:

QUESTION 63
You have an Azure subscription that contains an Azure OpenAI resource named OpenAI1 and a user named User1.
You need to ensure that User1 can upload datasets to OpenA11 and finetune the existing models. The solution must follow the principle of least privilege.
Which role should you assign to User1?

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

Correct Answer: C
Section:

QUESTION 64
DRAG DROP
You have an Azure subscription.
You are building a chatbot that will use an Azure OpenAI model.
You need to deploy the model.
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:

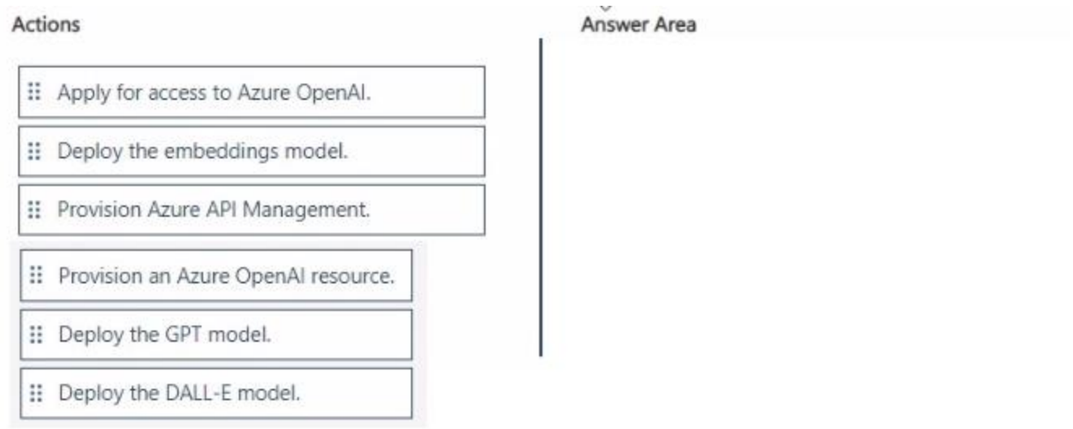

Correct Answer:

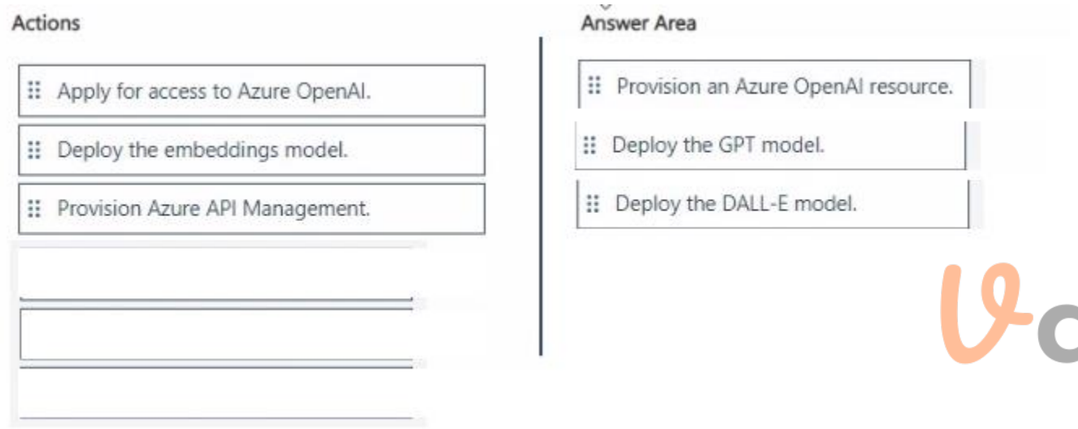

Section:
Explanation:

QUESTION 65
HOTSPOT
You are developing an application that will use the Azure AI Vision client library. The application has the following code.

```
def analyze_image(local_image):
  with open(local_image, "rb") as image_stream:
  image_analysis = client.analyze_image_in_stream(
    image=image_stream,
    visual_features=[
      VisualFeatureTypes.tags,
      VisualFeatureTypes.description
    ]
  )
  for caption in image_analysis.description.captions:
    print(f"\n{caption.text} with confidence {caption.confidence}")
  for tag in image_analysis.tags:
    print(f"\n{tag.name} with confidence {tag.confidence}")
```

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 an image 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 an image file from the local file system.	○	○

Section:
Explanation:

QUESTION 66
HOTSPOT
You have 1,000 scanned images of hand-written survey responses. The surveys do NOT have a consistent layout.
You have an Azure subscription that contains an Azure Al Document Intelligence resource named Aldoc1.
You open Document Intelligence Studio and create a new project.
You need to extract data from the survey responses. The solution must minimize development effort.
To where should you upload the images, and which type of model should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Upload to: An Azure Storage account
- An Azure Cosmos DB account
- An Azure Files share
- An Azure Storage account

Model type: Identity document (ID)
- Custom neural
- Custom template
- Identity document (ID)

Answer Area:

Answer Area

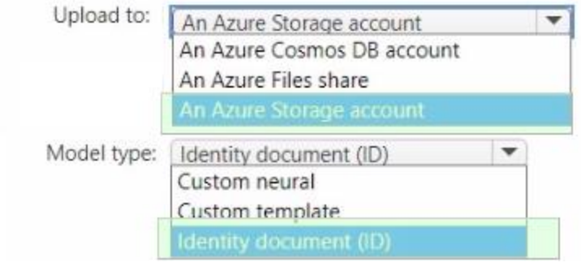

Section:
Explanation: