

Microsoft

Exam Questions AI-102

Designing and Implementing an Azure AI Solution



NEW QUESTION 1

DRAG DROP - (Topic 3)

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

The users report that the agent frequently responds with the following generic response: "Sorry, I don't understand that."

You need to improve the ability of the agent to respond to requests.

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

(Choose three.)

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

- enable active learning
- validate the utterances
- train and republish

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

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

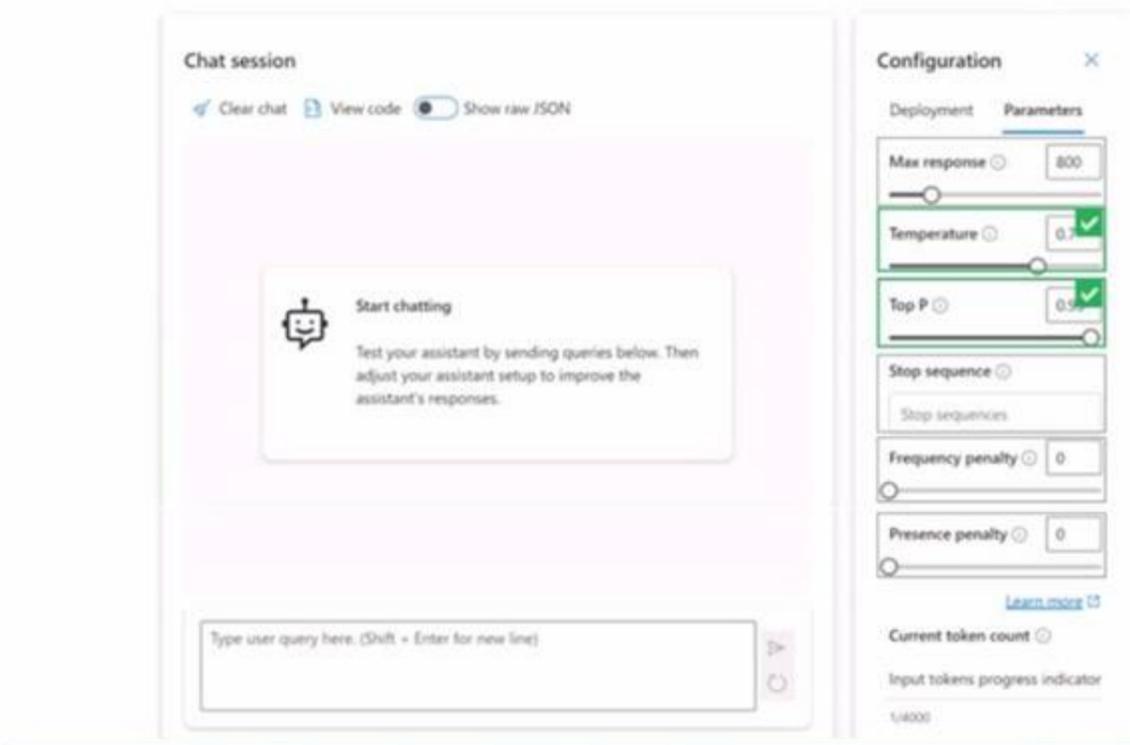
HOTSPOT - (Topic 3)

You build a chatbot by using Azure OpenAI Studio.

You need to ensure that the responses are more deterministic and less creative. Which two parameters should you configure? To answer, select the appropriate parameters in the answer area.

NOTE: Each correct answer is worth one point.

