

# Microsoft

## Exam Questions AI-102

Designing and Implementing an Azure AI Solution



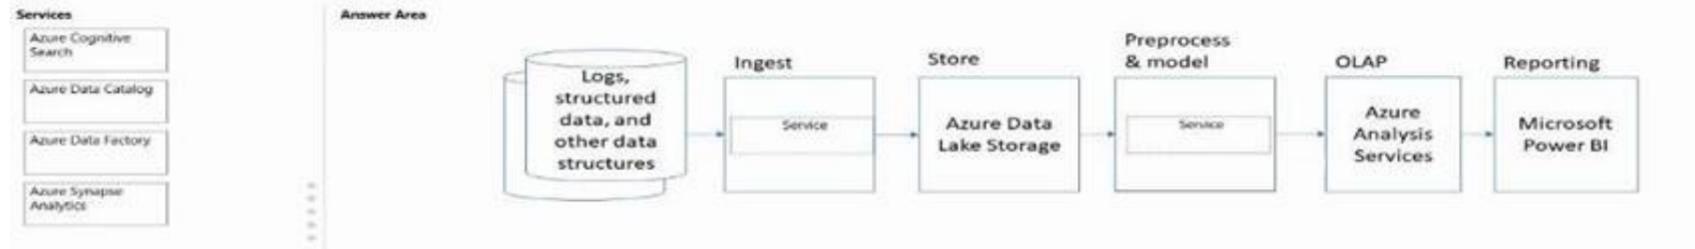
**NEW QUESTION 1**

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.

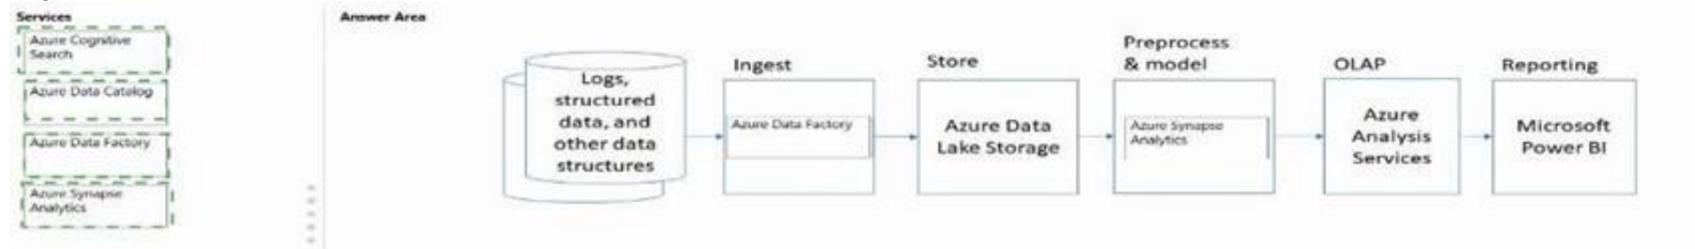
NOTE: Each correct match is worth one point.



- 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**

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 5**

HOTSPOT - (Topic 3)

You are building content for a video training solution.

You need to create narration to accompany the video content. The solution must use Custom Neural Voice.

What should you use to create a custom neural voice, and which service should you use to generate the narration? To answer, select the appropriate options in the answer area.

NOTE: Each correct answer is worth one point.

**Answer Area**

Custom neural voice:  ▼

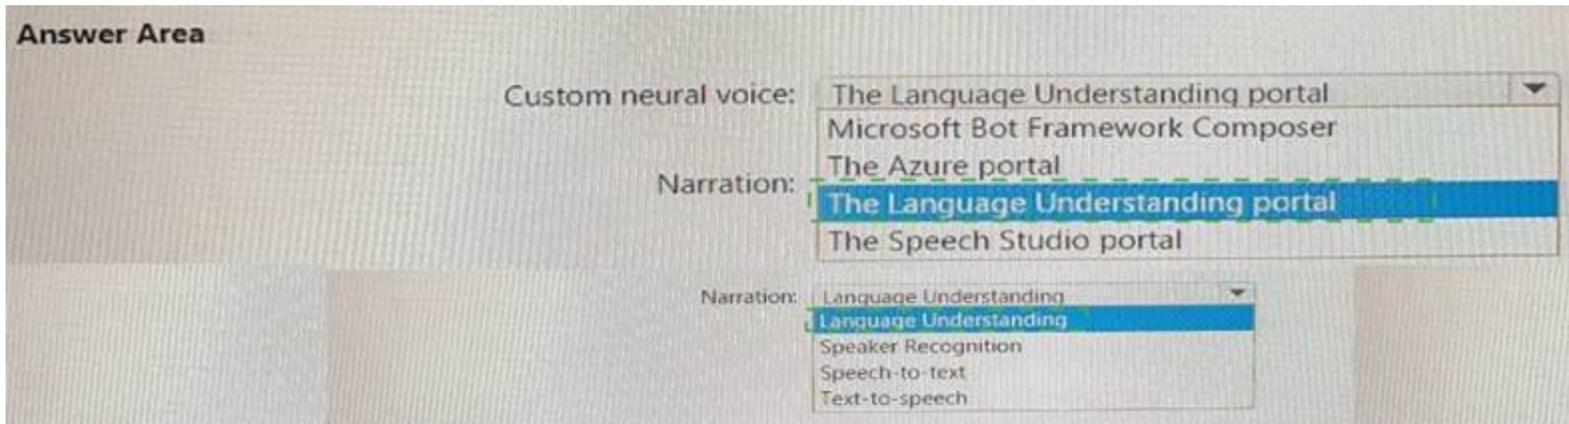
Narration:  ▼

Narration:  ▼

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 6**

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

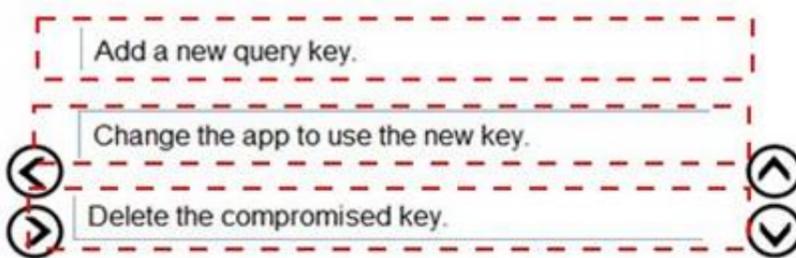
**Answer:** A

**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**



**NEW QUESTION 7**

- (Topic 3)

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

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

What should you add to the bot?

