



## Microsoft

### Exam Questions AI-102

Designing and Implementing an Azure AI Solution

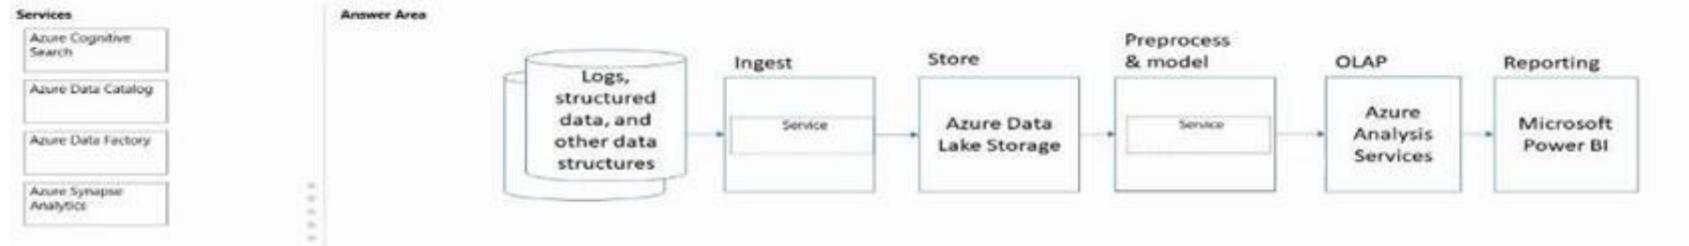
**NEW QUESTION 1**

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.

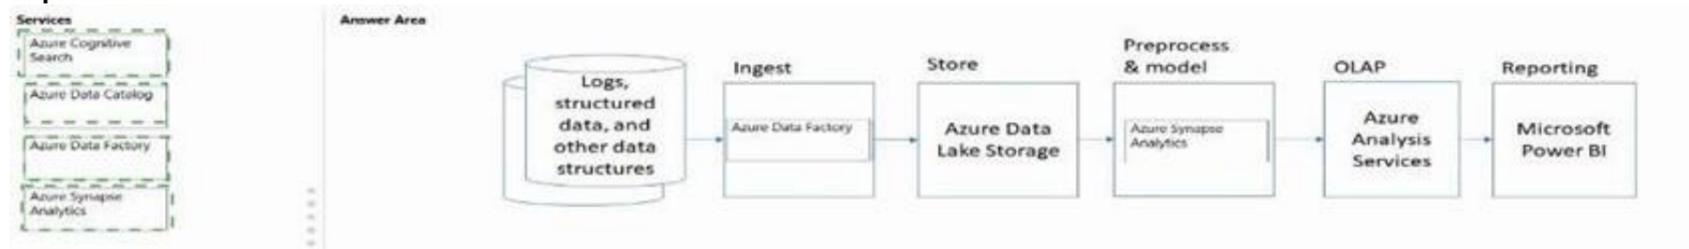
NOTE: Each correct match is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



**NEW QUESTION 2**

- (Topic 3)

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

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

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

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

Answer: D

**NEW QUESTION 3**

HOTSPOT - (Topic 3)

You build a chatbot by using Azure OpenAI Studio.

You need to ensure that the responses are more deterministic and less creative. Which two parameters should you configure? To answer, select the appropriate parameters in the answer area.

NOTE: Each correct answer is worth one point.

