

AI-102 Dumps

Designing and Implementing an Azure AI Solution

<https://www.certleader.com/AI-102-dumps.html>



NEW QUESTION 1

HOTSPOT - (Topic 3)

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.

Answer Area

JSON data: ▼

File projection
Object projection
Table projection

Scanned data: ▼

File projection
Object projection
Table projection

- A. Mastered
- B. Not Mastered

Answer: A

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.

NEW QUESTION 2

- (Topic 3)

Your company uses an Azure Cognitive Services solution to detect faces in uploaded images. The method to detect the faces uses the following code.

```
static async Task DetectFaces(string imagePath)
{
    HttpClient client = new HttpClient();
    DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", subscriptionKey);
    string requestParameter = "detectionModel=detection_01&returnFaceId=true&returnFaceLandmarks=false";
    string uri = endpoint + "/face/v1.0/detect?" + requestParameters;
    HttpResponseMessage response;
    byte[] byteData = GetImagesAsByteArray(imageFilePath);
    using (ByteArrayContent content = new ByteArrayContent(byteData))
    {
        Headers.ContentType = new MediaTypeHeaderValue("application/octet-stream");
        response = await PostAsync(uri, content);
        string contentString = await Content.ReadAsStringAsync();
        ProcessDetection(contentString);
    }
}
```

You discover that the solution frequently fails to detect faces in blurred images and in images that contain sideways faces.

- A. Use a different version of the Face API.
- B. Use the Computer Vision service instead of the Face service.
- C. Use the Identify method instead of the Detect method.
- D. Change the detection model. You need to increase the likelihood that the solution can detect faces in blurred images and images that contain sideways faces. What should you do?

Answer: D

Explanation:

Evaluate different models.

The best way to compare the performances of the detection models is to use them on a sample dataset. We recommend calling the Face - Detect API on a variety of images, especially images of many faces or of faces that are difficult to see, using each detection model. Pay attention to the number of faces that each model returns.

Reference:

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

NEW QUESTION 3

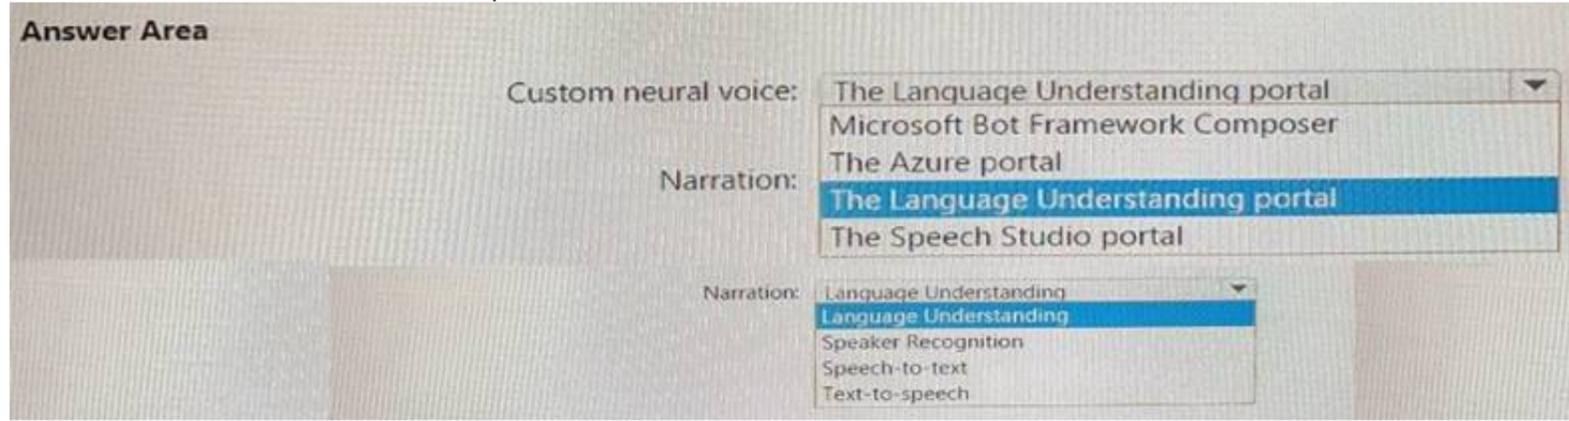
HOTSPOT - (Topic 3)

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.

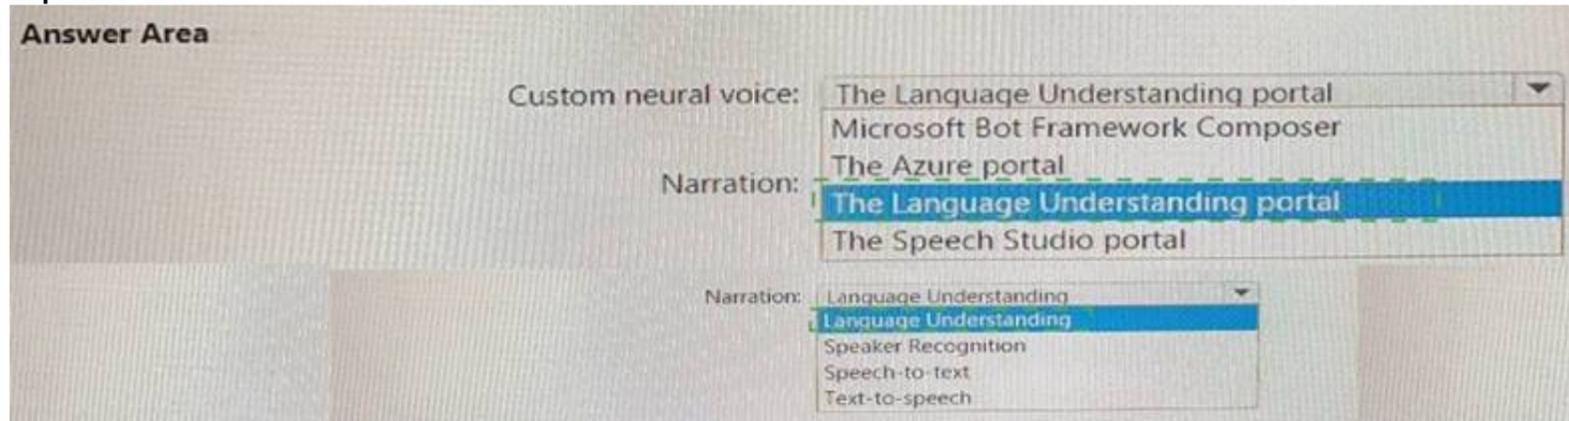
NOTE: Each correct answer is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 4

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 5

- (Topic 3)

You are building a chatbot.

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

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

Answer: B

NEW QUESTION 6

- (Topic 3)

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 create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1. You plan to create a new Azure Cognitive Search service named service1. You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet. Solution: You deploy service1 and a public endpoint, and you configure an IP firewall rule. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Reference:
<https://docs.microsoft.com/en-us/azure/private-link/private-link-overview>

NEW QUESTION 7

HOTSPOT - (Topic 3)

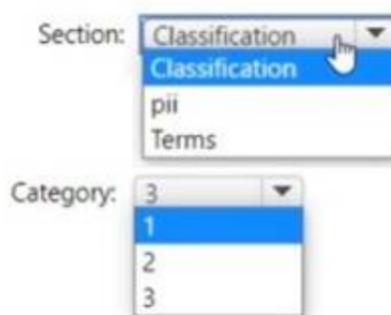
You are building a chatbot.

You need to use the Content Moderator service to identify messages that contain sexually explicit language.

Which section in the response from the service will contain the category score, and which category will be assigned to the message? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

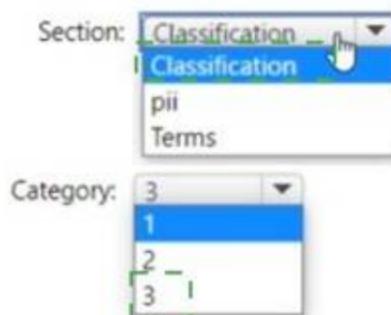


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 8

- (Topic 3)

You are building a retail kiosk system that will use a custom neural voice. You acquire audio samples and consent from the voice talent. You need to create a voice talent profile.

