

Microsoft.AI-102.vDec-2023.by.Karex.92q

Number: AI-102  
Passing Score: 800  
Time Limit: 120  
File Version: 12.0

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



## 01 - Implement Computer Vision Solutions

### QUESTION 1

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.microsoft.com/en-us/azure/media-services/videointer/customize-person-model-with-api>



### QUESTION 2

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 3

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 4

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 5**

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	Answer Area
detect	POST {Endpoint}/face/v1.0/ <input type="text"/>
findsimilars	Request Body
group	{
identify	"faceId": "c5c24a82-6845-4031-9d5d-978df9175426",
matchFace	"largeFaceListId": "sample_list",
matchPerson	"maxNumOfCandidatesReturned": 10,
verify	"mode": " <input type="text"/> "
	}

Correct Answer:

Values	Answer Area
<input type="text"/>	POST {Endpoint}/face/v1.0/ <input type="text" value="detect"/>
findsimilars	Request Body
group	{
identify	"faceId": "c5c24a82-6845-4031-9d5d-978df9175426",
matchFace	"largeFaceListId": "sample_list",
<input type="text"/>	"maxNumOfCandidatesReturned": 10,
verify	"mode": " <input type="text" value="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 6

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:

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

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.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
The code will return coordinates for the top-left corner of the rectangle that contains the brand logo of the displayed brands.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.	<input type="checkbox"/>	<input checked="" type="checkbox"/>

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 7

HOTSPOT

You develop an application that uses the Face API.

You need to add multiple images to a person group.

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

NOTE: Each correct selection is worth one point.

Hot Area:

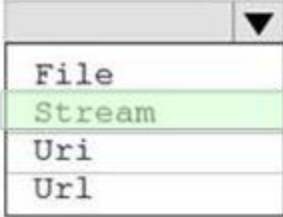
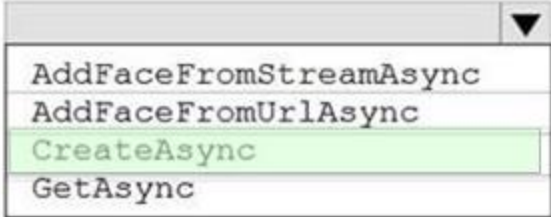
### Answer Area

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

 Vdumps

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);
        }
    }
});
```



#### 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 8

##### 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.	<input type="radio"/>	<input type="radio"/>
	The code will list tags and their associated confidence.	<input type="radio"/>	<input type="radio"/>
	The code will read a file from the local file system.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
The code will perform face recognition.	<input type="radio"/>	<input checked="" type="radio"/>
The code will list tags and their associated confidence.	<input checked="" type="radio"/>	<input type="radio"/>
The code will read a file from the local file system.	<input type="radio"/>	<input checked="" type="radio"/>

**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 9**

**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**

```
curl -H "Ocp-Apim-Subscription-Key: xxx" /
-o "sample.png" -H "Content-Type: application/json" /
?width=100&height=100&smartCropping=true" /
-d "{\"url\": \"https://upload.litwareinc.org/litware/bicycle.jpg\"}"
```

/vision/v3.1/

**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 10**

**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:**

Values	Answer Area
en-US	
false	
people,keywords	
people,search	
search	
true	
	<p><b>Cognitive Insights Widget</b></p> <p>https://www.videoindexer.ai/embed/insights/&lt;accountId&gt;/&lt;videoId&gt;/?widgets= [Value] controls= [Value]</p>
	<p><b>Player Widget</b></p> <p>https://www.videoindexer.ai/embed/player/&lt;accountId&gt;/&lt;videoId&gt;/? showcaptions= [Value] captions= [Value]</p>

**Correct Answer:**

Values	Answer Area
<input type="text" value="false"/>	<b>Cognitive Insights Widget</b> https://www.videoindexer.ai/embed/insights/<accountId>/<videoId>/?widgets= <input type="text" value="people,keywords"/> controls= <input type="text" value="search"/>
<input type="text" value="people,search"/>	<b>Player Widget</b> https://www.videoindexer.ai/embed/player/<accountId>/<videoId>/? showcaptions= <input type="text" value="true"/> captions= <input type="text" value="en-US"/>
<input type="text"/>	
<input type="text"/>	

**Section:**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-video-analyzer/video-analyzer-for-media-docs/video-indexer-embed-widgets>

**QUESTION 11**

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:**

**Actions**

- 
- 
- 
- 

**Answer Area**



**Correct Answer:**

### Actions

### Answer Area

  
  
  
  
◀ ▶ ⬆ ⬇

#### Section:

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/export-your-model>

#### QUESTION 12

##### 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.	<input type="radio"/>	<input type="radio"/>
The code will work for a group of 10,000 people.	<input type="radio"/>	<input type="radio"/>
AddFace can be called multiple times to add multiple face images to a person object.	<input type="radio"/>	<input type="radio"/>

Answer Area:

**Answer Area**

Statements	Yes	No
The code will add a face image to a person object in a person group.	<input checked="" type="radio"/>	<input type="radio"/>
The code will work for a group of 10,000 people.	<input type="radio"/>	<input checked="" type="radio"/>
AddFace can be called multiple times to add multiple face images to a person object.	<input checked="" type="radio"/>	<input type="radio"/>

**Section:**

**Explanation:**

Reference:

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

**QUESTION 13**

DRAG DROP

You have a Custom Vision resource named acvdev in a development environment.

You have a Custom Vision resource named acvprod in a production environment.

In acvdev, you build an object detection model named obj1 in a project named proj1.

You need to move obj1 to acvprod.

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

**Select and Place:**



**Actions**

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

**Answer Area**

Navigation icons: Left arrow, Right arrow, Up arrow, Down arrow.

**Correct Answer:**

**Actions**

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

**Answer Area**

Answer area containing three items in order:
 

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

 Navigation icons: Left arrow, Right arrow, Up arrow, Down arrow.

**Section:**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/copy-move-projects>

**QUESTION 14**

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:**

### Actions

- Train the classifier model.
- Upload and tag images.
- Initialize the training dataset.
- Train the object detection model.
- Create a project.

### Answer Area



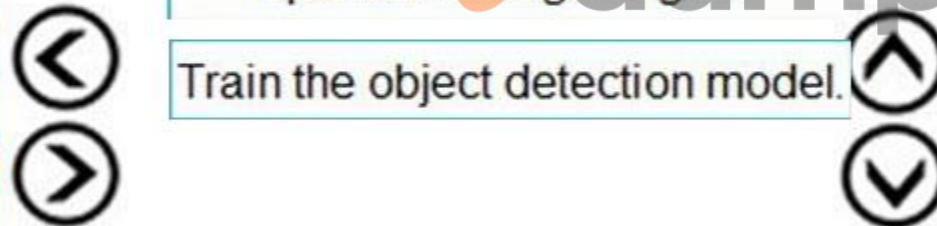
Correct Answer:

### Actions

- Train the classifier model.
- 
- Initialize the training dataset.
- 
- 

### Answer Area

- Create a project.
- Upload and tag images.
- Train the object detection model.



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 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)

 Vdumps

Answer 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)

vdumps

### 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

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**

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.

**Answer Area**

**Correct Answer:**

## Actions

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.
Upload the video to file storage.

## Answer Area

Upload the video to blob storage.
Index the video by using the Video Indexer API.
Extract the transcript from the Video Indexer API.
Translate the transcript by using the Translator API.



### 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>

### QUESTION 2

#### 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**