Answer Area

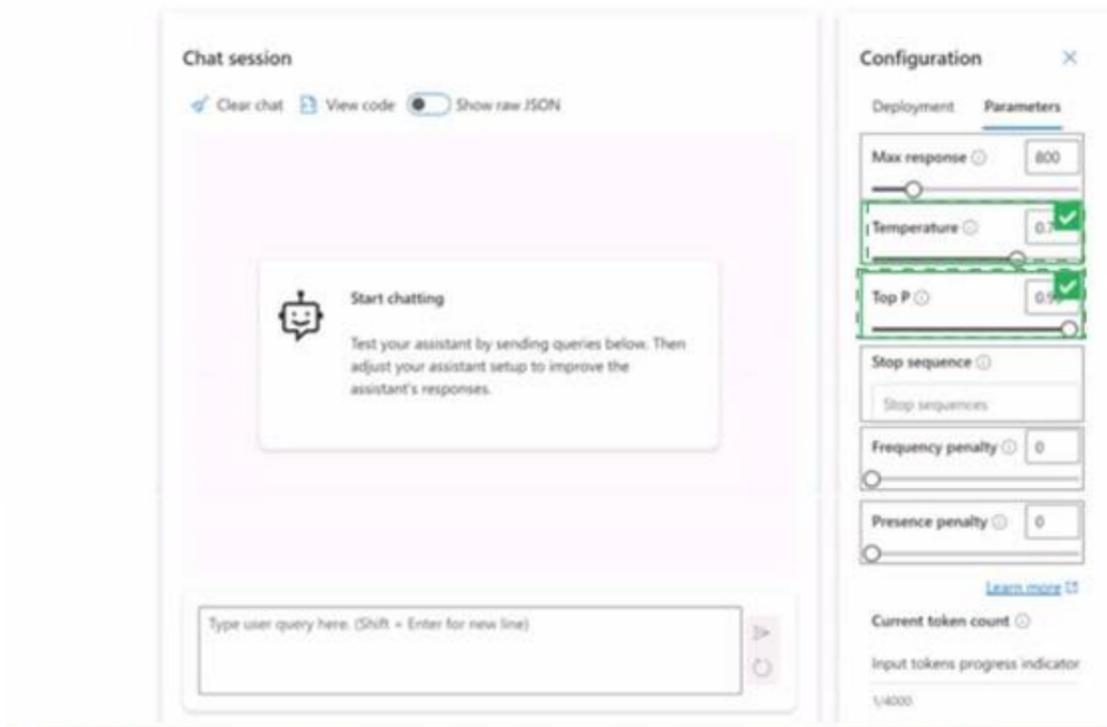
- A. Mastered

B. Not Mastered

Answer: A

Explanation:

Answer Area



**NEW QUESTION 4**

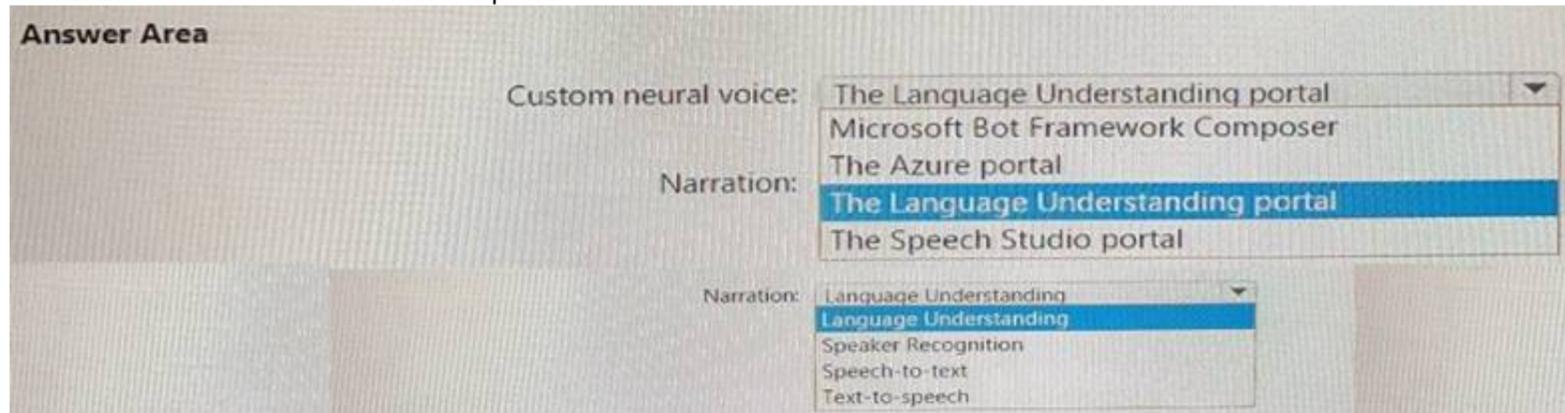
HOTSPOT - (Topic 3)

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.