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

**Answer:** D

**NEW QUESTION 8**

- (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 9**

- (Topic 3)

Your company needs to implement a relational database in Azure. The solution must minimize ongoing maintenance. Which Azure service should you use?

- A. SQL Server on Azure Virtual Machines
- B. Azure SOL Database
- C. Azure HDInsight
- D. Azure Cosmos DB

**Answer: B**

**NEW QUESTION 10**

- (Topic 3)

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

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

**Answer: D**

**NEW QUESTION 10**

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 style="display: flex; justify-content: space-around; align-items: center;"> <div style="border: 1px solid #ccc; width: 80%; height: 100%;"></div> <div style="text-align: center;"> <span style="font-size: 2em;">⤴</span>  <span style="font-size: 2em;">⤵</span> </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 12**

- (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 15**

- (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 18**

DRAG DROP - (Topic 3)

You have a collection of Microsoft Word documents and PowerPoint presentations in German.  
 You need to create a solution to translate the file to French. The solution must meet the following requirements:

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

You create a blob container for German files and a blob container for French files. You upload the original files to the container for German files.  
 Which three actions should you perform in sequence actions from the list of actions to the answer area and arrange them in the correct order.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**NEW QUESTION 20**

- (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 25**

- (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 28**

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.

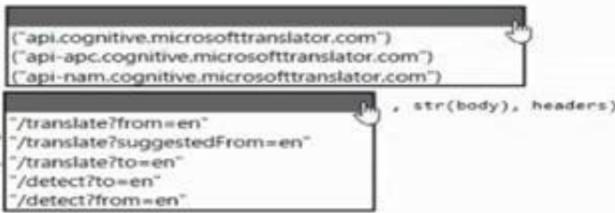
NOTE: Each correct selection is worth one point.

**Answer Area**

```

...
api_key = "FF956C6883821838691A8D200A4C606"
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
...