```
public static async Task<string> SuggestAltText(ComputerVisionClient client,
{
    List<VisualFeatureTypes?> features = new List<VisualFeatureTypes?>()
    {
        VisualFeatureTypes.Description
        VisualFeatureTypes.ImageType
        VisualFeatureTypes.Objects
        VisualFeatureTypes.Tags
    };
    ImageAnalysis results = await client.AnalyzeImageAsync(image, features);
    var c = results.Brands.DetectedBrands[0]
    var c = results.Description.Captions[0]
    var c = results.Metadata[0]
    var c = results.Objects[0]

    if(c.Confidence>0.5) return(c.Text);
}
```

image

Dictionary
stream
string

VisualFeatureTypes.Description
VisualFeatureTypes.ImageType
VisualFeatureTypes.Objects
VisualFeatureTypes.Tags

var c = results.Brands.DetectedBrands[0]
var c = results.Description.Captions[0]
var c = results.Metadata[0]
var c = results.Objects[0]

Answer Area:



**Answer Area**

```
public static async Task<string> SuggestAltText(ComputerVisionClient client,
{
    List<VisualFeatureTypes?> features = new List<VisualFeatureTypes?>()
    {
        VisualFeatureTypes.Description
        VisualFeatureTypes.ImageType
        VisualFeatureTypes.Objects
        VisualFeatureTypes.Tags
    };
    ImageAnalysis results = await client.AnalyzeImageAsync(image, features);
    var c = results.Brands.DetectedBrands[0]
    var c = results.Description.Captions[0]
    var c = results.Metadata[0]
    var c = results.Objects[0]

    if(c.Confidence>0.5) return(c.Text);
}
```

image

Dictionary
stream
string

VisualFeatureTypes.Description
VisualFeatureTypes.ImageType
VisualFeatureTypes.Objects
VisualFeatureTypes.Tags

var c = results.Brands.DetectedBrands[0]
var c = results.Description.Captions[0]
var c = results.Metadata[0]
var c = results.Objects[0]

Section:

Explanation:

Reference:

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

### 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.

Contoso identifies the following requirements for the chatbot:

Provide customers with answers to the FAQs.

Ensure that the customers can chat to a customer service agent.

Ensure that the members of a group named Management-Accountants can approve the FAQs.

Ensure that the members of a group named Consultant-Accountants can create and amend the FAQs.

Ensure that the members of a group named the Agent-CustomerServices can browse the FAQs.

Ensure that access to the customer service agents is managed by using Omnichannel for Customer Service. When the response confidence score is low, ensure that the chatbot can provide other response options to the customers.

#### Document Processing Requirements

Contoso identifies the following requirements for document processing:

The document processing solution must be able to process standardized financial documents that have the following characteristics:

- Contain fewer than 20 pages.

- Be formatted as PDF or JPEG files.

- Have a distinct standard for each office.

The document processing solution must be able to extract tables and text from the financial documents. The document processing solution must be able to extract information from receipt images.

Members of a group named Management-Bookkeeper must define how to extract tables from the financial documents. Members of a group named Consultant-Bookkeeper must be able to process the financial documents.

#### Knowledgebase Requirements

Contoso identifies the following requirements for the knowledgebase:

- Supports searches for equivalent terms

- Can transcribe jargon with high accuracy

- Can search content in different formats, including video

- Provides relevant links to external resources for further research



## QUESTION 1



You are developing the 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 2

### 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 3

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 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 Natural Language Processing Solutions

##### QUESTION 1

You are building a Language Understanding model for an e-commerce platform.

You need to construct an entity to capture billing addresses.

Which entity type should you use for the billing address?

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

**Correct Answer: B**

**Section:**



**Explanation:**

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

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

Incorrect Answers:

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

The geographical locations have subtypes:

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

D: Pattern.any is a variable-length placeholder used only in a pattern's template utterance to mark where the entity begins and ends.

E: A list entity represents a fixed, closed set of related words along with their synonyms. You can use list entities to recognize multiple synonyms or variations and extract a normalized output for them. Use the recommend option to see suggestions for new words based on the current list.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-entity-types>

**QUESTION 2**

You need to upload speech samples to a Speech Studio project for use in training.

How should you upload the samples?

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

**Correct Answer: B**

**Section:**

**Explanation:**

To upload your data, navigate to the Speech Studio . From the portal, click Upload data to launch the wizard and create your first dataset. You'll be asked to select a speech data type for your dataset, before allowing you to upload your data.

The default audio streaming format is WAV

Use this table to ensure that your audio files are formatted correctly for use with Custom Speech:



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

Reference:

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

### QUESTION 3

You are developing a method for an application that uses the Translator API.

The method will receive the content of a webpage, and then translate the content into Greek (el). The result will also contain a transliteration that uses the Roman alphabet.

You need to create the URI for the call to the Translator API.

You have the following URI.

<https://api.cognitive.microsofttranslator.com/translate?api-version=3.0>

Which three additional query parameters should you include in the URI? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

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

**Correct Answer: A, D, F**

**Section:**

**Explanation:**

### QUESTION 4

You have a chatbot that was built by using the Microsoft Bot Framework.

You need to debug the chatbot endpoint remotely.

Which two tools should you install on a local computer? Each correct answer presents part of the solution. (Choose two.) NOTE: Each correct selection is worth one point.

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

- D. Bot Framework CLI
- E. ngrok
- F. nginx

**Correct Answer: C, E**

**Section:**

**Explanation:**

Explanation/Reference:

Bot Framework Emulator is a desktop application that allows bot developers to test and debug bots, either locally or remotely. ngrok is a cross-platform application that "allows you to expose a web server running on your local machine to the internet." Essentially, what we'll be doing is using ngrok to forward messages from external channels on the web directly to our local machine to allow debugging, as opposed to the standard messaging endpoint configured in the Azure portal.

Reference:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-service-debug-emulator>

#### QUESTION 5

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

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

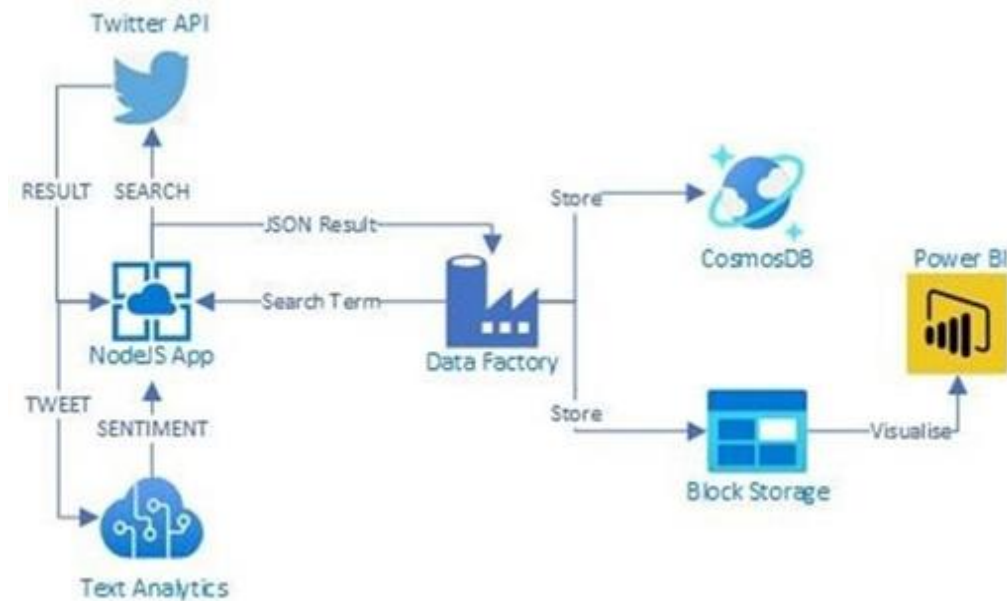
**Correct Answer: A**

**Section:**

**Explanation:**

Text Analytics Cognitive Service could be used to quickly determine the public perception for a specific topic, event or brand.

Example: A NodeJS app which pulls Tweets from Twitter using the Twitter API based on a specified search term. Then pass these onto Text Analytics for sentiment scoring before storing the data and building a visualisation in PowerBI. The Architecture looked something like this:



Reference:

<https://www.linkedin.com/pulse/measuring-public-perception-azure-cognitive-services-steve-dalai>

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/overview>

#### QUESTION 6



Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training.

Find contacts in London.

Who do I know in Seattle?

Search for contacts in Ukraine.

You need to implement the phrase list in Language Understanding.

Solution: You create a new intent for location.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: A**

**Section:**

**Explanation:**

An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance. Define a set of intents that corresponds to actions users want to take in your application.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-intent>

#### QUESTION 7

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training.

Find contacts in London.

Who do I know in Seattle?

Search for contacts in Ukraine.

You need to implement the phrase list in Language Understanding.

Solution: You create a new entity for the domain.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance. Define a set of intents that corresponds to actions users want to take in your application.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-intent>

#### QUESTION 8

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training.

Find contacts in London.

Who do I know in Seattle?

Search for contacts in Ukraine.

You need to implement the phrase list in Language Understanding.

Solution: You create a new pattern in the FindContact intent.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance. Define a set of intents that corresponds to actions users want to take in your application.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-intent>

#### QUESTION 9

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You develop an application to identify species of flowers by training a Custom Vision model.

You receive images of new flower species.

You need to add the new images to the classifier.

Solution: You add the new images, and then use the Smart Labeler tool.

Does this meet the goal?

A. Yes

B. No

**Correct Answer: B**

**Section:**

**Explanation:**

The model need to be extended and retrained.

Note: Smart Labeler to generate suggested tags for images. This lets you label a large number of images more quickly when training a Custom Vision model.

#### QUESTION 10

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You develop an application to identify species of flowers by training a Custom Vision model.

You receive images of new flower species.

You need to add the new images to the classifier.

Solution: You add the new images and labels to the existing model. You retrain the model, and then publish the model.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section:**

**Explanation:**

The model needs to be extended and retrained.

#### QUESTION 11

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You develop an application to identify species of flowers by training a Custom Vision model.

You receive images of new flower species.

You need to add the new images to the classifier.

Solution: You create a new model, and then upload the new images and labels.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

**Explanation:**

Explanation/Reference:

The model needs to be extended and retrained.



#### QUESTION 12

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

Which entity type should you use?

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

**Correct Answer: B**

**Section:**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-entity-types>

#### QUESTION 13

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

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

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

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

**Correct Answer: A, C**

**Section:**

**Explanation:**

Reference: <https://docs.microsoft.com/en-us/rest/api/cognitiveservices-qnamaker/qnamaker4.0/knowledgebase/create>

#### QUESTION 14

You are building a natural language model.

You need to enable active learning.

What should you do?

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

**Correct Answer: C**

**Section:**

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-how-to-review-endpoint-utterances#log-user-queries-to-enable-active-learning>

#### QUESTION 15

DRAG DROP

You are building a retail chatbot that will use a QnA Maker service.

You upload an internal support document to train the model. The document contains the following question: "What is your warranty period?"

Users report that the chatbot returns the default QnA Maker answer when they ask the following question: "How long is the warranty coverage?"

The chatbot returns the correct answer when the users ask the following question: "What is your warranty period?"

Both questions should return the same answer.

You need to increase the accuracy of the chatbot responses.

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

**Select and Place:**

## Actions

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

## Answer Area

Correct Answer:

## Actions

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

## Answer Area

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

**Section:**

**Explanation:**

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

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

Step 2: Retrain the model.

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

Step 3: Republish the model

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

Reference:

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

**QUESTION 16**



HOTSPOT

You are developing a service that records lectures given in English (United Kingdom).

You have a method named AppendToTranscriptFile that takes translated text and a language identifier.

You need to develop code that will provide transcripts of the lectures to attendees in their respective language. The supported languages are English, French, Spanish, and German.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
static async Task TranslateSpeechAsync()  
{  
    var config =SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");  
  
    var lang = new List<string>  
  
    config.SpeechRecognitionLanguage = "en-GB";  
    lang.ForEach(config.AddTargetLanguage);  
  
    using var audioConfig = AudioConfig.FromDefaultMicrophoneInput();  
    using var recognizer = new  
  
    var result = await recognizer.RecognizeOnceAsync();  
    if (result.Reason == ResultReason.TranslatedSpeech)
```

