

Microsoft

Exam Questions AI-102

Designing and Implementing an Azure AI Solution



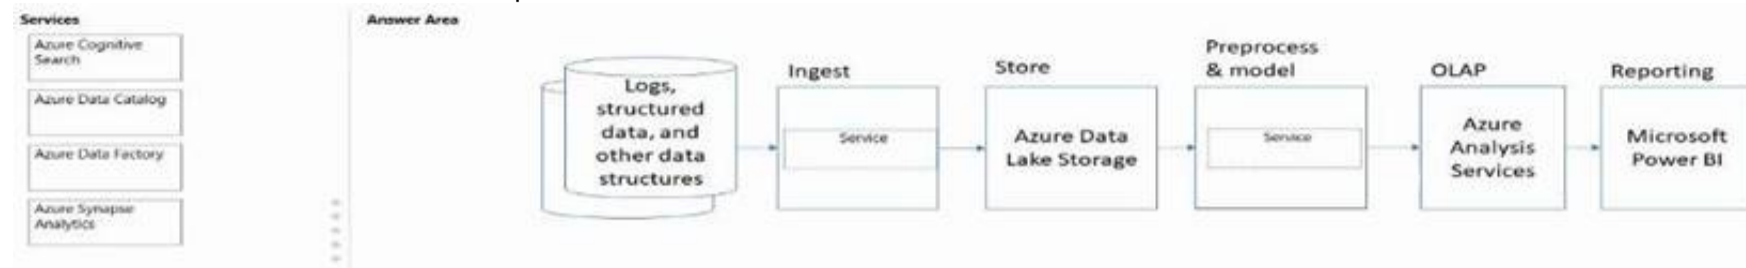
NEW QUESTION 1

DRAG DROP - (Topic 3)

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

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

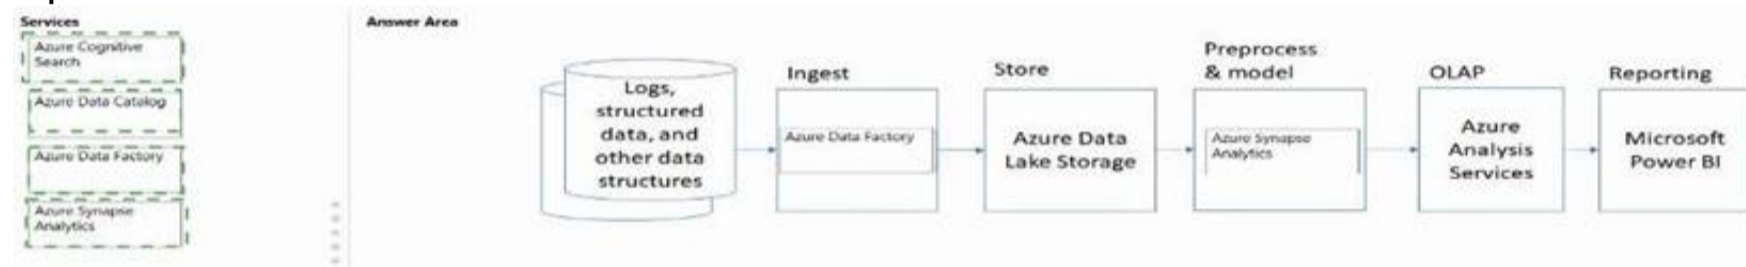
NOTE: Each correct match is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 2

- (Topic 3)

You need to enable speech capabilities for a chatbot.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable WebSockets for the chatbot app.
- B. Create a Speech service.
- C. Register a Direct Line Speech channel.
- D. Register a Cortana channel.
- E. Enable CORS for the chatbot app.
- F. Create a Language Understanding service.

Answer: ABC

Explanation:

You can use the Speech service to voice-enable a chat bot.

The Direct Line Speech channel uses the text-to-speech service, which has neural and standard voices.

You'll need to make a small configuration change so that your bot can communicate with the Direct Line Speech channel using web sockets.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/tutorial-voice-enable-your-bot-speech-sdk>

NEW QUESTION 3

- (Topic 3)

You have an Azure IoT hub that receives series data from machinery. You need to build an app that will perform the following actions:

- Perform anomaly detection across multiple correlated sensors
- Identify the root cause of process stops.
- Send incident alerts

The solution must minimize development time. Which Azure service should you use?

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

Answer: D

NEW QUESTION 4

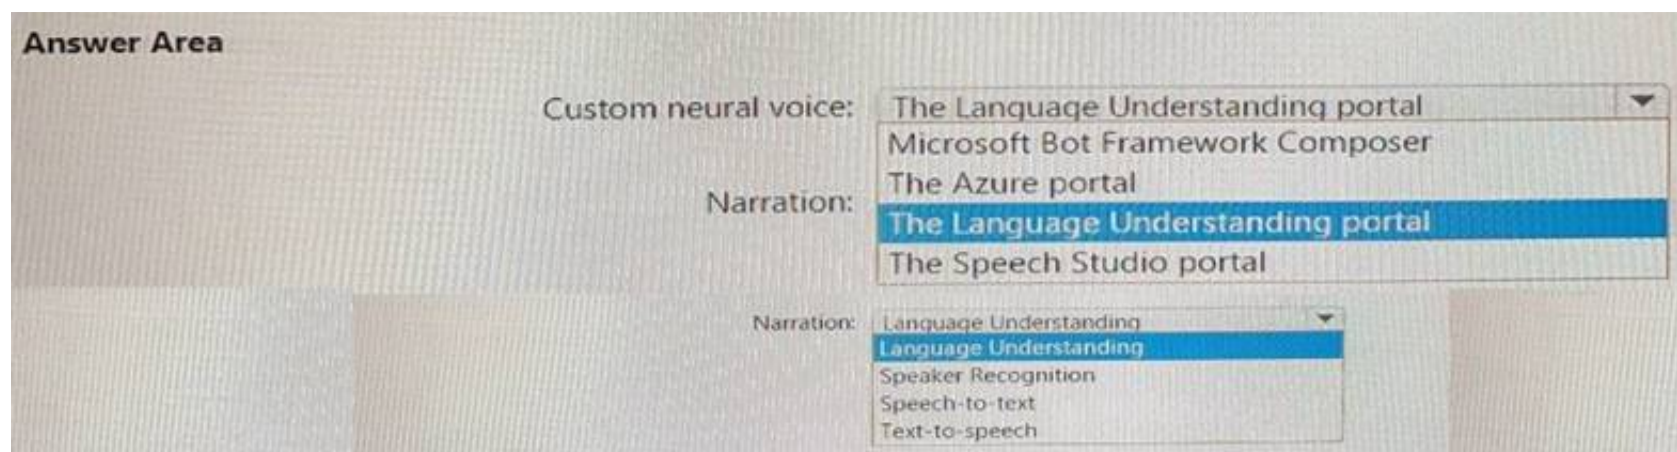
HOTSPOT - (Topic 3)

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.

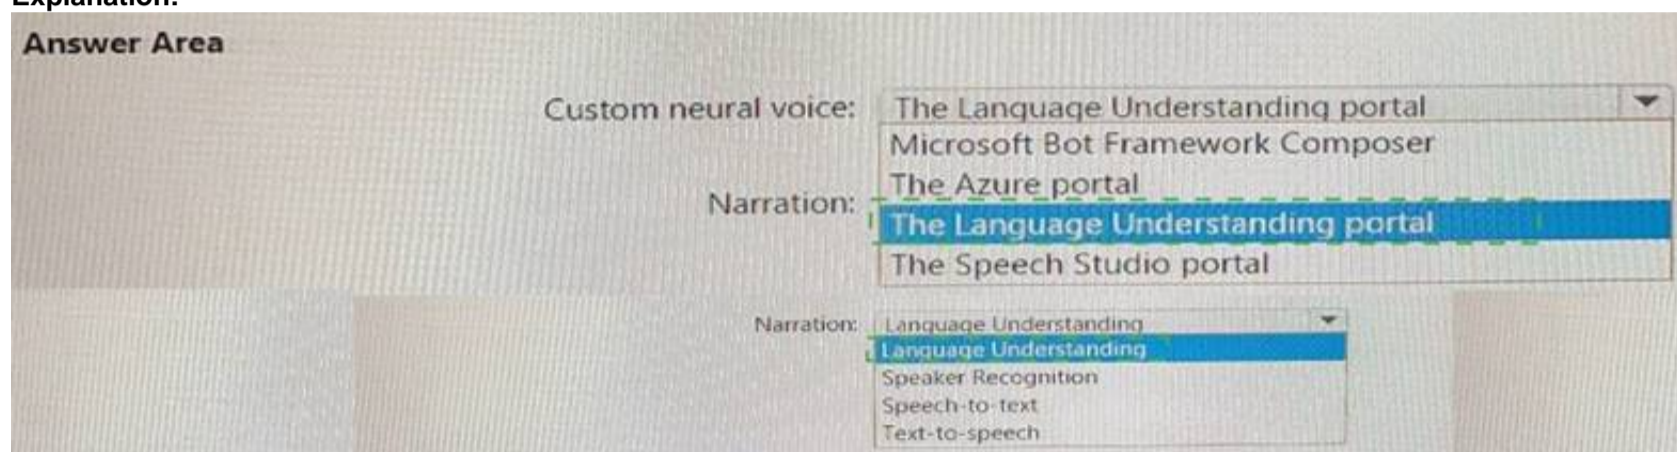
NOTE: Each correct answer is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 5

- (Topic 3)

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

You need to ensure that App1 will NOT return answers that include hate speech. What should you configure for Model1?