Answer Area

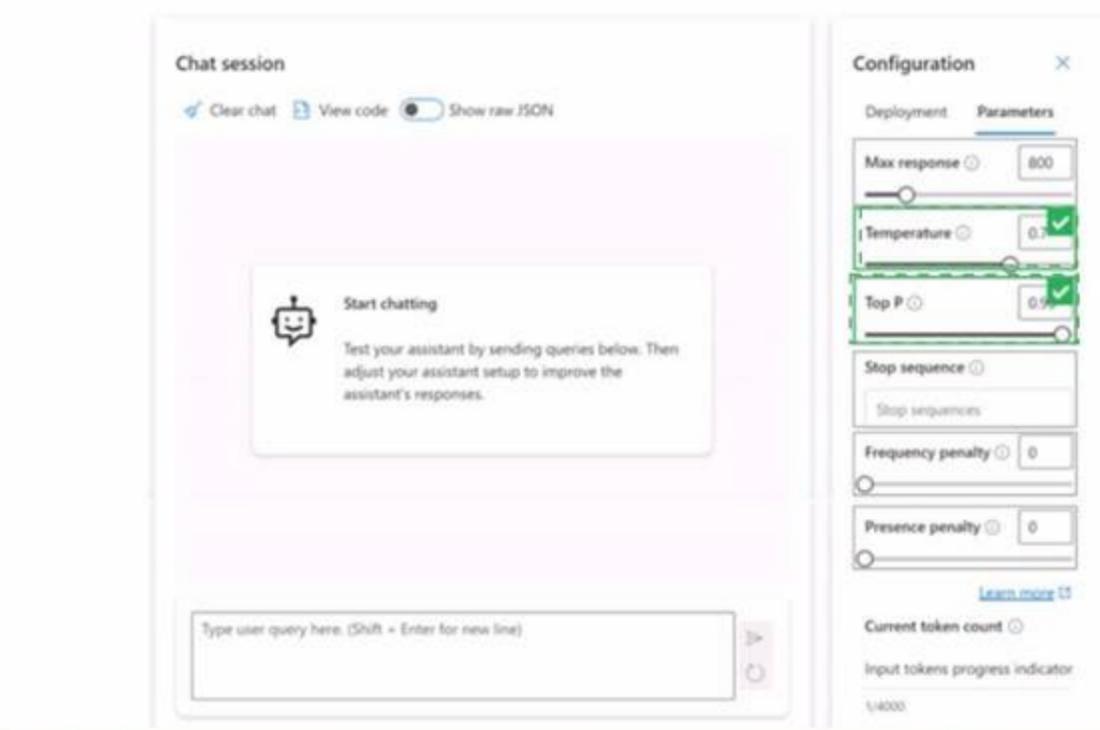


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 4

- (Topic 3)

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

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

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

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

Answer: D

Explanation:

Evaluate different models.

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

Reference:

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

NEW QUESTION 5

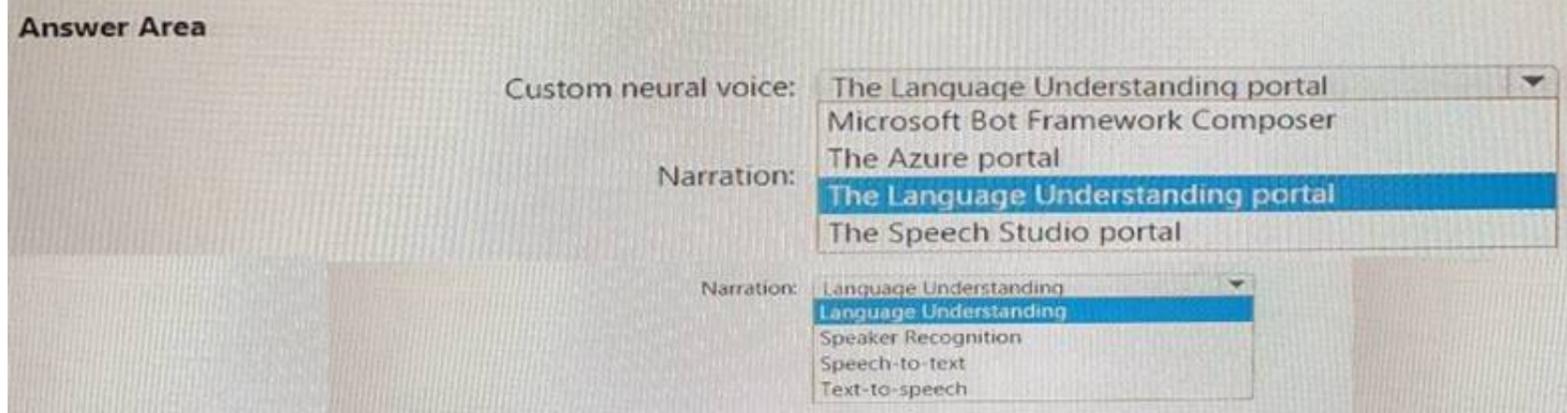
HOTSPOT - (Topic 3)

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.

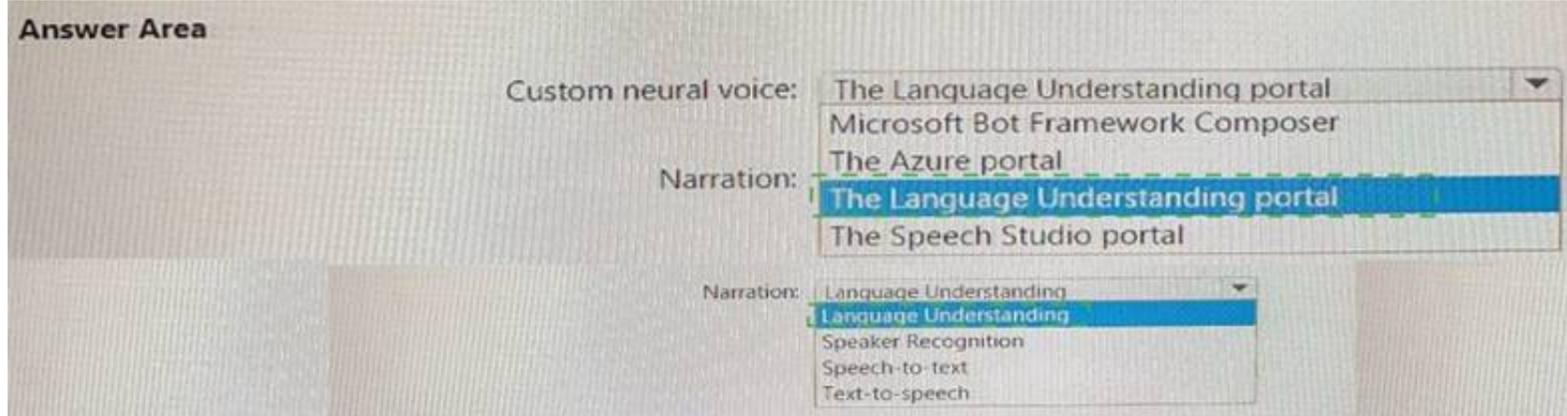
NOTE: Each correct answer is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 6

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 7

- (Topic 3)

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

uses Model!