["en-GB"]
["fr", "de", "es"]
["French", "Spanish", "German"]
{languages}

IntentRecognizer
SpeakerRecognizer
SpeechSynthesizer
TranslationRecognizer



Answer Area:

Answer Area

```
static async Task TranslateSpeechAsync()  
{  
    var config =SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");  
  
    var lang = new List<string>  
  
    config.SpeechRecognitionLanguage = "en-GB";  
    lang.ForEach(config.AddTargetLanguage);  
  
    using var audioConfig = AudioConfig.FromDefaultMicrophoneInput();  
    using var recognizer = new  
  
    var result = await recognizer.RecognizeOnceAsync();  
    if (result.Reason == ResultReason.TranslatedSpeech)
```

["en-GB"]
["fr", "de", "es"]
["French", "Spanish", "German"]
{languages}

IntentRecognizer
SpeakerRecognizer
SpeechSynthesizer
TranslationRecognizer

**Section:**

**Explanation:**

Box 1: {"fr", "de", "es"}

A common task of speech translation is to specify target translation languages, at least one is required but multiples are supported. The following code snippet sets both French and German as translation language targets.

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

Box 2: TranslationRecognizer

After you've created a SpeechTranslationConfig, the next step is to initialize a TranslationRecognizer.

Example code:

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



**QUESTION 17**

DRAG DROP

You train a Custom Vision model used in a mobile app.

You receive 1,000 new images that do not have any associated data.

You need to use the images to retrain the model. The solution must minimize how long it takes to retrain the model.

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

**Select and Place:**

Actions	Answer Area
Upload the images by category.	
Get suggested tags.	
Upload all the images.	⬅
Group the images locally into category folders.	➡
Review the suggestions and confirm the tags.	
Tag the images manually.	
	⬆
	⬇

**Correct Answer:**

Actions	Answer Area
	Group the images locally into category folders.
Get suggested tags.	Upload the images by category.
Upload all the images.	Tag the images manually.
Review the suggestions and confirm the tags.	