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

Answer: B

NEW QUESTION 6

- (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 7

- (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 8

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 9

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 10

HOTSPOT - (Topic 3)

You have an Azure Cognitive Search resource named Search 1 that is used by multiple apps. You need to secure Search 1. The solution must meet the following requirements:

- Prevent access to Search1 from the internet.
- Limit the access of each app to specific queries.

What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct answer is worth one point.

Answer Area

To prevent access from the internet: 


To limit access to queries: 


- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Answer Area

To prevent access from the internet: 

To limit access to queries: 

NEW QUESTION 10

- (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 11

- (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 enable customer-managed key (CMK) encryption.

Does this meet the goal?

- A. Yes
 B. No

Answer: B

Explanation:

Customer-managed key (CMK) encryption does not affect throttling.

Instead, you could migrate to a Cognitive Search service that uses a higher tier.

Note: 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 16

DRAG DROP - (Topic 3)

You plan to use containerized versions of the Anomaly Detector API on local devices for testing and in on-premises datacenters.

You need to ensure that the containerized deployments meet the following requirements:

? Prevent billing and API information from being stored in the command-line histories of the devices that run the container.
? Control access to the container images by using Azure role-based access control (Azure RBAC).
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.)
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions

Create a custom Dockerfile.

Pull the Anomaly Detector container image.

Distribute a docker run script.

Push the image to an Azure container registry.

Build the image.

Push the image to Docker Hub.

Answer Area

- A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Pull the Anomaly Detector container image. Step 2: Create a custom Dockerfile
Step 3: Build the image
Step 4: Push the image to an Azure container registry.
<https://docs.microsoft.com/en-us/azure/cognitive-services/containers/container-reuse-recipe>

NEW QUESTION 19

- (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 20

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
Stream processing has access to the most recent data received or data within a rolling time window.	<input type="radio"/>	<input type="radio"/>
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	<input type="radio"/>	<input type="radio"/>
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements

Stream processing has access to the most recent data received or data within a rolling time window.

Batch processing must occur immediately and have latency in the order of seconds or milliseconds.

Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.

Yes

No

NEW QUESTION 25

DRAG DROP - (Topic 3)

You have an app that manages feedback.

You need to ensure that the app can detect negative comments by using the Sentiment Analysis API in Azure Cognitive Service for Language. The solution must ensure that the managed feedback remains on your company's internal network.

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.

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

Actions

Provision the Language service resource in Azure.

Deploy a Docker container to an Azure container instance.

Deploy a Docker container to an on-premises server.

Identify the Language service endpoint URL and query the prediction endpoint.

Run the container and query the prediction endpoint.

>

<

Answer Area

<

>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

? Provision the Language service resource in Azure.

? Deploy a Docker container to an on-premises server.

? Run the container and query the prediction endpoint.

? According to the Microsoft documentation, the Language service is a cloud-based service that provides various natural language processing features, such as sentiment analysis, key phrase extraction, named entity recognition, etc. You can provision the Language service resource in Azure by following the steps in Create a Language resource. You will need to provide a name, a subscription, a resource group, a region, and a pricing tier for your resource. You will also get a key and an endpoint for your resource, which you will use to authenticate your requests to the Language service API.

? According to the Microsoft documentation, you can also use the Language service as a container on your own premises or in another cloud. This option gives you more control over your data and network, and allows you to use the Language service without an internet connection. You can deploy a Docker container to an on-premises server by following the steps in Deploy Language containers. You will need to have Docker installed on your server, pull the container image from the Microsoft Container Registry, and run the container with the appropriate parameters. You will also need to activate your container with your key and endpoint from your Azure resource.

? According to the Microsoft documentation, once you have deployed and activated your container, you can run it and query the prediction endpoint to get sentiment analysis results. The prediction endpoint is a local URL that follows this format: http://<container IP address>:<port>/text/analytics/v3.1-preview.4/sentiment. You can send HTTP POST requests to this endpoint with your text input in JSON format, and receive JSON responses with sentiment labels and scores for each document and sentence in your input.

NEW QUESTION 28

- (Topic 3)

You are examining the Language service output of an application.

The text analyzed is: Our tour guide took us up the Space Needle during our trip to Seattle last week.

The response contains the data shown in the following table.

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

Which Language service API is used to analyze the Text?

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

Answer: B

NEW QUESTION 33

- (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 35

HOTSPOT - (Topic 3)

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

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

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

## Logout
- forget me
- log out

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

## SelectNone
- none

@ ml DirectionalReference
@ prebuilt personName
```

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

NOTE: Each correct selection is worth one point.

SelectItem is [answer choice].

Choose {@DirectionalReference=top right} is [answer choice].

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SelectItem is [answer choice].

a domain

an entity

an intent

an utterance

Choose {@DirectionalReference=top right} is [answer choice].

a domain

an entity

an intent

an utterance

NEW QUESTION 37

DRAG DROP - (Topic 3)

You have a collection of Microsoft Word documents and PowerPoint presentations in German. You need to create a solution to translate the file to French. The solution must meet the following requirements:

- * Preserve the original formatting of the files.
- * Support the use of s custom glossary.

You create a blob container for German files and a blob container for French files. You upload the original files to the container for German files. Which three actions should you perform in sequence actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Perform an asynchronous translation by using the list of files to be translated.

Upload a glossary file to the container for German files.

Upload a glossary file to the container for French files.

Generate a list of files to be translated.

Define a document translation specification that has a French target.

Perform an asynchronous translation by using the document translation specification.

Answer Area

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions

Perform an asynchronous translation by using the list of files to be translated.

Upload a glossary file to the container for German files.

Upload a glossary file to the container for French files.

Generate a list of files to be translated.

Define a document translation specification that has a French target.

Perform an asynchronous translation by using the document translation specification.

Answer Area

Generate a list of files to be translated.

Define a document translation specification that has a French target.

Perform an asynchronous translation by using the document translation specification.

NEW QUESTION 38

- (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 41

HOTSPOT - (Topic 3)

You are developing an application that includes language translation. The application will translate text retrieved by using a function named get_text_to_be_translated. 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.

Answer Area

```

. . .
api_key = "FF956C68B83B21838691A8D200A4C606"
text = get_text_to_be_translated()
headers = {
    'Content-Type': 'application/json',
    'Ocp-Apim-Subscription-Key': api_key
}
body = {
    'Text': text
}
conn = httplib.HTTPSConnection
conn.request("POST",
response = conn.getresponse()
    conn.request("POST",
response = conn.getresponse()
response_data = response.r
. . .

```



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

. . .
api_key = "FF956C68B83B21838691A8D200A4C606"
text = get_text_to_be_translated()
headers = {
    'Content-Type': 'application/json',
    'Ocp-Apim-Subscription-Key': api_key
}
body = {
    'Text': text
}
conn = httplib.HTTPSConnection
conn.request("POST",
response = conn.getresponse()
    conn.request("POST",
response = conn.getresponse()
response_data = response.r
. . .

```



NEW QUESTION 44

- (Topic 3)

You have an Azure subscription that contains a Language service resource named ta1 and a virtual network named vnet1. You need to ensure that only resources in vnet1 can access ta1. What should you configure?

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

Answer: C

NEW QUESTION 49

- (Topic 3)

You have a Language Understanding resource named lu1. You build and deploy an Azure bot named bot1 that uses lu1. You need to ensure that bot1 adheres to the Microsoft responsible AI principle of inclusiveness. How should you extend bot1?

- A. Implement authentication for bot1.
- B. Enable active learning for lu1.
- C. Host lu1 in a container.
- D. Add Direct Line Speech to bot1.

Answer: D

Explanation:

Inclusiveness: AI systems should empower everyone and engage people.

Direct Line Speech is a robust, end-to-end solution for creating a flexible, extensible voice assistant. It is powered by the Bot Framework and its Direct Line Speech channel, that is optimized for voice-in, voice-out interaction with bots.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/direct-line- speech>

NEW QUESTION 53

- (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 56

DRAG DROP - (Topic 3)

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

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

Actions

Configure disk logging.

Export model1 to Host1.

Request approval to run the container.

Retrain the model.

Run the container.

Answer Area

>

<

^

v

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

According to the course AI-102T00: Designing and Implementing a Microsoft Azure AI Solution1, the correct sequence of actions to run model1 on Host1 is:
? Export model1 to Host1
? Run the container
? Configure disk logging
The explanation and references are as follows:
? Export model1 to Host1: This step is required to deploy the custom speech-to-text model to the Docker host. You can use the Azure portal or the Azure CLI to export the model as a container image2.
? Run the container: This step is required to start the container and run the model on the Docker host. You can use the Docker CLI to run the container image3.
? Configure disk logging: This step is optional but recommended to monitor the performance and health of the container. You can use the Docker CLI to configure disk logging for the container4.

NEW QUESTION 57

HOTSPOT - (Topic 3)

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

Answer Area

Statements	Yes	No
The output will include the following words: our and included.	<input type="radio"/>	<input type="radio"/>
The output will include the following words: Paris, Eiffel, and Tower.	<input type="radio"/>	<input type="radio"/>
The function will output all the key phrases from the input string to the console.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The output will include the following words: our and included.	<input checked="" type="radio"/>	<input type="radio"/>
The output will include the following words: Paris, Eiffel, and Tower.	<input checked="" type="radio"/>	<input type="radio"/>
The function will output all the key phrases from the input string to the console.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 59

- (Topic 3)

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

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

Answer: BC

NEW QUESTION 62

HOTSPOT - (Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

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

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

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

    using var audioConfig = AudioConfig.FromDefaultMicrophoneInput();
    using var recognizer = new (config, audioConfig);

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

Box 2: TranslationRecognizer

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

Example code:

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

NEW QUESTION 63

- (Topic 3)

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

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

Answer: A

NEW QUESTION 64

- (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 67

- (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 70

- (Topic 3)

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

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

Answer: C

NEW QUESTION 75

- (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 77

- (Topic 3)

You plan to perform predictive maintenance.

