

AI-102 Dumps

Designing and Implementing an Azure AI Solution

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



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

HOTSPOT - (Topic 3)

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

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

NOTE: Each correct selection is worth one point.

Answer Area

JSON data: ▼

- File projection
- Object projection
- Table projection

Scanned data: ▼

- File projection
- Object projection
- Table projection

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Object projection

Object projections are JSON representations of the enrichment tree that can be sourced from any node.

Box 2: File projection

File projections are similar to object projections and only act on the normalized_images collection.

NEW QUESTION 3

- (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 4

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

Relational data is stored in

a file system as unstructured data.
 a hierarchal folder structure.
 a tabular form of rows and columns.
 comma-separated value (CSV) files.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Relational data is stored in

a file system as unstructured data.
 a hierarchal folder structure.
 a Tabular form of rows and columns.
 comma-separated value (CSV) files.

NEW QUESTION 5

- (Topic 3)

You are building a chatbot.

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

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

Answer: B

NEW QUESTION 6

- (Topic 3)

You have an Azure subscription. The subscription contains an Azure OpenAI resource that hosts a GPT-3.5 Turbo model named Model1.

You configure Model1 to use the following system message: "You are an AI assistant that helps people solve mathematical puzzles. Explain your answers as if the request is by a 4- year-old."

Which type of prompt engineering technique is this an example of?

- A. few-shot learning
- B. affordance
- C. chain of thought
- D. priming

Answer: D

NEW QUESTION 7

DRAG DROP - (Topic 3)

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

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

? Prevent billing and API information from being stored in the command-line histories of the devices that run the container.

? Control access to the container images by using Azure role-based access control (Azure RBAC).

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

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

Actions

Answer Area

- Create a custom Dockerfile.
- Pull the Anomaly Detector container image.
- Distribute a `docker run` script.
- Push the image to an Azure container registry.
- Build the image.
- Push the image to Docker Hub.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Pull the Anomaly Detector container image. Step 2: Create a custom Dockerfile

Step 3: Build the image

Step 4: Push the image to an Azure container registry.

<https://docs.microsoft.com/en-us/azure/cognitive-services/containers/container-reuse-recipe>

NEW QUESTION 8

- (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 9

- (Topic 3)

What should you use to automatically delete blobs from Azure Blob Storage?

- A. the change feed
- B. a lifecycle management policy
- C. soft delete
- D. archive storage

Answer: D

NEW QUESTION 10

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

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

Find contacts in London. Who do I know in Seattle? Search for contacts in Ukraine.

You need to implement the phrase list in Language Understanding. Solution: You create a new intent for location.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

NEW QUESTION 12

- (Topic 3)

You have an Azure subscription that contains an Anomaly Detector resource. You deploy a Docker host server named Server 1 to the on-premises network. You need to host an instance of the Anomaly Detector service on Server 1. Which parameter should you include in the docker run command?

- A. Fluentd
- B. Billing
- C. Http Proxy
- D. Mounts

Answer: B

NEW QUESTION 16

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 17

DRAG DROP - (Topic 3)

You have a collection of Microsoft Word documents and PowerPoint presentations in German.

You need to create a solution to translate the file to French. The solution must meet the following requirements:

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

You create a blob container for German files and a blob container for French files. You upload the original files to the container for German files.

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

Actions

- Perform an asynchronous translation by using the list of files to be translated.
- Upload a glossary file to the container for German files.
- Upload a glossary file to the container for French files.
- Generate a list of files to be translated.
- Define a document translation specification that has a French target.
- Perform an asynchronous translation by using the document translation specification.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

- Perform an asynchronous translation by using the list of files to be translated.
- Upload a glossary file to the container for German files.
- Upload a glossary file to the container for French files.
- Generate a list of files to be translated.
- Define a document translation specification that has a French target.
- Perform an asynchronous translation by using the document translation specification.

Answer Area

Generate a list of files to be translated.