What should you upload to the profile?

- A. a five-minute wav or mp3 file of the voice talent describing the kiosk system
- B. a five-minute .flac audio file and the associated transcript as a w file
- C. a .wav or mp3 file of the voice talent consenting to the creation of a synthetic version of their voice
- D. a .zip file that contains 10-second .wav files and the associated transcripts as .txt files

Answer: D

NEW QUESTION 9

- (Topic 3)

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

Answer: B

Explanation:

The model needs to be extended and retrained.

NEW QUESTION 10

- (Topic 3)

You need to upload speech samples to a Speech Studio project. 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.

Answer: B

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>

NEW QUESTION 10

DRAG DROP - (Topic 3)

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.

Actions

- 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.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

- 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.

Answer Area

- Add a new query key.
- Change the app to use the new key.
- Delete the compromised key.

NEW QUESTION 13

- (Topic 3)

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.

Answer: A

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 new face for 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/video-indexer/customize-person-model-with-api>

NEW QUESTION 17

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

- A data analyst
- A data engineer
- A data scientist
- A database administrator

_____ is responsible for creating visuals and charts that help a company make informed decisions.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

- A data analyst
- A data engineer
- A data scientist
- A database administrator

_____ is responsible for creating visuals and charts that help a company make informed decisions.

NEW QUESTION 22

- (Topic 3)

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

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

Answer: A

NEW QUESTION 26

- (Topic 3)

You are building a Language Understanding solution.

You discover that many intents have similar utterances containing airport names or airport codes.

You need to minimize the number of utterances used to train the model. Which type of custom entity should you use?

- A. Pattern.any
- B. machine-learning
- C. list
- D. regular expression

Answer: C

NEW QUESTION 29

- (Topic 3)

You are designing a conversational interface for an app that will be used to make vacation requests. The interface must gather the following data:

- The start date of a vacation
- The end date of a vacation
- The amount of required paid time off

The solution must minimize dialog complexity. Which type of dialog should you use?

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

Answer: D

NEW QUESTION 31

- (Topic 3)

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

Answer: B

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>

NEW QUESTION 32

DRAG DROP - (Topic 3)

You plan to build a chatbot to support task tracking.

You create a Language Understanding service named lu1.

You need to build a Language Understanding model to integrate into the chatbot. The solution must minimize development time to build the model.

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

Actions

Answer Area

- Train the application.
- Publish the application.
- Add a new application.
- Add example utterances.
- Add the prebuilt domain ToDo.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

- * 1. Add a new application
- * 2. Add a prebuilt domain intent ToDo (it has already utterances so we can skip this step)
- * 3. Train
- * 4. Publish

NEW QUESTION 37

- (Topic 3)

You are developing an app that will use the Decision and Language APIs.

You need to provision resources for the app. The solution must ensure that each service is accessed by using a single endpoint and credential.

Which type of resource should you create?

- A. Language
- B. Speech
- C. Azure Cognitive Services
- D. Content Moderator

Answer: C

NEW QUESTION 38

- (Topic 3)

You build a conversational bot named bot1.

You need to configure the bot to use a QnA Maker application.

From the Azure Portal, where can you find the information required by bot1 to connect to the QnA Maker application?

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

Answer: C

Explanation:

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

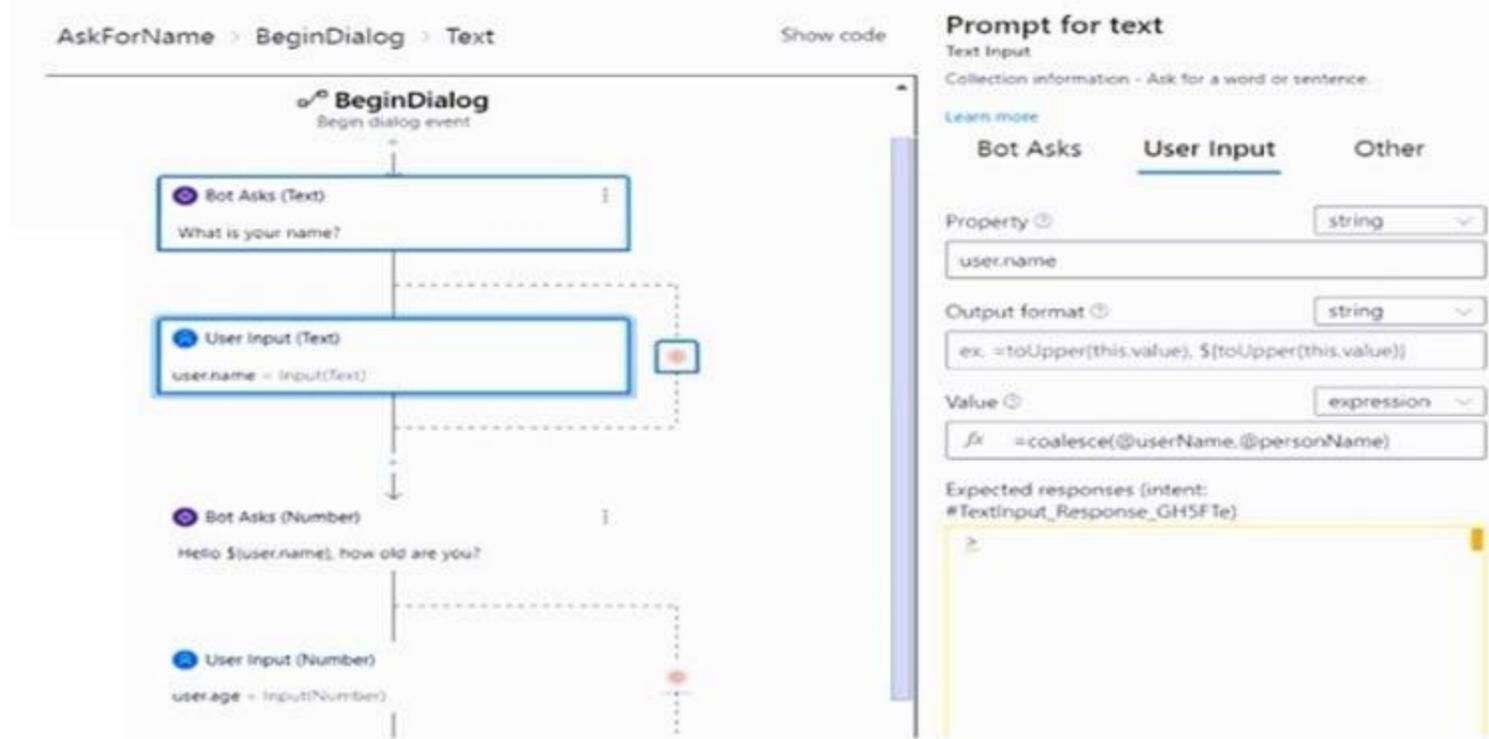
Reference:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-qna>

NEW QUESTION 43

HOTSPOT - (Topic 3)

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.

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"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
user.name is an entity.	<input checked="" 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 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"/>

NEW QUESTION 47

- (Topic 3)

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.

Answer: AB

NEW QUESTION 48

HOTSPOT - (Topic 3)

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.

```

* * *
var endpoint =  ;
                
                
                
                

var apiKey = "FF956C68883821B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);





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

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```

* * *
var endpoint =
    &quot;https://api.cognitive.microsofttranslator.com/translate&quot;;
    &quot;https://api.cognitive.microsofttranslator.com/transliterate&quot;;
    &quot;https://api-apc.cognitive.microsofttranslator.com/detect&quot;;
    &quot;https://api-nam.cognitive.microsofttranslator.com/detect&quot;;
    &quot;https://api-nam.cognitive.microsofttranslator.com/translate&quot;;

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);
* * *

```

NEW QUESTION 50

- (Topic 3)

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

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

Answer: A

NEW QUESTION 52

- (Topic 3)

You are building a chatbot by using Microsoft Bot Framework Composer.

You need to configure the chatbot to present a list of available options. The solution must ensure that an image is provided for each option.