You collect IoT sensor data from 100 industrial machines for a year. Each machine has 50 different sensors that generate data at one-minute intervals. In total, you have 5,000 time series datasets.

You need to identify unusual values in each time series to help predict machinery failures. Which Azure Cognitive Services service should you use?

- A. Anomaly Detector
- B. Cognitive Search
- C. Form Recognizer
- D. Custom Vision

Answer: A

NEW QUESTION 79


HOTSPOT - (Topic 3)

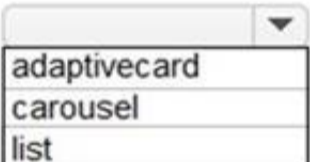
You are designing a conversation flow to be used in a chatbot.

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]

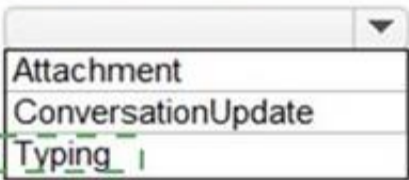
bot: I can help you with that! Let me see what I can find.
bot: Here's what I found.
bot:
[AttachmentLayout=  ]

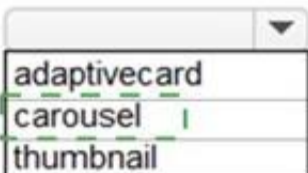
[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]
user: I like the first one.
bot: Sure, pulling up more information.
bot: [Attachment=cards\watchProfileCard.json  ]
user: That's nice! Thank you.
bot: Sure, you are most welcome!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


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
user=User1
bot=watchbot
user: I want a new watch.
bot: [  ][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 81

- (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 85

HOTSPOT - (Topic 3)

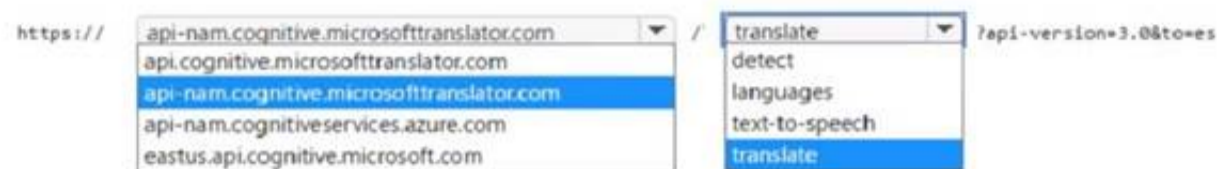
You are building an Azure web app named App1 that will translate text from English to Spanish.

You need to use the Text Translation REST API to perform the translation. The solution must ensure that you have data sovereignty in the United States.

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

NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

* 1. api-nam.cognitive.microsofttranslator.com

* 2. translate

<https://learn.microsoft.com/en-us/azure/cognitive-services/Translator/reference/v3-0-reference#base-urls>

Requests to Translator are, in most cases, handled by the datacenter that is closest to where the request originated. If there's a datacenter failure when using the global endpoint, the request may be routed outside of the geography.

To force the request to be handled within a specific geography, use the desired geographical endpoint. All requests are processed among the datacenters within the geography.

- United States

api-nam.cognitive.microsofttranslator.com

<https://learn.microsoft.com/en-us/azure/cognitive-services/translator/reference/rest-api-guide-translate>

Translate specified source language text into the target language text.

NEW QUESTION 87

- (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 92

- (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 97

- (Topic 3)

You plan create an index for an Azure Cognitive Search service by using the Azure portal. The Cognitive Search service will connect to an Azure SQL database

The Azure SQL database contains a table named UserMessages. Each row in User

Messages has a field named MessageCopy that contains the text of social media messages sent by a user

Users will perform full text searches against the MessageCopy field, and the values of the field will be shown to the users-

You need to configure the properties of the index for the MessageCopy field to support the solution.
 Which attributes should you enable for the field?

- A. Searchable and Retrievable
- B. Sortable and Retrievable
- C. Searchable and Facetable
- D. Filterable and Retrievable

Answer: A

NEW QUESTION 101

HOTSPOT - (Topic 3)

You run the following command.

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

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

Box 2: Yes

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

Box 3: Yes

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

NEW QUESTION 105

- (Topic 3)

You are building a bot on a local computer by using the Microsoft Bot Framework. The bot will use an existing Language Understanding model.

You need to translate the Language Understanding model locally by using the Bot Framework CLI. What should you do first?

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

Answer: B

Explanation:

You might want to manage the translation and localization for the language understanding content for your bot independently.

Translate command in the `@microsoft/bf-lu` library takes advantage of the Microsoft text translation API to automatically machine translate .lu files to one or more than 60+ languages supported by the Microsoft text translation cognitive service.

What is translated?

An .lu file and optionally translate Comments in the lu file LU reference link texts

List of .lu files under a specific path.

Reference:

<https://github.com/microsoft/botframework-cli/blob/main/packages/luis/docs/translate-command.md>

NEW QUESTION 106

HOTSPOT - (Topic 3)

You are building an app by using the Speech SDK. The app will translate speech from French to German by using natural language processing.

You need to define the source language and the output 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.

Answer Area

```
var speechTranslationConfig =
SpeechTranslationConfig.FromSubscription(speechKey, speechRegion);
speechTranslationConfig.
```

SpeechRecognitionLanguage

="fr"

SpeechRecognitionLanguage

AddTargetLanguage

SpeechSynthesisLanguage

TargetLanguages

VoiceName

```
speech_translation_config.
```

SpeechSynthesisLanguage

("de")

SpeechSynthesisLanguage

AddTargetLanguage

SpeechRecognitionLanguage

SpeechSynthesisLanguage

TargetLanguages

VoiceName

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Answer Area

```
var speechTranslationConfig =
SpeechTranslationConfig.FromSubscription(speechKey, speechRegion);
speechTranslationConfig.
```

SpeechRecognitionLanguage

="fr"

SpeechRecognitionLanguage

AddTargetLanguage

SpeechSynthesisLanguage

TargetLanguages

VoiceName

```
speech_translation_config.
```

SpeechSynthesisLanguage

("de")

SpeechSynthesisLanguage

AddTargetLanguage

SpeechRecognitionLanguage

SpeechSynthesisLanguage

TargetLanguages

VoiceName

NEW QUESTION 107

- (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 109

HOTSPOT - (Topic 3)

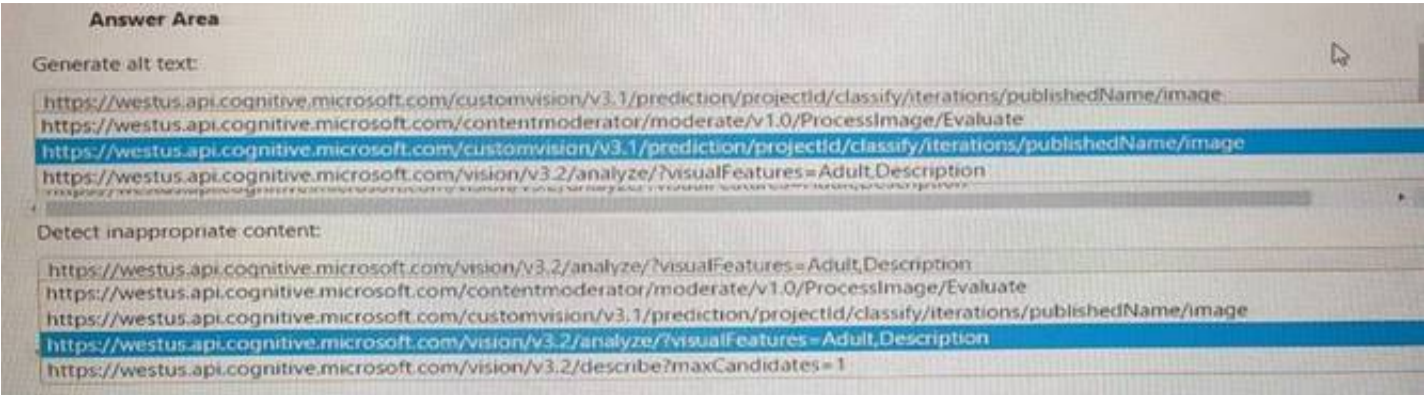
You are building an app that will enable users to upload images. The solution must meet the following requirements:

- 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.