Define a document translation specification that has a French target.

Perform an asynchronous translation by using the document translation specification.

NEW QUESTION 22

- (Topic 3)

You are building an AI solution that will use Sentiment Analysis results from surveys to calculate bonuses for customer service staff. You need to ensure that the solution meets the Microsoft responsible AI principles. What should you do?

- A. Add a human review and approval step before making decisions that affect the staffs financial situation
- B. Include the Sentiment Analysis results when surveys return a low confidence score.
- C. Use all the surveys, including surveys by customers who requested that their account be deleted and their data be removed.
- D. Publish the raw survey data to a central location and provide the staff with access to the location.

Answer: A

NEW QUESTION 26

- (Topic 3)

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

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

Answer: C

NEW QUESTION 28

- (Topic 3)

You need to store event log data that is semi-structured and received as the logs occur. What should you use?

- A. Azure Table storage
- B. Azure Queue storage
- C. Azure Files

Answer: A

NEW QUESTION 30

- (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 35

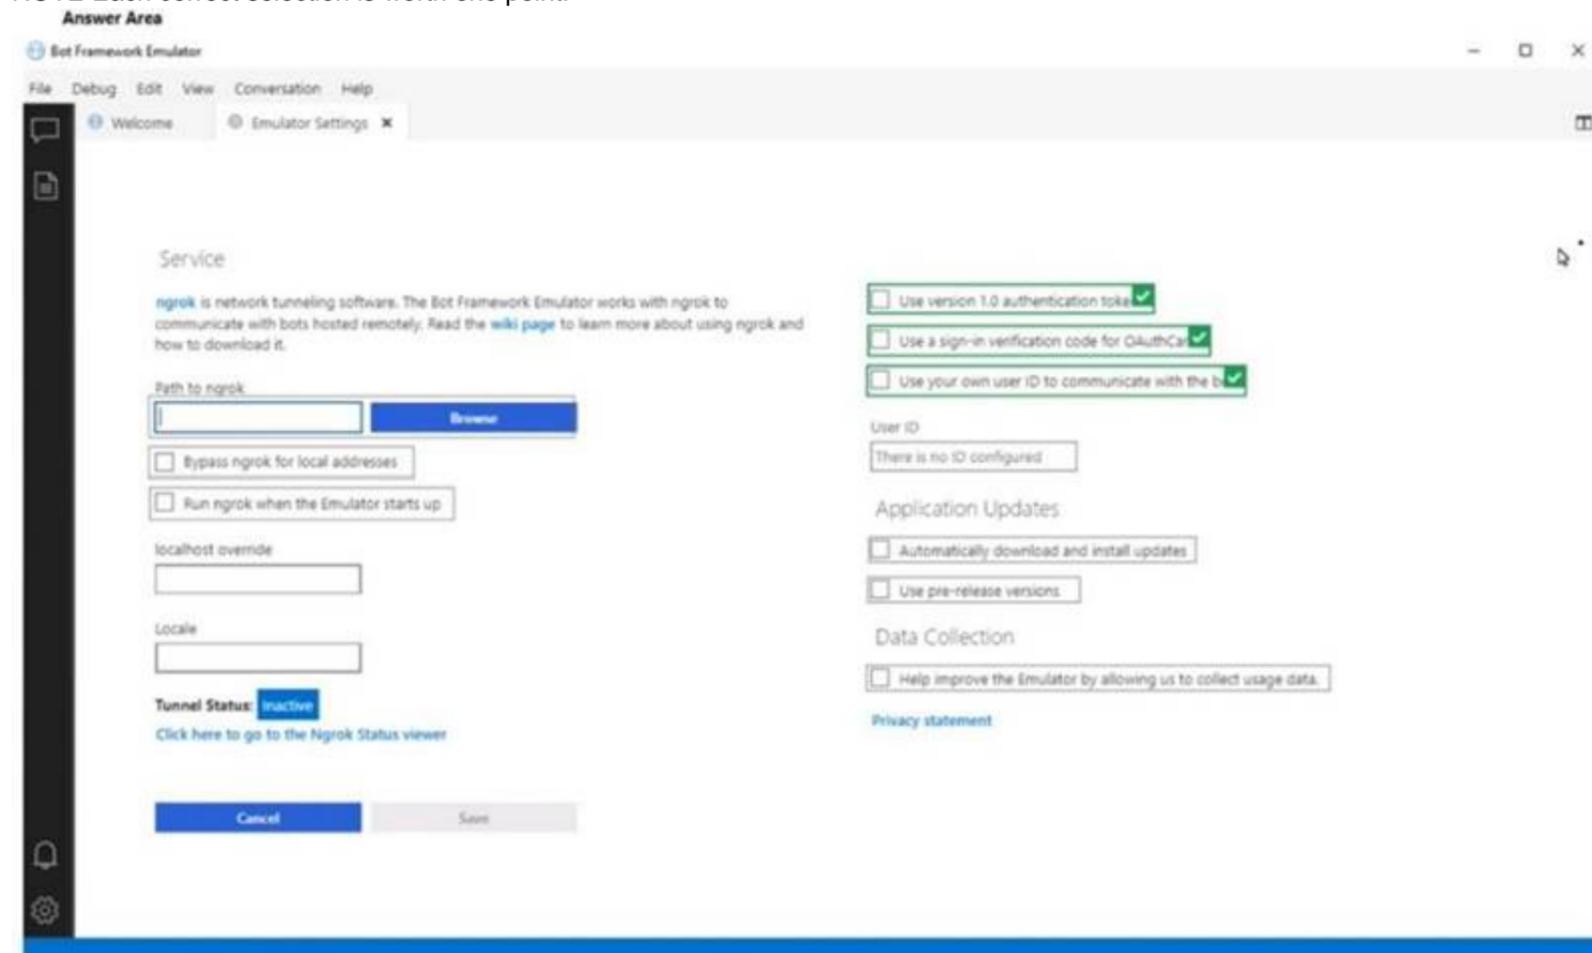
HOTSPOT - (Topic 3)