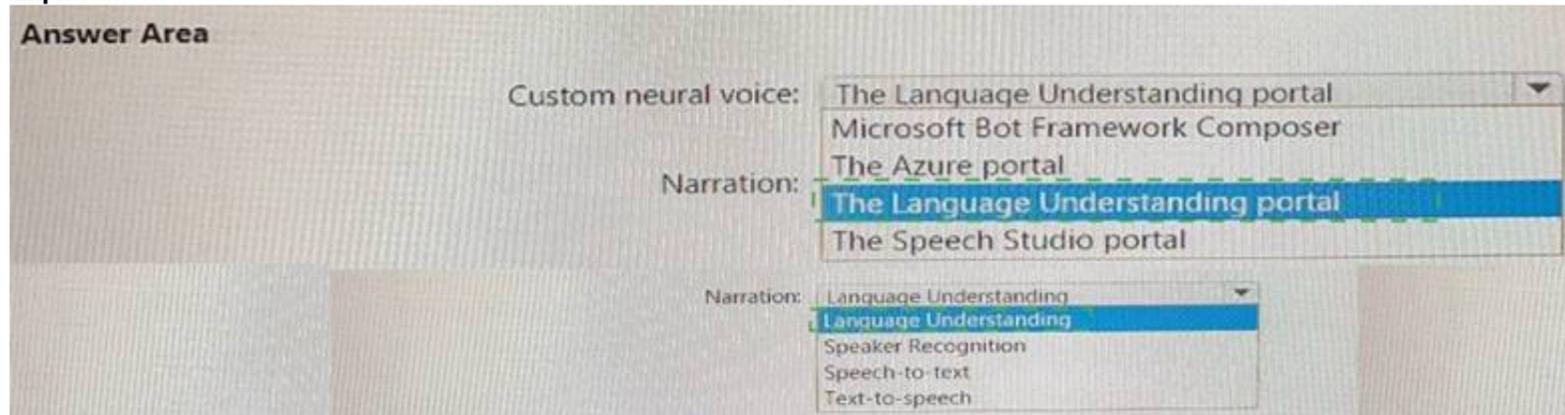


A. Mastered

B. Not Mastered

Answer: A

Explanation:



**NEW QUESTION 5**

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

**Answer Area**

Relational data is stored in

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

Relational data is stored in

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

**NEW QUESTION 6**

- (Topic 3)

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

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

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

You need to add the new images to the classifier.

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

- A. Yes
- B. No

**Answer:** B

**Explanation:**

The model needs to be extended and retrained.

**NEW QUESTION 7**

- (Topic 3)

You need to recommend a non-relational data store that is optimized for storing and retrieving text files, videos, audio streams, and virtual disk images. The data store must store data, some metadata, and a unique ID for each file. Which type of data store should you recommend?

- A. columnar
- B. key/value
- C. document
- D. object

**Answer:** D

**NEW QUESTION 8**

- (Topic 3)

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

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

**Answer:** B

**Explanation:**

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

The default audio streaming format is WAV

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

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

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

**NEW QUESTION 9**

DRAG DROP - (Topic 3)

You have a web app that uses Azure Cognitive Search.

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

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

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

**Actions**

- Add a new query key.
- Regenerate the secondary admin key.
- Change the app to use the secondary admin key.
- Change the app to use the new key.
- Regenerate the primary admin key.
- Delete the compromised key.

**Answer Area**



- A. Mastered
- B. Not Mastered

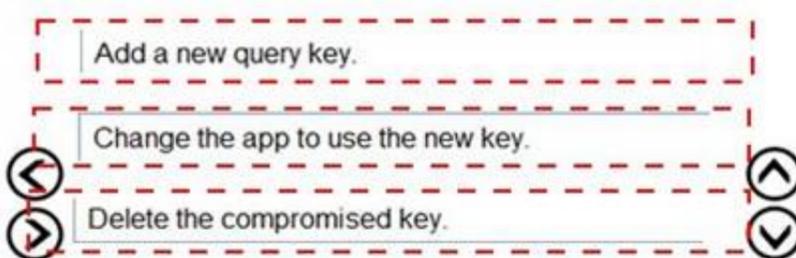
**Answer:** A

**Explanation:**

**Actions**

- Add a new query key.
- Regenerate the secondary admin key.
- Change the app to use the secondary admin key.
- Change the app to use the new key.
- Regenerate the primary admin key.
- Delete the compromised key.

**Answer Area**



**NEW QUESTION 10**

- (Topic 3)

You have a Video Indexer service that is used to provide a search interface over company videos on your company's website.

You need to be able to search for videos based on who is present in the video. What should you do?

- A. Create a person model and associate the model to the videos.
- B. Create person objects and provide face images for each object.
- C. Invite the entire staff of the company to Video Indexer.
- D. Edit the faces in the videos.
- E. Upload names to a language model.