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

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

Answer: B

NEW QUESTION 8

- (Topic 3)

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

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

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

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

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

Solution: You deploy service1 and a public endpoint, and you configure an IP firewall rule. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Reference:

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

NEW QUESTION 9

HOTSPOT - (Topic 3)

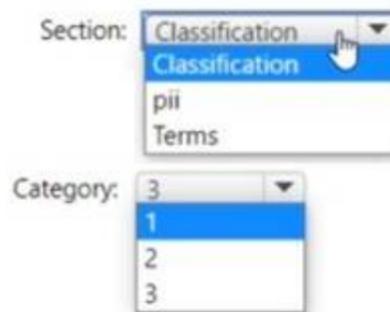
You are building a chatbot.

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

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

NOTE: Each correct selection is worth one point.

Answer Area

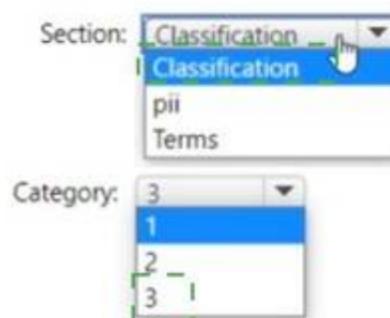


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 10

- (Topic 3)

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

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

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

You need to add the new images to the classifier.

Solution: You create a new model, and then upload the new images and labels. Does this meet the goal?

- A. Yes

B. No

Answer: B

Explanation:

The model needs to be extended and retrained.

NEW QUESTION 10

- (Topic 3)

You use the Microsoft Bot Framework Composer to build a chatbot that enables users to purchase items.

You need to ensure that the users can cancel in-progress transactions. The solution must minimize development effort.

What should you add to the bot?

- A. a language generator
- B. a custom event
- C. a dialog trigger
- D. a conversation activity

Answer: D

NEW QUESTION 15

- (Topic 3)

You have a Video Indexer service that is used to provide a search interface over company videos on your company's website.

You need to be able to search for videos based on who is present in the video. What should you do?

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

Answer: A

Explanation:

Video Indexer supports multiple Person models per account. Once a model is created, you can use it by providing the model ID of a specific Person model when uploading/indexing or reindexing a video. Training a new face for a video updates the specific custom model that the video was associated with.

Note: Video Indexer supports face detection and celebrity recognition for video content. The celebrity recognition feature covers about one million faces based on commonly requested data source such as IMDB, Wikipedia, and top LinkedIn influencers. Faces that aren't recognized by the celebrity recognition feature are detected but left unnamed. Once you label a face with a name, the face and name get added to your account's Person model. Video Indexer will then recognize this face in your future videos and past videos.

Reference:

<https://docs.microsoft.com/en-us/azure/media-services/video-indexer/customize-person-model-with-api>

NEW QUESTION 19

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 20

- (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 22

- (Topic 3)

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

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

Answer: A

NEW QUESTION 27

- (Topic 3)

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

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

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

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

Answer: D

NEW QUESTION 30

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.

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	Yes	No
Stream processing has access to the most recent data received or data within a rolling time window.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	<input type="checkbox"/>	<input checked="" type="checkbox"/>

NEW QUESTION 32

DRAG DROP - (Topic 3)

You plan to build a chatbot to support task tracking.

You create a Language Understanding service named lu1.

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

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

Actions

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 34

- (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 38

- (Topic 3)

You build a conversational bot named bot1.

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

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

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

Answer: C

Explanation:

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


Reference:

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

NEW QUESTION 41

HOTSPOT - (Topic 3)

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



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

NOTE: Each correct selection is worth one point.

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:

Answer Area

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

NEW QUESTION 43

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].

▼
a domain
an entity
an intent
an utterance

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

▼
a domain
an entity
an intent
an utterance

- 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 48

HOTSPOT - (Topic 3)

You are developing an application that includes language translation.

The application will translate text retrieved by using a function named getTextToBeTranslated. The text can be in one of many languages. The content of the text