```



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

**Answer Area**

```

...
api_key = "FF956C6883821838691A8D200A4C606"
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 31**

- (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 33**

- (Topic 3)

You train a Conversational Language Understanding model to understand the natural language input of users.

You need to evaluate the accuracy of the model before deploying it.

What are two methods you can use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. From the language authoring REST endpoint, retrieve the model evaluation summary.
- B. From Language Studio, enable Active Learning, and then validate the utterances logged for review.
- C. From Language Studio, select Model performance.
- D. From the Azure portal, enable log collection in Log Analytics, and then analyze the logs.

**Answer:** AC

**NEW QUESTION 37**

- (Topic 3)

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:**

"When an AI application relies on personal data, such as a facial recognition system that takes images of people to recognize them; you should make it clear to the user how their data is used and retained, and who has access to it." from: <https://docs.microsoft.com/en-us/learn/paths/prepare-for-ai-engineering/>

**NEW QUESTION 39**

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 41**

- (Topic 3)

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

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

**Answer:** D

**Explanation:**

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

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

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

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

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

#### NEW QUESTION 45

- (Topic 3)

You plan to provision a QnA Maker service in a new resource group named RG1. In RG1, you create an App Service plan named AP1.

Which two Azure resources are automatically created in RG1 when you provision the QnA Maker service? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Language Understanding
- B. Azure SQL Database
- C. Azure Storage
- D. Azure Cognitive Search
- E. Azure App Service

**Answer:** DE

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/set-up-qnamaker-service-azure?tabs=v1#delete-azure-resources>

"When you create a QnAMaker resource, you host the data in your own Azure subscription. Azure Search is used to index your data." & "When you create a QnAMaker resource, you host the runtime in your own Azure subscription. App Service is the compute engine that runs the QnA Maker queries for you."

#### NEW QUESTION 50

- (Topic 3)

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

? Supports chit-chat, knowledge base, and multilingual models

? Performs sentiment analysis on user messages

? Selects the best language model automatically

What should you integrate into the chatbot?

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

**Answer:** C

#### Explanation:

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