**Section:**

**Explanation:**

Reference:

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

**QUESTION 18**

HOTSPOT

You are developing an application that includes language translation.

The application will translate text retrieved by using a function named `getTextToBeTranslated`. The text can be in one of many languages. The content of the text must remain within the Americas Azure geography.

You need to develop code to translate the text to a single language.

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

NOTE: Each correct selection is worth one point.

**Hot Area:**

 Vdumps

### Answer Area

```
...  
var endpoint =  ;  
var apiKey = "FF956C68883B21B38691ABD200A4C606";  
var text = getTextToBeTranslated();  
var body = '[{"Text":"' + text + '"}]';  
var client = new HttpClient();  
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);  
  
  
  
var uri = endpoint + &quot;?from=en&quot;;  
var uri = endpoint + &quot;?suggestedFrom=en&quot;;  
var uri = endpoint + &quot;?to=en&quot;;  
HttpResponseMessage response;  
var content = new StringContent(body, Encoding.UTF8, "application/json");  
var response = await client.PutAsync(uri, content);  
...
```

Vdumps

Answer Area:

### Answer Area

```
...
var endpoint =  ;
var apiKey = "FF956C68883B21B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);
var uri = ;
var uri = ;
var uri = ;
HttpResponseMessage response;
var content = new StringContent(body, Encoding.UTF8, "application/json");
var response = await client.PutAsync(uri, content);
...
```

 dumps

Section:

Explanation:

### QUESTION 19

HOTSPOT

You run the following command.

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

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

Hot Area:



## Answer Area

Statements	Yes	No
Going to <code>http://localhost:5000/status</code> will query the Azure endpoint to verify whether the API key used to start the container is valid.	<input type="radio"/>	<input type="radio"/>
The container logging provider will write log data.	<input type="radio"/>	<input type="radio"/>
Going to <code>http://localhost:5000/swagger</code> will provide the details to access the documentation for the available endpoints.	<input type="radio"/>	<input type="radio"/>

Answer Area:

## Answer Area

Statements	Yes	No
Going to <code>http://localhost:5000/status</code> will query the Azure endpoint to verify whether the API key used to start the container is valid.	<input checked="" type="radio"/>	<input type="radio"/>
The container logging provider will write log data.	<input checked="" type="radio"/>	<input type="radio"/>
Going to <code>http://localhost:5000/swagger</code> will provide the details to access the documentation for the available endpoints.	<input checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

Box 1: Yes

`http://localhost:5000/status` : Also requested with GET, this verifies if the api-key used to start the container is valid without causing an endpoint query.

Box 2: Yes

The command saves container and LUIS logs to output mount at `C:\output`, located on container host

Box 3: Yes

`http://localhost:5000/swagger` : The container provides a full set of documentation for the endpoints and a Try it out feature. With this feature, you can enter your settings into a web-based HTML form and make the query without having to write any code. After the query returns, an example CURL command is provided to demonstrate the HTTP headers and body format that's required.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-container-howto>

**QUESTION 20**

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

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

**Correct Answer: A**

**Section:**

**Explanation:**

Active learning is a technique of machine learning in which the machine learned model is used to identify informative new examples to label. In LUIS, active learning refers to adding utterances from the endpoint traffic whose current predictions are unclear to improve your model.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-glossary>

**QUESTION 21**

DRAG DROP

You are building a Language Understanding model for purchasing tickets.

You have the following utterance for an intent named PurchaseAndSendTickets.

Purchase [2 audit business] tickets to [Paris] [next Monday] and send tickets to [email@domain.com]

You need to select the entity types. The solution must use built-in entity types to minimize training data whenever possible.

Which entity type should you use for each label? To answer, drag the appropriate entity types to the correct labels. Each entity type may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

**Select and Place:**

Entity Types	Answer Area
Email	
List	Paris: <input type="text"/>
Regex	email@domain.com: <input type="text"/>
GeographyV2	2 audit business: <input type="text"/>
Machine learned	

**Correct Answer:**

## Entity Types

## Answer Area

List	Paris:	GeographyV2
Regex	email@domain.com:	Email
	2 audit business:	Machine learned

### Section:

#### Explanation:

Box 1: GeographyV2

The prebuilt geographyV2 entity detects places. Because this entity is already trained, you do not need to add example utterances containing GeographyV2 to the application intents.

Box 2: Email

Email prebuilt entity for a LUIS app: Email extraction includes the entire email address from an utterance. Because this entity is already trained, you do not need to add example utterances containing email to the application intents.

Box 3: Machine learned

The machine-learning entity is the preferred entity for building LUIS applications.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-geographyv2>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-email>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/reference-entity-machine-learned-entity>

## 02 - Implement Natural Language Processing Solutions

### Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

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

To start the case study

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

When you are ready to answer a question, click the Question button to return to the question.

### Overview

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

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

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

### Existing Environment

### Applications

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

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

#### Azure Resources

You have the following resources:

- An Azure Active Directory (Azure AD) tenant

An Azure Active Directory (Azure AD) tenant

- The tenant supports internal authentication.

- All employees belong to a group named AllUsers.

- Senior managers belong to a group named LeadershipTeam.

An Azure Functions resource

- A function app posts to Azure Event Grid when stock levels of a product change between OK, Low Stock, and Out of Stock. The function app uses the Azure Cosmos DB change feed. An Azure Cosmos DB account

- The account uses the Core (SQL) API.

- The account stores data for the Product Management app and the Inventory Tracking app.

- An Azure Storage account

An Azure Storage account

- The account contains blob containers for assets related to products.

- The assets include images, videos, and PDFs.

An Azure Cognitive Services resource named wwics

A Video Indexer resource named wwivi

#### Requirements

##### Business Goals

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

##### Planned Changes

Wide World Importers plans to start the following projects:

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

##### Business Requirements

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

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

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

Minimize costs.

##### Governance and Security Requirements

Wide World Importers identifies the following governance and security requirements:

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

Azure Cognitive Services must be inaccessible directly from the internet.

##### Accessibility Requirements

Wide World Importers identifies the following accessibility requirements:

All images must have relevant alt text.

All videos must have transcripts that are associated to the video and included in product descriptions. Product descriptions, transcripts, and alt text must be available in English, Spanish, and Portuguese.

##### Product Creation Requirements

Wide World Importers identifies the following requirements for improving the Product Management app:

Minimize how long it takes for employees to create products and add assets.

Remove the need for manual translations.

##### Smart E-Commerce Requirements

Wide World Importers identifies the following requirements for the smart e-commerce project:

Ensure that the Cognitive Search solution meets a Service Level Agreement (SLA) of 99.9% availability for searches and index writes. Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products. Support autocompletion and autosuggestion based on all product name variants.

Store all raw insight data that was generated, so the data can be processed later.

Update the stock level field in the product index immediately upon changes.

Update the product index hourly.

##### Shopping On-the-Go Requirements



Wide World Importers identifies the following requirements for the shopping on-the-go chatbot:

Answer common questions.

Support interactions in English, Spanish, and Portuguese.

Replace an existing FAQ process so that all Q&A is managed from a central location.