must remain within the Americas Azure geography.
 You need to develop code to translate the text to a single language.
 How should you complete the code? To answer, select the appropriate options in the answer area.
 NOTE: Each correct selection is worth one point.

```

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

    var apiKey = "FF956C68883B21B38691ABD200A4C606";
    var text = getTextToBeTranslated();
    var body = '[{"Text":"' + text + '"}]';
    var client = new HttpClient();
    client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```

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

    var apiKey = "FF956C68883B21B38691ABD200A4C606";
    var text = getTextToBeTranslated();
    var body = '[{"Text":"' + text + '"}]';
    var client = new HttpClient();
    client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);

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

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

NEW QUESTION 52

HOTSPOT - (Topic 3)

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

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

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

For the second argument, you call the function and specify the following string. Our tour of Paris included a visit to the Eiffel Tower.
 For each of the following statements, select Yes if the statement is true. Otherwise, select No.
 NOTE: Each correct selection is worth one point.

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 type="radio"/>	<input checked="" type="radio"/>
The function will output all the key phrases from the input string to the console.	<input type="radio"/>	<input checked="" type="radio"/>

NEW QUESTION 54

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 a 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	Answer Area
Perform an asynchronous translation by using the list of files to be translated.	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/>
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.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 59

- (Topic 3)

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

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

Which service should you use?

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

Answer: A

NEW QUESTION 62

- (Topic 3)

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

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

Answer: C

NEW QUESTION 64

HOTSPOT - (Topic 3)

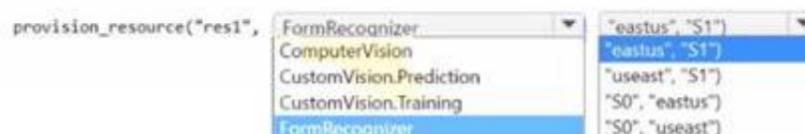
You plan to provision Azure Cognitive Services resources by using the following method.

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

You need to create a Standard tier resource that will convert scanned receipts into text. How should you call the method? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

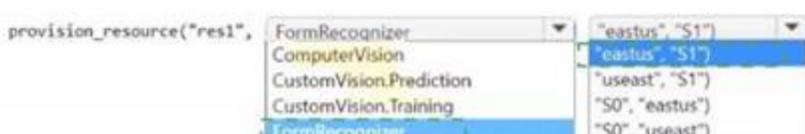


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



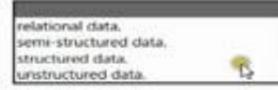
NEW QUESTION 65

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

Varying fields for each entity in a JSON document is an example of



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Varying fields for each entity in a JSON document is an example of



NEW QUESTION 69

- (Topic 3)

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

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

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 72

DRAG DROP - (Topic 3)

You are developing an application that will recognize faults in components produced on a factory production line. The components are specific to your business. You need to use the Custom Vision API to help detect common faults. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Create a project Create a new project.

Step 2: Upload and tag the images

Choose training images. Then upload and tag the images.

Step 3: Train the classifier model. Train the classifier

NEW QUESTION 76

- (Topic 3)

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

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

Answer: C

NEW QUESTION 81

- (Topic 3)

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

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

Answer: B

NEW QUESTION 86

- (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 91

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	Answer Area
Configure disk logging.	
Export model1 to Host1.	⬆
Request approval to run the container.	⬆
Retrain the model.	
Run the container.	

- 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 96

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 99

- (Topic 3)

You have a 20-GB file named file1, avi that is stored on a local drive. You need to index file1, avi by using the Azure Video indexer website. What should you do first?

- A. Upload File1 avi to an Azure storage queue.
- B. upload File1.avi to the www.youtube.com seepage
- C. Upload file1 avi to the Azure video indexer website.
- D. Upload file1 avi to Microsoft OneDrive.

Answer: D

Explanation:

This is because the Azure Video Indexer website allows you to upload videos from a URL or from your file system, but there are some limitations and considerations for each option1.

If you upload from your file system, the size of the file is limited to 2 GB, which is less than the size of file1.avi (20 GB). Therefore, this option is not feasible.

If you upload from a URL, the size of the file is limited to 30 GB, which is enough for file1.avi. However, the URL must be publicly accessible and valid, and the file must be accessible. You cannot use URLs from streaming services such as YouTube1. Therefore, options A and B are not valid.

The best option is to upload file1.avi to a cloud storage service such as Microsoft OneDrive, and then generate a shareable link for the file. You can then paste the link in the

Azure Video Indexer website and upload the video from the URL. This way, you can avoid the file size limitation and ensure that the file is accessible and valid2.

NEW QUESTION 102

- (Topic 3)

You need to build a chatbot that meets the following requirements:

- ? Supports chit-chat, knowledge base, and multilingual models
- ? Performs sentiment analysis on user messages
- ? Selects the best language model automatically

What should you integrate into the chatbot?

- A. QnA Maker, Language Understanding, and Dispatch
- B. Translator, Speech, and Dispatch
- C. Language Understanding, Text Analytics, and QnA Maker
- D. Text Analytics, Translator, and Dispatch

Answer: C

Explanation:

Language Understanding: An AI service that allows users to interact with your applications, bots, and IoT devices by using natural language.

QnA Maker is a cloud-based Natural Language Processing (NLP) service that allows you to create a natural conversational layer over your data. It is used to find the most appropriate answer for any input from your custom knowledge base (KB) of information.

Text Analytics: Mine insights in unstructured text using natural language processing (NLP)—no machine learning expertise required. Gain a deeper understanding of customer opinions with sentiment analysis. The Language Detection feature of the Azure Text Analytics REST API evaluates text input

Reference:

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

NEW QUESTION 104

DRAG DROP - (Topic 3)

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

In acvdev, you build an object detection model named obj1 in a project named proj1. You need to move obj1 to acvprod.

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 106

- (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 G8 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 107

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

A JSON document is an example of

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A JSON document is an example of

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

NEW QUESTION 111

- (Topic 3)

You have the following data sources:

? Finance: On-premises Microsoft SQL Server database

? Sales: Azure Cosmos DB using the Core (SQL) API

? Logs: Azure Table storage

? HR: Azure SQL database

You need to ensure that you can search all the data by using the Azure Cognitive Search REST API. What should you do?

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

Answer: B

Explanation:

On-premises Microsoft SQL Server database cannot be used as an index data source. Note: Indexer in Azure Cognitive Search: : Automate aspects of an indexing operation by

configuring a data source and an indexer that you can schedule or run on demand. This feature is supported for a limited number of data source types on Azure.

Indexers crawl data stores on Azure.

? Azure Blob Storage

? Azure Data Lake Storage Gen2 (in preview)

? Azure Table Storage

? Azure Cosmos DB

? Azure SQL Database

? SQL Managed Instance

? SQL Server on Azure Virtual Machines

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-indexer-overview#supported-data-sources>

NEW QUESTION 112

HOTSPOT - (Topic 3)

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

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 117

- (Topic 3)

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

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

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

Answer: B

NEW QUESTION 120

- (Topic 3)

What is a characteristic of a non-relational database?

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

Answer: C

NEW QUESTION 125

HOTSPOT - (Topic 3)

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

