

Microsoft

Exam Questions AI-102

Designing and Implementing an Azure AI Solution



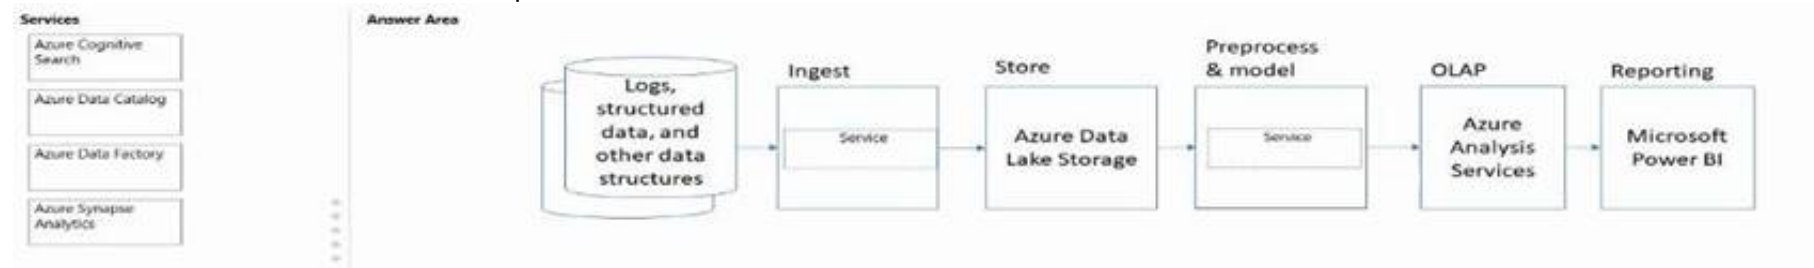
NEW QUESTION 1

DRAG DROP - (Topic 3)

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

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


NOTE: Each correct match is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 2

- (Topic 3)

You need to enable speech capabilities for a chatbot.

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

NOTE: Each correct selection is worth one point.

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

Answer: ABC

Explanation:

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

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

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

Reference:

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

NEW QUESTION 3

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
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	<input type="radio"/>	<input type="radio"/>
All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 4

- (Topic 3)
You have an Azure IoT hub that receives series data from machinery. You need to build an app that will perform the following actions:

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

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

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

Answer: D

NEW QUESTION 5

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 6

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 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 Modell!

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)

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

What should you upload to the profile?

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

Answer: D

NEW QUESTION 9

DRAG DROP - (Topic 3)

You have a web app that uses Azure Cognitive Search.

When reviewing billing for the app, you discover much higher than expected charges. You suspect that the query key is compromised.

You need to prevent unauthorized access to the search endpoint and ensure that users only have read only access to the documents collection. The solution must minimize app downtime.

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

Actions

Add a new query key.

Regenerate the secondary admin key.

Change the app to use the secondary admin key.

Change the app to use the new key.

Regenerate the primary admin key.

Delete the compromised key.

Answer Area

⏪

⏩

⏴

⏵

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Add a new query key.

Regenerate the secondary admin key.

Change the app to use the secondary admin key.

Change the app to use the new key.

Regenerate the primary admin key.

Delete the compromised key.

Answer Area

⏪

⏩

⏴

⏵

Add a new query key.

Change the app to use the new key.

Delete the compromised key.

NEW QUESTION 10

- (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 10

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

A data analyst

A data engineer

A data scientist

A database administrator

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A data analyst

A data engineer

A data scientist

A database administrator

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

NEW QUESTION 15

- (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 17

- (Topic 3)

You are building a Language Understanding solution.

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

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

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

Answer: C

NEW QUESTION 21

DRAG DROP - (Topic 3)

You have an app that manages feedback.

You need to ensure that the app can detect negative comments by using the Sentiment Analysis API in Azure Cognitive Service for Language. The solution must

ensure that the managed feedback remains on your company's internal network.

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

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

Actions	Answer Area
Provision the Language service resource in Azure.	<div> <div>></div> <div><</div> <div>↑</div> <div>↓</div> </div>
Deploy a Docker container to an Azure container instance.	
Deploy a Docker container to an on-premises server.	
Identify the Language service endpoint URL and query the prediction endpoint.	
Run the container and query the prediction endpoint.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

? Provision the Language service resource in Azure.

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

? Run the container and query the prediction endpoint.