Provide all employees with the ability to edit Q&As. Only senior managers must be able to publish updates. Support purchases by providing information about relevant products to customers. Product displays must include images and warnings when stock levels are low or out of stock.

Product JSON Sample

You have the following JSON sample for a product.

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



#### QUESTION 1

##### HOTSPOT

You are planning the product creation project.

You need to build the REST endpoint to create the multilingual product descriptions.

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

NOTE: Each correct selection is worth one point.

##### Hot Area:



### Answer Area

		?api-version=3.0&to=es&to=pt
api.cognitive.microsofttranslator.com	/detect	
api-nam.cognitive.microsofttranslator.com	/languages	
westus.tts.speech.microsoft.com	/text-to-speech	
wwics.cognitiveservices.azure.com/translator	/translate	

Answer Area:

### Answer Area

		?api-version=3.0&to=es&to=pt
api.cognitive.microsofttranslator.com	/detect	
api-nam.cognitive.microsofttranslator.com	/languages	
westus.tts.speech.microsoft.com	/text-to-speech	
wwics.cognitiveservices.azure.com/translator	/translate	

Section:

Explanation:

Box 1: api.cognitive.microsofttranslator.com

Translator 3.0: Translate. Send a POST request to:

<https://api.cognitive.microsofttranslator.com/translate?api-version=3.0>

Box 2: /translate

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/translator/reference/v3-0-translate>

### 01 - Implement Knowledge Mining Solutions

#### QUESTION 1

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 2**

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 3**

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 4

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 5

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/coqnitve-services/text-analytics/overview>

**QUESTION 6**

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($"{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.	<input type="radio"/>	<input type="radio"/>
The output will contain the following words: the, cat, sat, on, and mat.	<input type="radio"/>	<input type="radio"/>
The output will contain the confidence level for key phrases.	<input type="radio"/>	<input type="radio"/>

**Answer Area:**

## Answer Area

Statements	Yes	No
The call will output key phrases from the input string to the console.	<input checked="" type="radio"/>	<input type="radio"/>
The output will contain the following words: the, cat, sat, on, and mat.	<input type="radio"/>	<input checked="" type="radio"/>
The output will contain the confidence level for key phrases.	<input type="radio"/>	<input checked="" type="radio"/>

### 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 7

#### HOTSPOT

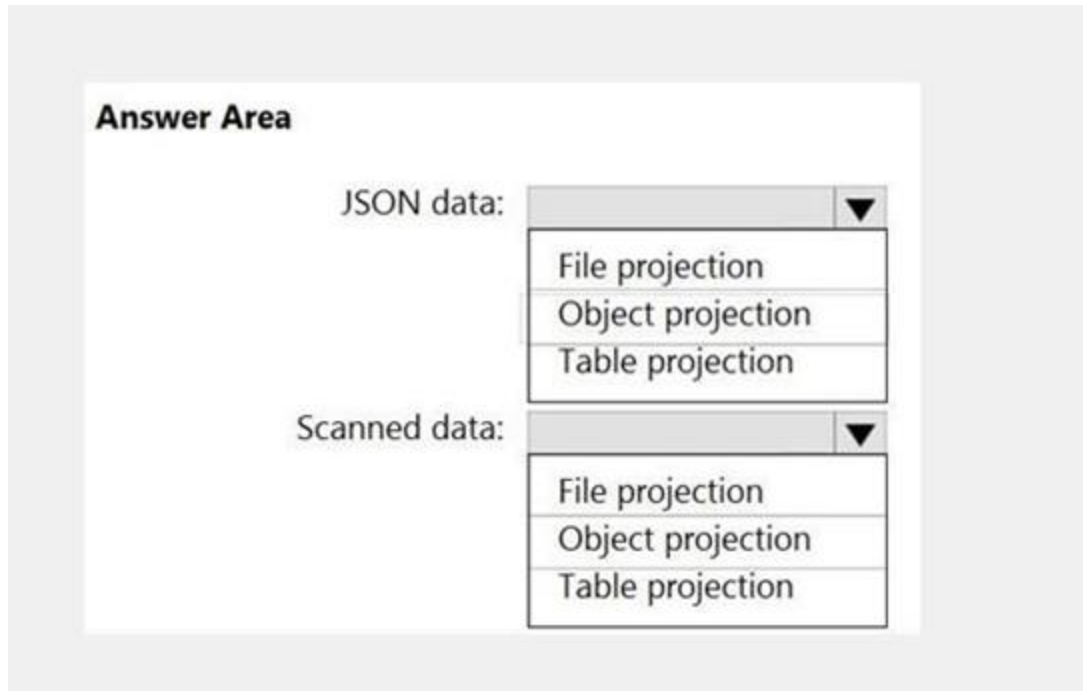
You are creating an enrichment pipeline that will use Azure Cognitive Search. The knowledge store contains unstructured JSON data and scanned PDF documents that contain text.

Which projection type should you use for each data type? To answer, select the appropriate options in the answer area.

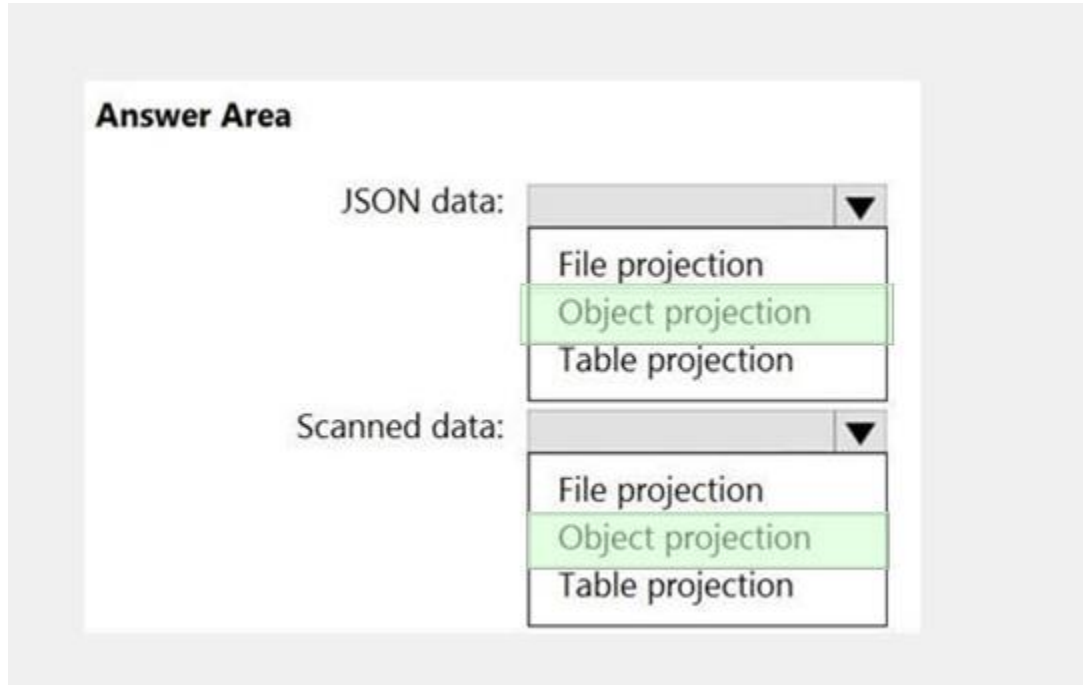
NOTE: Each correct selection is worth one point.