You have a chatbot.

You need to test the bot by using the Bot Framework Emulator. The solution must ensure that you are prompted for credentials when you sign in to the bot.

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

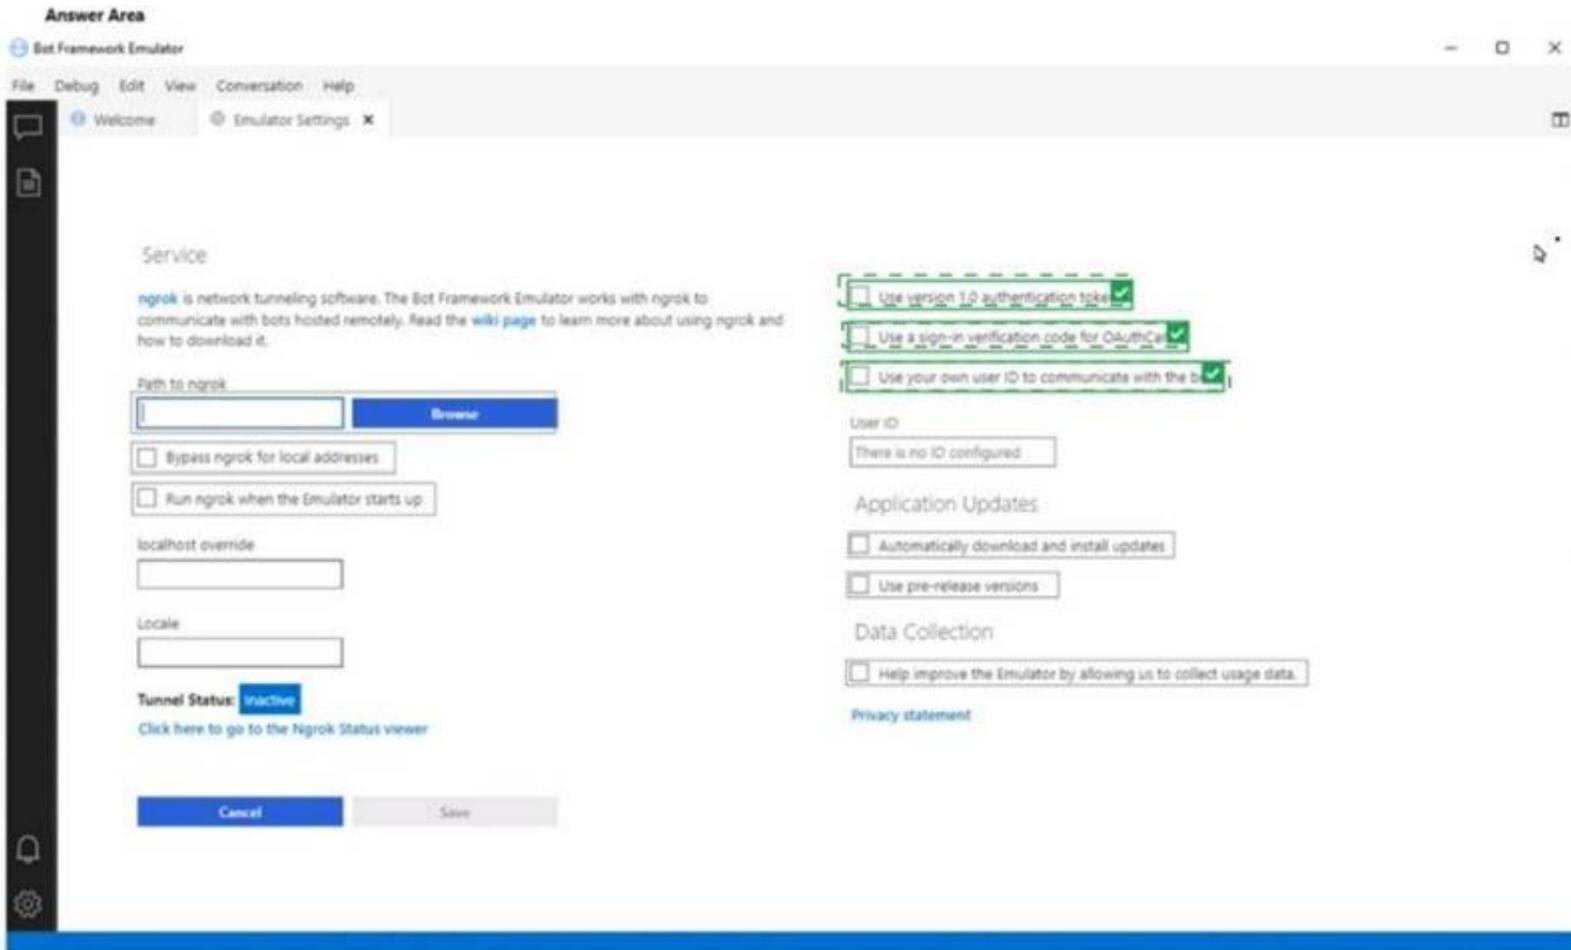
NOTE Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 36

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 40

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 42

- (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 47

HOTSPOT - (Topic 3)

You develop a test method to verify the results retrieved from a call to the Computer Vision API. The call is used to analyze the existence of company logos in images. The call returns a collection of brands named brands.

You have the following code segment.

```
foreach (var brand in brands)
{
    if (brand.Confidence >= .75)
        Console.WriteLine($"Logo of {brand.Name} between {brand.Rectangle.X},
{brand.Rectangle.Y} and {brand.Rectangle.W}, {brand.Rectangle.H}");
}
```

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 51

- (Topic 3)

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

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

Which type of dialog should you use?