Which two features should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

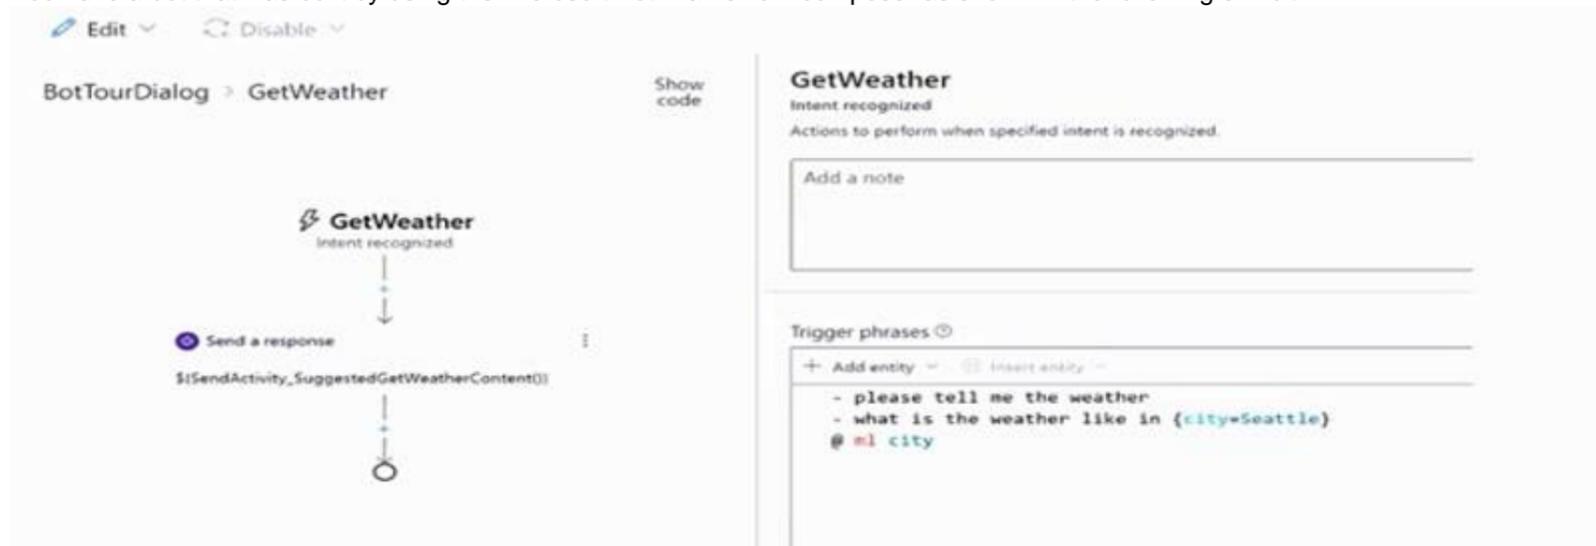
- A. an Azure function
- B. an adaptive card
- C. an entity
- D. a dialog
- E. an utterance

Answer: BD

NEW QUESTION 54

HOTSPOT - (Topic 3)

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.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 59

- (Topic 3)

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 log=true to the prediction endpoint query.
- D. Enable sentiment analysis.

Answer: C

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>

NEW QUESTION 62

- (Topic 3)

You have a collection of 50,000 scanned documents that contain text. You plan to make the text available through Azure Cognitive Search. You need to configure an enrichment pipeline to perform optical character recognition (OCR) and text analytics. The solution must minimize costs. What should you attach to the skillset?

- A. a new Computer Vision resource
- B. a free (Limited enrichments) Cognitive Services resource
- C. an Azure Machine Learning pipeline
- D. a new Cognitive Services resource that uses the SO pricing tier

Answer: A

Explanation:

The Computer Vision API uses text recognition APIs to extract and recognize text information from images. Read uses the latest recognition models, and is optimized for large, text-heavy documents and noisy images.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/cognitive-search-with-skillsets>

NEW QUESTION 66

DRAG DROP - (Topic 3)

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.

Actions

Answer Area

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



- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 71

- (Topic 3)

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

Answer: B

NEW QUESTION 73

- (Topic 3)

You train a Conversational Language Understanding model to understand the natural language input of users. You need to evaluate the accuracy of the model before deploying it. What are two methods you can use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. From the language authoring REST endpoint, retrieve the model evaluation summary.
- B. From Language Studio, enable Active Learning, and then validate the utterances logged for review.
- C. From Language Studio, select Model performance.
- D. From the Azure portal, enable log collection in Log Analytics, and then analyze the logs.

Answer: AC

NEW QUESTION 74

DRAG DROP - (Topic 3)

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.

Actions

Answer Area

- 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.



- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
Use the ExportProject endpoint on acvdev.	Use the GetProjects endpoint on acvdev.
Use the GetProjects endpoint on acvdev.	Use the ExportProject endpoint on acvdev.
Use the ImportProject endpoint on acvprod.	Use the ImportProject endpoint on acvprod.
Use the ExportIteration endpoint on acvdev.	
Use the GetIterations endpoint on acvdev.	
Use the UpdateProject endpoint on acvprod.	

NEW QUESTION 77

- (Topic 3)

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

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

Answer: A

Explanation:

Reference:

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

NEW QUESTION 80

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

A JSON document is an example of

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A JSON document is an example of

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

NEW QUESTION 81

- (Topic 3)

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

Answer: A

Explanation:

An ML entity can be composed of smaller sub-entities, each of which can have its own properties. For example, Address could have the following structure:

Address: 4567 Main Street, NY, 98052, USA Building Number: 4567

Street Name: Main Street State: NY

Zip Code: 98052 Country: USA
<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-entity-types>

NEW QUESTION 83

- (Topic 3)

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

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

You have a chatbot that uses question answering in Azure Cognitive Service for Language Users report that the responses of the chatbot lack formality when answering spurious questions

You need to ensure that the chatbot provides formal responses to spurious questions. Solution: From Language Studio, you change the chitchat source to qna_chitchit_friendly.tsv. and then retrain and republish the model. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 88

- (Topic 3)

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)

Answer: AD

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/>

NEW QUESTION 91

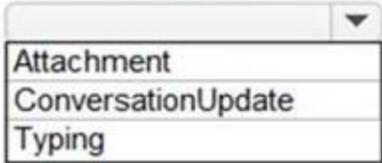
HOTSPOT - (Topic 3)

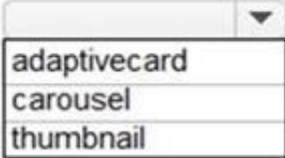
You are designing a conversation flow to be used in a chatbot.

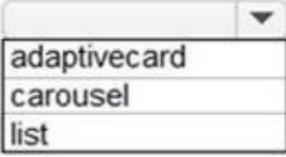
You need to test the conversation flow by using the Microsoft Bot Framework Emulator. How should you complete the .chat file? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

user=User1
bot=watchbot
user: I want a new watch.
bot: [  ][Delay=3000]

bot: I can help you with that! Let me see what I can find.
bot: Here's what I found.
bot:
[AttachmentLayout=  ]

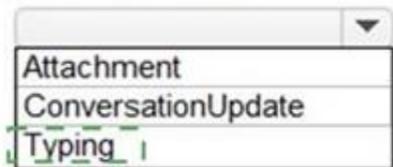
[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]
user: I like the first one.
bot: Sure, pulling up more information.
bot: [Attachment=cards\watchProfileCard.json  ]
user: That's nice! Thank you.
bot: Sure, you are most welcome!
    