#### 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 8**

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:**

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

**Answer Area:**

Statements	Yes	No
CompanyDescription is available for indexing.	<input checked="" type="radio"/>	<input type="radio"/>
The definition calls a web API as part of the enrichment process.	<input checked="" type="radio"/>	<input type="radio"/>
The enrichment step is called only for the first organization under "/document/organizations".	<input type="radio"/>	<input checked="" type="radio"/>

**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 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:

Actions	Answer Area
Add a new query key.	
Regenerate the secondary admin key.	
Change the app to use the secondary admin key.	⬅️
Change the app to use the new key.	➡️
Regenerate the primary admin key.	
Delete the compromised key.	

Correct Answer:

Actions	Answer Area
	Add a new query key.
Regenerate the secondary admin key.	Change the app to use the new key.
Change the app to use the secondary admin key.	Delete the compromised key.
Regenerate the primary admin key.	



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 autocomplete 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 suggesterName in the body.
- B. Add a suggester that 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 searchFields 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 autocomplete 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:

## Exam H

### QUESTION 1

You build a Language Understanding model by using the Language Understanding portal. You export the model as a JSON file as shown in the following sample.

```
{
  "text": "average amount of rain by month at chicago last year",
  "intent": "Weather.CheckWeatherValue",
  "entities": [
    {
      "entity": "Weather.WeatherRange",
      "startPos": 0,
      "endPos": 6,
      "children": []
    },
    {
      "entity": "Weather.WeatherCondition",
      "startPos": 18,
      "endPos": 21,
      "children": []
    },
    {
      "entity": "Weather.Historic",
      "startPos": 23,
      "endPos": 30,
      "children": []
    }
  ]
}
```

To what does the Weather.Historic entity correspond in the utterance?

- A. by month
- B. chicago
- C. rain
- D. location

**Correct Answer: A**

**Section:**

### QUESTION 2

You need to develop an extract solution for the receipt images. The solution must meet the document processing requirements and the technical requirements. You upload the receipt images to the Form Recognizer API for analysis, and the API returns the following JSON.



```
"documentResults":[
  {
    "docType":"prebuilt:receipt",
    "pageRange":[
      1,
      1
    ],
    "fields":{
      "ReceiptType":{
        "type":"string",
        "valueString":"Itemized",
        "confidence":0.672
      },
      "MerchantName":{
        "type":"string",
        "valueString":"Tailwind",
        "text":"Tailwind",
        "boundingBox":[],
        "page":1,
        "confidence":0.913,
        "elements":[
          "#/readResults/0/lines/0/words/0"
        ]
      }
    }
  },
  ...
]
```

Vdumps

Which expression should you use to trigger a manual review of the extracted information by a member of the ConsultantBookkeeper group?

- A. `documentResults.docType == "prebuilt:receipt"`
- B. `documentResults.fields.confidence < 0.7`
- C. `documentResults.fields.ReceiptType.confidence > 0.7`
- D. `documentResults.fields.MerchantName.confidence < 0.7`

**Correct Answer: D**

**Section:**

**Explanation:**

Need to specify the field name, and then use `< 0.7` to handle trigger if confidence score is less than 70%. Scenario:

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.

Reference:

<https://docs.microsoft.com/en-us/azure/applied-ai-services/form-recognizer/api-v2-0/reference-sdk-api-v2-0>

**QUESTION 3**

HOTSPOT

You build a QnA Maker resource to meet the chatbot requirements.

Which RBAC role should you assign to each group? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Management-Accountants	<input type="checkbox"/>	Owner
	<input type="checkbox"/>	Contributor
	<input type="checkbox"/>	Cognitive Services User
	<input type="checkbox"/>	Cognitive Services QnA Maker Read
	<input type="checkbox"/>	Cognitive Services QnA Maker Editor

Consultant-Accountants	<input type="checkbox"/>	Owner
	<input type="checkbox"/>	Contributor
	<input type="checkbox"/>	Cognitive Services User
	<input type="checkbox"/>	Cognitive Services QnA Maker Read
	<input type="checkbox"/>	Cognitive Services QnA Maker Editor


  

Agent-CustomerServices	<input type="checkbox"/>	Owner
	<input type="checkbox"/>	Contributor
	<input type="checkbox"/>	Cognitive Services User
	<input type="checkbox"/>	Cognitive Services QnA Maker Read
	<input type="checkbox"/>	Cognitive Services QnA Maker Editor

*Vdumps*

Answer Area:





**Section:**

**Explanation:**

Box 1: Cognitive Service User

Ensure that the members of a group named Management-Accountants can approve the FAQs.

Approve=publish.

Cognitive Service User (read/write/publish): API permissions: All access to Cognitive Services resource except for ability to:

1. Add new members to roles.
2. Create new resources.

Box 2: Cognitive Services QnA Maker Editor

Ensure that the members of a group named Consultant-Accountants can create and amend the FAQs.

QnA Maker Editor: API permissions:

1. Create KB API
2. Update KB API
3. Replace KB API
4. Replace Alterations
5. "Train API" [in new service model v5]

Box 3: Cognitive Services QnA Maker Read

Ensure that the members of a group named the Agent-CustomerServices can browse the FAQs.

QnA Maker Read: API Permissions:

1. Download KB API
2. List KBs for user API



3. Get Knowledge base details

4. Download AlterationsGenerate Answer

Reference: <https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/concepts/role-based-access-control>