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

Answer: A

NEW QUESTION 55

DRAG DROP - (Topic 3)

You plan to build a chatbot to support task tracking.

You create a Conversational Language Understanding service named lu1.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 59

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="radio"/>	<input type="radio"/>
The definition calls a web API as part of the enrichment process.	<input type="radio"/>	<input type="radio"/>
The enrichment step is called only for the first organization under "/document/organizations".	<input type="radio"/>	<input type="radio"/>

- 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 61

DRAG DROP - (Topic 3)

You are developing a call to the Face API. The call must find similar faces from an existing list named employeefaces. The employeefaces list contains 60,000 images.

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: LargeFaceListID

LargeFaceList: Add a face to a specified large face list, up to 1,000,000 faces.

Note: Given query face's faceId, to search the similar-looking faces from a faceId array, a face list or a large face list. A "faceListId" is created by FaceList - Create containing persistedFaceIds that will not expire. And a "largeFaceListId" is created by LargeFaceList - Create containing persistedFaceIds that will also not expire.

NEW QUESTION 65

- (Topic 3)

You are building a chatbot.

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

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

Answer: B

NEW QUESTION 67

- (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 70

- (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 72

- (Topic 3)

Which Azure Storage service implements the key/value model?

- A. Azure Files
- B. Azure Blob
- C. Azure Table
- D. Azure Queue

Answer: C

NEW QUESTION 77

- (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 79

- (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 83

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.

The screenshot shows the Microsoft Bot Framework Composer interface. On the left, a dialog flow is visible starting with a 'BeginDialog' event. The flow includes:

- A 'Bot Asks (Text)' step with the prompt 'What is your name?'.
- A 'User Input (Text)' step with the configuration 'username = Input(Text)'.
- A 'Bot Asks (Number)' step with the prompt 'Hello \$(user.name), how old are you?'.
- A 'User Input (Number)' step with the configuration 'user.age = Input(Number)'.

 On the right, the configuration panel for the selected 'Text input' prompt is shown. It includes:

- Property: string (dropdown), value: user.name
- Output Format: string (dropdown), value: ex. =toUpper(this.value), \$(toUpper(this.value))
- Value: expression (dropdown), value: fx =coalesce(@user.Name.@personName)
- Expected responses (intent): #TextInput_Response_GH5FTe

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No
 User.name is a property. Box 2: Yes
 Box 3: Yes
 The coalesce() function evaluates a list of expressions and returns the first non-null (or non-empty for string) expression.

NEW QUESTION 84

HOTSPOT - (Topic 3)

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

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

▼ \

`http://contoso.blob.core.windows.net`

`https://contoso.cognitiveservices.azure.com`

`mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase`

`mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment`

```
Eula=accept \
```

Billing=

▼ \

`http://contoso.blob.core.windows.net`

`https://contoso.cognitiveservices.azure.com`

`mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase`

`mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment`

```
ApiKey=xxxxxxxxxxxxxxxxxxxxxxxx
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 89

HOTSPOT - (Topic 3)

You have an Azure subscription that contains an Azure AI Document Intelligence resource named DM. You build an app named App1 that analyzes PDF files for handwritten content by using DM. You need to ensure that App1 will recognize the handwritten content. How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