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

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

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

NEW QUESTION 26

- (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 31

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

var apiKey = "FF956C68B83B21B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);


    var uri = endpoint + &quot;?from=en&quot;;
    var uri = endpoint + &quot;?suggestedFrom=en&quot;;
    var uri = endpoint + &quot;?to=en&quot;;


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

```

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

```

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

var apiKey = "FF956C68B83B21B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);


    var uri = endpoint + &quot;?from=en&quot;;
    var uri = endpoint + &quot;?suggestedFrom=en&quot;;
    var uri = endpoint + &quot;?to=en&quot;;


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

```

NEW QUESTION 35

- (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 37

- (Topic 3)

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

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

Answer: A

NEW QUESTION 39

- (Topic 3)

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

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

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

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

Answer: BD

NEW QUESTION 44

- (Topic 3)

You plan to build an app that will generate a list of tags for uploaded images. The app must meet the following requirements:

- Generate tags in a user's preferred language.
- Support English, French, and Spanish.
- Minimize development effort

You need to build a function that will generate the tags for the app. Which Azure service endpoint should you use?

- A. Custom Vision image classification
- B. Content Moderator Image Moderation
- C. Custom Translator
- D. Computer Vision Image Analysis

Answer: A

NEW QUESTION 49

HOTSPOT - (Topic 3)

You are developing an application that includes language translation.

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

```

...
api_key = "FF956C68B83821838691A8D200A4C606"
text = get_text_to_be_translated()
headers = {
    'Content-Type': 'application/json',
    'Ocp-Apim-Subscription-Key': api_key
}
body = {
    'Text': text
}
conn = httpLib.HTTPSConnection
conn.request("POST",
response = conn.getresponse()
    conn.request("POST",
response = conn.getrespons
response_data = response.r
...

```

☐ ("api.cognitive.microsofttranslator.com")
☐ ("api-apc.cognitive.microsofttranslator.com")
☐ ("api-nam.cognitive.microsofttranslator.com")

☐ "/translate?from=en"
☐ "/translate?suggestedFrom=en"
☐ "/translate?to=en"
☐ "/detect?to=en"
☐ "/detect?from=en"

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

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

```



NEW QUESTION 53

- (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 55

- (Topic 3)

You are developing a monitoring system that will analyze engine sensor data, such as rotation speed, angle, temperature, and pressure. The system must generate an alert in response to atypical values. What should you include in the solution?

- A. Application Insights in Azure Monitor
- B. metric alerts in Azure Monitor
- C. Multivariate Anomaly Detection
- D. Univariate Anomaly Detection

Answer: C

NEW QUESTION 58

- (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 61

DRAG DROP - (Topic 3)

You have a Docker host named Host1 that contains a container base image.

You have an Azure subscription that contains a custom speech-to-text model named model1.

You need to run model1 on Host1.

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

Actions

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

➡

⬅

Answer Area

- 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:

Passing Certification Exams Made Easy

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

- ? 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 63

- (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 66

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 68

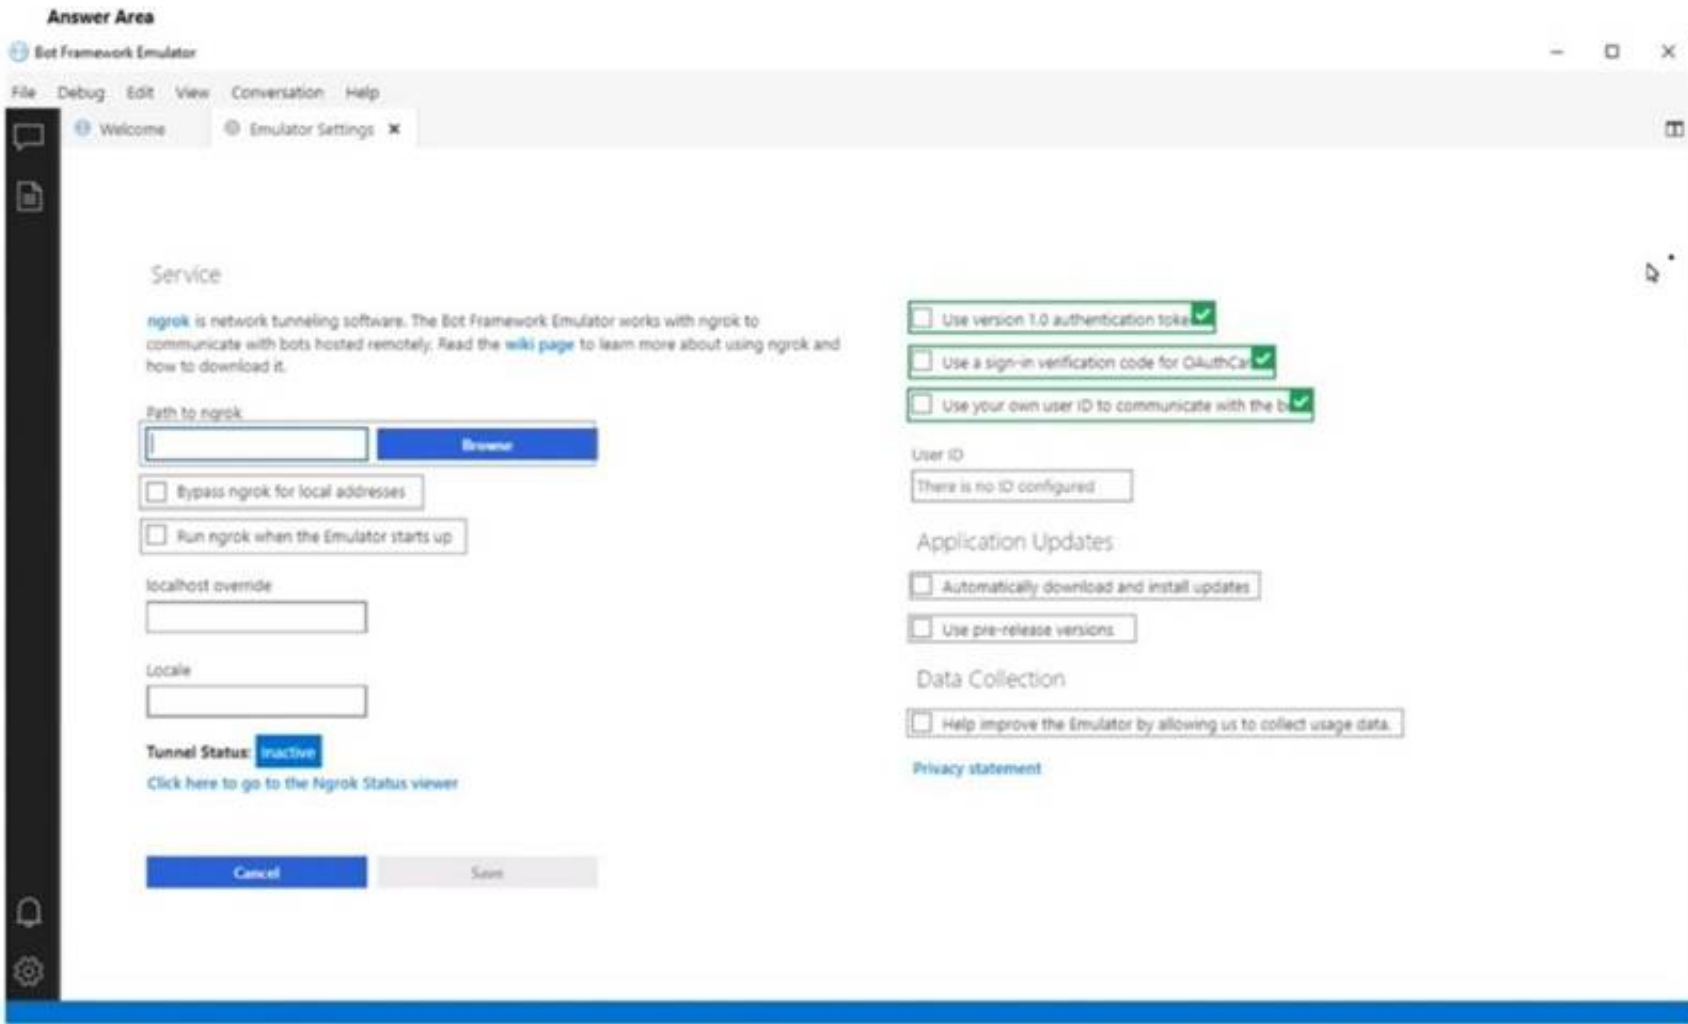
HOTSPOT - (Topic 3)

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 70

- (Topic 3)

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

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

Answer: A

NEW QUESTION 72

- (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 73

- (Topic 3)

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

A. isolation

B. atomicity

C. consistency

D. durability

Answer: B

NEW QUESTION 75

- (Topic 3)

You are building a Conversational Language Understanding model.

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

- Set all the lights to on.
- Turn off the lights in the living room.
- What is the current thermostat temperature?
- Lower the temperature of the thermostat by five degrees. Which three elements should you add to the model?

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

A. a location Intent

B. a change setting entity

C. a device intent

D. a change setting intent

E. a query setting intent

F. a device entity

Answer: BCF

NEW QUESTION 79

- (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 80

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

A block of code that runs in a database is called

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A block of code that runs in a database is called

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

NEW QUESTION 85

DRAG DROP - (Topic 3)

You are building an app that will scan confidential documents and use the Language service to analyze the contents. You provision an Azure Cognitive Services resource. You need to ensure that the app can make requests to the Language service endpoint. The solution must ensure that confidential documents remain on-premises. 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

Pull an image from Docker Hub.

Provision an on-premises Kubernetes cluster that has internet connectivity.

Provision an Azure Kubernetes Service (AKS) resource.

Run the container and specify an App ID and Client Secret.

Provision an on-premises Kubernetes cluster that is isolated from the internet.

Pull an image from the Microsoft Container Registry (MCR).

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Pull an image from Docker Hub.

Provision an on-premises Kubernetes cluster that has internet connectivity.

Provision an Azure Kubernetes Service (AKS) resource.