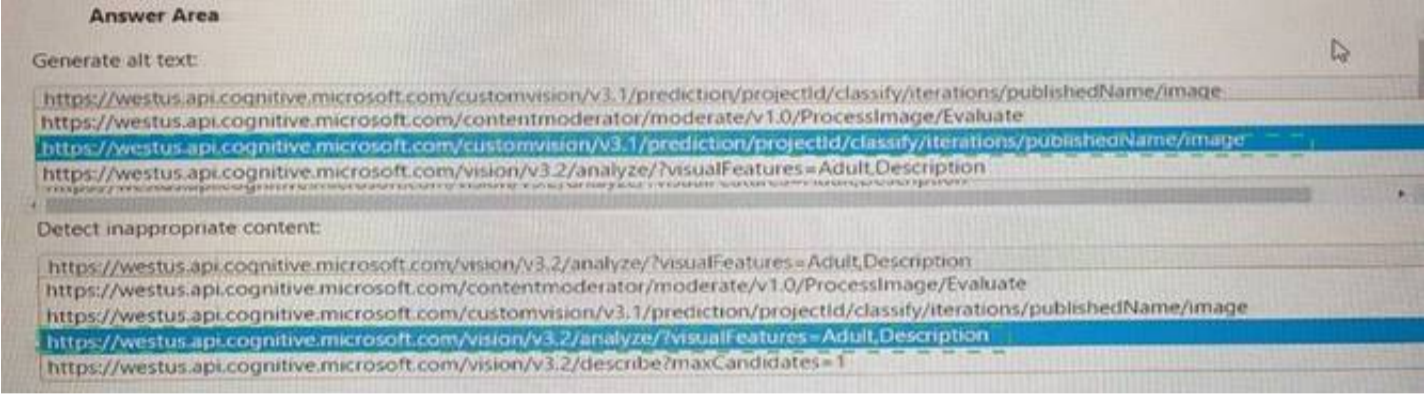
NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 113

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

- 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 116

DRAG DROP - (Topic 3)

You have 100 chatbots that each has its own Language Understanding model. Frequently, you must add the same phrases to each model. You need to programmatically update the Language Understanding models to include the new phrases. How should you complete the code? 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
AddPhraseListAsync	var phraselistId = await client.Features. []
Phraselist	(appId, versionId, new []
PhraselistCreateObject	{
Phrases	EnabledForAllModels = false,
SavePhraselistAsync	IsExchangeable = true,
UploadPhraseListAsync	Name = "PL1",
	Phrases = "item1,item2,item3,item4,item5"
	});

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: AddPhraseListAsync Example: Add phraselist feature

```
var phraselistId = await client.Features.AddPhraseListAsync(applId, versionId, new
```

```
PhraselistCreateObject
```

```
{  
    EnabledForAllModels = false, IsExchangeable = true, Name = "QuantityPhraselist", Phrases = "few,more,extra"  
});
```

```
Box 2: PhraselistCreateObject
```

NEW QUESTION 120

- (Topic 3)

You are building a chatbot.

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

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

Answer: B

NEW QUESTION 125

- (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 127

- (Topic 3)

You have a Language service resource that performs the following:

- Sentiment analysis
- Named Entity Recognition (NER)
- Personally Identifiable Information (PII) identification

You need to prevent the resource from persisting input data once the data is analyzed. Which query parameter in the Language service API should you configure?

- A. loggingOptOut
- B. piiCategories
- C. showStats
- D. Model-version

Answer: A

NEW QUESTION 129

- (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 add indexes.

Does this meet the goal?

- A. Yes

B. No

Answer: B

Explanation:

Instead, you could migrate to a Cognitive Search service that uses a higher tier.

Note: 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>

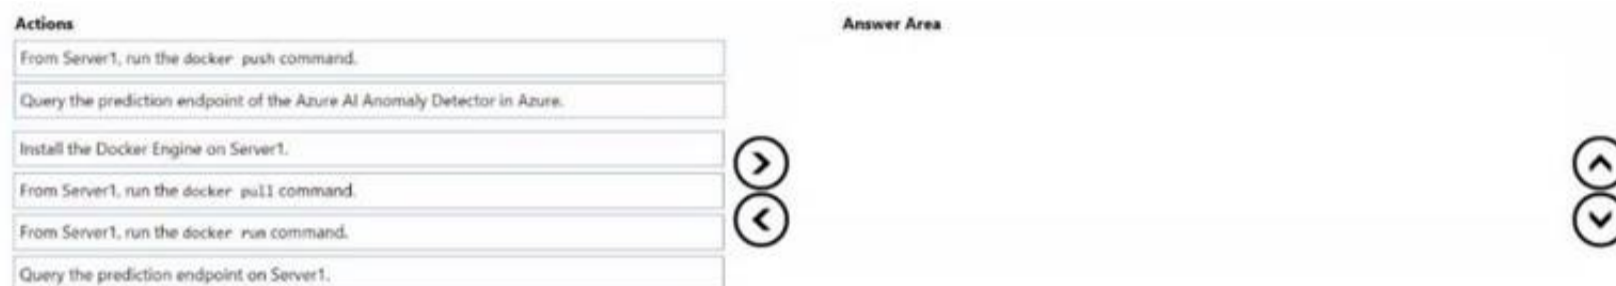
NEW QUESTION 132

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.

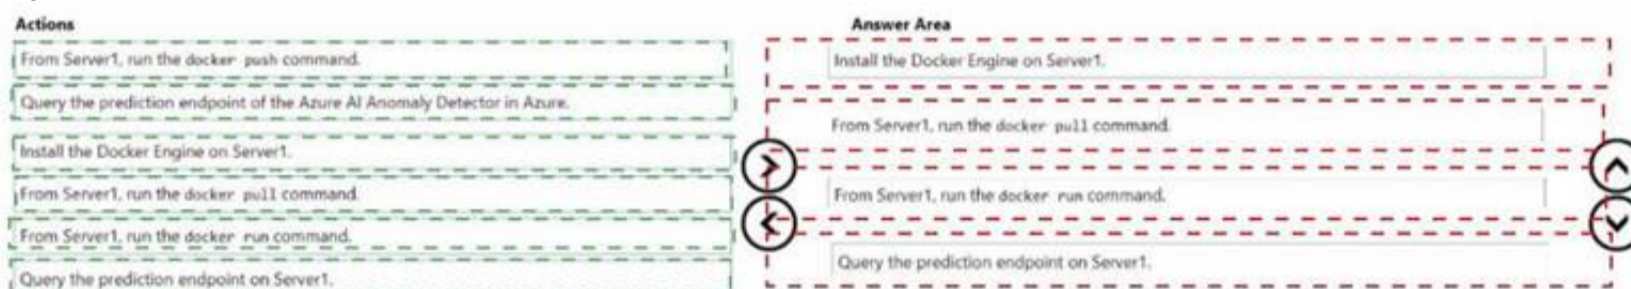


A. Mastered

B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 133

- (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. readInputStreamAsync

B. describeImageInStreamAsync

C. toggleImageInStreamAsync

D. analyzeImageByDomainInStreamAsync

Answer: D

NEW QUESTION 137

- (Topic 3)

You are building a chatbot that will use question answering in Azure Cognitive Service for Language.

You have a PDF named Doc1.pdf that contains a product catalogue and a price list You upload Doc1.pdf and train the model.

During testing, users report that the chatbot responds correctly to the following question: What is the price of < product>?

The chatbot fails to respond to the following question: How much does <product* cost? You need to ensure that the chatbot responds correctly to both questions.

Solution: From Language Studio, you create an entity for cost, and then retrain and republish the model.

Does this meet the goal?

A. Yes

B. No

Answer: B

NEW QUESTION 142

- (Topic 3)

You have an Azure subscription that contains an Azure App Service app named App1. You provision a multi-service Azure Cognitive Services resource named CSAccount1.

You need to configure App1 to access CSAccount1. The solution must minimize administrative effort.

What should you use to configure App1?

- A. a system-assigned managed identity and an X.509 certificate
- B. the endpoint URI and an OAuth token
- C. the endpoint URI and a shared access signature (SAS) token
- D. the endpoint URI and subscription key

Answer: D

NEW QUESTION 145

- (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 add the new images, and then use the Smart Labeler tool. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

The model need to be extended and retrained.

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

NEW QUESTION 147

- (Topic 3)

You need to measure the public perception of your brand on social media messages. Which Azure Cognitive Services service should you use?

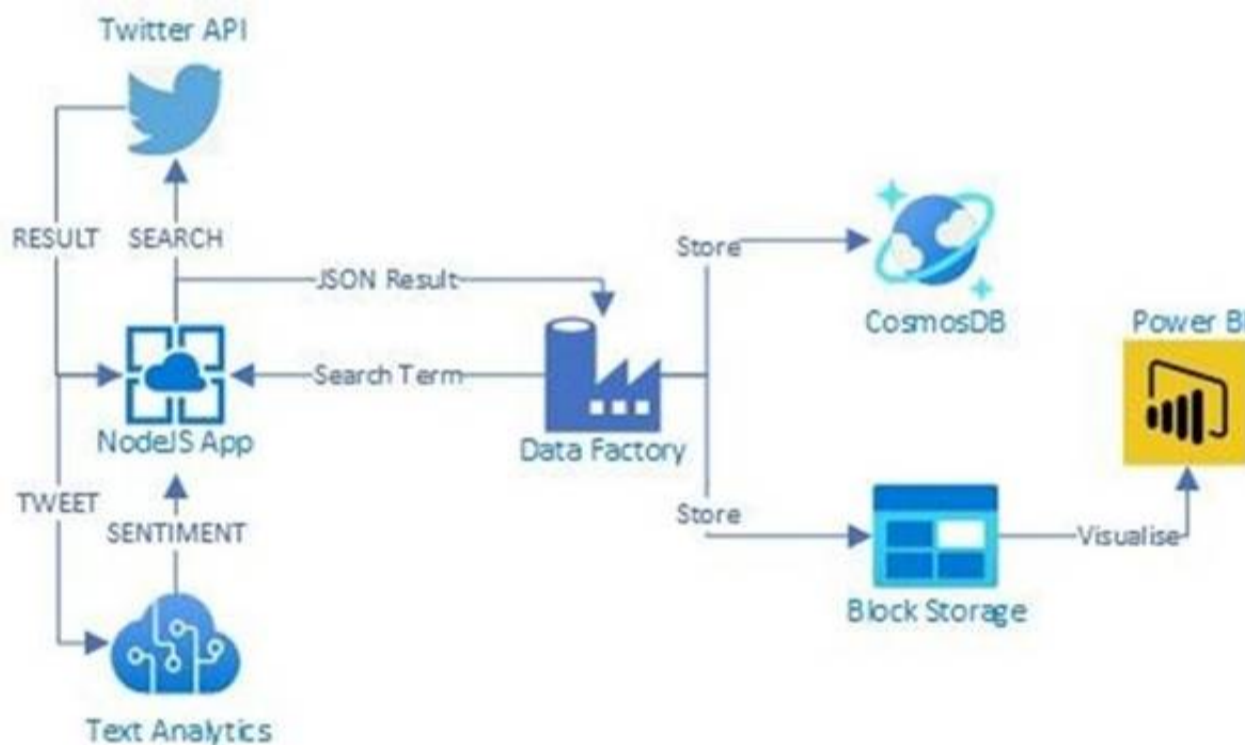
- A. Text Analytics
- B. Content Moderator
- C. Computer Vision
- D. Form Recognizer

Answer: A

Explanation:

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

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



Reference:

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

NEW QUESTION 149

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.