**Answer:** A

**Explanation:**

Video Indexer supports multiple Person models per account. Once a model is created, you can use it by providing the model ID of a specific Person model when uploading/indexing or reindexing a video. Training a new face for a video updates the specific custom model that the video was associated with.

Note: Video Indexer supports face detection and celebrity recognition for video content. The celebrity recognition feature covers about one million faces based on commonly requested data source such as IMDB, Wikipedia, and top LinkedIn influencers. Faces that aren't recognized by the celebrity recognition feature are detected but left unnamed. Once you label a face with a name, the face and name get added to your account's Person model. Video Indexer will then recognize this face in your future videos and past videos.

Reference:

<https://docs.microsoft.com/en-us/azure/media-services/video-indexer/customize-person-model-with-api>

**NEW QUESTION 10**

- (Topic 3)

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

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

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

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

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

Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

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>

**NEW QUESTION 13**

HOTSPOT - (Topic 3)

You are building a call handling system that will receive calls from French-speaking and German-speaking callers. The system must perform the following tasks;

- Capture inbound voice messages as text.
- Replay messages in English on demand.

Which Azure Cognitive Services should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

● ● ● ● ●

To capture messages: Speech-to-text ▼  
Speaker Recognition  
Speech-to-text  
Text-to-speech  
Translator

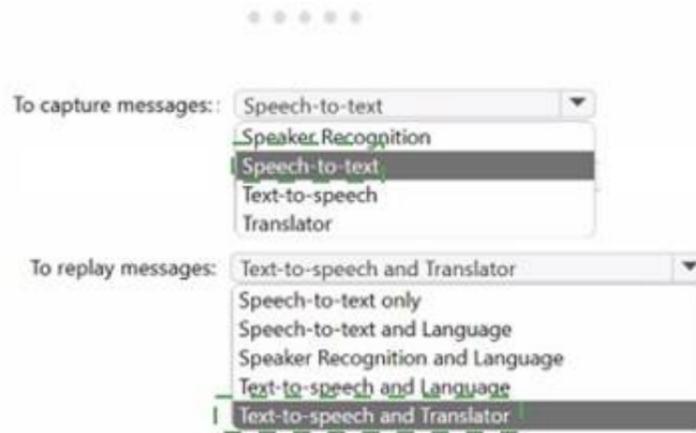
To replay messages: Text-to-speech and Translator ▼  
Speech-to-text only  
Speech-to-text and Language  
Speaker Recognition and Language  
Text-to-speech and Language  
Text-to-speech and Translator

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area



**NEW QUESTION 17**

- (Topic 3)

You are building a Language Understanding solution.

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

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

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

**Answer: C**

**NEW QUESTION 18**

- (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 SQL Database
- C. Azure HDInsight
- D. Azure Cosmos DB

**Answer: B**

**NEW QUESTION 23**

- (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 pattern in the FindContact intent.

Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance.

Define a set of intents that corresponds to actions users want to take in your application. Reference:

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

**NEW QUESTION 28**

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
Stream processing has access to the most recent data received or data within a rolling time window.	<input type="radio"/>	<input type="radio"/>
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	<input type="radio"/>	<input type="radio"/>
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

**Answer Area**

Statements	Yes	No
Stream processing has access to the most recent data received or data within a rolling time window.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	<input type="checkbox"/>	<input checked="" type="checkbox"/>

**NEW QUESTION 31**

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; flex-direction: column; align-items: center;"> <div style="margin-bottom: 10px;"><input type="checkbox"/></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 32**

DRAG DROP - (Topic 3)

You plan to build a chatbot to support task tracking.

You create a Language Understanding service named lu1.

You need to build a Language Understanding model to integrate into the chatbot. The solution must minimize development time to build the model.

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

Actions	Answer Area
Train the application.	<div style="display: flex; flex-direction: column; align-items: center;"> <div style="margin-bottom: 10px;"><input type="checkbox"/></div> </div>
Publish the application.	
Add a new application.	
Add example utterances.	
Add the prebuilt domain ToDo.	