```
{
  "status": "succeeded",
  "createdDateTime": "2023-09-14T21:01:02Z",
  "lastUpdatedDateTime": "2023-09-14T21:01:03Z",
  "analyzeResult": {
    "apiVersion": "2023-07-31",
    "modelId": "prebuilt-healthInsuranceCard.us",
    "stringIndexType": "utf16CodeUnit",
    "content": "Blood Pressure 118/72",
    "pages": [
      {
        ...
        "words": [
          {
            "content": "Blood",
            "polygon": [ ... ],
            "confidence": 0.766,
            "span": { ... }
          },
          {
            "content": "Pressure",
            "polygon": [ ... ],
            "confidence": 0.716,
            "span": { ... }
          },
          {
            "content": "118/72",
            "polygon": [ ... ],
            "confidence": 0.761,
            "span": { ... }
          }
        ]
      }
    ],
    "documents": [
      {
        "docType": "healthInsuranceCard.us",
        "boundingRegions": [ ... ]
      }
    ],
    "fields": {},
    "confidence": 1,
    "spans": [ - ]
  }
}
```

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 129

- (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 StartRecognizeReceiptsFromUri 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 132

- (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 135

- (Topic 3)

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

You need to construct an entity to capture billing addresses. Which entity type should you use?

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

Answer: B

Explanation:

Reference:

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

ML Entity with Structure

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

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

Street Name: Main Street State: NY

Zip Code: 98052 Country: USA

NEW QUESTION 138

- (Topic 3)

A customer uses Azure Cognitive Search.

The customer plans to enable a server-side encryption and use customer-managed keys (CMK) stored in Azure.

What are three implications of the planned change? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. The index size will increase.
- B. Query times will increase.
- C. A self-signed X.509 certificate is required.
- D. The index size will decrease.

- E. Query times will decrease.
- F. Azure Key Vault is required.

Answer: ABF

Explanation:

"Customer-managed keys require an additional billable service, Azure Key Vault, which can be in a different region, but under the same subscription, as Azure Cognitive Search. Enabling CMK encryption will increase index size and degrade query performance."
same document also lists Azure Key Vault as a requirement: <https://docs.microsoft.com/en-us/azure/search/search-security-overview#data-protection>

NEW QUESTION 140

- (Topic 3)

You create five bots by using Microsoft Bot Framework Composer.

You need to make a single bot available to users that combines the bots. The solution must support dynamic routing to the bots based on user input.

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

NOTE: Each correct selection is worth one point.

- A. Create an Orchestrator model.
- B. Change the Recognizer/Dispatch type.
- C. Create a composer extension.
- D. Enable WebSockets.
- E. Create a custom recognizer JSON file.
- F. Install the Orchestrator package.

Answer: ABF

NEW QUESTION 141

- (Topic 3)

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

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

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

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

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

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

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Reference:

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

NEW QUESTION 143

- (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 144

HOTSPOT - (Topic 3)

You are building an Azure Cognitive Search custom skill. You have the following custom skill schema definition.

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes
 Once you have defined a skillset, you must map the output fields of any skill that directly contributes values to a given field in your search index.
 Box 2: Yes
 The definition is a custom skill that calls a web API as part of the enrichment process.
 Box 3: No
 For each organization identified by entity recognition, this skill calls a web API to find the description of that organization.

NEW QUESTION 148

- (Topic 3)

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

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

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

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

- A. Yes
- B. No

Answer: A

NEW QUESTION 151

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

```
{
  "faceId": "18c51a87-3a69-47a8-aedc-a54745f708a1",
  : "employeefaces",
  "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 faceld, to search the similar-looking faces from a faceld array, a face list or a large face list. A "faceListId" is created by FaceList - Create containing persistedFacelds that will not expire. And a "largeFaceListId" is created by LargeFaceList - Create containing persistedFacelds that will also not expire.