Features

Custom named entity recognition (NER)

Key phrase extraction

Language detection

Named Entity Recognition (NER)

Personally identifiable information (PII) detection

Sentiment analysis

Answer Area

Identify code names for internal product development:

Identify messages that include credit card numbers:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Features

Custom named entity recognition (NER)

Key phrase extraction

Language detection

Named Entity Recognition (NER)

Personally identifiable information (PII) detection

Sentiment analysis

Answer Area

Identify code names for internal product development:

Custom named entity recognition (NER)

Identify messages that include credit card numbers:

Personally identifiable information (PII) detection

NEW QUESTION 153

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.

Services

Speaker Recognition

Speech to Text

Text Analytics

Text to Speech

Translator

Answer Area

Detect the incoming language:

Respond in the callers' own language:

- 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 154

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.

AskForName > BeginDialog > Text

Show code

BeginDialog
Begin dialog event

Bot Asks (Text)
What is your name?

User Input (Text)
user.name = Input(Text)

Bot Asks (Number)
Hello \$(user.name), how old are you?

User Input (Number)
user.age = Input(Number)

Prompt for text
Text input
Collection information - Ask for a word or sentence.
[Learn more](#)

Bot Asks

User Input

Other

Property ⓘ
string
user.name

Output Format ⓘ
string
ex. =toUpper(this.value), \$(toUpper(this.value))

Value ⓘ
expression
fx =coalesce(@user.Name.@personName)

Expected responses (intent:
#TextInput_Response_GH5FTe)

>

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
	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:

Box 1: No

User.name is a property. Box 2: Yes

Box 3: Yes

The coalesce() function evaluates a list of expressions and returns the first non-null (or non-empty for string) expression.

NEW QUESTION 156

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

elastic pool

MySQL server

PostgreSQL server

virtual machine

 that hosts SQL Server.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Passing Certification Exams Made Easy

visit - <https://www.surepassexam.com>

Answer Area

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



that hosts SQL Server.

NEW QUESTION 157

- (Topic 3)

What is the primary purpose of a data warehouse?

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

Answer: C

NEW QUESTION 162

DRAG DROP - (Topic 3)

You Build a bot in JavaScript.

From the Azure Command-Line interface (CLI), you run the following command. az bot prepare-deploy

You need to deploy the bot to Azure.

Which three Azure CLI commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the client order.

Commands

- az ad app credential
- az ad app create
- az deployment group create
- az webapp deployment source config-zip
- az ad app update

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

To deploy the bot to Azure, you should run the following three Azure CLI commands in sequence:

? az deployment group create This command will create the Azure resources for your bot using an ARM template and a parameters file. You need to specify the resource group name, the template file path, and the parameters file path. For example:

az deployment group create --resource-group myResourceGroup --template-file "deploymentTemplates\template-with-preexisting-rg.json" --parameters "deploymentTemplates\parameters-for-template-BotApp-with-rg.json"

This command will also output the app ID and password of your bot, which you will need for the next command1.

? az webapp deployment source config-zip This command will deploy your bot code

to the app service that you created in the previous step. You need to specify the resource group name, the app service name, and the zip file path of your bot code. For example:

az webapp deployment source config-zip --resource-group myResourceGroup --name myBotAppService --src "code.zip"

This command will also output the URL of your bot endpoint, which you will need for the next command2.

? az ad app update This command will update your bot registration with the endpoint

URL of your bot. You need to specify the app ID of your bot and the endpoint URL. For example:

az ad app update --id myBotAppId --set replyUrls="https://myBotAppService.azurewebsites.net/api/messages"

This command will complete the deployment process and make your bot ready to be tested3.

NEW QUESTION 165

- (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 a network security group (NSG) for vnet1.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

<https://docs.microsoft.com/en-us/azure/virtual-network/network-security-groups-overview#network-security-groups>

<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-overview>

NEW QUESTION 169

- (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

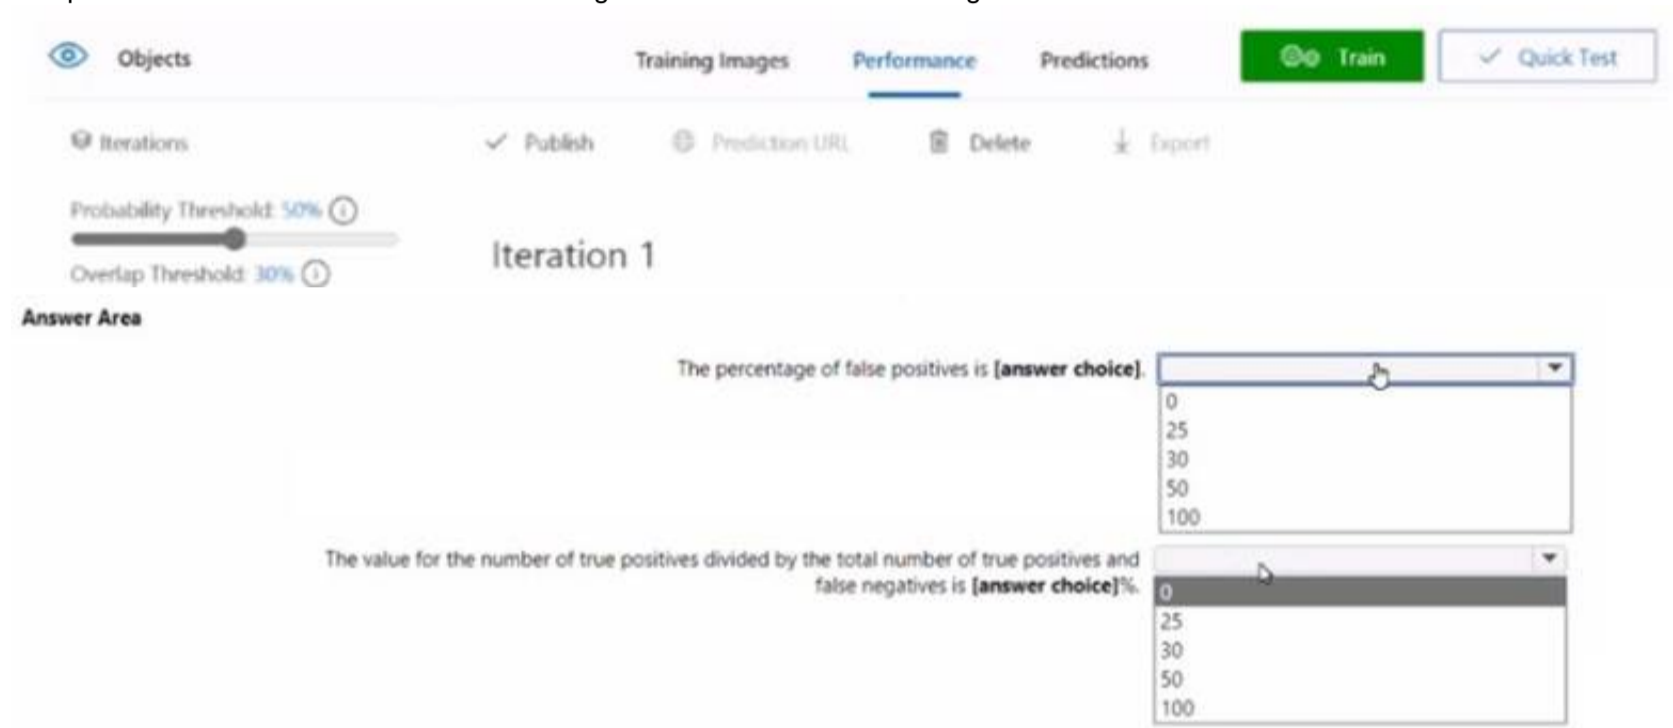
Answer: B

NEW QUESTION 170

HOTSPOT - (Topic 3)

You are building a model to detect objects in images.

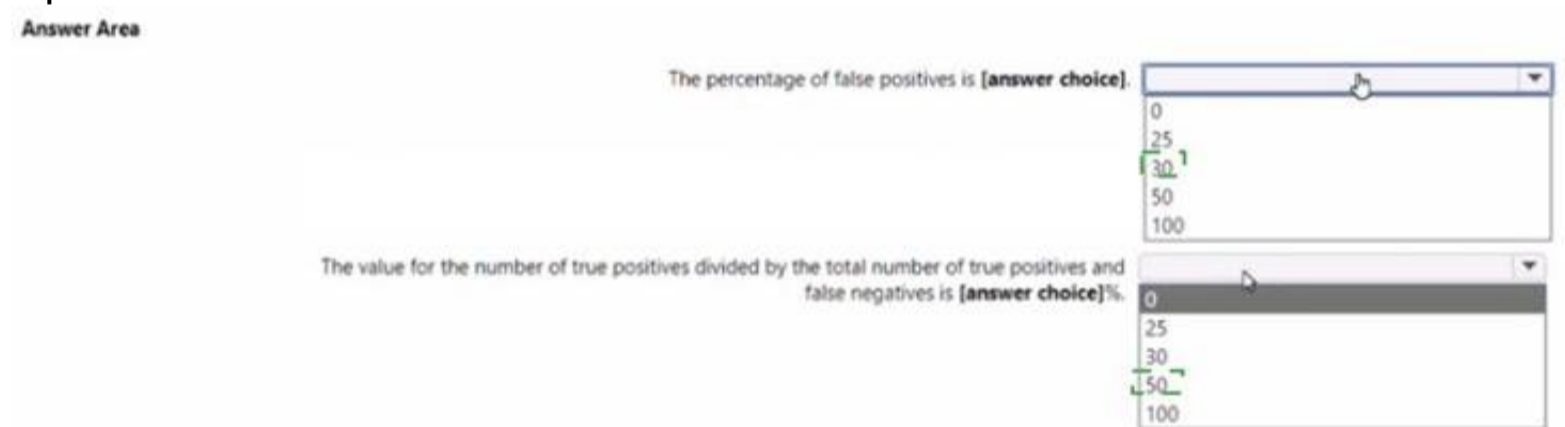
The performance of the model based on training data is shown in the following exhibit.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 174

- (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>

sol/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 179

- (Topic 3)

You have data saved in the following format.