```

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

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

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

```

NEW QUESTION 90

DRAG DROP - (Topic 3)

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

Actions

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

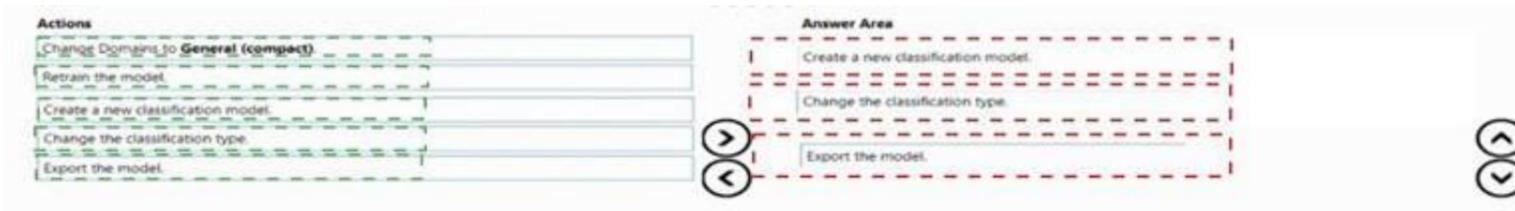
Answer Area

⏪
⏩

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 91

DRAG DROP - (Topic 3)

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

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

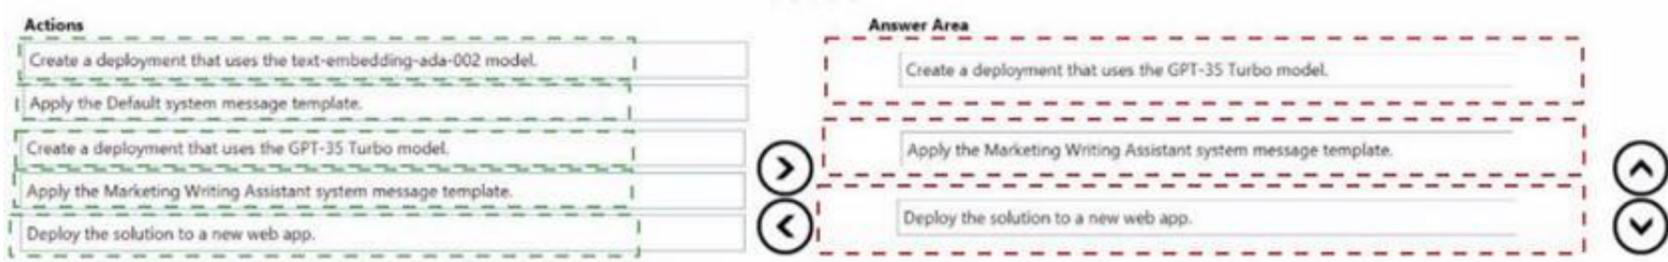
Which three actions should you perform in sequence in Azure OpenAI Studio? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 95

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 96

HOTSPOT - (Topic 3)

You are building an app that will share user images.

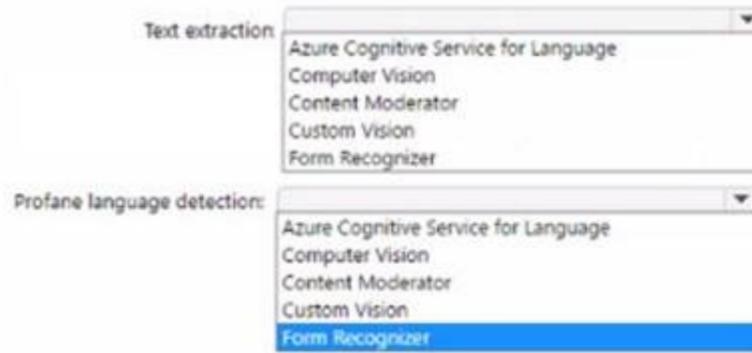
You need to configure the app to meet the following requirement's

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

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

NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Extraction: b) Computer Vision

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

Profane Language Detection: c) Content Moderator

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

NEW QUESTION 101

DRAG DROP - (Topic 3)

You have a Language Understanding solution that runs in a Docker container.

You download the Language Understanding container image from the Microsoft Container Registry (MCR).

You need to deploy the container image to a host computer.

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

Actions	Answer Area
From the Language Understanding portal, retrain the model.	
From the host computer, run the container and specify the input directory.	
From the Language Understanding portal, export the solution as a package file.	
From the host computer, move the package file to the Docker input directory.	
From the host computer, build the container and specify the output directory.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
From the Language Understanding portal, retrain the model.	
From the host computer, run the container and specify the input directory.	
From the Language Understanding portal, export the solution as a package file.	From the Language Understanding portal, export the solution as a package file.
From the host computer, move the package file to the Docker input directory.	From the host computer, move the package file to the Docker input directory.
From the host computer, build the container and specify the output directory.	From the host computer, build the container and specify the output directory.

NEW QUESTION 102

DRAG DROP - (Topic 3)

You are building a transcription service for technical podcasts.

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

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

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

Actions

- Create a Speaker Recognition model.
- Create a Conversational Language Understanding model.
- Create a Custom Speech project.
- Create a speech-to-text model.
- Upload training datasets.
- Train the model.
- Deploy the model.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

NEW QUESTION 103

- (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 106

- (Topic 3)

You have a mobile app that manages printed forms.

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

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

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

Answer: A

NEW QUESTION 111

- (Topic 3)

You are developing a method that uses the Computer Vision client library. The method will perform optical character recognition (OCR) in images. The method has the following code.