NEW QUESTION 152

HOTSPOT - (Topic 3)

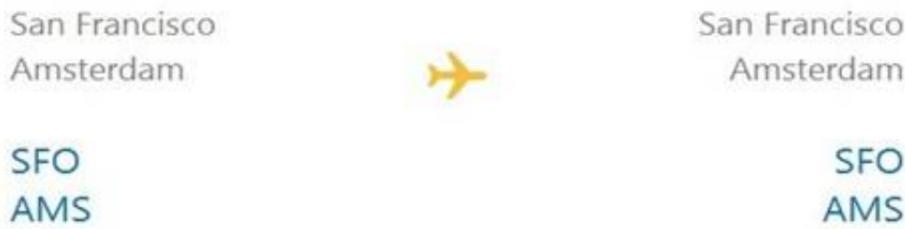
You are building a chatbot that will provide information to users as shown in the following exhibit.

Passengers

Sarah Hum
 Jeremy Goldberg
 Evan Litvak

2 Stops

Tue, May 30, 2017 10:25 PM



Non-Stop

Fri, Jun 2, 2017 11:55 PM



Total **\$4,032.54**

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

NOTE: Each correct selection is worth one point.

Answer Area

The chatbot is showing [answer choice].

- an Adaptive Card
- a Hero Card
- a Thumbnail Card

The card includes [answer choice].

- an action set
- an image
- an image group
- media

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Adaptive card Box 2: an image

NEW QUESTION 154

- (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 156

- (Topic 3)

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

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

Which type of detection should you use?

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

Answer: A

Explanation:

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

NEW QUESTION 157

- (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 158

DRAG DROP - (Topic 3)

You are developing a webpage that will use the Video Indexer service to display videos of internal company meetings.

You embed the Player widget and the Cognitive Insights widget into the page. You need to configure the widgets to meet the following requirements:

? Ensure that users can search for keywords.

? Display the names and faces of people in the video.

? Show captions in the video in English (United States).

How should you complete the URL for each widget? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
 Graphical user interface, text, application, Word, email Description automatically generated

NEW QUESTION 161

DRAG DROP - (Topic 3)

You are building a Language Understanding model for purchasing tickets.

You have the following utterance for an intent named PurchaseAndSendTickets. Purchase [2 audit business] tickets to [Paris] [next Monday] and send tickets to [email@domain.com]

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
 Box 1: GeographyV2
 The prebuilt geographyV2 entity detects places. Because this entity is already trained, you do not need to add example utterances containing GeographyV2 to the application intents.
 Box 2: Email
 Email prebuilt entity for a LUIS app: Email extraction includes the entire email address from an utterance. Because this entity is already trained, you do not need to add example utterances containing email to the application intents.
 Box 3: Machine learned
 The machine-learning entity is the preferred entity for building LUIS applications.

NEW QUESTION 165

- (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 166

- (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 169

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.

The screenshot shows a 'Features' list on the left with the following items: Custom named entity recognition (NER), Key phrase extraction, Language detection, Named Entity Recognition (NER), Personally Identifiable Information (PII) detection, and Sentiment analysis. The 'Answer Area' on the right contains two empty text boxes: 'Identify code names for internal product development:' and 'Identify messages that include credit card numbers:'.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The screenshot shows the same 'Features' list as above. In the 'Answer Area', the first requirement 'Identify code names for internal product development:' has 'Custom named entity recognition (NER)' assigned to it. The second requirement 'Identify messages that include credit card numbers:' has 'Personally Identifiable Information (PII) detection' assigned to it.

NEW QUESTION 170

HOTSPOT - (Topic 3)

You have a collection of press releases stored as PDF files.
 You need to extract text from the files and perform sentiment analysis.

Which service should you use for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

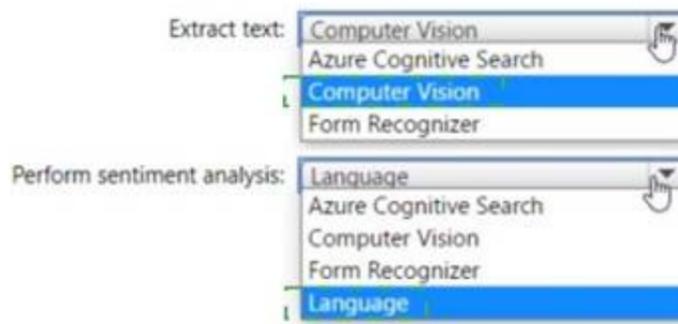
The screenshot shows the 'Answer Area' with two dropdown menus. The first dropdown is labeled 'Extract text:' and has 'Computer Vision' selected. The second dropdown is labeled 'Perform sentiment analysis:' and has 'Language' selected.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 172

- (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 entity for the domain.

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 174

- (Topic 3)

You are building a bot by using Microsoft Bot Framework.

You need to configure the bot to respond to spoken requests. The solution must minimize development effort.

What should you do?

- A. Deploy the bot to Azure and register the bot with a Direct Une Speech channel
- B. Integrate the bot with Cortana by using the Bot Framework SDK.
- C. Create an Azure function that will call the Speech service and connect the bot to the function.
- D. Deploy the bot to Azure and register the bot with a Microsoft Teams channel.