```

- A. Mastered
- B. Not Mastered

Answer: A

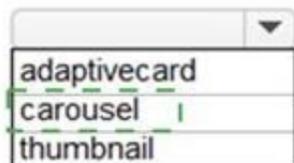
Explanation:

```
user=User1
bot=watchbot
user: I want a new watch.
bot: [ Attachment ] [Delay=3000]
```



```
bot: I can help you with that! Let me see what I can find.
bot: Here's what I found.
bot:
```

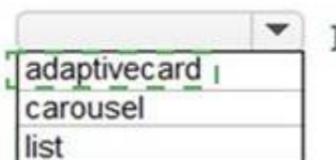
```
[AttachmentLayout=
```



```
]
[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]
```

```
user: I like the first one.
bot: Sure, pulling up more information.
```

```
bot: [Attachment=cards\watchProfileCard.json
```



```
user: That's nice! Thank you.
```

```
bot: Sure, you are most welcome!
```

NEW QUESTION 96

HOTSPOT - (Topic 3)

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<input type="radio"/>	<input type="radio"/>
Normalizing a database reduces data redundancy.	<input type="radio"/>	<input type="radio"/>
Normalization improves data integrity.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<input type="radio"/>	<input checked="" type="radio"/>
Normalizing a database reduces data redundancy.	<input checked="" type="radio"/>	<input type="radio"/>
Normalization improves data integrity.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 101

- (Topic 3)

You create a bot by using the Microsoft Bot Framework SDK.

You need to configure the bot to respond to events by using custom text responses. What should you use?

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

Answer: B

NEW QUESTION 103

- (Topic 3)

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

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

Answer: B

NEW QUESTION 107

- (Topic 3)

What is a characteristic of a non-relational database?

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

Answer: C

NEW QUESTION 110

- (Topic 3)

You have receipts that are accessible from a URL.

You need to extract data from the receipts by using Form Recognizer and the SDK. The solution must use a prebuilt model.

Which client and method should you use?

- A. the FormRecognizerClient client and the StartRecognizeReceiptsFromUri method
- B. the FormTrainingClient client and the StartRecognizeContentFromUri method
- C. the FormRecognizerClient client and the StartRecognizeReceiptsFromUri method
- D. the FormTrainingClient client and the StartRecognizeReceiptsFromUri method

Answer: D

Explanation:

To analyze receipts from a URL, use the StartRecognizeReceiptsFromUri method Example code:

```
private static async Task AnalyzeReceipt( FormRecognizerClient recognizerClient, string receiptUri)  
{
```

```
    RecognizedFormCollection receipts = await recognizerClient.StartRecognizeReceiptsFromUri(new Uri(receiptUri)).WaitForCompletionAsync();
```

Reference:

<https://docs.microsoft.com/en-us/azure/applied-ai-services/form-recognizer/quickstarts/client-library>

NEW QUESTION 112

- (Topic 3)

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

Answer: B

NEW QUESTION 114

- (Topic 3)

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

Answer: B

Explanation:

Reference:

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

ML Entity with Structure

An ML entity can be composed of smaller sub-entities, each of which can have its own properties. For example, Address could have the following structure:

Address: 4567 Main Street, NY, 98052, USA Building Number: 4567

Street Name: Main Street State: NY

Zip Code: 98052 Country: USA

NEW QUESTION 118

- (Topic 3)

Your company wants to reduce how long it takes for employees to log receipts in expense reports. All the receipts are in English.

You need to extract top-level information from the receipts, such as the vendor and the transaction total. The solution must minimize development effort.

Which Azure Cognitive Services service should you use?

- A. Custom Vision
- B. Personalizer
- C. Form Recognizer
- D. Computer Vision

Answer: B

Explanation:

Azure Form Recognizer is a cognitive service that lets you build automated data processing software using machine learning technology. Identify and extract text, key/value pairs, selection marks, tables, and structure from your documents—the service outputs structured data that includes the relationships in the original file, bounding boxes, confidence and more.

Form Recognizer is composed of custom document processing models, prebuilt models for invoices, receipts, IDs and business cards, and the layout model.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/form-recognizer>

NEW QUESTION 123

- (Topic 3)

You are building a flight booking bot by using the Microsoft Bot Framework SDK.

The bot will ask users for their departure date. The bot must repeat the question until a valid date is given, or the users cancel the transaction.

Which type of dialog should you use?

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

Answer: A

NEW QUESTION 124

- (Topic 3)

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 create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1.

You plan to create a new Azure Cognitive Search service named service1.

You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet.

Solution: You deploy service1 and a public endpoint to a new virtual network, and you configure Azure Private Link.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/cognitive-services-virtual-networks?tabs=portal#use-private-endpoints>

NEW QUESTION 129

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

A block of code that runs in a database is called

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A block of code that runs in a database is called

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

NEW QUESTION 132

- (Topic 3)

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.

Answer: A

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>

NEW QUESTION 137

- (Topic 3)

You have an Azure Cognitive Search solution and a collection of blog posts that include a category field. You need to index the posts. The solution must meet the following requirements:

- Include the category field in the search results.
- Ensure that users can search for words in the category field.
- Ensure that users can perform drill down filtering based on category. Which index attributes should you configure for the category field?

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

Answer: B

NEW QUESTION 141

HOTSPOT - (Topic 3)

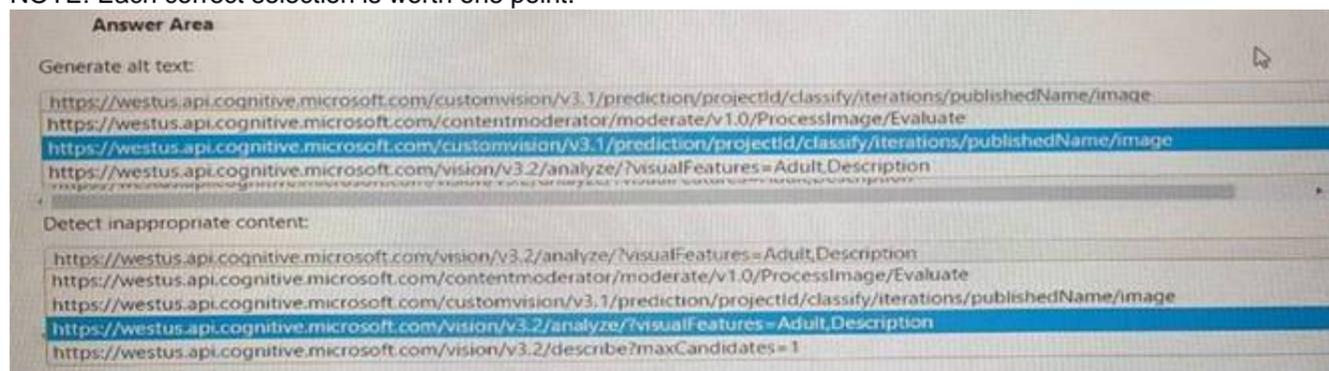
You are building an app that will enable users to upload images. The solution must meet the following requirements:

- 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.

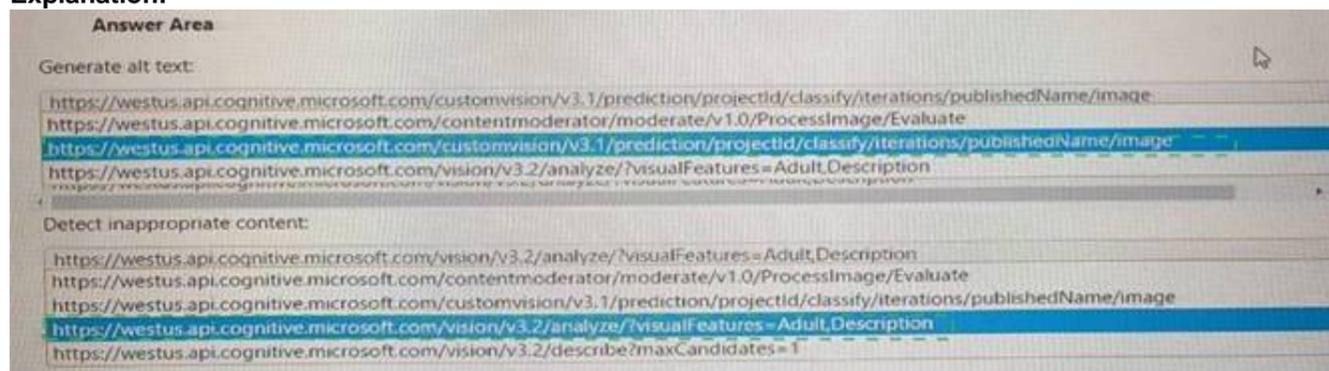
NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 143

- (Topic 3)

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 create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1.

You plan to create a new Azure Cognitive Search service named service1.

You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet.

Solution: You deploy service1 and a private endpoint to vnet1. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.

The service could be an Azure service such as:

- ? Azure Storage
- ? Azure Cosmos DB
- ? Azure SQL Database
- ? Your own service using a Private Link Service.

Reference:

<https://docs.microsoft.com/en-us/azure/private-link/private-endpoint-overview>

NEW QUESTION 146

DRAG DROP - (Topic 3)

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.

Values	Answer Area
<input type="text" value="faceListId"/>	<pre>{ "faceId": "18c51a87-3a69-47a8-aedc-a54745f708a1", <input type="text"/>: "employeefaces", "maxNumOfCandidatesReturned": 1, "mode": <input type="text"/> }</pre>
<input type="text" value="LargeFaceListId"/>	
<input type="text" value="matchFace"/>	
<input type="text" value="matchPerson"/>	

- A. Mastered
- B. Not Mastered

Answer: A

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.

NEW QUESTION 147

HOTSPOT - (Topic 3)

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

Passengers

- Sarah Hum
- Jeremy Goldberg
- Evan Litvak

2 Stops

Tue, May 30, 2017 10:25 PM

San Francisco Amsterdam		San Francisco Amsterdam
SFO AMS		SFO AMS

Non-Stop

Fri, Jun 2, 2017 11:55 PM

San Francisco Amsterdam		San Francisco Amsterdam
SFO AMS		SFO AMS

Total **\$4,032.54**

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer Area

The chatbot is showing **[answer choice]**.

	▼
an Adaptive Card	
a Hero Card	
a Thumbnail Card	

The card includes **[answer choice]**.

	▼
an action set	
an image	
an image group	
media	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Adaptive card Box 2: an image

NEW QUESTION 149

- (Topic 3)

You have the following C# method for creating Azure Cognitive Services resources programmatically.

