

# Microsoft

## Exam Questions AI-102

Designing and Implementing an Azure AI Solution



**NEW QUESTION 1**

- (Exam Topic 2)

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:**

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 2**

- (Exam Topic 2)

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

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

**Answer:** A

**Explanation:**

Reference:

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

**NEW QUESTION 3**

- (Exam Topic 2)

You train a Custom Vision model to identify a company's products by using the Retail domain. You plan to deploy the model as part of an app for Android phones. You need to prepare the model for deployment.

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 the model domain.
- Retrain the model.
- Test the model.
- Export the model.

**Answer Area**



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Graphical user interface, text, application Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/export-your-model>

**NEW QUESTION 4**

- (Exam Topic 2)

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 5**

- (Exam Topic 2)

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:**

Graphical user interface, text Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-security-api-keys>

**NEW QUESTION 6**

- (Exam Topic 2)

You are building a Language Understanding model for an e-commerce platform. You need to construct an entity to capture billing addresses.

Which entity type should you use for the billing address?

- A. machine learned
- B. Regex
- C. geographyV2
- D. Pattern.any
- E. list

**Answer:** B

**Explanation:**

A regular expression entity extracts an entity based on a regular expression pattern you provide. It ignores case and ignores cultural variant. Regular expression is best for structured text or a predefined sequence of alphanumeric values that are expected in a certain format. For example:

Entity	Regular expression	Example
Flight Number	flight [A-Z]{2} [0-9]{4}	flight AS 1234
Credit Card Number	[0-9]{16}	5478789865437632

Incorrect answers

C: The prebuilt geographyV2 entity detects places. Because this entity is already trained, you do not need to add example utterances containing GeographyV2 to the application intents. GeographyV2 entity is supported in English culture.

The geographical locations have subtypes:

Subtype	Purpose
poi	point of interest
city	name of city
countryRegion	name of country or region
continent	name of continent
state	name of state or province

D: Pattern.any is a variable-length placeholder used only in a pattern's template utterance to mark where the entity begins and ends.

E: A list entity represents a fixed, closed set of related words along with their synonyms. You can use list entities to recognize multiple synonyms or variations and extract a normalized output for them. Use the recommend option to see suggestions for new words based on the current list.

Reference:

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

**NEW QUESTION 7**

- (Exam Topic 2)

You are building a retail chatbot that will use a QnA Maker service.

You upload an internal support document to train the model. The document contains the following question: "What is your warranty period?"

Users report that the chatbot returns the default QnA Maker answer when they ask the following question: "How long is the warranty coverage?"

The chatbot returns the correct answer when the users ask the following question: "What is your warranty period?"

Both questions should return the same answer.

You need to increase the accuracy of the chatbot responses.

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

Actions	Answer Area
Add a new question and answer (QnA) pair.	
Retrain the model.	
Add additional questions to the document.	
Republish the model.	
Add alternative phrasing to the question and answer (QnA) pair.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1: Add alternative phrasing to the question and answer (QnA) pair.

Add alternate questions to an existing QnA pair to improve the likelihood of a match to a user query. Step 2: Retrain the model.

Periodically select Save and train after making edits to avoid losing changes. Step 3: Republish the model

Note: A knowledge base consists of question and answer (QnA) pairs. Each pair has one answer and a pair contains all the information associated with that answer.

Reference:  
<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/edit-knowledge-base>

**NEW QUESTION 8**

- (Exam Topic 2)

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

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

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

You need to add the new images to the classifier.

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

- A. Yes
- B. No

**Answer:** B

**Explanation:**

The model needs to be extended and retrained.

**NEW QUESTION 9**

- (Exam Topic 2)

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-
  
xxxxxxxxxxxx/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>

Reference:

<https://docs.microsoft.com/en-us/rest/api/deviceupdate/resourcemanager/accounts/create> <https://www.analyticsvidhya.com/blog/2020/12/microsoft-azure-cognitive-services-api-for-ai-development/>

**NEW QUESTION 10**

- (Exam Topic 2)

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 10**

- (Exam Topic 2)

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:**

Box 1: No

Box 2: Yes

The ComputerVision.analyzeImageInStreamAsync operation extracts a rich set of visual features based on the image content.