Answer: B

NEW QUESTION 176

- (Topic 3)

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

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

Answer: B

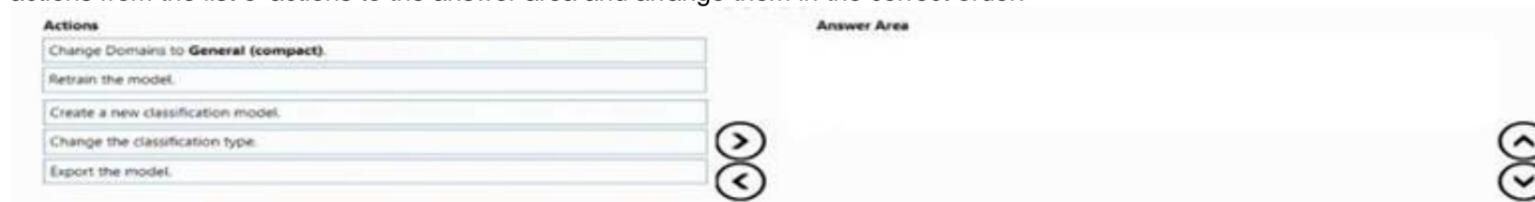
NEW QUESTION 179

DRAG DROP - (Topic 3)

You have a Custom Vision service project that performs object detection. The project uses the General domain for classification and contains a trained model.

You need to export the model for use on a network that is disconnected from the internet. Which three actions should you perform in sequence? To answer, move the appropriate

actions from the list o' actions to the answer area and arrange them in the correct order.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 183

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.

- 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 184

- (Topic 3)

You successfully run the following HTTP request.

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

What is the result of the request?

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

Answer: B

Explanation:

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

NEW QUESTION 185

DRAG DROP - (Topic 3)

You are building a transcription service for technical podcasts.

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

NEW QUESTION 187

- (Topic 3)

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

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

You create an Azure Cognitive Services resource.

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

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

Answer: B

NEW QUESTION 192

- (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 container
- 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 container
- 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 folder
- 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 196

- (Topic 3)

You are building a chatbot.

You need to ensure that the bot will recognize the names of your company's products and codenames. The solution must minimize development effort.

Which Azure Cognitive Service for Language service should you include in the solution?

- A. custom text classification
- B. entity linking

- C. custom Named Entity Recognition (NER)
- D. key phrase extraction

Answer: C

NEW QUESTION 201

- (Topic 3)

Which scenario is an example of a streaming workload?

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

Answer: D

NEW QUESTION 203

- (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 204

- (Topic 3)

You need to implement a table projection to generate a physical expression of an Azure Cognitive Search index.

Which three properties should you specify in the skillset definition JSON configuration table node? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

- A. tableName
- B. generatedKeyName
- C. dataSource
- D. dataSourceConnection
- E. source

Answer: ABE

Explanation:

Defining a table projection.

Each table requires three properties:

tableName: The name of the table in Azure Storage.

generatedKeyName: The column name for the key that uniquely identifies this row. source: The node from the enrichment tree you are sourcing your enrichments from. This node is usually the output of a shaper, but could be the output of any of the skills.

Reference:

<https://docs.microsoft.com/en-us/azure/search/knowledge-store-projection-overview>

NEW QUESTION 208

- (Topic 3)

You are building a language model by using a Language Understanding service. You create a new Language Understanding resource.

You need to add more contributors. What should you use?

- A. a conditional access policy in Azure Active Directory (Azure AD)
- B. the Access control (1AM) page for the authoring resources in the Azure portal
- C. the Access control (1AM) page for the prediction resources in the Azure portal

Answer: B

Explanation:

Reference:

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

NEW QUESTION 211

- (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 214

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 checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 219

DRAG DROP - (Topic 3)

You are developing a photo application that will find photos of a person based on a sample image by using the Face API.

You need to create a POST request to find the photos.

How should you complete the request? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

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

NOTE: Each correct selection is worth one point.

Values

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

Answer Area

POST {Endpoint}/face/v1.0/
 Request Body

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: findsimilar

<https://docs.microsoft.com/en-us/rest/api/faceapi/face/find-similar>

Box 2: matchPerson

Find similar has two working modes, "matchPerson" and "matchFace". "matchPerson" is the default mode that it tries to find faces of the same person as possible by using internal same-person thresholds. It is useful to find a known person's other photos. Note that an empty list will be returned if no faces pass the internal thresholds. "matchFace" mode ignores same-person thresholds and returns ranked similar faces anyway, even the similarity is low. It can be used in the cases like searching celebrity-looking faces.

NEW QUESTION 222

- (Topic 3)

You have a chatbot that was built by using the Microsoft Bot Framework. You need to debug the chatbot endpoint remotely.

Which two tools should you install on a local computer? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

- A. Fiddler
- B. Bot Framework Composer
- C. Bot Framework Emulator
- D. Bot Framework CLI
- E. ngrok
- F. nginx

Answer: CE

Explanation:

Bot Framework Emulator is a desktop application that allows bot developers to test and debug bots, either locally or remotely.

ngrok is a cross-platform application that "allows you to expose a web server running on your local machine to the internet." Essentially, what we'll be doing is using ngrok to forward messages from external channels on the web directly to our local machine to allow debugging, as opposed to the standard messaging endpoint configured in the Azure portal.