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

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

Reference:

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

#### NEW QUESTION 55

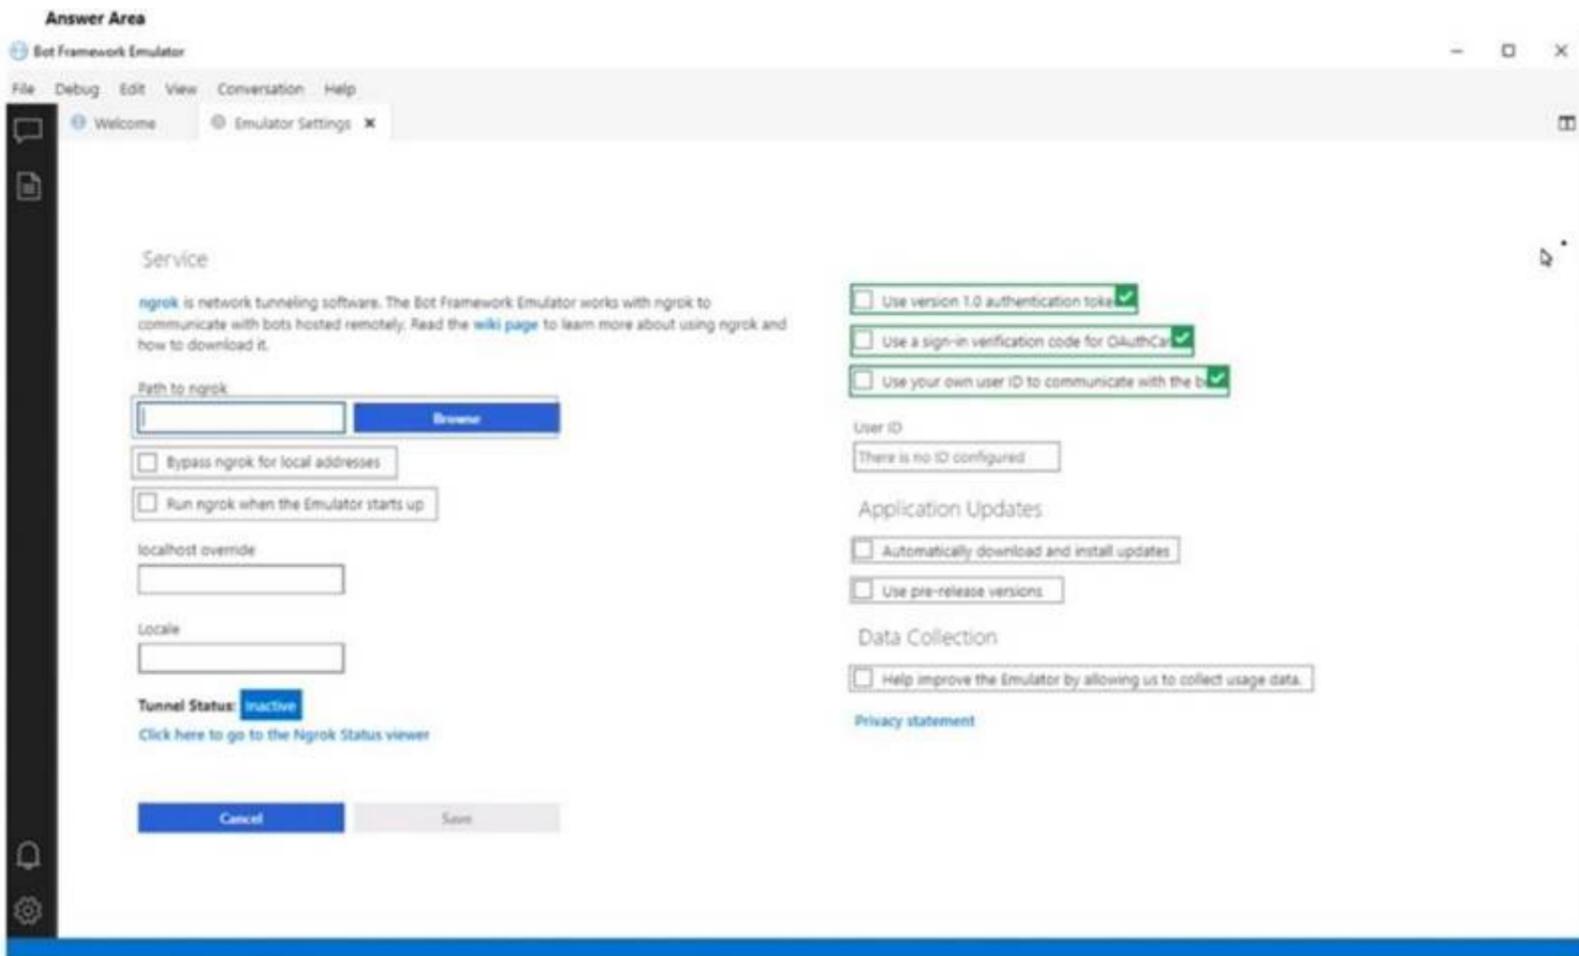
HOTSPOT - (Topic 3)

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.

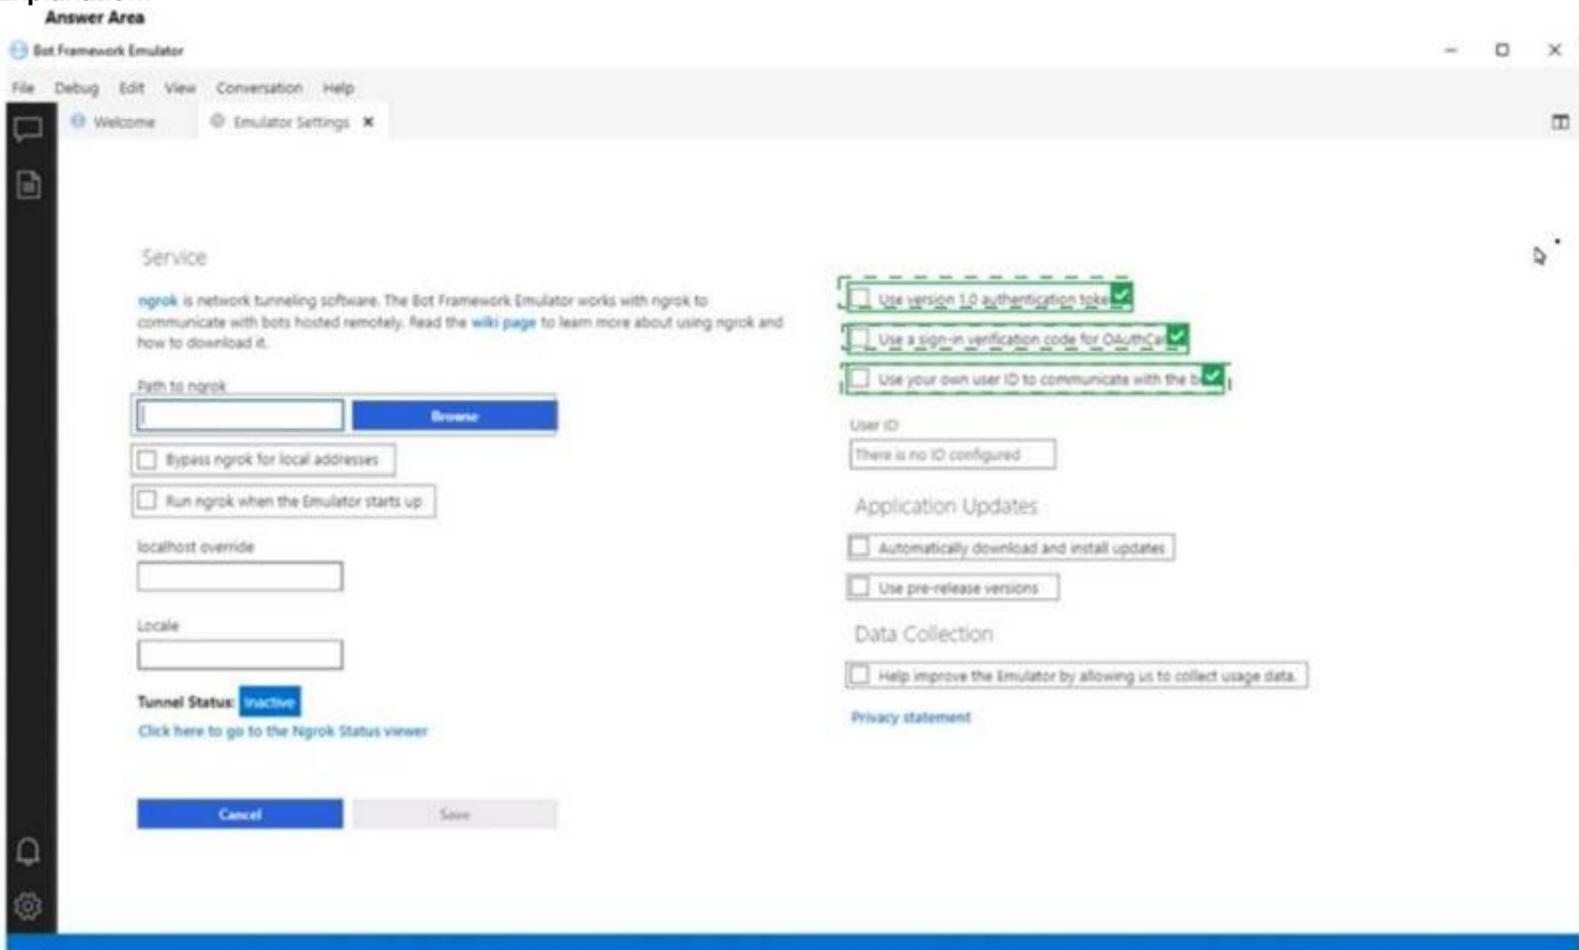
NOTE Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 56**

DRAG DROP - (Topic 3)

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

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

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

**Actions**

- Use the ExportProject endpoint on acvdev.
- Use the GetProjects endpoint on acvdev.
- Use the ImportProject endpoint on acvprod.
- Use the ExportIteration endpoint on acvdev.
- Use the GetIterations endpoint on acvdev.
- Use the UpdateProject endpoint on acvprod.

**Answer Area**

⏪

⏩

⏴

⏵

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Actions**

- Use the ExportProject endpoint on acvdev.
- Use the GetProjects endpoint on acvdev.
- Use the ImportProject endpoint on acvprod.
- Use the ExportIteration endpoint on acvdev.
- Use the GetIterations endpoint on acvdev.
- Use the UpdateProject endpoint on acvprod.

**Answer Area**

- Use the GetProjects endpoint on acvdev.
- Use the ExportProject endpoint on acvdev.
- Use the ImportProject endpoint on acvprod.

⏪

⏩

⏴

⏵

**NEW QUESTION 58**

- (Topic 3)  
 What should you use to build a Microsoft Power Bi paginated report?

- A. Power BI Report Builder
- B. Charculator
- C. Power BI Desktop
- D. the Power BI service

**Answer:** A

**NEW QUESTION 61**

- (Topic 3)  
 You use the Custom Vision service to build a classifier. After training is complete, you need to evaluate the classifier. Which two metrics are available for review? Each correct answer presents a complete solution. (Choose two.)  
 NOTE: Each correct selection is worth one point.

- A. recall
- B. F-score
- C. weighted accuracy
- D. precision
- E. area under the curve (AUC)

**Answer:** AD

**Explanation:**

Custom Vision provides three metrics regarding the performance of your model: precision, recall, and AP.  
 Reference:  
<https://www.tallan.com/blog/2020/05/19/azure-custom-vision/>

**NEW QUESTION 63**

- (Topic 3)  
 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 64

- (Topic 3)

You have receipts that are accessible from a URL.

You need to extract data from the receipts by using Form Recognizer and the SDK. The solution must use a prebuilt model.

Which client and method should you use?

- A. the FormRecognizerClient client and the StartRecognizeReceiptsFromUri method
- B. the FormTrainingClient client and the StartRecognizeContentFromUri method
- C. the FormRecognizerClient client and the StartRecognizeReceiptsFromUri method
- D. the FormTrainingClient client and the StartRecognizeReceiptsFromUri method

**Answer:** D

#### Explanation:

To analyze receipts from a URL, use the StartRecognizeReceiptsFromUri method Example code:

```
private static async Task AnalyzeReceipt( FormRecognizerClient recognizerClient, string receiptUri)
```

```
{
```

```
    RecognizedFormCollection receipts = await recognizerClient.StartRecognizeReceiptsFromUri(new Uri(receiptUri)).WaitForCompletionAsync();
```

```
Reference:
```

```
https://docs.microsoft.com/en-us/azure/applied-ai-services/form-recognizer/quickstarts/client-library
```

#### NEW QUESTION 66

- (Topic 3)

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

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

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

**Answer:** B

#### Explanation:

Reference:

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

ML Entity with Structure

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

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

Street Name: Main Street State: NY

Zip Code: 98052 Country: USA

#### NEW QUESTION 67

- (Topic 3)

Your company wants to reduce how long it takes for employees to log receipts in expense reports. All the receipts are in English.

You need to extract top-level information from the receipts, such as the vendor and the transaction total. The solution must minimize development effort.

Which Azure Cognitive Services service should you use?

- A. Custom Vision
- B. Personalizer
- C. Form Recognizer
- D. Computer Vision

**Answer:** B

#### Explanation:

Azure Form Recognizer is a cognitive service that lets you build automated data processing software using machine learning technology. Identify and extract text, key/value pairs, selection marks, tables, and structure from your documents—the service outputs structured data that includes the relationships in the original file, bounding boxes, confidence and more.

Form Recognizer is composed of custom document processing models, prebuilt models for invoices, receipts, IDs and business cards, and the layout model.

Reference:

```
https://docs.microsoft.com/en-us/azure/cognitive-services/form-recognizer
```

#### NEW QUESTION 72

- (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 77**

- (Topic 3)

You have an Azure Cognitive Search solution and a collection of blog posts that include a category field. You need to index the posts. The solution must meet the following requirements:

- Include the category field in the search results.
- Ensure that users can search for words in the category field.
- Ensure that users can perform drill down filtering based on category. Which index attributes should you configure for the category field?

- A. searchable, facetable, and retrievable
- B. retrievable, filterable, and sortable
- C. retrievable, facetable, and key
- D. searchable, sortable, and retrievable

**Answer: B**

**NEW QUESTION 81**

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	Answer Area
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.	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Actions	Answer Area
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.	Provision an on-premises Kubernetes cluster that is isolated from the internet.
Pull an image from the Microsoft Container Registry (MCR).	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.	Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

**NEW QUESTION 86**

HOTSPOT - (Topic 3)

You are building an app that will enable users to upload images. The solution must meet the following requirements:

- Automatically suggest alt text for the images.
- Detect inappropriate images and block them.
- Minimize development effort.

You need to recommend a computer vision endpoint for each requirement.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Generate alt text:

- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>

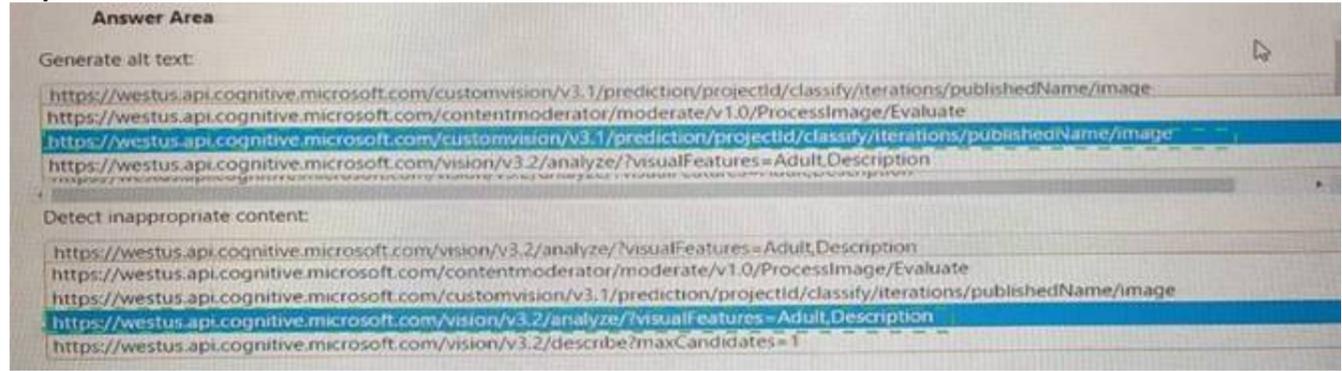
Detect inappropriate content:

- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1>

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**



**NEW QUESTION 87**

DRAG DROP - (Topic 3)

You develop an app in O named App1 that performs speech-to-speech translation. You need to configure App1 to translate English to German. How should you complete the speechTransiationConf ig object? 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.

<p><b>Values</b></p> <ul style="list-style-type: none"> <li>addTargetLanguage</li> <li>speechSynthesisLanguage</li> <li>speechRecognitionLanguage</li> <li>voiceName</li> </ul>	<p><b>Answer Area</b></p> <pre>var translationConfig = SpeechTranslationConfig.FromSubscription(SPEECH_SUBSCRIPTION_KEY, SPEECH_SERVICE_REGION); translationConfig. [ ] = "en-US"; translationConfig. [ ] ("de");</pre>
---	---

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

<p><b>Values</b></p> <ul style="list-style-type: none"> <li>addTargetLanguage</li> <li>speechSynthesisLanguage</li> <li>speechRecognitionLanguage</li> <li>voiceName</li> </ul>	<p><b>Answer Area</b></p> <pre>var translationConfig = SpeechTranslationConfig.FromSubscription(SPEECH_SUBSCRIPTION_KEY, SPEECH_SERVICE_REGION); translationConfig. [speechRecognitionLanguage] = "en-US"; translationConfig. [addTargetLanguage] ("de");</pre>
---	---

**NEW QUESTION 91**

- (Topic 3)

You are building a chatbot that will use question answering in Azure Cognitive Service for Language. You have a PDF named Docl.pdf that contains a product catalogue and a price list You upload Docl.pdf and train the model. During testing, users report that the chatbot responds correctly to the following question: What is the price of < product>? The chatbot fails to respond to the following question: How much does <product\* cost? You need to ensure that the chatbot responds correctly to both questions. Solution: From Language Studio, you create an entity for cost, and then retrain and republish the model. Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**NEW QUESTION 94**

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.

<p><b>Workload Types</b></p> <ul style="list-style-type: none"> <li>Batch</li> <li>Streaming</li> </ul>	<p><b>Answer Area</b></p> <table border="0"> <tr> <td style="border: 1px solid gray; padding: 2px;">Workload type</td> <td>Data for a product catalog will be loaded every 12 hours to a data warehouse.</td> </tr> <tr> <td style="border: 1px solid gray; padding: 2px;">Workload type</td> <td>Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time.</td> </tr> <tr> <td style="border: 1px solid gray; padding: 2px;">Workload type</td> <td>Updates to inventory data will be loaded to a data warehouse every 1 million transactions.</td> </tr> </table>	Workload type	Data for a product catalog will be loaded every 12 hours to a data warehouse.	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.
Workload type	Data for a product catalog will be loaded every 12 hours to a data warehouse.						
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 95**

- (Topic 3)

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

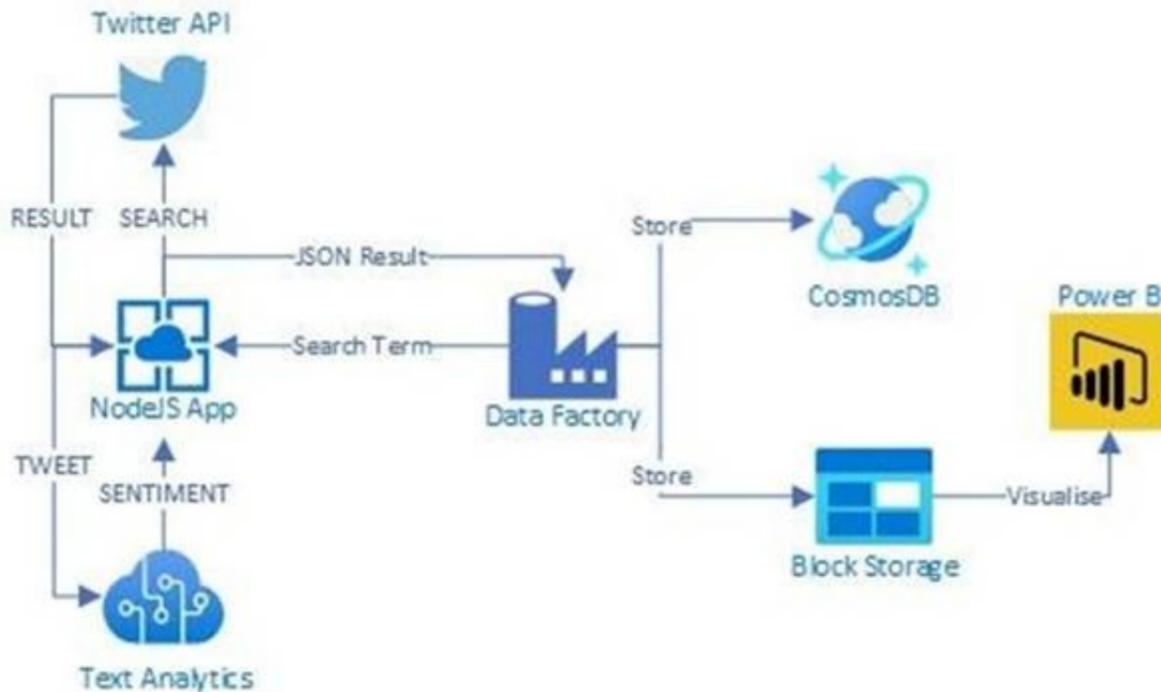
- A. Text Analytics
- B. Content Moderator
- C. Computer Vision
- D. Form Recognizer

**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 100**

- (Topic 3)

You are building a bot by using Microsoft Bot Framework.

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

What should you do?

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

**Answer:** B

**NEW QUESTION 104**

DRAG DROP - (Topic 3)

You have an Azure subscription that contains an Azure OpenA1 resource named AH. You plan to build an app named App1 that will write press releases by using AM.

You need to deploy an Azure OpenA1 model for Appl. The solution must minimize development effort.

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

Actions	Answer Area
Create a deployment that uses the text-embedding-ada-002 model.	
Apply the Default system message template.	
Create a deployment that uses the GPT-35 Turbo model.	
Apply the Marketing Writing Assistant system message template.	
Deploy the solution to a new web app.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Actions	Answer Area
Create a deployment that uses the text-embedding-ada-002 model.	Create a deployment that uses the GPT-35 Turbo model.
Apply the Default system message template.	Apply the Marketing Writing Assistant system message template.
Create a deployment that uses the GPT-35 Turbo model.	Deploy the solution to a new web app.
Apply the Marketing Writing Assistant system message template.	
Deploy the solution to a new web app.	

**NEW QUESTION 105**

- (Topic 3)

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

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

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

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

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

Solution: You deploy service1 and a public endpoint, and you configure a network security group (NSG) for vnet1.

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

<https://docs.microsoft.com/en-us/azure/virtual-network/network-security-groups-overview#network-security-groups>

<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-overview>

**NEW QUESTION 106**

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

**Answer Area**

When provisioning an Azure Cosmos DB \_\_\_\_\_, you need to specify which type of API you will use.

account  
 container  
 database  
 item

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

When provisioning an Azure Cosmos DB \_\_\_\_\_, you need to specify which type of API you will use.

account  
 container  
 database  
 item

**NEW QUESTION 110**

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.

Objects Training Images **Performance** Predictions Train Quick Test

Iterations Publish Prediction URL Delete Export

Probability Threshold: 50% Overlap Threshold: 30%

Iteration 1

**Answer Area**

The percentage of false positives is [answer choice].

The value for the number of true positives divided by the total number of true positives and false negatives is [answer choice]%

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

The percentage of false positives is [answer choice].

The value for the number of true positives divided by the total number of true positives and false negatives is [answer choice]%

**NEW QUESTION 114**

HOTSPOT - (Topic 3)

You have a library that contains thousands of images.

You need to tag the images as photographs, drawings, or clipart.

Which service endpoint and response property should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Service endpoint:

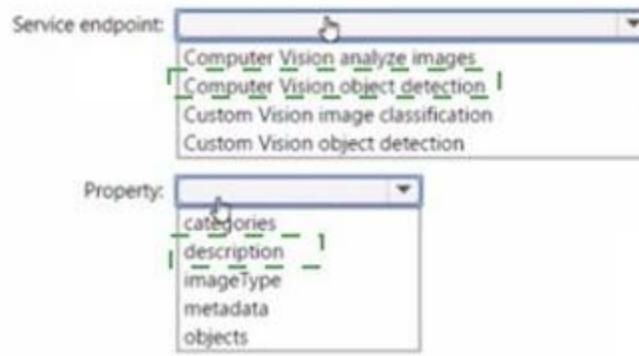
Property:

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**



**NEW QUESTION 116**

DRAG DROP - (Topic 3)

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

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

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

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

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 119**

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.

Statements	Yes	No
The code will create and maintain the <code>UserProfile</code> object in the underlying storage layer.	<input type="radio"/>	<input type="radio"/>
The code will create and maintain the <code>ConversationData</code> object in the underlying storage layer.	<input type="radio"/>	<input type="radio"/>
The <code>UserProfile</code> and <code>ConversationData</code> 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 120**

HOTSPOT - (Topic 3)

You develop an application that uses the Face API. You need to add multiple images to a person group.

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**