Run the container and specify an App ID and Client Secret.

Provision an on-premises Kubernetes cluster that is isolated from the internet.

Pull an image from the Microsoft Container Registry (MCR).

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Answer Area

Provision an on-premises Kubernetes cluster that is isolated from the internet.

Pull an image from the Microsoft Container Registry (MCR).

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

NEW QUESTION 88

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 93

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 96

- (Topic 3)

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

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

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

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

D. create_resource(client, "res1", "CustomVision.Prediction", "S0", "westus")

Answer: B

Explanation:

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

NEW QUESTION 100

- (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 105

DRAG DROP - (Topic 3)

Match the types of workloads to the appropriate scenarios.

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

NOTE: Each correct match is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 110

- (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 112

- (Topic 3)

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

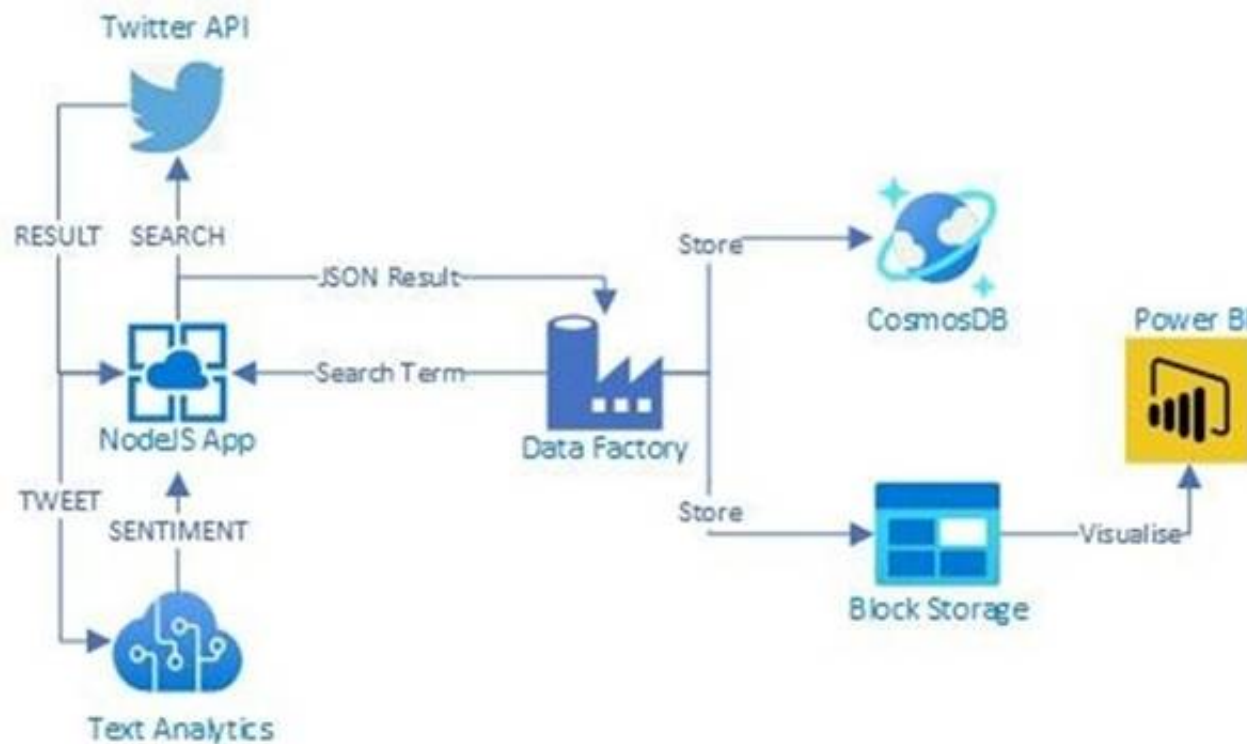
- A. Text Analytics
- B. Content Moderator
- C. Computer Vision
- D. Form Recognizer

Answer: A

Explanation:

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

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



Reference:

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

NEW QUESTION 115

DRAG DROP - (Topic 3)

You are building a customer support chatbot.

You need to configure the bot to identify the following:

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

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

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

NOTE: Each correct selection is worth one point.

Features

- Custom named entity recognition (NER)
- Key phrase extraction
- Language detection
- Named Entity Recognition (NER)
- Personally Identifiable Information (PII) detection
- Sentiment analysis

Answer Area

Identify code names for internal product development:

Identify messages that include credit card numbers:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Features

- Custom named entity recognition (NER)
- Key phrase extraction
- Language detection
- Named Entity Recognition (NER)
- Personally Identifiable Information (PII) detection
- Sentiment analysis

Answer Area

Identify code names for internal product development: Custom named entity recognition (NER)

Identify messages that include credit card numbers: Personally Identifiable Information (PII) detection

NEW QUESTION 120

DRAG DROP - (Topic 3)

You need to develop an automated call handling system that can respond to callers in their own language. The system will support only French and English.

