

AI-102^{Q&As}

Designing and Implementing a Microsoft Azure AI Solution

Pass Microsoft AI-102 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/ai-102.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

You have a chatbot that uses question answering in Azure Cognitive Service for Language.

Users report that the responses of the chatbot lack formality when answering spurious questions.

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

Solution: From Language Studio, you change the chitchat source to qna_chitchat_friendly.tsv, and then retrain and republish the model.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

<https://learn.microsoft.com/en-us/azure/cognitive-services/language-service/question-answering/how-to/chit-chat>

QUESTION 2

DRAG DROP

You are building a Language Understanding model for purchasing tickets.

You have the following utterance for an intent named PurchaseAndSendTickets.

Purchase [2 audit business] tickets to [Paris] [next Monday] and send tickets to [email@domain.com]

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

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

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

Select and Place:

Entity Types

Answer Area

Email

List

Regex

GeographyV2

Machine learned

Paris:

email@domain.com:

2 audit business:

Correct Answer:

Entity Types

Answer Area

List

Regex

Paris:

GeographyV2

email@domain.com:

Email

2 audit business:

Machine learned

Box 1: GeographyV2

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

Box 2: Email

Email prebuilt entity for a LUIS app: Email extraction includes the entire email address from an utterance. Because this entity is already trained, you do not need to add example utterances containing email to the application intents.

Box 3: Machine learned

The machine-learning entity is the preferred entity for building LUIS applications.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-geographyv2>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-email>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/reference-entity-machine-learned-entity>

QUESTION 3

You are developing the knowledgebase.

You use Azure Video Analyzer for Media (previously Video indexer) to obtain transcripts of webinars.

You need to ensure that the solution meets the knowledgebase requirements.

What should you do?

- A. Create a custom language model
- B. Configure audio indexing for videos only
- C. Enable multi-language detection for videos
- D. Build a custom Person model for webinar presenters

Correct Answer: A

<https://learn.microsoft.com/en-us/azure/azure-video-indexer/customize-language-model-overview> Azure Video Indexer supports automatic speech recognition through integration with the Microsoft Custom Speech Service. You can customize the Language model by uploading adaptation text, namely text from the domain whose vocabulary you'd like the engine to adapt to. Once you train your model, new words appearing in the adaptation text will be recognized, assuming default pronunciation, and the Language model will learn new probable sequences of words.

QUESTION 4

HOTSPOT

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.

Hot Area:

Answer Area

```
...
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;;

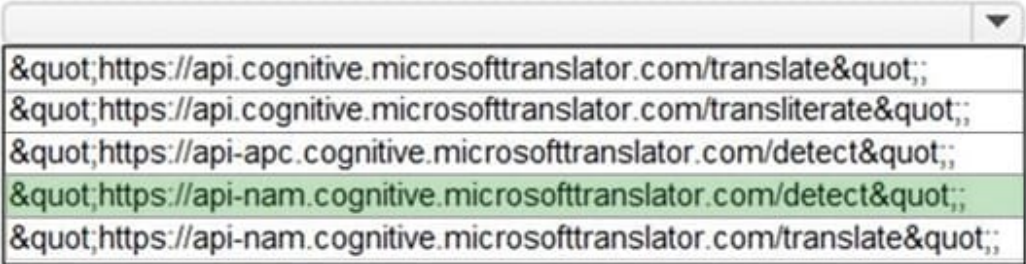
var apiKey = "FF956C68B83B21B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);


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

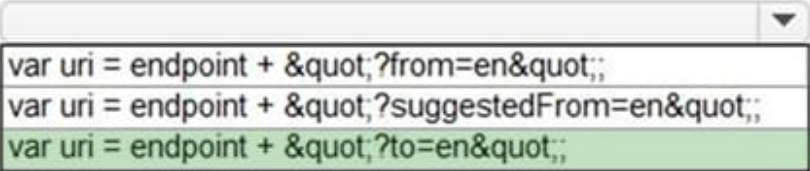
HttpResponseMessage response;
var content = new StringContent(body, Encoding.UTF8, "application/json");
var response = await client.PutAsync(uri, content);
...
```

Correct Answer:

Answer Area

```
...
var endpoint = ;

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



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

QUESTION 5

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

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

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

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

You need to reduce the likelihood that search query requests are throttled. Solution: You enable customer-managed key (CMK) encryption.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

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

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

Note: A simple fix to most throttling issues is to throw more resources at the search service (typically replicas for query-based throttling, or partitions for indexing-based throttling). However, increasing replicas or partitions adds cost, which is

why it is important to know the reason why throttling is occurring at all. Reference: <https://docs.microsoft.com/en-us/azure/search/search-performance-analysis>

QUESTION 6

You are building a solution that will detect anomalies in sensor data from the previous 24 hours. You need to ensure that the solution scans the entire dataset, at the same time, for anomalies. Which type of detection should you use?

A. batch

B. streaming

C. change points

Correct Answer: A

<https://learn.microsoft.com/en-us/azure/cognitive-services/anomaly-detector/overview#univariate-anomaly-detection>

Batch detection Use your time series to detect any anomalies that might exist throughout your data. This operation generates a model using your entire time series data, with each point analyzed with the same model.

QUESTION 7

HOTSPOT

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.

Hot Area:

Answer Area

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	

Correct Answer:

Answer Area

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	

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

QUESTION 8

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. Migrate the data in HR to Azure Blob storage.

- B. Migrate the data in HR to the on-premises SQL server.
- C. Export the data in Finance to Azure Data Lake Storage.
- D. Ingest the data in Logs into Azure Sentinel.

Correct Answer: C

<https://learn.microsoft.com/en-us/azure/search/search-indexer-overview>

QUESTION 9

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: AD

QUESTION 10

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.

Correct Answer: B

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>

QUESTION 11

HOTSPOT

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.

Hot Area:

Answer Area

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

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

Eula=accept \

Billing=

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

ApiKey=xxxxxxxxxxxxxxxxxxxxxx

Correct Answer:

Answer Area

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

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

Eula=accept \

Billing=

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

ApiKey=xxxxxxxxxxxxxxxxxxxxxx

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://cognitiveservices.azure.com>

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

{ENDPOINT_URI} is the endpoint for accessing the Text Analytics API: [https://cognitiveservices.a](https://cognitiveservices.azure.com) 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-containers?tabs=sentiment>

QUESTION 12

HOTSPOT

You are developing a text processing solution.

You develop the following method.

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

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

You call the method by using the following code.

```
GetKeyPhrases(textAnalyticsClient, "the cat sat on the mat");
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Box 1: Yes

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

Box 2: No

\\the\\ is not a key phrase.

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

"wonderful staff".

Box 3: No

Key phrase extraction does not have confidence levels.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-keyword-extraction>

QUESTION 13

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

Correct Answer: A

<https://learn.microsoft.com/en-us/azure/search/search-howto-large-index#run-indexers-in-parallel> If you partition your data, you can create multiple indexer-data-source combinations that pull from each data source and write to the same search index. Because each indexer is distinct, you can run them at the same time, populating a search index more quickly than if you ran them sequentially.

QUESTION 14

HOTSPOT

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

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

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

Hot Area:

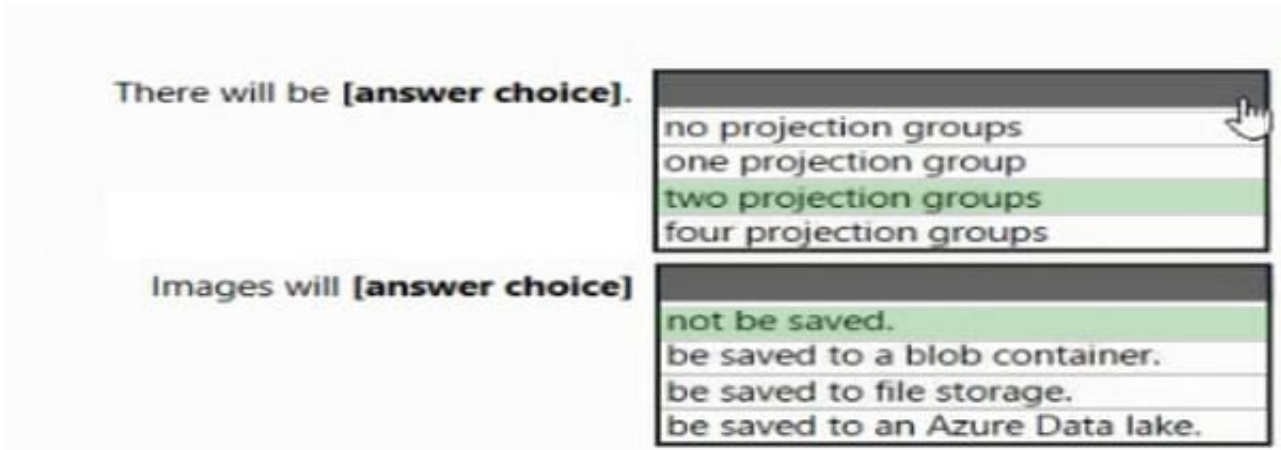
There will be [answer choice].