Reference:

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

NEW QUESTION 227

HOTSPOT - (Topic 3)

You need to create a new resource that will be used to perform sentiment analysis and optical character recognition (OCR). The solution must meet the following requirements:

? Use a single key and endpoint to access multiple services.

? Consolidate billing for future services that you might use.

? Support the use of Computer Vision in the future.

How should you complete the HTTP request to create the new resource? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

▼ <https://management.azure.com/subscriptions/xxxxxxx-xxxx-xxxx-xxxx-xxxx-xxxxxxxxxx/resourceGroups/RG1/providers/Microsoft.CognitiveServices/accounts/CS1?api-version=2017-04-18>

PATCH
POST
PUT

```
{
  "location": "West US",
  "kind": "CognitiveServices",
  "sku": {
    "name": "S0"
  },
  "properties": {},
  "identity": {
    "type": "SystemAssigned"
  }
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: PUT

Sample Request: PUT <https://management.azure.com/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/test-rg/providers/Microsoft.DeviceUpdate/accounts/contoso?api-version=2020-03-01-preview>

NEW QUESTION 228

- (Topic 3)

What are two benefits of platform as a service (PaaS) relational database offerings in Azure, such as Azure SQL Database? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Answer: AD

NEW QUESTION 230

DRAG DROP - (Topic 3)

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 232

- (Topic 3)

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

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

NEW QUESTION 234

HOTSPOT - (Topic 3)

You are building a model that will be used in an iOS app.

You have images of cats and dogs. Each image contains either a cat or a dog.

You need to use the Custom Vision service to detect whether the images is of a cat or a dog.

How should you configure the project in the Custom Vision portal? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Project Types:
 Classification Types:
 Domains:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Classification

Box 2: Multiclass

A multiclass classification project is for classifying images into a set of tags, or target labels. An image can be assigned to one tag only.

Box 3: General

General: Optimized for a broad range of image classification tasks. If none of the other specific domains are appropriate, or if you're unsure of which domain to choose, select one of the General domains.

NEW QUESTION 235

- (Topic 3)

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

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

Answer: B

NEW QUESTION 238

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 243

- (Topic 3)

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



What should you create in Cover BI Desktop;1

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

Answer: C

NEW QUESTION 247

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	Answer Area
Cassandra API	API Graph data
Gremlin API	API JSON documents
MongoDB API	API Key/value data
Table API	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

APIs	Answer Area
Cassandra API	
Gremlin API	Gremlin API Graph data
MongoDB API	MongoDB API JSON documents
Table API	Table API Key/value data

NEW QUESTION 248

HOTSPOT - (Topic 3)

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

Which Azure Cognitive Services API should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

https://

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 249

- (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 251

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 253

- (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 257

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.	<div style="display: flex; justify-content: space-around;"> ⬅ ⬆ </div>
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 261

- (Topic 3)

You build a bot by using the Microsoft Bot Framework SDK and the Azure Bot Service.

You plan to deploy the bot to Azure.

You register the bot by using the Bot Channels Registration service.

Which two values are required to complete the deployment? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. botId
- B. tenanId
- C. appld
- D. objecId
- E. appSecrec

Answer: CE

Explanation:

Reference:

<https://github.com/MicrosoftDocs/bot-docs/blob/live/articles/bot-service-quickstart-registration.md>

NEW QUESTION 263

- (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. /formrecognizer/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 265

- (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 add an additional component to meet the technical requirements and the chatbot requirements. What should you add?

- A. Dispatch
- B. chatdown
- C. Language Understanding
- D. Microsoft Translator

Answer: A

Explanation:

Scenario: All planned projects must support English, French, and Portuguese.

If a bot uses multiple LUIS models and QnA Maker knowledge bases (knowledge bases), you can use the Dispatch tool to determine which LUIS model or QnA Maker knowledge base best matches the user input. The dispatch tool does this by creating a single LUIS app to route user input to the correct model.

Reference:

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

NEW QUESTION 266

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 269

- (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 270

- (Topic 2)

You need to develop an extract solution for the receipt images. The solution must meet the document processing requirements and the technical requirements.

You upload the receipt images to the Form Recognizer API for analysis, and the API returns the following JSON.

```

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

```

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

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

Answer: D

Explanation:

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

Reference:

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

NEW QUESTION 275

- (Topic 2)

You are developing the knowledgebase by using Azure Cognitive Search.

You need to meet the knowledgebase requirements for searching equivalent terms. What should you include in the solution?

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

Answer: A

Explanation:

Within a search service, synonym maps are a global resource that associate equivalent terms, expanding the scope of a query without the user having to actually provide the term. For example, assuming "dog", "canine", and "puppy" are mapped synonyms, a query on "canine" will match on a document containing "dog".

Create synonyms: A synonym map is an asset that can be created once and used by many indexes.

Reference:

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

NEW QUESTION 277

.....

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](#)