```
static void create_resource(CognitiveServicesManagementClient client, string
resource_name, string kind, string account_tier, string location)
{
    CognitiveServicesAccount parameters =
        new CognitiveServicesAccount(null, null, kind, location, resource_name,
new CognitiveServicesAccountProperties(), new Sku(account_tier));
    var result = client.Accounts.Create(resource_group_name, account_tier,
parameters);
}
```

You need to call the method to create a free Azure resource in the West US Azure region. The resource will be used to generate captions of images automatically. Which code should you use?

- A. create_resource(client, "res1", "ComputerVision", "F0", "westus")
- B. create_resource(client, "res1", "CustomVision.Prediction", "F0", "westus")
- C. create_resource(client, "res1", "ComputerVision", "S0", "westus")
- D. create_resource(client, "res1", "CustomVision.Prediction", "S0", "westus")

Answer: B

Explanation:

<https://azure.microsoft.com/en-us/pricing/details/cognitive-services/computer-vision/>

NEW QUESTION 154

- (Topic 3)

You are building a solution that will detect anomalies in sensor data from the previous 24 hours.

You need to ensure that the solution scans the entire dataset, at the same time, for anomalies.

Which type of detection should you use?

- A. batch
- B. streaming
- C. change point

Answer: A

Explanation:

Batch anomaly detection is a type of anomaly detection that scans the entire dataset at once for outliers and unusual patterns. Batch anomaly detection is suitable for offline analysis of historical data, such as sensor data from the previous 24 hours. Batch anomaly detection can use various techniques, such as statistical methods, machine learning methods, or hybrid methods, to identify anomalies in the data.

NEW QUESTION 155

DRAG DROP - (Topic 3)

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.

Actions	Answer Area
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.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
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.	Install the Docker Engine on Server1.
From Server1, run the docker pull command.	From Server1, run the docker pull command.
From Server1, run the docker run command.	From Server1, run the docker run command.
Query the prediction endpoint on Server1.	Query the prediction endpoint on Server1.

NEW QUESTION 156

- (Topic 3)

You have an app that analyzes images by using the Computer Vision API.

You need to configure the app to provide an output for users who are vision impaired. The solution must provide the output in complete sentences.

Which API call should you perform?

- A. readInStreamAsync
- B. describeImageInStreamAsync
- C. toggleImageInStreamAsync
- D. analyzeImageByDomainInStreamAsync

Answer: D

NEW QUESTION 159

- (Topic 3)

You are building an app that will share user images.

You need to configure the app to perform the following actions when a user uploads an image:

- Categorize the image as either a photograph or a drawing.
- Generate a caption for the image.

The solution must minimize development effort.

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

NOTE: Each correct selection is worth one point.

- A. object detection in Computer Vision
- B. content tags in Computer Vision
- C. image descriptions in Computer Vision
- D. image type detection in Computer Vision
- E. image classification in Custom Vision

Answer: CD

Explanation:

? According to the Microsoft documentation, Computer Vision is a cloud-based service that provides developers with access to advanced algorithms for processing images and returning information. By uploading an image or specifying an image URL, Computer Vision algorithms can analyze visual content in different ways based on inputs and user choices.

? According to the Microsoft documentation, image type detection is one of the features of Computer Vision that can categorize an image as either a photograph or a drawing. You can use the image type detection feature by calling the Analyze Image API with the visualFeatures parameter set to ImageType. The API will return a JSON response with an imageType field that indicates whether the image is a photo or a clipart.

? According to the Microsoft documentation, image descriptions is another feature of Computer Vision that can generate a caption for an image. You can use the image descriptions feature by calling the Analyze Image API with the visualFeatures parameter set to Description. The API will return a JSON response with a description field that contains a list of captions for the image, each with a confidence score.

? Therefore, by using these two features of Computer Vision, you can achieve your app requirements with minimal development effort. You don't need to use any other services, such as object detection, content tags, or Custom Vision, which are designed for different purposes.

NEW QUESTION 162

DRAG DROP - (Topic 3)

Match the types of workloads to the appropriate scenarios.

To answer, drag the appropriate workload type from the column on the left to its scenario on the right. Each workload type may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Workload Types	Answer Area
Batch	Workload type: Data for a product catalog will be loaded every 12 hours to a data warehouse.
Streaming	Workload type: Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time.
	Workload type: Updates to inventory data will be loaded to a data warehouse every 1 million transactions.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Workload Types	Answer Area
Batch	Batch: Data for a product catalog will be loaded every 12 hours to a data warehouse.
Streaming	Streaming: Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time.
	Batch: Updates to inventory data will be loaded to a data warehouse every 1 million transactions.

NEW QUESTION 165

DRAG DROP - (Topic 3)

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.

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

- A. Mastered
- B. Not Mastered

Answer: A

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.

NEW QUESTION 170

DRAG DROP - (Topic 3)

You are building a customer support chatbot.

You need to configure the bot to identify the following:

- Code names for internal product development
 - Messages that include credit card numbers
- The solution must minimize development effort.

Which Azure Cognitive Service for Language feature should you use for each requirement? To answer, drag the appropriate features to the correct requirements.

Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content

NOTE: Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 174

DRAG DROP - (Topic 3)

You need to develop an automated call handling system that can respond to callers in their own language. The system will support only French and English. Which Azure Cognitive Services service should you use to meet each requirement? To answer, drag the appropriate services to the correct requirements. Each service may be used once, more than once, or not at all. You may need to drag the split bat between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Text Analytics

The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.

Box 2: Translator

Translator is a cloud-based neural machine translation service that is part of the Azure Cognitive Services family of REST APIs. Translator can be used with any operating system and powers many Microsoft products and services used by thousands of businesses worldwide to perform language translation and other language-related operations.

NEW QUESTION 177

- (Topic 3)

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 SOL query. What should you create?

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

Answer: B

NEW QUESTION 179

HOTSPOT - (Topic 3)

You plan to deploy a containerized version of an Azure Cognitive Services service that will be used for text analysis. You configure <https://contoso.cognitiveservices.azure.com> as the endpoint URI for the service, and you pull the latest version of the Text Analytics Sentiment Analysis container. You need to run the container on an Azure virtual machine by using Docker. How should you complete the command? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \
```

http://contoso.blob.core.windows.net
<https://contoso.cognitiveservices.azure.com>
mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment

```
Eula=accept \
```