- no projection groups
- one projection group
- two projection groups
- four projection groups

Images will [answer choice]

- not be saved.
- be saved to a blob container.
- be saved to file storage.
- be saved to an Azure Data lake.

Correct Answer:



QUESTION 15

HOTSPOT

You are developing an application that includes language translation.

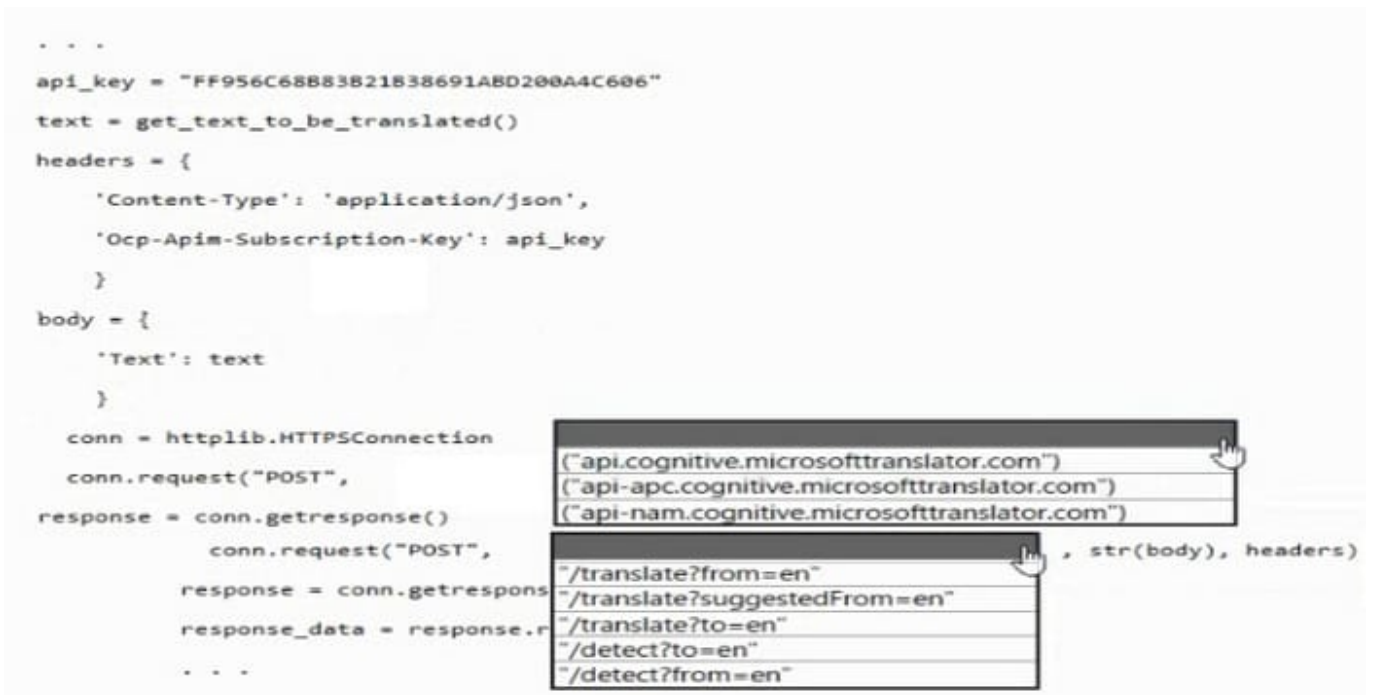
The application will translate text retrieved by using a function named `get_text_to_be_translated`. The text can be in one of many languages. The content of the text must remain within the Americas Azure geography.

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.

Hot Area:



Correct Answer:

```
...
api_key = "FF956C68883821838691ABD200A4C606"
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
...
, str(body), headers)
```



[Latest AI-102 Dumps](#)

[AI-102 Exam Questions](#)

[AI-102 Braindumps](#)