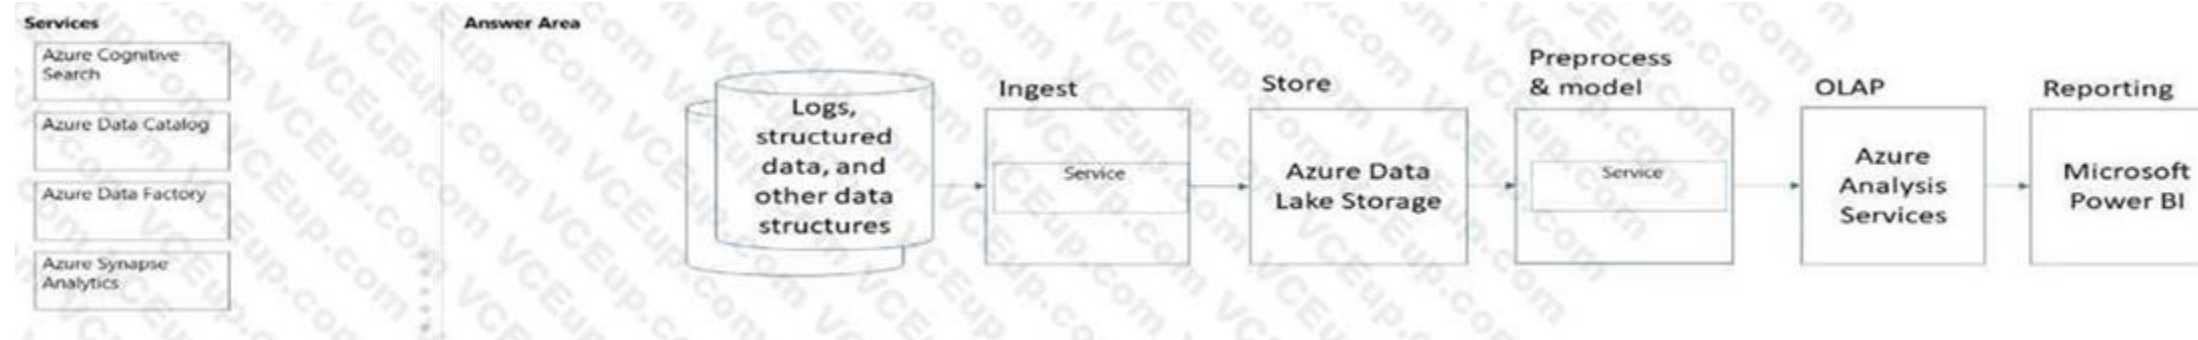
#### QUESTION 4

DRAG DROP

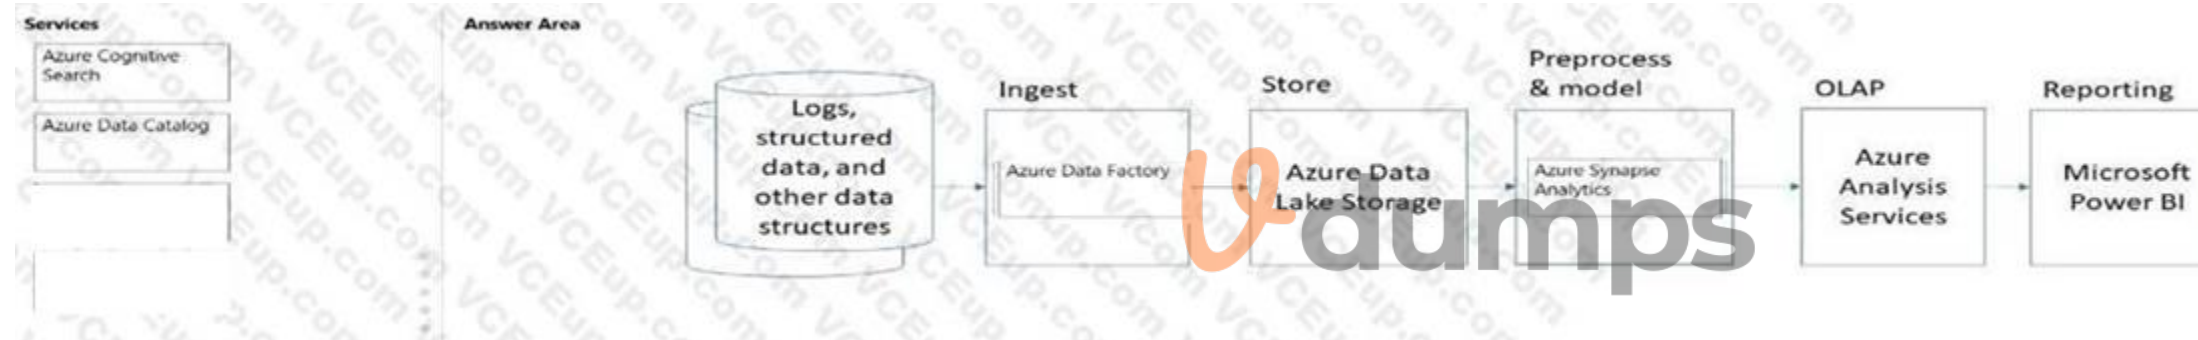
Match the Azure services to the appropriate locations in the architecture.

To answer, drag the appropriate service from the column on the left to its location on the right. Each service 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 5

HOTSPOT

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:

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	<input type="radio"/>	<input type="radio"/>
All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	<input type="radio"/>	<input type="radio"/>

Answer Area:

**Answer Area**

**Statements**

Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.  Yes  No

Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.  Yes  No

All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.  Yes  No

**Section:**

**Explanation:**

**QUESTION 6**

HOTSPOT

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**

Normalization involves eliminating relationships between database tables.  Yes  No

Normalizing a database reduces data redundancy.  Yes  No

Normalization improves data integrity.  Yes  No

**Answer Area:**

**Answer Area**

**Statements**

Normalization involves eliminating relationships between database tables.  Yes  No

Normalizing a database reduces data redundancy.  Yes  No

Normalization improves data integrity.  Yes  No

**Section:**

**Explanation:**

**QUESTION 7**

HOTSPOT

Select the answer that correctly completes the sentence.

**Hot Area:**

**Answer Area**

A block of code that runs in a database is called

- a stored procedure.
- a table.
- a view.
- an index.

**Answer Area:**



**Section:**

**Explanation:**

**QUESTION 8**

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 SQL Database
- C. Azure HDInsight
- D. Azure Cosmos DB

**Correct Answer: B**

**Section:**

**QUESTION 9**

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

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

**Correct Answer: B**

**Section:**

**QUESTION 10**

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 11**

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 12

What is a primary characteristic of a relational database?

- A. data is queried and manipulated by using a variant of the SQL 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 13

What are 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 14

What should you use to build a Microsoft Power BI paginated report?

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

**Correct Answer: A**

**Section:**



**QUESTION 15**

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 16**

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 17**

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 BI?

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

**Correct Answer: C**

**Section:**

**QUESTION 18**

Your company has a reporting 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 19**

You are building a Conversational Language Understanding model.



You need to ensure that the model will support the following sample utterances:

- Set all the lights to on.
- Turn off the lights in the living room.
- What is the current thermostat temperature?
- Lower the temperature of the thermostat by five degrees.

Which three elements should you add to the model?

Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

**Correct Answer: B, C, F**

**Section:**

#### QUESTION 20

You are building a social media extension that will convert text to speech. The solution must meet the following requirements:

- Support messages of up to 400 characters.
- Provide users with multiple voice options.
- Minimize costs.

You create an Azure Cognitive Services resource.

Which Speech API endpoint provides users with the available voice options?



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

**Correct Answer: C**

**Section:**

#### QUESTION 21

You are building a chatbot that will use question answering in Azure Cognitive Service for Language. You have a PDF named Doc1.pdf that contains a product catalogue and a price list

You upload Doc1.pdf and train the model.

During testing, users report that the chatbot responds correctly to the following QWhat is the price of < product>? The chatbot fails to respond to the following QHow much does <product\* cost?

You need to ensure that the chatbot responds correctly to both questions.

Solution: From Language Studio, you create an entity for cost, and then retrain and republish the model. Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section:**

#### QUESTION 22

**HOTSPOT**

You are building an app that will enable users to upload images. The solution must meet the following requirements:

- Automatically suggest alt text for the images.
- Detect inappropriate images and block them.
- Minimize development effort.

You need to recommend a computer vision endpoint for each requirement.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Hot Area:**

**Answer Area**

Generate alt text:

- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1>

Detect inappropriate content:

- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1>

**Answer Area:**

**Answer Area**

Generate alt text:

- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1>

Detect inappropriate content:

- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1>

**Section:**

**Explanation:**

**QUESTION 23**

**HOTSPOT**



You are building content for a video training solution.

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

Hot Area:

**Answer Area**

Custom neural voice:

Narration:

Narration:

Answer Area:

**Answer Area**

Custom neural voice:

Narration:

Narration:

Section:

Explanation:

**QUESTION 24**

HOTSPOT

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



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

**Hot Area:**

**Answer Area**

If a user asks "what is the weather like in New York", the bot will [answer choice].