```
Parallel.For(0, PersonCount, async i =>
{
    Guid personId = persons[i].PersonId;
    string personImageDir = $"/path/to/person/{i}/images";
    foreach (string imagePath in Directory.GetFiles(personImageDir, "*.jpg"))
    {
        using (
            

|        |
|--------|
| File   |
| Stream |
| Uri    |
| Url    |


            t = File.OpenRead(imagePath))
        {
            await faceClient.PersonGroupPerson.
            

|                        |
|------------------------|
| AddFaceFromStreamAsync |
| AddFaceFromUrlAsync    |
| CreateAsync            |
| GetAsync               |


            (personGroupId, personId, t);
        }
    }
});
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Stream

The File.OpenRead(String) method opens an existing file for reading.

Example: Open the stream and read it back. using (FileStream fs = File.OpenRead(path))

Box 2: AddFaceFromStreamAsync

Step 5 on <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-add-faces>

**NEW QUESTION 125**

- (Topic 3)

You build a language model by using Conversational Language Understanding. 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 m Ukraine.

You need to implement the phrase list in Conversational Language Understanding. Solution: You create a new utterance for each phrase in the FindContact intent.

- A. Yes
- B. No

**Answer:** B

**NEW QUESTION 128**

- (Topic 3)

Which scenario is an example of a streaming workload?

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

**Answer:** D

**NEW QUESTION 130**

- (Topic 3)

You are examining the Text Analytics output of an application.

The text analyzed is: "Our tour guide took us up the Space Needle during our trip to Seattle last week."

The response contains the data shown in the following table.

Text	Category	ConfidenceScore
Tour guide	PersonType	0.45
Space Needle	Location	0.38
Trip	Event	0.78
Seattle	Location	0.78
Last week	DateTime	0.80

Which Text Analytics API is used to analyze the text?

- A. Sentiment Analysis
- B. Named Entity Recognition
- C. Entity Linking
- D. Key Phrase Extraction

**Answer:** B

**NEW QUESTION 132**

- (Topic 3)

You build a chatbot that uses the Azure OpenAI GPT 3.5 model.

You need to improve the quality of the responses from the chatbot. The solution must minimize development effort.

What are two ways to achieve the goal? Each correct answer presents a complete solution. NOTE: Each correct answer is worth one point.

- A. Fine-tune the model.
- B. Provide grounding content.
- C. Add sample request/response pairs.
- D. Retrain the language model by using your own data.
- E. Train a custom large language model (LLM).

**Answer:** BC

**NEW QUESTION 133**

- (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 137**

HOTSPOT - (Topic 3)

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

- ? Use a single key and endpoint to access multiple services.
- ? Consolidate billing for future services that you might use.
- ? Support the use of Computer Vision in the future.

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

NOTE: Each correct selection is worth one point.

**Answer Area**

PATCH
POST
PUT

  
xxxx-xxxx-
  
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/providers/Microsoft.DeviceUpdate/accounts/contoso?api-version=2020-03-01-preview

**NEW QUESTION 139**

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

**Answer Area**

A relational database is appropriate for scenarios that involve a high volume of

<ul style="list-style-type: none"> <li>changes to relationships between entities.</li> <li>geographically distributed writes.</li> <li>transactional writes.</li> <li>writes that have varying data structures.</li> </ul>
--

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

A relational database is appropriate for scenarios that involve a high volume of

<ul style="list-style-type: none"> <li>changes to relationships between entities.</li> <li>geographically distributed writes.</li> <li>transactional writes.</li> <li>writes that have varying data structures.</li> </ul>
--

**NEW QUESTION 142**

DRAG DROP - (Topic 3)

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

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

NOTE: Each correct match is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 144**

HOTSPOT - (Topic 3)

You are building an app that will process incoming email and direct email messages to either French or English language support teams. Which Azure Cognitive Services API should you use? To answer, select the appropriate options in the answer area.

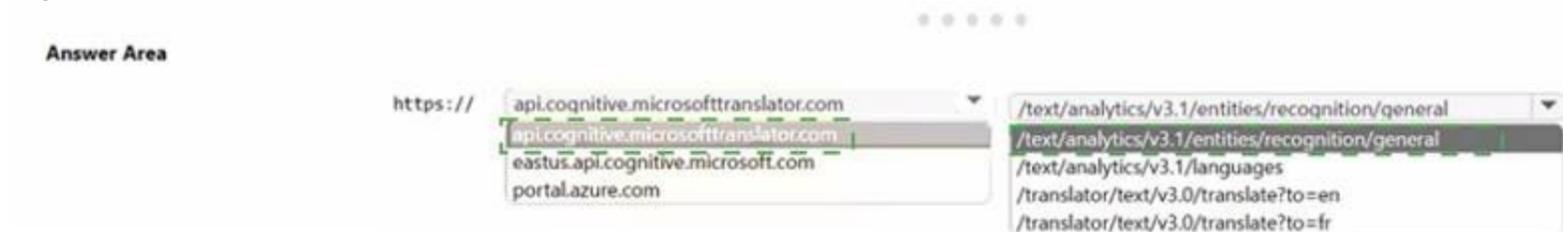
NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



**NEW QUESTION 147**

- (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 150**

HOTSPOT - (Topic 3)

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

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

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

**NEW QUESTION 153**

- (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 157**

- (Topic 3)

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

- A. Revoke

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

**Answer: B**

**NEW QUESTION 160**

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 162**

- (Topic 3)

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

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

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

questions that are outside the scope of the knowledge base.

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

Solution: From Language Studio, you modify the question and answer pairs for the custom intents, and then retrain and republish the model.

Does this meet the goal?

A. Yes

B. No

**Answer: B**

**NEW QUESTION 167**

HOTSPOT - (Topic 2)

You build a QnA Maker resource to meet the chatbot requirements.

Which RBAC role should you assign to each group? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Management-Accountants	<div style="border: 1px solid black; padding: 5px;"> <div style="border-bottom: 1px solid black; display: flex; justify-content: flex-end; align-items: center;">▼</div> <ul style="list-style-type: none"> <li>Owner</li> <li>Contributor</li> <li>Cognitive Services User</li> <li>Cognitive Services QnA Maker Read</li> <li>Cognitive Services QnA Maker Editor</li> </ul> </div>
Consultant-Accountants	<div style="border: 1px solid black; padding: 5px;"> <div style="border-bottom: 1px solid black; display: flex; justify-content: flex-end; align-items: center;">▼</div> <ul style="list-style-type: none"> <li>Owner</li> <li>Contributor</li> <li>Cognitive Services User</li> <li>Cognitive Services QnA Maker Read</li> <li>Cognitive Services QnA Maker Editor</li> </ul> </div>
Agent-CustomerServices	<div style="border: 1px solid black; padding: 5px;"> <div style="border-bottom: 1px solid black; display: flex; justify-content: flex-end; align-items: center;">▼</div> <ul style="list-style-type: none"> <li>Owner</li> <li>Contributor</li> <li>Cognitive Services User</li> <li>Cognitive Services QnA Maker Read</li> <li>Cognitive Services QnA Maker Editor</li> </ul> </div>

A. Mastered

B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Cognitive Service User

Ensure that the members of a group named Management-Accountants can approve the FAQs.