Which Azure Cognitive Services service should you use to meet each requirement? To answer, drag the appropriate services to the correct requirements. Each service may be used once, more than once, or not at all. You may need to drag the split bat between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Services

Answer Area

Speaker Recognition

Speech to Text

Text Analytics

Text to Speech

Translator

Detect the incoming language:

Respond in the callers' own language:

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Box 1: Text Analytics

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

Box 2: Translator

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

NEW QUESTION 122

- (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 124

HOTSPOT - (Topic 3)

You plan to deploy a containerized version of an Azure Cognitive Services service that will be used for text analysis.

You configure <https://contoso.cognitiveservices.azure.com> as the endpoint URI for the service, and you pull the latest version of the Text Analytics Sentiment Analysis container.

You need to run the container on an Azure virtual machine by using Docker.

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

NOTE: Each correct selection is worth one point.

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

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

Eula=accept \

Billing=

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

ApiKey=xxxxxxxxxxxxxxxxxxxxxx

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Box 1: mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment

To run the Sentiment Analysis v3 container, execute the following docker run command. docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \

mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment \ Eula=accept \

Billing={ENDPOINT_URI} \

ApiKey={API_KEY} is the endpoint for accessing the Text Analytics API. https://<your- custom-subdomain>.cognitiveservices.azure.com

Box 2: https://contoso.cognitiveservices.azure.com

{ENDPOINT_URI} is the endpoint for accessing the Text Analytics API: https://<your- custom-subdomain>.cognitiveservices.a The endpoint for accessing the Text Analytics API. zure.com

NEW QUESTION 128

- (Topic 3)

What is the primary purpose of a data warehouse?

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

Answer: C

NEW QUESTION 133

HOTSPOT - (Topic 3)

You are building an app that will share user images.

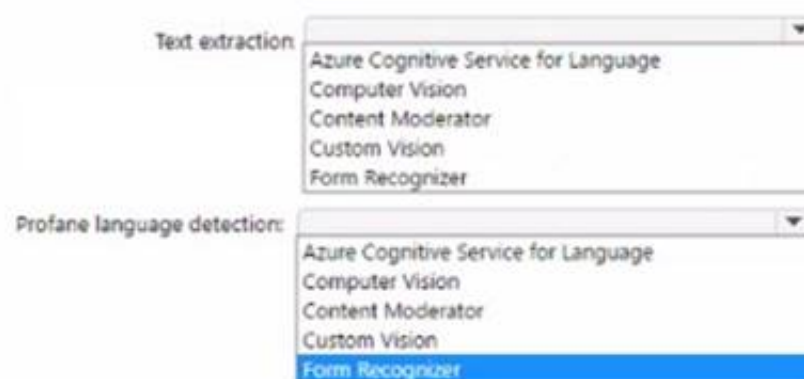
You need to configure the app to meet the following requirement's

- 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 apocopate opinions in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area



Text extraction: Computer Vision

Profane language detection: Content Moderator

- 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 135

- (Topic 3)