- Identify New York as a city entity
- change to a different dialog
- identify New York as a city entity
- identify New York as a state entity
- respond with the weather in Seattle

The GetWeather dialog uses a [answer choice] trigger.

- Language Understanding Intent recogn...
- Custom events
- Dialog events
- Language Understanding Intent recognized
- QnA Intent recognized

**Answer Area:**

**Answer Area**

If a user asks "what is the weather like in New York", the bot will [answer choice].

- Identify New York as a city entity
- change to a different dialog
- identify New York as a city entity
- identify New York as a state entity
- respond with the weather in Seattle

The GetWeather dialog uses a [answer choice] trigger.

- Language Understanding Intent recogn...
- Custom events
- Dialog events
- Language Understanding Intent recognized
- QnA intent recognized

**Section:**


**Explanation:**

**QUESTION 25**

**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**

Statements	Yes	No
user.name is an entity.	<input type="radio"/>	<input type="radio"/>
The dialog asks for a user name and a user age and assigns appropriate values to the user.name and user.age properties.	<input type="radio"/>	<input type="radio"/>
The chatbot attempts to take the first non-null entity value for userName Or personName and assigns the value to user.name.	<input type="radio"/>	<input type="radio"/>

**Answer Area:**

**Answer Area**

Statements	Yes	No
user.name is an entity.	<input type="radio"/>	<input checked="" type="radio"/>
The dialog asks for a user name and a user age and assigns appropriate values to the user.name and user.age properties.	<input checked="" type="radio"/>	<input type="radio"/>
The chatbot attempts to take the first non-null entity value for userName Or personName and assigns the value to user.name.	<input checked="" type="radio"/>	<input type="radio"/>

**Section:**

**Explanation:**

**QUESTION 26**

**DRAG DROP**

You plan to build a chatbot to support task tracking.

You create a Conversational Language Understanding service named lu1.

You need to build a Conversational Language Understanding model to Integrate into the chatbot. The solution must minimize development time to build the model. Which four actions should you perform In sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

**Select and Place:**



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

**Correct Answer:**

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

**Section:**

**Explanation:**

**QUESTION 27**

You have an Azure Cognitive Search solution and an enrichment pipeline that performs Sentiment Analysis on social media posts.

You need to define a knowledge store that will include the social media posts and the Sentiment Analysis results.

Which two fields should you include in the definition? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. storageContainer
- B. tables
- C. storageConnectionString
- D. files
- E. objects

**Correct Answer: D, E**

**Section:**

**QUESTION 28**

**HOTSPOT**

You are building a call handling system that will receive calls from French-speaking and German-speaking callers. The system must perform the following tasks;

\* Capture inbound voice messages as text.

\* Replay messages in English on demand.

Which Azure Cognitive Services 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

.....

To capture messages:

- Speaker Recognition
- Speech-to-text**
- Text-to-speech
- Translator

To replay messages:

- Speech-to-text only
- Speech-to-text and Language
- Speaker Recognition and Language
- Text-to-speech and Language
- Text-to-speech and Translator**

Answer Area:

Answer Area

.....

To capture messages:

- Speaker Recognition
- Speech-to-text**
- Text-to-speech
- Translator

To replay messages:

- Speech-to-text only
- Speech-to-text and Language
- Speaker Recognition and Language
- Text-to-speech and Language
- Text-to-speech and Translator**

Section:

Explanation:

**QUESTION 29**

HOTSPOT

You are building an app that will process incoming email and direct email messages to either French or English language support teams.

Which Azure Cognitive Services API 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

.....

https://

- api.cognitive.microsofttranslator.com**
- eastus.api.cognitive.microsofttranslator.com
- portal.azure.com

- /text/analytics/v3.1/entities/recognition/general**
- /text/analytics/v3.1/languages
- /translator/text/v3.0/translate?to=en
- /translator/text/v3.0/translate?to=fr

Answer Area:

Answer Area



Section:

Explanation:

QUESTION 30

You are building a Chatbot by using the Microsoft Bot Framework SDK. The bot will be used to accept food orders from customers and allow the customers to customize each food item. You need to configure the bot to ask the user for additional input based on the type of item ordered. The solution must minimize development effort. Which two types of dialogs should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. adaptive
- B. action
- C. waterfall
- D. prompt
- E. input

Correct Answer: B, C

Section:

QUESTION 31

You are building an app that will include one million scanned magazine articles. Each article will be stored as an image file. You need to configure the app to extract text from the images. The solution must minimize development effort. What should you include in the solution?

- A. Computer Vision Image Analysis
- B. the Read API in Computer Vision
- C. Form Recognizer
- D. Azure Cognitive Service for Language

Correct Answer: B

Section:

QUESTION 32

You have an Azure subscription that contains an AI enrichment pipeline in Azure Cognitive Search and an Azure Storage account that has 10 GB of scanned documents and images. You need to index the documents and images in the storage account. The solution must minimize how long it takes to build the index. What should you do?

- A. From the Azure portal, configure parallel indexing.
- B. Create a text-based indexer by using the REST API.
- C. From the Azure portal, configure scheduled indexing.
- D. Configure field mappings by using the REST API.

Correct Answer: A



**Section:**

**QUESTION 33**

**HOTSPOT**

You have an app named App1 that uses Azure AI Document Intelligence to analyze medical You have an app named App1 that uses Azure AI 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.	<input type="radio"/>	<input type="radio"/>
The text content was recognized with greater than 70 percent confidence.	<input type="radio"/>	<input type="radio"/>
The form elements were recognized with greater than 70 percent confidence.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
The chosen model is suitable for the intended use case.	<input type="radio"/>	<input checked="" type="radio"/>
The text content was recognized with greater than 70 percent confidence.	<input checked="" type="radio"/>	<input type="radio"/>
The form elements were recognized with greater than 70 percent confidence.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

**QUESTION 34**

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,
    "count": 1,
    "maxValue": 100,
    "size": 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
      ...
    }
  }
}
```

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,  
    "count": 1,  
    "maxValue": 100,  
    "size": 600,  
  },  
  "properties": {  
    "model": {  
      "format": "OpenAI",  
      ...  
    }  
  }  
}
```



Section:

Explanation:

### QUESTION 35

You have an Azure subscription. The subscription contains an Azure OpenAI resource that hosts a GPT-4 model named Model1 and an app named App1. App1 uses Model1!

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 36**

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 37**

DRAG DROP

You have a monitoring solution that uses the Azure AI Anomaly Detector service.

You provision a server named Server1 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:**

**Actions**

- From Server1, run the docker push command.
- Query the prediction endpoint of the Azure AI Anomaly Detector in Azure.
- Install the Docker Engine on Server1.
- From Server1, run the docker pull command.
- From Server1, run the docker run command.
- Query the prediction endpoint on Server1.

**Answer Area**



**Correct Answer:**

### Actions

From Server1, run the docker push command.

Query the prediction endpoint of the Azure AI Anomaly Detector in Azure.

### Answer Area

Install the Docker Engine on Server1.

From Server1, run the docker pull command.

From Server1, run the docker run command.

Query the prediction endpoint on Server1.

### Section:

### Explanation:

Install the Docker Engine on Server1.

From Server1, run 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 38

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:**