Box 3: No

Images will be read from a stream. Reference:

<https://docs.microsoft.com/en-us/java/api/com.microsoft.azure.cognitiveservices.vision.computervision.compute>

**NEW QUESTION 14**

- (Exam Topic 2)

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

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

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

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

You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet. Solution: You deploy service1 and a public endpoint, and you configure an IP firewall rule.

Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Reference:

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

**NEW QUESTION 18**

- (Exam Topic 2)

You plan to perform predictive maintenance.

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

You need to identify unusual values in each time series to help predict machinery failures.

Which Azure Cognitive Services service should you use?

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

**Answer: A**

**NEW QUESTION 19**

- (Exam Topic 2)

You need to upload speech samples to a Speech Studio project. How should you upload the samples?

- A. Combine the speech samples into a single audio file in the .wma format and upload the file.
- B. Upload a .zip file that contains a collection of audio files in the .wav format and a corresponding text transcript file.
- C. Upload individual audio files in the FLAC format and manually upload a corresponding transcript in Microsoft Word format.
- D. Upload individual audio files in the .wma format.

**Answer: B**

**Explanation:**

To upload your data, navigate to the Speech Studio . From the portal, click Upload data to launch the wizard and create your first dataset. You'll be asked to select a speech data type for your dataset, before allowing you to upload your data.

The default audio streaming format is WAV

Use this table to ensure that your audio files are formatted correctly for use with Custom Speech:

Property	Value
File format	RIFF (WAV)
Sample rate	8,000 Hz or 16,000 Hz
Channels	1 (mono)
Maximum length per audio	2 hours
Sample format	PCM, 16-bit
Archive format	.zip
Maximum archive size	2 GB

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/how-to-custom-speech-test-and-train>

**NEW QUESTION 20**

- (Exam Topic 2)

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:**

Step 1: Add prebuilt domain models as required.

Prebuilt models provide domains, intents, utterances, and entities. You can start your app with a prebuilt model or add a relevant model to your app later.

Note: Language Understanding (LUIS) provides prebuilt domains, which are pre-trained models of intents and entities that work together for domains or common categories of client applications.

The prebuilt domains are trained and ready to add to your LUIS app. The intents and entities of a prebuilt domain are fully customizable once you've added them to your app.

Step 2: Enable active learning

To enable active learning, you must log user queries. This is accomplished by calling the endpoint query with the log=true querystring parameter and value.

Step 3: Train and republish the Language Understanding model

The process of reviewing endpoint utterances for correct predictions is called Active learning. Active learning captures endpoint queries and selects user's endpoint utterances that it is unsure of. You review these utterances to select the intent and mark entities for these real-world utterances. Accept these changes into your example utterances then train and publish. LUIS then identifies utterances more accurately.

Reference:

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

**NEW QUESTION 22**

- (Exam Topic 2)

You are building a natural language model. You need to enable active learning.

What should you do?

- A. Add show-all-intents=true to the prediction endpoint query.
- B. Enable speech priming.
- C. Add log=true to the prediction endpoint query.
- D. Enable sentiment analysis.

**Answer:** C

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-how-to-review-endpoint-utterances#log-user>

**NEW QUESTION 23**

- (Exam Topic 2)

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/Comp>

**NEW QUESTION 25**

- (Exam Topic 2)

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.

**Answer Area**

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

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

Eula=accept \

Billing=

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

ApiKey=xxxxxxxxxxxxxxxxxxxxxx

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: [mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment](https://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  
 Reference:  
<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-install-co>

**NEW QUESTION 27**

- (Exam Topic 2)

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**

**Answer Area**

- 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.