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

Which format was used?

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

Answer: A

NEW QUESTION 184

HOTSPOT - (Topic 3)

You are building a chatbot by using the Microsoft Bot Framework SDK.

You use an object named UserProfile to store user profile information and an object named ConversationData to store information related to a conversation.

You create the following state accessors to store both objects in state.

var userStateAccessors = _userState.CreateProperty<UserProfile>(nameof(UserProfile)); var conversationStateAccessors =

_conversationState.CreateProperty<ConversationData>(nameof(ConversationData)); The state storage mechanism is set to Memory Storage.

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
The code will create and maintain the UserProfile object in the underlying storage layer.	<input type="radio"/>	<input type="radio"/>
The code will create and maintain the ConversationData object in the underlying storage layer.	<input type="radio"/>	<input type="radio"/>
The UserProfile and ConversationData objects will persist when the Bot Framework runtime terminates.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

You create property accessors using the CreateProperty method that provides a handle to the BotState object. Each state property accessor allows you to get or set the value of the associated state property.

Box 2: Yes

Box 3: No

Before you exit the turn handler, you use the state management objects' SaveChangesAsync() method to write all state changes back to storage.

NEW QUESTION 187

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

Create a Speaker Recognition model.

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.

Answer Area

>

<

⬆

⬇

- 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 189

- (Topic 3)

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: DE

NEW QUESTION 194

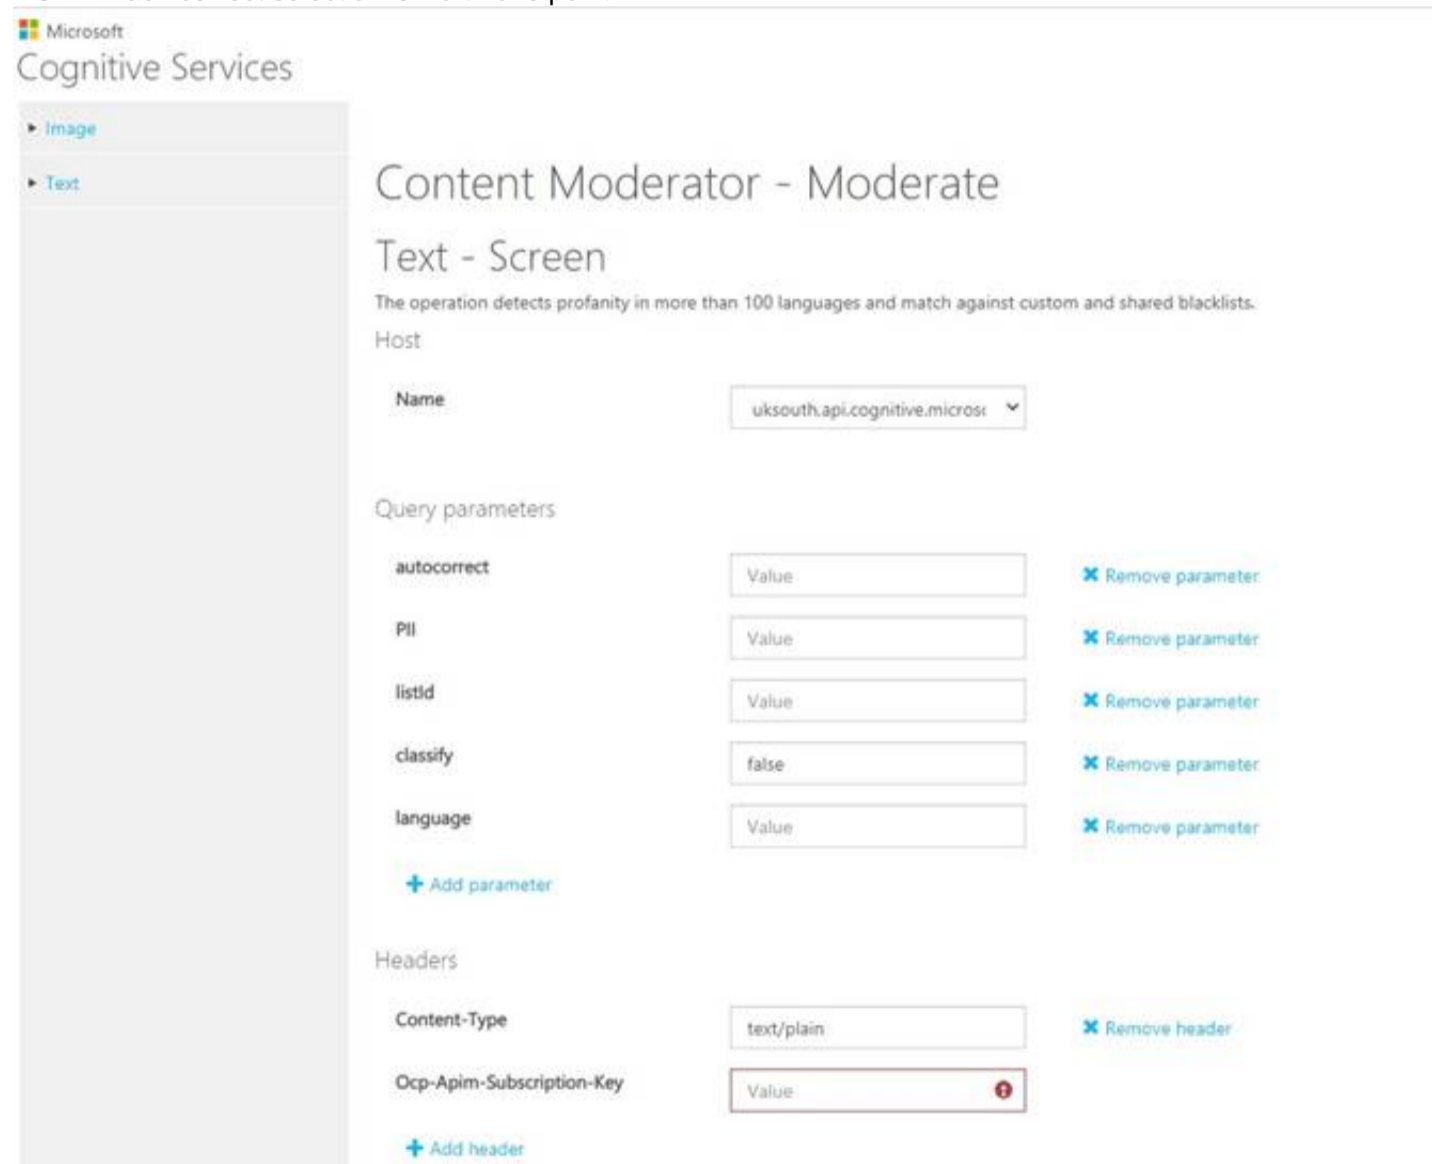
HOTSPOT - (Topic 3)

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.