Billing=

http://contoso.blob.core.windows.net
<https://contoso.cognitiveservices.azure.com>
mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment

```
ApiKey=xxxxxxxxxxxxxxxxxxxxxx
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: `mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment`
To run the Sentiment Analysis v3 container, execute the following docker run command. `docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \ mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment \ Eula=accept \ Billing={ENDPOINT_URI} \ ApiKey={API_KEY}` is the endpoint for accessing the Text Analytics API. `https://<your- custom-subdomain>.cognitiveservices.azure.com`
Box 2: `https://contoso.cognitiveservices.azure.com`
{ENDPOINT_URI} is the endpoint for accessing the Text Analytics API: `https://<your- custom-subdomain>.cognitiveservices.azure.com` The endpoint for accessing the Text Analytics API. `zure.com`

NEW QUESTION 182

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

In an infrastructure as a service (IaaS) instance of Microsoft SQL Server on Azure, you manage the _____ that hosts SQL Server.

elastic pool
MySQL server
PostgreSQL server
virtual machine

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

In an infrastructure as a service (IaaS) instance of Microsoft SQL Server on Azure, you manage the _____ that hosts SQL Server.

elastic pool
MySQL server
PostgreSQL server
virtual machine

NEW QUESTION 184

DRAG DROP - (Topic 3)

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

Actions

- Change Domains to **General (compact)**.
- Retrain the model.
- Create a new classification model.
- Change the classification type.
- Export the model.

Answer Area

⏪ ⏩

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 186

DRAG DROP - (Topic 3)

You have an Azure subscription that contains an Azure OpenAI resource named AH. You plan to build an app named App1 that will write press releases by using AM.

You need to deploy an Azure OpenAI model for App1. The solution must minimize development effort.

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 191

HOTSPOT - (Topic 3)

You are developing a streaming Speech to Text solution that will use the Speech SDK and MP3 encoding.

You need to develop a method to convert speech to text for streaming MP3 data. How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

var audioFormat = 
var speechConfig = SpeechConfig.FromSubscription("18c51a87-3a69-47a8-aedc-a54745f708a1", "westus");
var audioConfig = AudioConfig.FromStreamInput(pushStream, audioFormat);
using (var recognizer = new )
{
    var result = await recognizer.RecognizeOnceAsync();
    var text = result.Text;
}
    
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 196

HOTSPOT - (Topic 3)

You are building an app that will share user images.

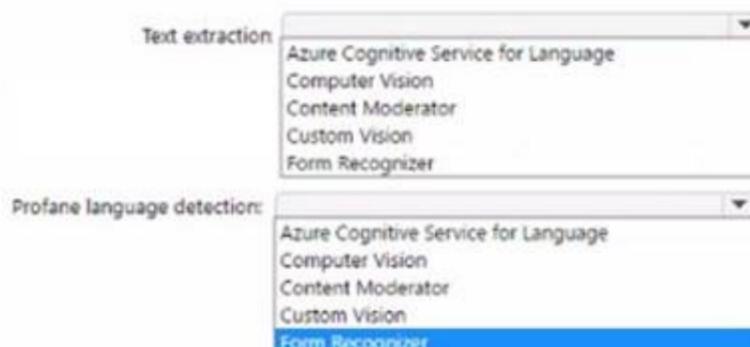
You need to configure the app to meet the following requirement's

- uploaded images must be scanned and any text must be extracted from the images.
- Extracted text must be analyzed for the presence of profane language.
- The solution must minimize development effort.

what should you use for each requirement? To answer, select the appropriate opinions in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Extraction: b) Computer Vision

Computer Vision is a service that can analyze images and extract text from them using optical character recognition (OCR) or read API1. OCR can detect and extract printed or handwritten text from images in various languages and formats, such as PDF, TIFF, or JPEG2. Read API can perform asynchronous batch processing of text from multiple images, and can also handle text in tables, mixed languages, and rotated text3.

Profane Language Detection: c) Content Moderator

Content Moderator is a service that can detect and filter out potentially offensive or unwanted content in text, images, and videos. For text, Content Moderator can identify profane or abusive language, personal data, and custom terms that you define. You can also use Content Moderator to review and moderate the content manually or automatically using workflows.

NEW QUESTION 201

- (Topic 3)

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

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

Answer: B

NEW QUESTION 202

- (Topic 3)

You successfully run the following HTTP request.