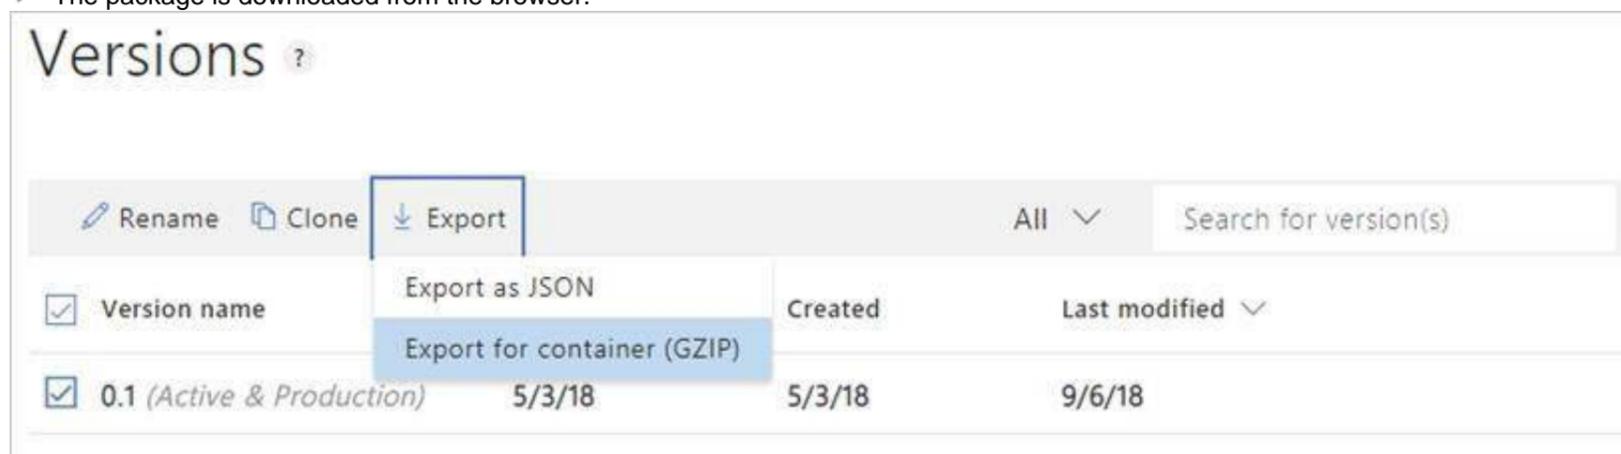
- A. Mastered
- B. Not Mastered

**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. Reference:

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

**NEW QUESTION 31**

- (Exam Topic 2)