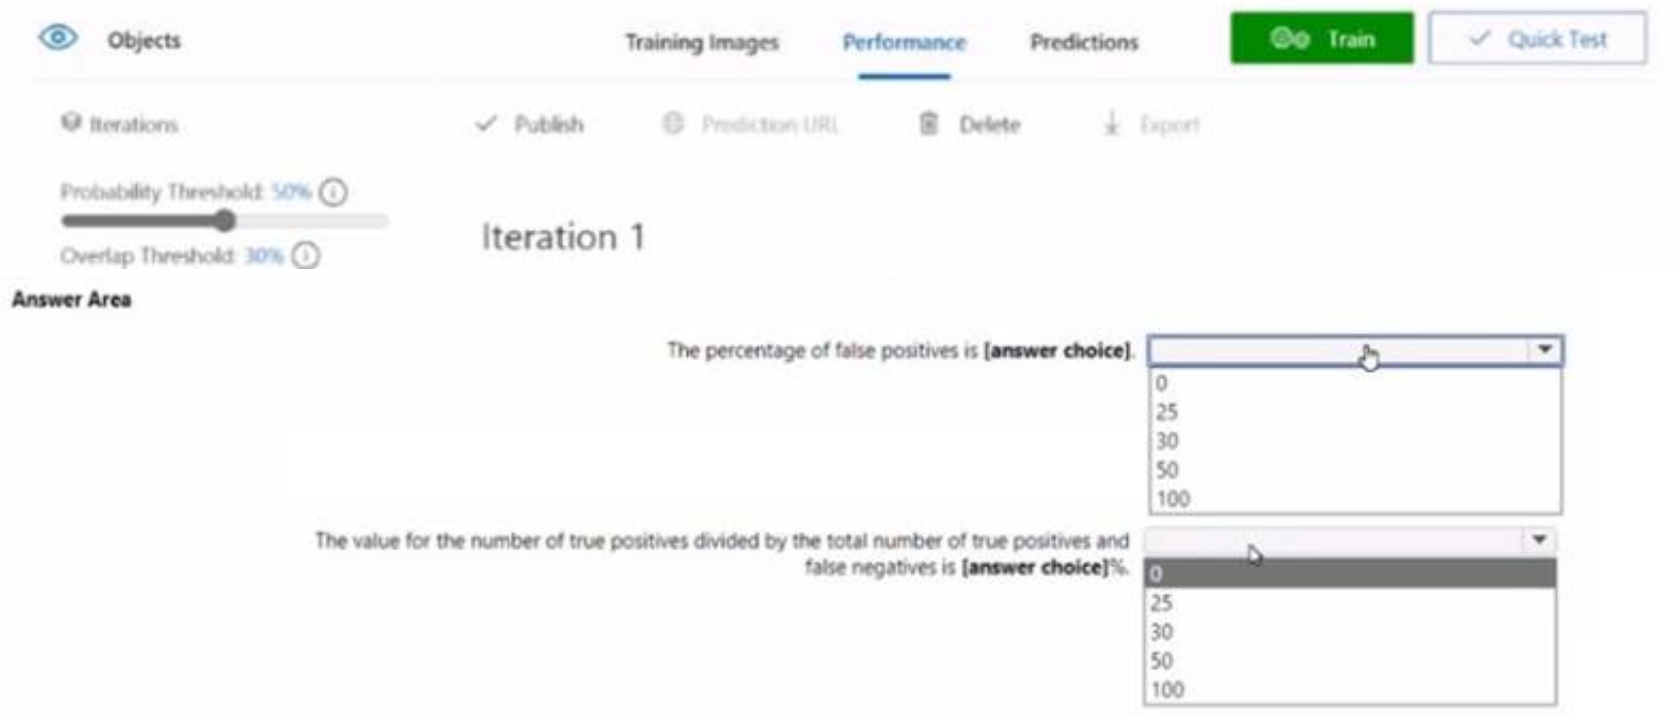
You have an app named App1 that uses an Azure Cognitive Services model to identify anomalies in a time series data stream. You need to run App1 in a location that has limited connectivity. The solution must minimize costs. What should you use to host the model?

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

Answer: B

NEW QUESTION 139

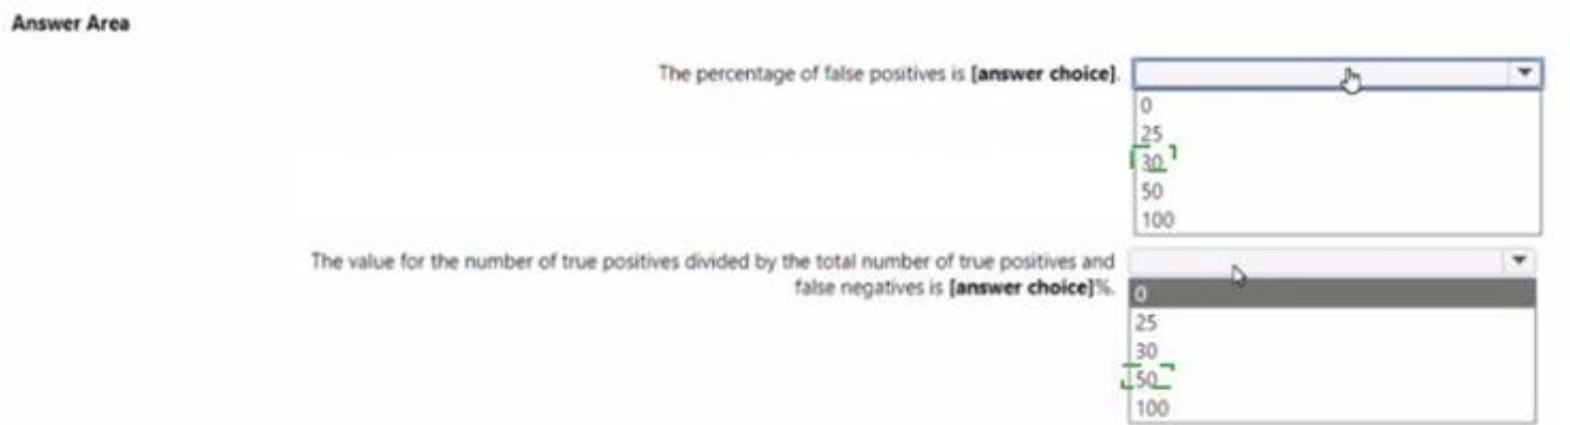
HOTSPOT - (Topic 3)
You are building a model to detect objects in images.
The performance of the model based on training data is shown in the following exhibit.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 144

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

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.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 148

- (Topic 3)
You have data saved in the following format.