```
public static async Task ReadFileUrl(ComputerVisionClient client, string urlFile)
{
    const int numberOfCharsInOperationId = 36;

    var txtHeaders = await client.ReadAsync(urlFile, language: "en");

    string opLocation = txtHeaders.OperationLocation;
    string operationId = opLocation.Substring(opLocation.Length -
numberOfCharsInOperationId);

    ReadOperationResult results;

    results = await client.GetReadResultAsync(Guid.Parse(operationId));

    var textUrlFileResults = results.AnalyzeResult.ReadResults;
    foreach (ReadResult page in textUrlFileResults)
    {
        foreach (Line line in page.Lines)
        {
            Console.WriteLine(line.Text);
        }
    }
}
```

During testing, you discover that the call to the GetReadResultAsync method occurs before the read operation is complete. You need to prevent the GetReadResultAsync method from proceeding until the read operation is complete. Which two actions should you perform? Each correct answer presents part of the solution. (Choose two.)
NOTE: Each correct selection is worth one point.

- A. Remove the Guid.Parse(operationId) parameter.
- B. Add code to verify the results.Status value.
- C. Add code to verify the status of the txtHeaders.Status value.
- D. Wrap the call to GetReadResultAsync within a loop that contains a delay.

Answer: BD

Explanation:

Example code : do

```
{
results = await client.GetReadResultAsync(Guid.Parse(operationId));
}
```

while ((results.Status == OperationStatusCodes.Running || results.Status == OperationStatusCodes.NotStarted));

Reference:

<https://github.com/Azure-Samples/cognitive-services-quickstart-code/blob/master/dotnet/ComputerVision/ComputerVisionQuickstart.cs>

NEW QUESTION 115

- (Topic 3)

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

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

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

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

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

- A. Yes
- B. No

Answer: A

Explanation:

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

Reference:

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

NEW QUESTION 116

- (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 119

HOTSPOT - (Topic 3)

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

```

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

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

Answer Area

There will be [answer choice].
 no projection groups
 one projection group
 two projection groups
 four projection groups

Images will [answer choice]
 not be saved.
 be saved to a blob container.
 be saved to file storage.
 be saved to an Azure Data lake.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

There will be [answer choice].
 no projection groups
 one projection group
 two projection groups
 four projection groups

Images will [answer choice]
 not be saved.
 be saved to a blob container.
 be saved to file storage.
 be saved to an Azure Data lake.

NEW QUESTION 121

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

PATCH
POST
PUT

xxxx-xxxx-

xxxxxxxxxxxxx/resourceGroups/RG1/providers/Microsoft.CognitiveServices/
accounts/CS1?api-version=2017-04-18

{
 "location": "West US",
 "kind": "

CognitiveServices
ComputerVision
TextAnalytics

",
 "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 126

HOTSPOT - (Topic 3)

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

```

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

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

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

```

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

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 127

- (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 131

- (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 136

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

is responsible for identifying which business rules must be applied to the data of a company.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

is responsible for identifying which business rules must be applied to the data of a company.