Approve=publish.

Cognitive Service User (read/write/publish): API permissions: All access to Cognitive Services resource except for ability to:

- \* 1. Add new members to roles.
- \* 2. Create new resources.

Box 2: Cognitive Services QnA Maker Editor

Ensure that the members of a group named Consultant-Accountants can create and amend the FAQs.

QnA Maker Editor: API permissions:

- \* 1. Create KB API
- \* 2. Update KB API
- \* 3. Replace KB API
- \* 4. Replace Alterations
- \* 5. "Train API" [in new service model v5]

Box 3: Cognitive Services QnA Maker Read

Ensure that the members of a group named the Agent-CustomerServices can browse the FAQs.

QnA Maker Read: API Permissions:

- \* 1. Download KB API
- \* 2. List KBs for user API
- \* 3. Get Knowledge base details
- \* 4. Download Alterations Generate Answer

**NEW QUESTION 171**

- (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 173**

DRAG DROP - (Topic 1)

You are planning the product creation project.

You need to recommend a process for analyzing videos.

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

Actions	Answer Area
Index the video by using the Video Indexer API.	
Upload the video to blob storage.	
Analyze the video by using the Computer Vision API.	
Extract the transcript from Microsoft Stream.	
Send the transcript to the Language Understanding API as an utterance.	
Extract the transcript from the Video Indexer API.	
Translate the transcript by using the Translator API.	
Upload the video to file storage.	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Scenario: All videos must have transcripts that are associated to the video and included in product descriptions.

Product descriptions, transcripts, and all text must be available in English, Spanish, and Portuguese.

Step 1: Upload the video to blob storage

Given a video or audio file, the file is first dropped into a Blob Storage. T

Step 2: Index the video by using the Video Indexer API.

When a video is indexed, Video Indexer produces the JSON content that contains details of the specified video insights. The insights include: transcripts, OCRs, faces, topics, blocks, etc.

Step 3: Extract the transcript from the Video Indexer API. Step 4: Translate the transcript by using the Translator API.

**NEW QUESTION 175**

- (Topic 1)

You are developing the smart e-commerce project.

You need to implement autocompletion as part of the Cognitive Search solution. Which three actions should you perform? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

- A. Make API queries to the autocomplete endpoint and include suggesterName in the body.
- B. Add a suggester that has the three product name fields as source fields.
- C. Make API queries to the search endpoint and include the product name fields in the searchFields query parameter.
- D. Add a suggester for each of the three product name fields.
- E. Set the searchAnalyzer property for the three product name variants.
- F. Set the analyzer property for the three product name variants.

**Answer:** ABF

**Explanation:**

Scenario: Support autocompletion and autosuggestion based on all product name variants. A: Call a suggester-enabled query, in the form of a Suggestion request or Autocomplete

request, using an API. API usage is illustrated in the following call to the Autocomplete REST API.

POST /indexes/myxboxgames/docs/autocomplete?search&api-version=2020-06-30

```
{
  "search": "minecraf", "suggesterName": "sg"
}
```

B: In Azure Cognitive Search, typeahead or "search-as-you-type" is enabled through a suggester. A suggester provides a list of fields that undergo additional tokenization, generating prefix sequences to support matches on partial terms. For example, a suggester that includes a City field with a value for "Seattle" will have prefix combinations of "sea", "seat", "seatt", and "seattl" to support typeahead.

F: Use the default standard Lucene analyzer ("analyzer": null) or a language analyzer (for example, "analyzer": "en.Microsoft") on the field.

Reference:

<https://docs.microsoft.com/en-us/azure/search/index-add-suggesters>

**NEW QUESTION 179**

HOTSPOT - (Topic 1)

You are developing the shopping on-the-go project.

You are configuring access to the QnA Maker resources.

Which role should you assign to AllUsers and LeadershipTeam? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

AllUsers:	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span>▼</span> </div> <ul style="list-style-type: none"> <li>Cognitive Service User</li> <li>Contributor</li> <li>Owner</li> <li>QnA Maker Editor</li> <li>QnA Maker Read</li> </ul> </div>
LeadershipTeam:	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> <span>▼</span> </div> <ul style="list-style-type: none"> <li>Cognitive Service User</li> <li>Contributor</li> <li>Owner</li> <li>QnA Maker Editor</li> <li>QnA Maker Read</li> </ul> </div>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: QnA Maker Editor

Scenario: Provide all employees with the ability to edit Q&As.

The QnA Maker Editor (read/write) has the following permissions: Create KB API Update KB API Replace KB API Replace Alterations "Train API" [in new service model v5]

Box 2: Contributor

Scenario: Only senior managers must be able to publish updates. Contributor permission: All except ability to add new members to roles

**NEW QUESTION 182**

DRAG DROP - (Topic 1)

You are developing the smart e-commerce project.

You need to design the skillset to include the contents of PDFs in searches.

How should you complete the skillset design diagram? To answer, drag the appropriate services to the correct stages. Each service 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.

Services	Answer Area			
	Source	Cracking	Preparation	Destination
Azure Blob storage				
Custom Vision API				
Azure Files				
Language Understanding API				
Translator API				
Computer Vision API				
Azure Cosmos DB				

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Azure Blob storage

At the start of the pipeline, you have unstructured text or non-text content (such as images, scanned documents, or JPEG files). Data must exist in an Azure data storage service that can be accessed by an indexer.

Box 2: Computer Vision API

Scenario: Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products.

The Computer Vision Read API is Azure's latest OCR technology (learn what's new) that extracts printed text (in several languages), handwritten text (English only), digits, and currency symbols from images and multi-page PDF documents.

Box 3: Translator API

Scenario: Product descriptions, transcripts, and all text must be available in English, Spanish, and Portuguese.

Box 4: Azure Files

Scenario: Store all raw insight data that was generated, so the data can be processed later.

**NEW QUESTION 184**

.....

## **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](#)