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

Which format was used?

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

Answer: A

NEW QUESTION 150

HOTSPOT - (Topic 3)
You are building a chatbot by using the Microsoft Bot Framework SDK.
You use an object named UserProfile to store user profile information and an object named ConversationData to store information related to a conversation.
You create the following state accessors to store both objects in state.
var userStateAccessors = _userState.CreateProperty<UserProfile>(nameof(UserProfile)); var conversationStateAccessors =
_conversationState.CreateProperty<ConversationData>(nameof(ConversationData)); The state storage mechanism is set to Memory Storage.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
You create property accessors using the CreateProperty method that provides a handle to the BotState object. Each state property accessor allows you to get or set the value of the associated state property.
Box 2: Yes
Box 3: No
Before you exit the turn handler, you use the state management objects' SaveChangesAsync() method to write all state changes back to storage.

NEW QUESTION 151

HOTSPOT - (Topic 3)
You are building a chatbot.
You need to use the Content Moderator API to identify aggressive and sexually explicit language.
Which three settings should you configure? To answer, select the appropriate settings in the answer area.
NOTE: Each correct selection is worth one point.

Microsoft

Cognitive Services

Image

Text

Content Moderator - Moderate

Text - Screen

The operation detects profanity in more than 100 languages and match against custom and shared blacklists.

Host

Name

uksouth.api.cognitive.microso

Query parameters

autocorrect

Value

Remove parameter

Pii

Value

Remove parameter

listId

Value

Remove parameter

classify

false

Remove parameter

language

Value

Remove parameter

+ Add parameter

Headers

Content-Type

text/plain

Remove header

Ocp-Apim-Subscription-Key

Value

+ Add header

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Microsoft

Cognitive Services

Image

Text

Content Moderator - Moderate

Text - Screen

The operation detects profanity in more than 100 languages and match against custom and shared blacklists.

Host

Name

uksouth.api.cognitive.microso

Query parameters

autocorrect

Value

Remove parameter

Pii

Value

Remove parameter

listId

Value

Remove parameter

classify

false

Remove parameter

language

Value

Remove parameter

+ Add parameter

Headers

Content-Type

text/plain

Remove header

Ocp-Apim-Subscription-Key

Value

+ Add header

NEW QUESTION 152

- (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 154

- (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 157

- (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 160

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: findsimilars

<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 161

- (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 163

DRAG DROP - (Topic 3)

You plan to use a Language Understanding application named app1 that is deployed to a container.

App1 was developed by using a Language Understanding authoring resource named lu1. App1 has the versions shown in the following table.

Version	Trained date	Published date
V1.2	None	None
V1.1	2020-10-01	None
V1.0	2020-09-01	2020-09-15

You need to create a container that uses the latest deployable version of app1.

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

- Run a container that has version set as an environment variable.
- Export the model by using the Export as JSON option.
- Select v1.1 of app1.
- Run a container and mount the model file.
- Select v1.0 of app1.
- Export the model by using the Export for containers (GZIP) option.
- Select v1.2 of app1.

Answer Area

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Step 1: Export the model using the Export for containers (GZIP) option. Export versioned app's package from LUIS portal
The versioned app's package is available from the Versions list page.
? Sign on to the LUIS portal.
? Select the app in the list.
? Select Manage in the app's navigation bar.
? Select Versions in the left navigation bar.
? Select the checkbox to the left of the version name in the list.
? Select the Export item from the contextual toolbar above the list.
? Select Export for container (GZIP).
? The package is downloaded from the browser.



Step 2: Select v1.1 of app1.
A trained or published app packaged as a mounted input to the container with its associated App ID.
Step 3: Run a contain and mount the model file.
Run the container, with the required input mount and billing settings.

NEW QUESTION 168

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

https://management.azure.com/subscriptions/xxxxxxx-xxxx-

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 173

DRAG DROP - (Topic 3)
You need to analyze video content to identify any mentions of specific company names.
Which three actions should you perform in sequence? To answer move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Add the specific company names to the exclude list.

Sign in to the Custom Vision website.

From Content model customization, select **Language**.

Sign in to the Azure Video Analyzer for Media website.

From Content model customization, select **Brands**.

Add the specific company names to the include list.

>

<

Answer Area

>

<

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Add the specific company names to the exclude list.

Sign in to the Custom Vision website.

From Content model customization, select **Language**.

Sign in to the Azure Video Analyzer for Media website.

From Content model customization, select **Brands**.

Add the specific company names to the include list.

>

<

Answer Area

Sign in to the Azure Video Analyzer for Media website.

From Content model customization, select **Brands**.

Add the specific company names to the include list.