NEW QUESTION 139

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 141

DRAG DROP - (Topic 3)

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

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

NOTE: Each correct match is worth one point.

APIs

- Cassandra API
- Gremlin API
- MongoDB API
- Table API

Answer Area

- API Graph data
- API JSON documents
- API Key/value data

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

APIs

- Cassandra API
- Gremlin API
- MongoDB API
- Table API

Answer Area

- Gremlin API Graph data
- MongoDB API JSON documents
- Table API Key/value data

NEW QUESTION 146

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a

graph.
table.
partition key.
document.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a

graph.
table.
partition key.
document.

NEW QUESTION 149

- (Topic 3)

You have an Azure Cognitive Search instance that indexes purchase orders by using Form Recognizer. You need to analyze the extracted information by using Microsoft Power BI. The solution must minimize development effort. What should you add to the indexer?

- A. a table projection
- B. a projection group
- C. an object projection
- D. a file projection

Answer: C

NEW QUESTION 152

HOTSPOT - (Topic 3)

You have a Computer Vision resource named contoso1 that is hosted in the West US Azure region. You need to use contoso1 to make a different size of a product photo by using the smart cropping feature.

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

westus.api.cognitive.microsoft.com generateThumbnail
<https://docs.microsoft.com/en-us/rest/api/computervision/3.1/generate-thumbnail/generate-thumbnail#examples>
 POST
<https://westus.api.cognitive.microsoft.com/vision/v3.1/generateThumbnail?width=500&height=500&smartCropping=True>
 Ocp-Apim-Subscription-Key: {API key}

NEW QUESTION 154

- (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 156

- (Topic 3)

What is a primary characteristic of a relational database?

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

Answer: C

NEW QUESTION 159

- (Topic 3)

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

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

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

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

Answer: C

NEW QUESTION 160

HOTSPOT - (Topic 3)

You are developing a text processing solution. You develop the following method.

```
static void GetKeyPhrases(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.ExtractKeyPhrases(text);
    Console.WriteLine("Key phrases:");

    foreach (string keyphrase in response.Value)
    {
        Console.WriteLine($"{keyphrase}");
    }
}
```

You call the method by using the following code. GetKeyPhrases(textAnalyticsClient, "the cat sat on the mat");

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

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

The Key Phrase Extraction API evaluates unstructured text, and for each JSON document, returns a list of key phrases.

Box 2: No

'the' is not a key phrase.

This capability is useful if you need to quickly identify the main points in a collection of documents. For example, given input text "The food was delicious and there were wonderful staff", the service returns the main talking points: "food" and "wonderful staff".

Box 3: No

Key phrase extraction does not have confidence levels.

NEW QUESTION 161

- (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 164

- (Topic 2)

You are developing the chatbot.

You create the following components:

- * A QnA Maker resource
- * A chatbot by using the Azure Bot Framework SDK.

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

- A. QnADialogResponseOptions.CardNoMatchText
- B. Qna MakerOptions-ScoreThreshold
- C. Qna Maker Options StrickFilters
- D. QnaMakerOptions.RankerType

Answer: D

Explanation:

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

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

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

Reference:

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

NEW QUESTION 168

- (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 172

- (Topic 2)

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

- A. an indexer for Azure Blob storage attached to a skillset that contains the language detection skill and the text translation skill

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

Answer: C

Explanation:

The wiki contains text in English, French and Portuguese.

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

The Document Extraction skill extracts content from a file within the enrichment pipeline. This allows you to take advantage of the document extraction step that normally happens before the skillset execution with files that may be generated by other skills.

Note: The Translator Text API will be used to determine the from language. The Language detection skill is not required.

Reference:

<https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-document-extraction> <https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-text-translation>

NEW QUESTION 177

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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