```
POST https://management.azure.com/subscriptions/18c51a87-3a69-47a8-aedc-a54745f708a1/resourceGroups/RG1/providers/Microsoft.CognitiveServices/accounts/contoso/regenerateKey?api-version=2017-04-18
Body{"keyName": "Key2"}
```

What is the result of the request?

- A. A key for Azure Cognitive Services was generated in Azure Key Vault.
- B. A new query key was generated.
- C. The primary subscription key and the secondary subscription key were rotated.
- D. The secondary subscription key was reset.

Answer: B

Explanation:

<https://docs.microsoft.com/en-us/rest/api/searchmanagement/2021-04-01-preview/query-keys/create>

NEW QUESTION 203

DRAG DROP - (Topic 3)

You are building a transcription service for technical podcasts.

Testing reveals that the service fails to transcribe technical terms accurately. You need to improve the accuracy of the service.

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

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions	Answer Area
Create a Speaker Recognition model.	<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;"> > < </div> <div style="text-align: center;"> ↑ ↓ </div> </div>
Create a Conversational Language Understanding model.	
Create a Custom Speech project.	
Create a speech-to-text model.	
Upload training datasets.	
Train the model.	
Deploy the model.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<https://learn.microsoft.com/en-us/azure/cognitive-services/speech-service/custom-speech-overview#how-does-it-work>

With Custom Speech, you can upload your own data, test and train a custom model, compare accuracy between models, and deploy a model to a custom endpoint.

- Create a project and choose a model. Use a Speech resource that you create in the Azure portal. If you will train a custom model with audio data, choose a Speech resource region with dedicated hardware for training audio data.
- Upload test data. Upload test data to evaluate the speech to text offering for your applications, tools, and products.
- Train a model. Provide written transcripts and related text, along with the corresponding audio data. Testing a model before and after training is optional but recommended.
- Deploy a model. Once you're satisfied with the test results, deploy the model to a custom endpoint. With the exception of batch transcription, you must deploy a custom endpoint to use a Custom Speech model.

NEW QUESTION 206

HOTSPOT - (Topic 3)

You are building a chatbot.

You need to use the Content Moderator API to identify aggressive and sexually explicit language.

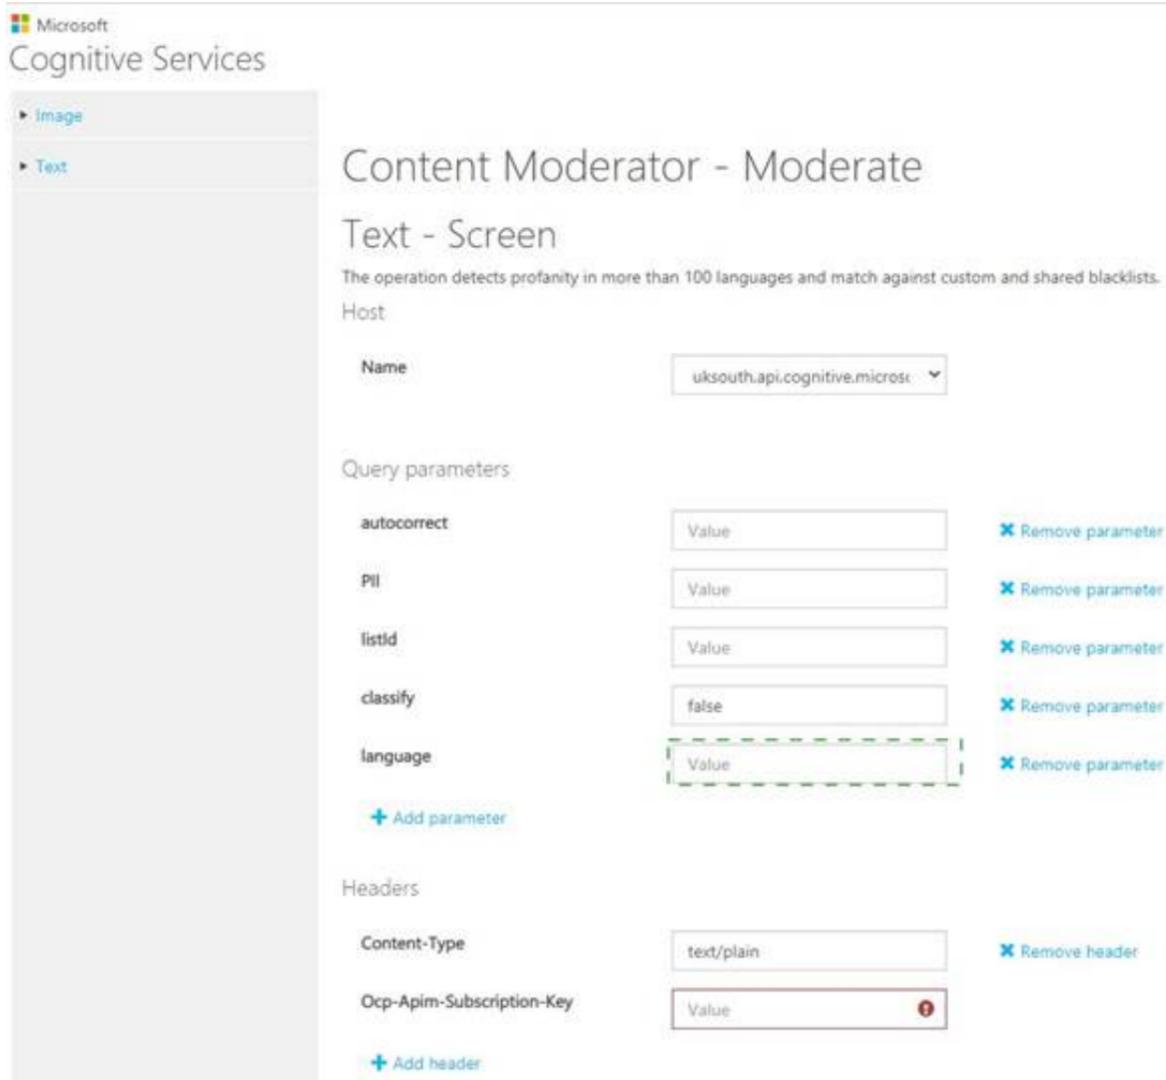
Which three settings should you configure? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 207

- (Topic 3)

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 container
- B. Create a Cognitive Search service for each containe
- C. Within each indexer definition, schedule the same runtime execution pattern.
- D. Split the data into multiple blob container
- E. Create an indexer for each containe
- F. Increase the search unit
- G. Within each indexer definition, schedule a sequential execution pattern.
- H. Create a Cognitive Search service for each type of document.
- I. Split the data into multiple virtual folder
- J. Create an indexer for each folde
- K. Increase the search unit
- L. Within each indexer definition, schedule the same runtime execution pattern.

Answer: D

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-howto-indexing-azure-blob-storage>

NEW QUESTION 212

- (Topic 3)

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

Answer: D

NEW QUESTION 216

HOTSPOT - (Topic 3)

A1 You have an Azure OpenAI resource named AH that hosts three deployments of the GPT 3.5 model. Each deployment is optimized for a unique workload.

You plan to deploy three apps. Each app will access AM by using the REST API and will use the deployment that was optimized for the apps intended workload.

You need to provide each app with access to AH and the appropriate deployment. The solution must ensure that only the apps can access AM.

What should you use to provide access to AM. and what should each app use to connect to its appropriate deployment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Provide access to AI1 by using:

Connect to the deployment by using:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Provide access to AI1 by using:

Connect to the deployment by using:

NEW QUESTION 217

- (Topic 3)

You have a mobile app that manages printed forms.

You need the app to send images of the forms directly to Forms Recognizer to extract relevant information. For compliance reasons, the image files must not be stored in the cloud.

In which format should you send the images to the Form Recognizer API endpoint?

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

Answer: A

NEW QUESTION 221

- (Topic 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 = txtHeaders.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.

Answer: BD

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>

NEW QUESTION 224

- (Topic 3)

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

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

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

You discover that some search query requests to the Cognitive Search service are being throttled.

You need to reduce the likelihood that search query requests are throttled. Solution: You migrate to a Cognitive Search service that uses a higher tier. Does this meet the goal?

A. Yes

B. No

Answer: A

Explanation:

A simple fix to most throttling issues is to throw more resources at the search service (typically replicas for query-based throttling, or partitions for indexing-based throttling). However, increasing replicas or partitions adds cost, which is why it is important to know the reason why throttling is occurring at all.

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-performance-analysis>

NEW QUESTION 228

- (Topic 3)

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.

Answer: CDF

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>

NEW QUESTION 230

DRAG DROP - (Topic 3)

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.)

Actions

Answer Area

- 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.

- A. Mastered
- B. Not Mastered

Answer: A

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.

NEW QUESTION 233

HOTSPOT - (Topic 3)

You create a knowledge store for Azure Cognitive Search by using the following JSON.

```

"knowledgeStore" : {
  "storageConnectionString": "DefaultEndpointsProtocol=https;AccountName=<Acct Name>;AccountKey=<Acct Key>;",
  "projections": [
    {
      "tables": [
        {
          "tableName": "unrelatedDocument",
          "generatedKeyName": "Documentid",
          "source": "/document/pbiShape"
        },
        {
          "tableName": "unrelatedKeyPhrases",
          "generatedKeyName": "KeyPhraseid",
          "source": "/document/pbiShape/keyPhrases"
        }
      ]
    },
    {
      "storageContainer": "unrelatedocrlayout",
      "source": null,
      "sourceContext": "/document/normalized_images/*/layoutText",
      "inputs": [
        {
          "name": "ocrLayoutText",
          "source": "/document/normalized_images/*/layoutText"
        }
      ]
    }
  ],
  "files": []
}
    