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

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

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

## Logout
- forget me
- log out

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

## SelectNone
- none

@m1 DirectionalReference
@prebuilt personName
```

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

NOTE: Each correct selection is worth one point.

SelectItem is [answer choice]

- a domain
- an entity
- an intent
- an utterance

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

- a domain
- an entity
- an intent
- an utterance

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

Reference:

<https://github.com/solliancenet/tech-immersion-data-ai/blob/master/ai-exp1/README.md>

**NEW QUESTION 32**

- (Exam Topic 2)

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: Push the image to an Azure container registry.  
 To push an image to an Azure Container registry, you must first have an image.  
 Step 4: Distribute the docker run script  
 Use the docker run command to run the containers. Reference:  
<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-intro>

**NEW QUESTION 35**

- (Exam Topic 2)

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.  
 Reference:  
<https://docs.microsoft.com/en-us/azure/search/cognitive-search-output-field-mapping>

**NEW QUESTION 40**

- (Exam Topic 2)

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

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

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

Actions	Answer Area
Use the <code>ExportProject</code> endpoint on acvdev.	
Use the <code>GetProjects</code> endpoint on acvdev.	
Use the <code>ImportProject</code> endpoint on acvprod.	⬅
Use the <code>ExportIteration</code> endpoint on acvdev.	➡
Use the <code>GetIterations</code> endpoint on acvdev.	
Use the <code>UpdateProject</code> endpoint on acvprod.	⬆
	⬇

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Text Description automatically generated

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/copy-move-projects>

**NEW QUESTION 44**

- (Exam Topic 2)

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

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

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

NOTE: Each correct selection is worth one point.

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

**Answer:** CE

**Explanation:**

Reference:

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

**NEW QUESTION 45**

- (Exam Topic 2)

You are developing a new sales system that will process the video and text from a public-facing website. You plan to notify users that their data has been processed by the sales system.

Which responsible AI principle does this help meet?

- A. transparency
- B. fairness
- C. inclusiveness
- D. reliability and safety

**Answer:** D

**Explanation:**

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/responsible-ai>

### NEW QUESTION 50

- (Exam Topic 2)

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 52

- (Exam Topic 2)

A customer uses Azure Cognitive Search.

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

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

NOTE: Each correct selection is worth one point.

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

**Answer: ABE**

#### Explanation:

Reference:

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

### NEW QUESTION 55

- (Exam Topic 2)

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

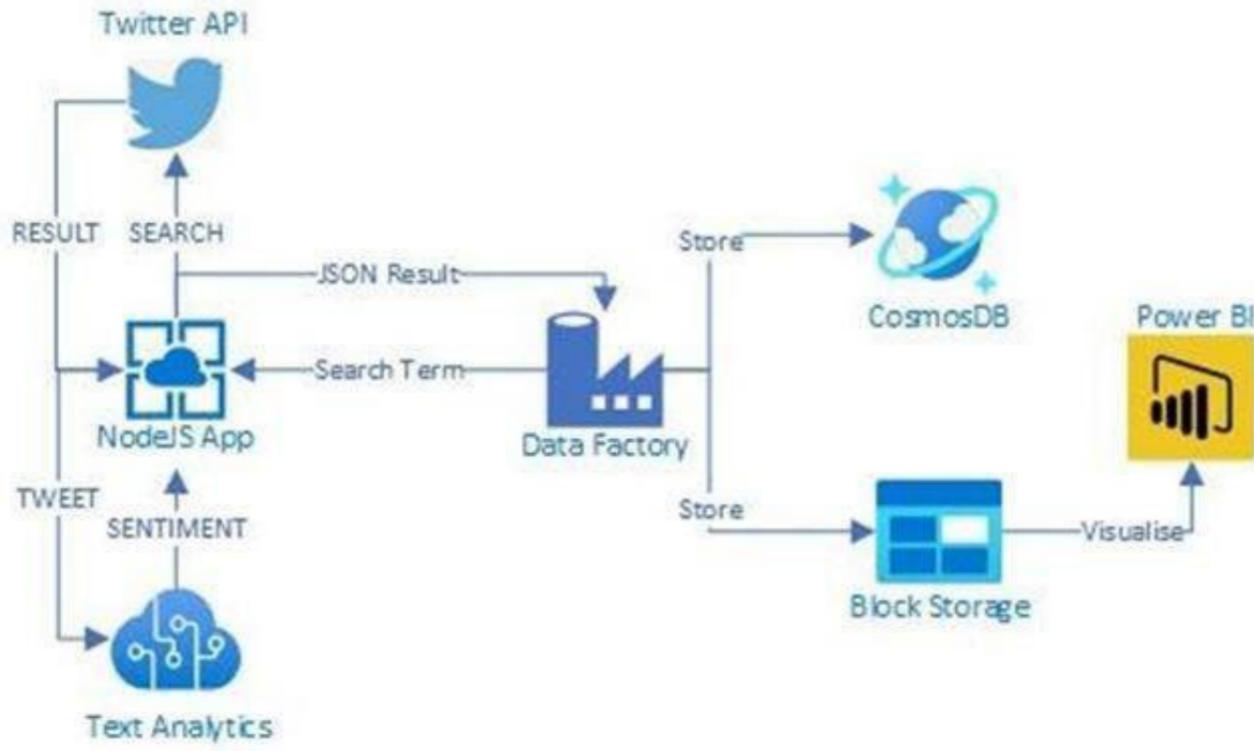
- A. Text Analytics
- B. Content Moderator
- C. Computer Vision
- D. Form Recognizer

**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 57**

.....

## **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](#)**