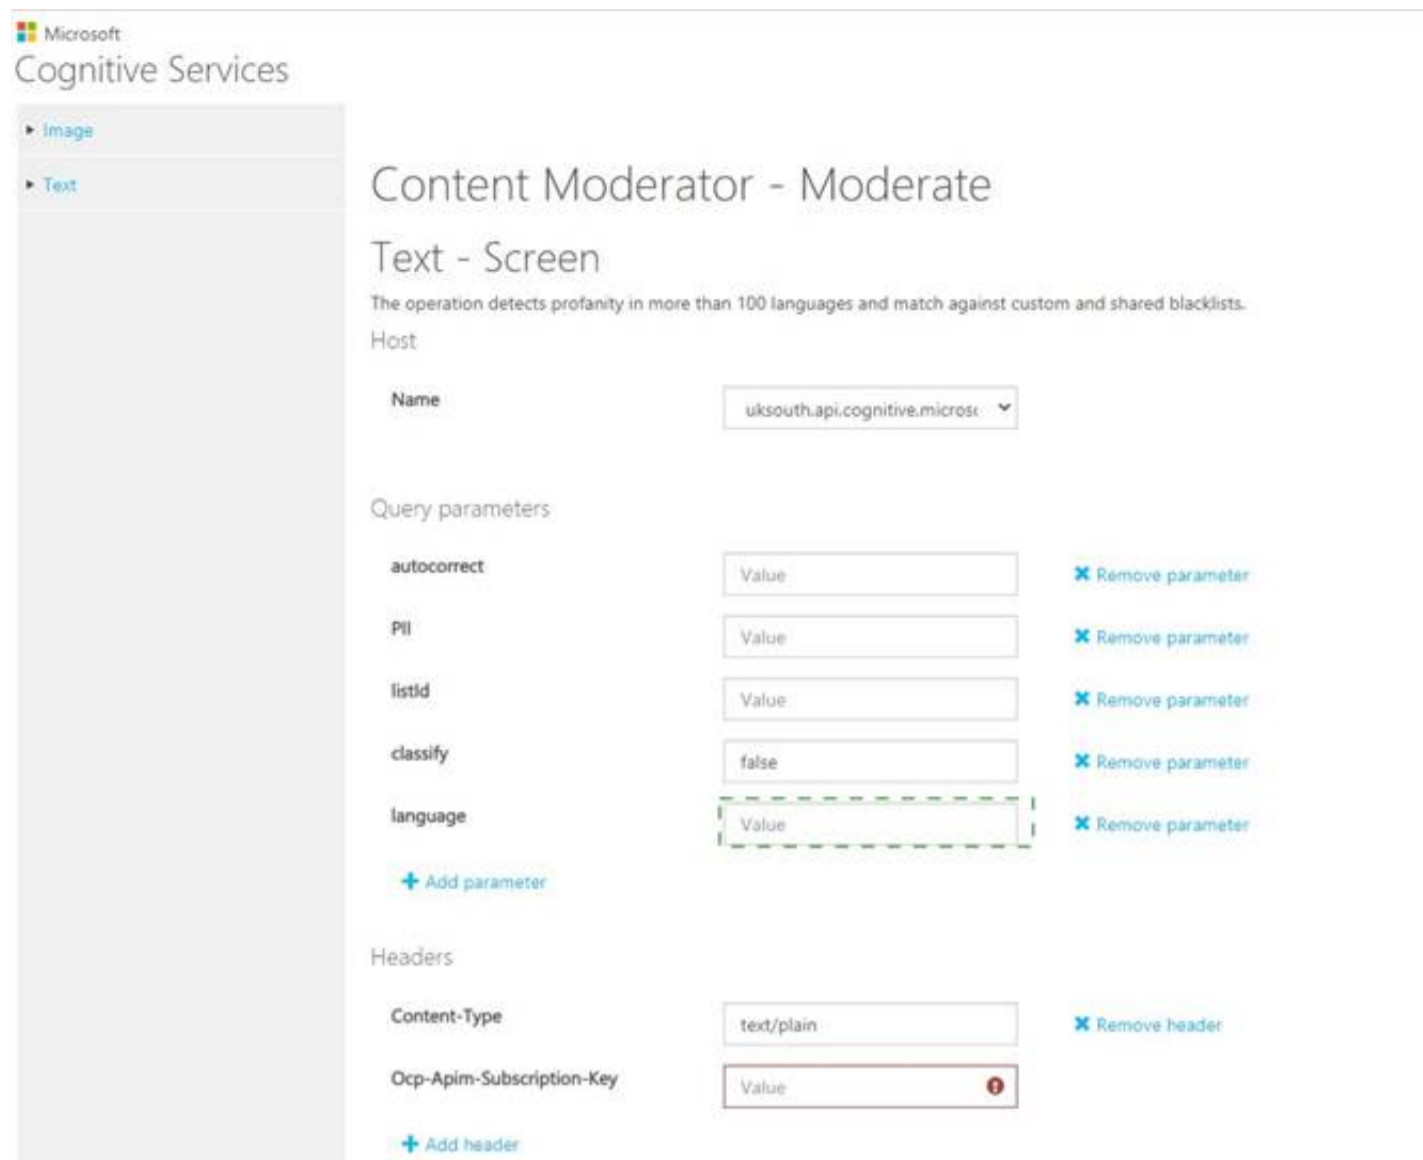
NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 195

- (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 198

- (Topic 3)

You are examining the Text Analytics output of an application.

The text analyzed is: "Our tour guide took us up the Space Needle during our trip to Seattle last week."

The response contains the data shown in the following table.

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

Which Text Analytics API is used to analyze the text?

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

Answer: B

NEW QUESTION 200

- (Topic 3)

You build a chatbot that uses the Azure OpenAI GPT 3.5 model.

You need to improve the quality of the responses from the chatbot. The solution must minimize development effort.

What are two ways to achieve the goal? Each correct answer presents a complete solution. NOTE: Each correct answer is worth one point.

- A. Fine-tune the model.
- B. Provide grounding content.
- C. Add sample request/response pairs.
- D. Retrain the language model by using your own data.
- E. Train a custom large language model (LLM).

Answer: BC

NEW QUESTION 205

- (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 209

- (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 210

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
Azure Databricks is an Apache Spark-based analytics platform.	<input type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input type="radio"/>	<input type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Azure Databricks is an Apache Spark-based analytics platform.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input type="radio"/>	<input type="radio"/>

NEW QUESTION 215

- (Topic 3)

You are building a multilingual chatbot.

You need to send a different answer for positive and negative messages.

Which two Text Analytics APIs should you use? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

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

Answer: BD

Explanation:

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

NEW QUESTION 216

HOTSPOT - (Topic 3)

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

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

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 217

DRAG DROP - (Topic 3)

You build a bot by using the Microsoft Bot Framework SDK. You need to test the bot interactively on a local machine. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the 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

Register the bot with the Azure Bot Service.

Open the Bot Framework Composer.

Build and run the bot.

Open the Bot Framework Emulator.

Connect to the bot endpoint.

>

<

Answer Area

↑

↓

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Register the bot with the Azure Bot Service.

Open the Bot Framework Composer.

Build and run the bot.

Open the Bot Framework Emulator.

Connect to the bot endpoint.

Answer Area

Build and run the bot.

Open the Bot Framework Emulator.

Connect to the bot endpoint.

NEW QUESTION 220

- (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 224

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.

Actions

Add the specific company names to the exclude list.

Sign in to the Custom Vision website.

From Content model customization, select Language.

Sign in to the Azure Video Analyzer for Media website.

From Content model customization, select Brands.

Add the specific company names to the include list.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Add the specific company names to the exclude list.

Sign in to the Custom Vision website.

From Content model customization, select Language.

Sign in to the Azure Video Analyzer for Media website.

From Content model customization, select Brands.

Add the specific company names to the include list.

Answer Area

Sign in to the Azure Video Analyzer for Media website.

From Content model customization, select Brands.

Add the specific company names to the include list.

NEW QUESTION 227

HOTSPOT - (Topic 3)
You are building a solution that students will use to find references for essays. You use the following code to start building the solution.

```
using Azure;
using System;
using Azure.AI.TextAnalytics;

private static readonly AzureKeyCredential credentials = new AzureKeyCredential("<key>");
private static readonly Uri endpoint = new Uri("<endpoint>");

static void EntityLinker(TextAnalyticsClient client)
{
    var response = client.RecognizeLinkedEntities(
        "Our tour guide took us up the Space Needle during our trip to Seattle last week.");
}
```

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
The code will detect the language of documents.	<input type="radio"/>	<input type="radio"/>
The <code>ur1</code> attribute returned for each linked entity will be a Bing search link.	<input type="radio"/>	<input type="radio"/>
The <code>matches</code> attribute returned for each linked entity will provide the location in a document where the entity is referenced.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The code will detect the language of documents.	<input type="radio"/>	<input checked="" type="radio"/>
The <code>ur1</code> attribute returned for each linked entity will be a Bing search link.	<input checked="" type="radio"/>	<input type="radio"/>
The <code>matches</code> attribute returned for each linked entity will provide the location in a document where the entity is referenced.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 232

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

- 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 233

- (Topic 3)

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

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

Answer: B

NEW QUESTION 237

HOTSPOT - (Topic 3)

You have an Azure subscription that contains an Azure OpenAI resource named AH. You build a chatbot that will use AI1 to provide generative answers to specific questions. You need to ensure that the responses are more creative and less deterministic.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 239

DRAG DROP - (Topic 3)

Match the Azure Cosmos DB APIs to the appropriate data structures.

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

NOTE: Each correct match is worth one point.

APIs

Cassandra API

Gremlin API

MongoDB API

Table API

Answer Area

API	Graph data
API	JSON documents
API	Key/value data

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

APIs

Cassandra API

Gremlin API

MongoDB API

Table API

Answer Area

Gremlin API

MongoDB API

Table API

Graph data

JSON documents

Key/value data

NEW QUESTION 241

- (Topic 3)
You build a custom Form Recognizer model.
You receive sample files to use for training the model as shown in the following table.

Name	Type	Size
File1	PDF	20 MB
File2	MP4	100 MB
File3	JPG	20 MB
File4	PDF	100 MB
File5	GIF	1 MB
File6	JPG	40 MB

Which three files can you use to train the model? Each correct answer presents a complete solution. (Choose three.)
NOTE: Each correct selection is worth one point.

- A. File1
- B. File2
- C. File3
- D. File4
- E. File5
- F. File6

Answer: ACF

Explanation:

Input requirements
Form Recognizer works on input documents that meet these requirements:
Format must be JPG, PNG, PDF (text or scanned), or TIFF. Text-embedded PDFs are best because there's no possibility of error in character extraction and location.
File size must be less than 50 MB.
Reference:
<https://docs.microsoft.com/en-us/azure/cognitive-services/form-recognizer/overview>

NEW QUESTION 246

- (Topic 3)
You have an Azure Cognitive Services model named Model that identifies the intent of text input.
You develop an app in C# named App1. You need to configure App1 to use Model1. Which package should you add to App1?