>

<

NEW QUESTION 178

- (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 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 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 chat scenarios, which may result in a poor user experience and a loss of engagement. Instead, you should choose a chat personality that matches the tone and style of your chatbot, such as Professional or Caring. You can also edit the 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 182

- (Topic 3)

You are building an Azure WebJob that will create knowledge bases from an array of URLs.
You instantiate a QnAMakerClient object that has the relevant API keys and assign the object to a variable named client.
You need to develop a method to create the knowledge bases.
Which two actions should you include in the method? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

Answer: AC

Explanation:

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

NEW QUESTION 184

- (Topic 3)

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

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

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

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

Answer: ACF

Explanation:

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

NEW QUESTION 185

- (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?

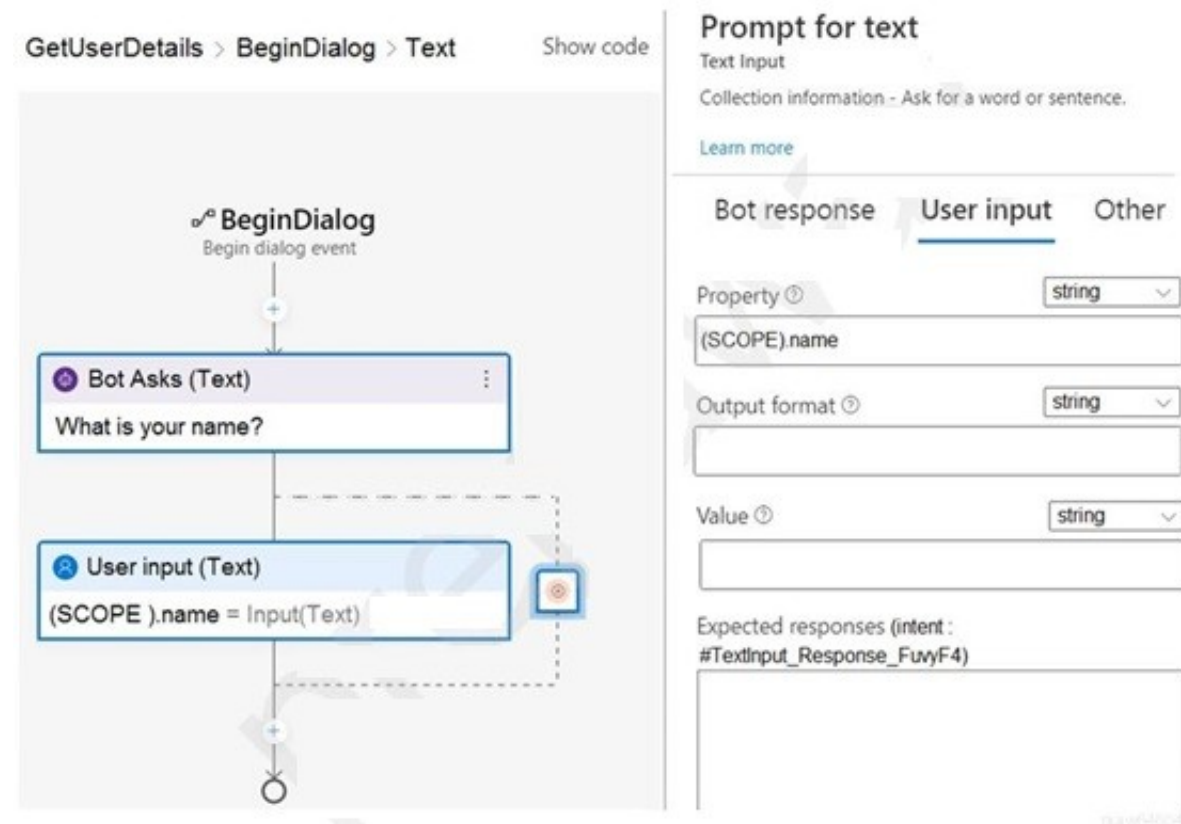
- A. a table projection
- B. a projection group
- C. an object projection
- D. a file projection

Answer: C

NEW QUESTION 190

- (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 193

- (Topic 3)

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

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

Answer: B

NEW QUESTION 194

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 198

- (Topic 3)
You are developing a method for an application that uses the Translator API. The method will receive the content of a webpage, and then translate the content into Greek (el). The result will also contain a transliteration that uses the Roman alphabet. You need to create the URI for the call to the Translator API. You have the following URI. <https://api.cognitive.microsofttranslator.com/translate?api-version=3.0> Which three additional query parameters should you include in the URI? Each correct answer presents part of the solution. (Choose three.)
NOTE: Each correct selection is worth one point.

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

F. toScript=Latn

Answer: ADF

Explanation:

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

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

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

Reference:

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

NEW QUESTION 202

- (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. tenantId
- C. appId
- D. objectId
- E. appSecret

Answer: CE

Explanation:

Reference:

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

NEW QUESTION 206

- (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 210

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 214

- (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 From 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 218

- (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 219

.....

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](#)