```

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

Answer Area

There will be [answer choice].

- no projection groups
- one projection group
- two projection groups
- four projection groups

Images will [answer choice]

- not be saved.
- be saved to a blob container.
- be saved to file storage.
- be saved to an Azure Data lake.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

There will be [answer choice].

- no projection groups
- one projection group
- two projection groups
- four projection groups

Images will [answer choice]

- not be saved.
- be saved to a blob container.
- be saved to file storage.
- be saved to an Azure Data lake.

NEW QUESTION 234

- (Topic 3)

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

Answer: AD

NEW QUESTION 235

DRAG DROP - (Topic 3)

You need to analyze video content to identify any mentions of specific company names.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 237

- (Topic 3)

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

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

You have a chatbot that uses question answering in Azure Cognitive Service for Language Users report that the responses of the chatbot lack formality when answering spurious questions

You need to ensure that the chatbot provides formal responses to spurious questions. Solution: From Language Studio, you remove all the chit-chat question and answer pairs, and then retrain and republish the model Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Removing all the chit-chat question and answer pairs from the project will not ensure that the chatbot provides formal responses to spurious questions. It will only make the chatbot unable to handle any chit-chat scenarios, which may result in a poor user experience and a loss of engagement. Instead, you should choose a chit-chat personality that matches the tone and style of your chatbot, such as Professional or Caring. You can also edit the chit-chat questions and answers to suit your specific needs, or add new ones that are not in the predefined data set. This way, you can ensure that the chatbot responds appropriately to spurious questions, while still maintaining a conversational and engaging interaction with the user.

NEW QUESTION 242

- (Topic 3)

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

- A. Content Moderator
- B. Form Recognizer

- C. Computer Vision
- D. Language service

Answer: D

NEW QUESTION 247

HOTSPOT - (Topic 3)

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.

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)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Classification

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.

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.

NEW QUESTION 250

- (Topic 3)

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

Answer: B

NEW QUESTION 251

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

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

<ul style="list-style-type: none"> changes to relationships between entities. geographically distributed writes. transactional writes. writes that have varying data structures.
--

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

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

NEW QUESTION 256

- (Topic 3)

You are building an Azure WebJob 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
- C. Knowledgebas
- D. CreateAsync method.
- E. Create a list of QnADTO objects that represents data from the WebJob.
- F. Create a CreateKbDTO object.

Answer: AC

Explanation:

Reference:

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

NEW QUESTION 259

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

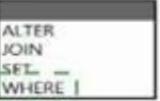
The  clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

The  clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

NEW QUESTION 263

- (Topic 3)

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

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

Answer: B

NEW QUESTION 266

DRAG DROP - (Topic 3)

You have a factory that produces cardboard packaging for food products. The factory has intermittent internet connectivity.

The packages are required to include four samples of each product.

You need to build a Custom Vision model that will identify defects in packaging and provide the location of the defects to an operator. The model must ensure that each package contains the four products.

Which project type and domain should you use? To answer, drag the appropriate options to the correct targets. Each option may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content

NOTE: Each correct selection is worth one point.

Options	Answer Area
Food	Project type: <input type="text"/>
General	Domain: <input type="text"/>
General (compact)	
Image classification	
Logo	
Object detection	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Options	Answer Area
Food	Project type: Object detection
General	Domain: Food
General (compact)	
Image classification	
Logo	
Object detection	

NEW QUESTION 269

- (Topic 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

Answer: ADF

Explanation:

C: textType is an optional parameter. It defines whether the text being translated is plain text or HTML text (used for web pages).

D: to is a required parameter. It specifies the language of the output text. The target language must be one of the supported languages included in the translation scope.

F: toScript is an optional parameter. It specifies the script of the translated text. We use Latin (Roman alphabet) script.

Reference:

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

NEW QUESTION 274

- (Topic 2)

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/v3.2/read/analyzeResults
- B. /formrecognizer/v2.0/prebuilt/receipt/analyze
- C. /vision/v3.2/read/analyze
- D. /vision/v3.2/describe
- E. /formercognizer/v2.0/custom/models{modelId}/ analyze

Answer: BC

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/v3.2/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-v2- preview/operations/GetAnalyzeReceiptResult>

<https://docs.microsoft.com/en-us/rest/api/computervision/3.1/read/read>

NEW QUESTION 275

- (Topic 2)

You are developing the chatbot.

You create the following components:

* A QnA Maker resource

* A chatbot by using the Azure Bot Framework SDK.

You need to integrate the components to meet the chatbot requirements. Which property should you use?

- A. QnADialogResponseOptions.CardNoMatchText
- B. Qna MakerOptions-ScoreThreshold
- C. Qna Maker Op t ions StrickFilters
- D. QnaMakerOptions.RankerType

Answer: D

Explanation:

Scenario: When the response confidence score is low, ensure that the chatbot can provide other response options to the customers.

When no good match is found by the ranker, the confidence score of 0.0 or "None" is returned and the default response is "No good match found in the KB". You can override this default response in the bot or application code calling the endpoint. Alternately, you can also set the override response in Azure and this changes the default for all knowledge bases deployed in a particular QnA Maker service.

Choosing Ranker type: By default, QnA Maker searches through questions and answers. If you want to search through questions only, to generate an answer, use the RankerType=QuestionOnly in the POST body of the GenerateAnswer request.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/concepts/best- practices>

NEW QUESTION 277

- (Topic 2)

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

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

Answer: C

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.

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>

NEW QUESTION 282

- (Topic 2)

You are developing the knowledgebase.

You use Azure Video Analyzer for Media (previously Video indexer) to obtain transcripts of webinars.

You need to ensure that the solution meets the knowledgebase requirements. What should you do?

- A. Create a custom language model
- B. Configure audio indexing for videos only
- C. Enable multi-language detection for videos
- D. Build a custom Person model for webinar presenters

Answer: B

Explanation:

Can search content in different formats, including video

Audio and video insights (multi-channels). When indexing by one channel, partial result for those models will be available.

Keywords extraction: Extracts keywords from speech and visual text.

Named entities extraction: Extracts brands, locations, and people from speech and visual text via natural language processing (NLP).

NEW QUESTION 285

HOTSPOT - (Topic 1)

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.

Answer Area

		<code>?api-version=3.0&to=es&to=pt</code>
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	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: api-nam.cognitive.microsofttranslator.com

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

Box 2: /translate

NEW QUESTION 290

HOTSPOT - (Topic 1)

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

You are configuring access to the QnA Maker resources.

Which role should you assign to AllUsers and LeadershipTeam? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

AllUsers:	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">Cognitive Service User</div> <div style="padding: 2px;">Contributor</div> <div style="padding: 2px;">Owner</div> <div style="padding: 2px;">QnA Maker Editor</div> <div style="padding: 2px;">QnA Maker Read</div> </div>
LeadershipTeam:	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">Cognitive Service User</div> <div style="padding: 2px;">Contributor</div> <div style="padding: 2px;">Owner</div> <div style="padding: 2px;">QnA Maker Editor</div> <div style="padding: 2px;">QnA Maker Read</div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: QnA Maker Editor

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

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

Update KB API Replace KB API Replace Alterations "Train API" [in

new service model v5]

Box 2: Contributor

Scenario: Only senior managers must be able to publish updates. Contributor permission: All except ability to add new members to roles

NEW QUESTION 292

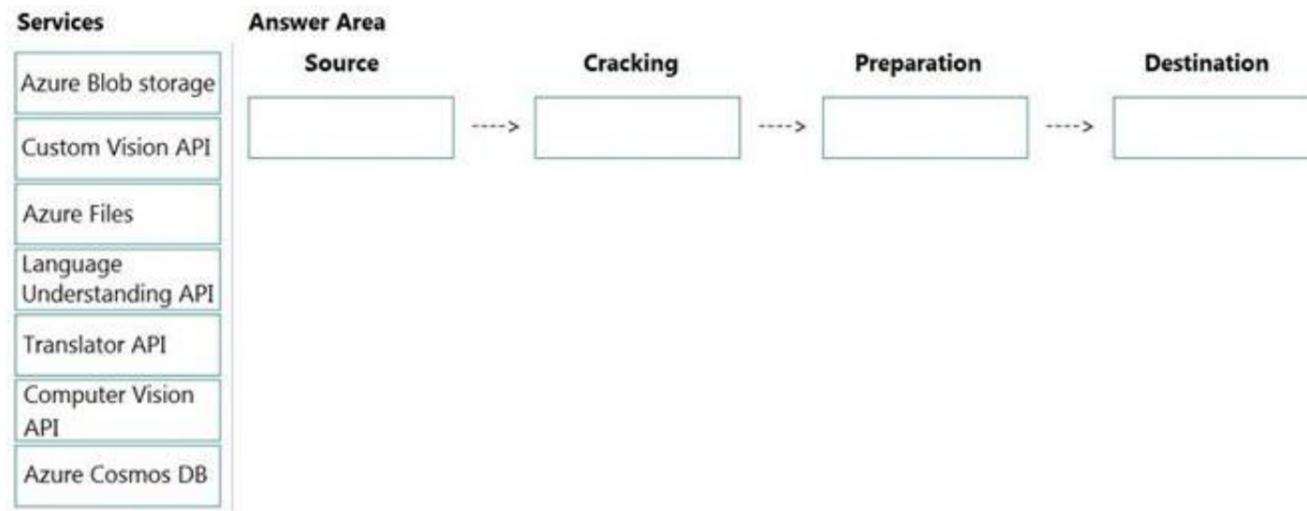
DRAG DROP - (Topic 1)

You are developing the smart e-commerce project.

You need to design the skillset to include the contents of PDFs in searches.

How should you complete the skillset design diagram? To answer, drag the appropriate services to the correct stages. Each service 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.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Blob storage

At the start of the pipeline, you have unstructured text or non-text content (such as images, scanned documents, or JPEG files). Data must exist in an Azure data storage service that can be accessed by an indexer.

Box 2: Computer Vision API

Scenario: Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products.

The Computer Vision Read API is Azure's latest OCR technology (learn what's new) that extracts printed text (in several languages), handwritten text (English only), digits, and currency symbols from images and multi-page PDF documents.

Box 3: Translator API

Scenario: Product descriptions, transcripts, and all text must be available in English, Spanish, and Portuguese.

Box 4: Azure Files

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

NEW QUESTION 297

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AI-102 Exam with Our Prep Materials Via below:

<https://www.certleader.com/AI-102-dumps.html>