- A. Azure.AI.Language.Conversations
- B. SpeechServicesToolkit
- C. Universal.Microsoft.CognitiveServices.Speech
- D. Xamarin.Cognitive.Speech

Answer: C

NEW QUESTION 251

HOTSPOT - (Topic 3)
You have a Computer Vision resource named contoso1 that is hosted in the West US Azure region.
You need to use contoso1 to make a different size of a product photo by using the smart cropping feature.
How should you complete the API URL? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

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

https://api.projectoxford.ai

https://contoso1.cognitiveservices.azure.com

https://westus.api.cognitive.microsoft.com

/vision/v3.1/

areaOfInterest

detect

generateThumbnail

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

westus.api.cognitive.microsoft.com/generateThumbnail

<https://docs.microsoft.com/en-us/rest/api/computervision/3.1/generate-thumbnail/generate-thumbnail#examples>

POST

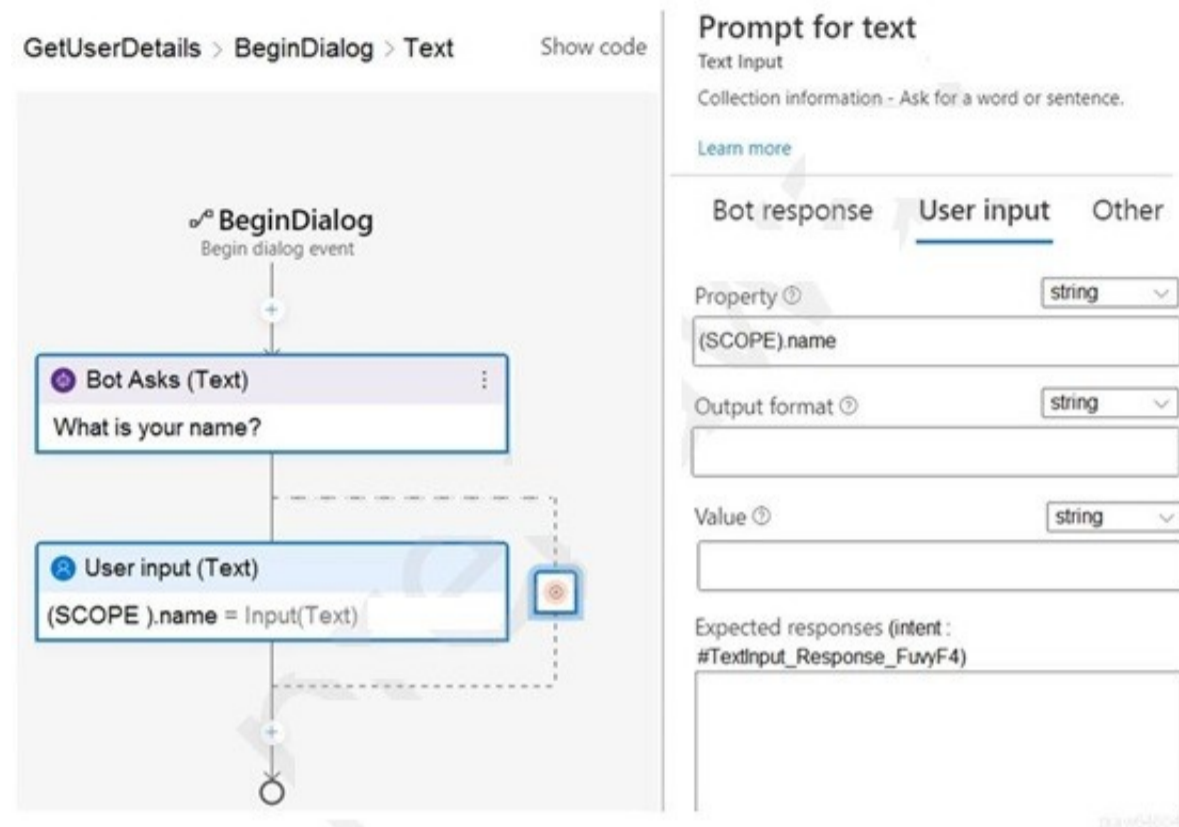
<https://westus.api.cognitive.microsoft.com/vision/v3.1/generateThumbnail?width=500&height=500&smartCropping=True>

Ocp-Apim-Subscription-Key: {API key}

NEW QUESTION 252

- (Topic 3)

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



The chatbot contains a dialog named GetUserDetails. GetUserDetails contains a TextInput control that prompts users for their name.

The user input will be stored in a property named name.

You need to ensure that you can dispose of the property when the last active dialog ends. Which scope should you assign to name?

- A. dialog
- B. user
- C. turn
- D. conversation

Answer: A

Explanation:

The dialog scope associates properties with the active dialog. Properties in the dialog scope are retained until the dialog ends.

Reference:

<https://docs.microsoft.com/en-us/composer/concept-memory?tabs=v2x>

NEW QUESTION 253

- (Topic 3)

What is a primary characteristic of a relational database?

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

Answer: C

NEW QUESTION 255

- (Topic 3)

You have a file share that contains 5,000 images of scanned invoices.

You need to analyze the images. The solution must extract the following data:

- Invoice items
 - Sales amounts
 - Customer details
- What should you use?

- A. Custom Vision
- B. Computer Vision
- C. Immersive Reader
- D. Form Recognizer

Answer: C

NEW QUESTION 256

HOTSPOT - (Topic 3)
You plan to deploy an Azure OpenAI resource by using an Azure Resource Manager (ARM) template.
You need to ensure that the resource can respond to 600 requests per minute.
How should you complete the template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity": 600,
    "count": 1,
    "maxValue": 100,
    "size": 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
      ...
    }
  }
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Answer Area

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity": 600,
    "count": 1,
    "maxValue": 100,
    "size": 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
      ...
    }
  }
}
```

NEW QUESTION 259

- (Topic 3)
What is used to define a query in a stream processing jobs in Azure Stream Analytics?

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

Answer: A

NEW QUESTION 264

- (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 random

questions that are outside the scope of the knowledge base.

You need to ensure that the chatbot provides formal responses to these spurious questions.

Solution: From Language Studio, you modify the question and answer pairs for the custom intents, and then retrain and republish the model.

Does this meet the goal?

A. Yes

B. No

Answer: B

NEW QUESTION 269

DRAG DROP - (Topic 3)

You have a chatbot that uses a QnA Maker application.

You enable active learning for the knowledge base used by the QnA Maker application. You need to integrate user input into the model.

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

Actions

Answer Area

Add a task to the Azure resource.

Approve and reject suggestions.

Publish the knowledge base.

Modify the automation task logic app to run an Azure Resource Manager template that creates the Azure Cognitive Services resource.

For the knowledge base, select Show active learning suggestions.

Save and train the knowledge base.

Select the properties of the Azure Cognitive Services resource.



A. Mastered

B. Not Mastered

Answer: A

Explanation:

Step 1: For the knowledge base, select Show active learning suggestions.

In order to see the suggested questions, on the Edit knowledge base page, select View Options, then select Show active learning suggestions.

Step 2: Approve and reject suggestions.

Each QnA pair suggests the new question alternatives with a check mark, , to accept the question or an x to reject the suggestions. Select the check mark to add the question.

Step 3: Save and train the knowledge base.

Select Save and Train to save the changes to the knowledge base. Step 4: Publish the knowledge base.

Select Publish to allow the changes to be available from the GenerateAnswer API.

When 5 or more similar queries are clustered, every 30 minutes, QnA Maker suggests the alternate questions for you to accept or reject.

NEW QUESTION 272

- (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 277

- (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 281

HOTSPOT - (Topic 2)

You are developing the knowledgebase by using Azure Cognitive Search. You need to build a skill that will be used by indexers.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: "categories": ["Locations", "Persons", "Organizations"], Locations, Persons, Organizations are in the outputs.
Scenario: Contoso plans to develop a searchable knowledgebase of all the intellectual property
Note: The categories parameter is an array of categories that should be extracted. Possible category types: "Person", "Location", "Organization", "Quantity", "Datetime", "URL", "Email". If no category is provided, all types are returned.
Box 2: {"name": " entities"}
The include wikis, so should include entities in the outputs.
Note: entities is an array of complex types that contains rich information about the entities extracted from text, with the following fields
name (the actual entity name. This represents a "normalized" form) wikipediaId
wikipediaLanguage
wikipediaUrl (a link to Wikipedia page for the entity) etc.

NEW QUESTION 285

- (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 290

HOTSPOT - (Topic 2)
You build a QnA Maker resource to meet the chatbot requirements.
Which RBAC role should you assign to each group? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Management-Accountants	<div>▼</div> <div>Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor</div>
Consultant-Accountants	<div>▼</div> <div>Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor</div>
Agent-CustomerServices	<div>▼</div> <div>Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor</div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Cognitive Service User
Ensure that the members of a group named Management-Accountants can approve the FAQs.
Approve=publish.
Cognitive Service User (read/write/publish): API permissions: All access to Cognitive Services resource except for ability to:
* 1. Add new members to roles.
* 2. Create new resources.
Box 2: Cognitive Services QnA Maker Editor

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

QnA Maker Editor: API permissions:

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

Box 3: Cognitive Services QnA Maker Read

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

QnA Maker Read: API Permissions:

- * 1. Download KB API
- * 2. List KBs for user API
- * 3. Get Knowledge base details
- * 4. Download Alterations Generate Answer

NEW QUESTION 295

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><div></div><div>Cognitive Service User</div><div>Contributor</div><div>Owner</div><div>QnA Maker Editor</div><div>QnA Maker Read</div></div>
LeadershipTeam:	<div><div></div><div>Cognitive Service User</div><div>Contributor</div><div>Owner</div><div>QnA Maker Editor</div><div>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 299

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.

Services

Azure Blob storage

Custom Vision API

Azure Files

Language Understanding API

Translator API

Computer Vision API

Azure Cosmos DB

Answer Area

Source

Cracking

Preparation

Destination

- 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 301

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

AI-102 Practice Exam Features:

- * AI-102 Questions and Answers Updated Frequently
- * AI-102 Practice Questions Verified by Expert Senior Certified Staff
- * AI-102 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * AI-102 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AI-102 Practice Test Here](#)