- A. Mastered
- B. Not Mastered

Answer: A

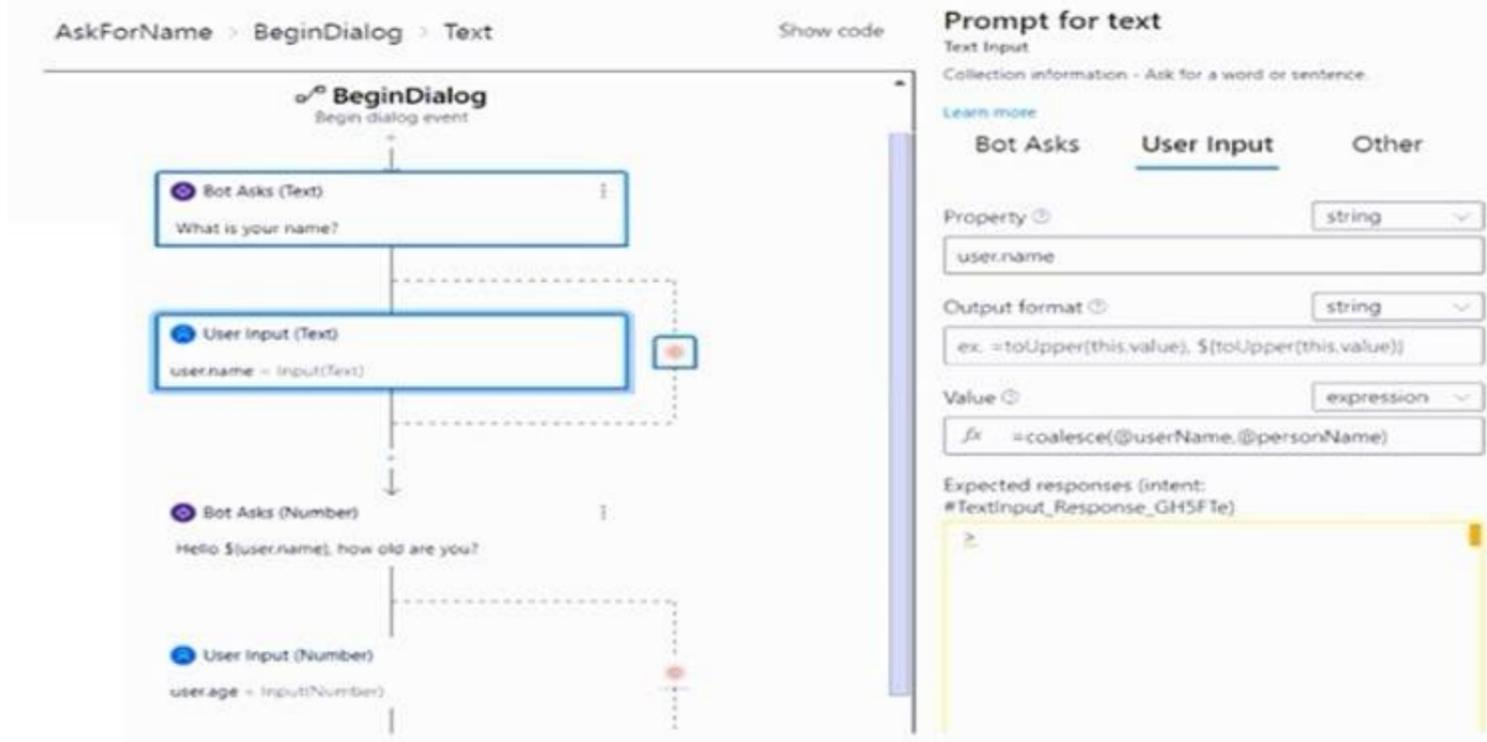
**Explanation:**

- \* 1. Add a new application
- \* 2. Add a prebuilt domain intent ToDo (it has already utterances so we can skip this step)
- \* 3. Train
- \* 4. Publish

**NEW QUESTION 34**

HOTSPOT - (Topic 3)

You are building a chatbot by using the Microsoft Bot Framework Composer. You have the dialog design shown in the following exhibit.



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

NOTE: Each correct selection is worth one point.

**Answer Area**

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

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

**Answer Area**

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

**NEW QUESTION 39**

HOTSPOT - (Topic 3)

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

@ ml DirectionalReference
@ prebuilt personName
```

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

SelectItem is [answer choice].

▼
a domain
an entity
an intent
an utterance

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

▼
a domain
an entity
an intent
an utterance

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

**NEW QUESTION 42**

- (Topic 3)

What are two uses of data visualization? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Communicate the significance of data.
- B. Represent trends and patterns over time.
- C. Implement machine learning to predict future values.
- D. Enforce business logic across reports.

**Answer:** AB

**NEW QUESTION 45**

- (Topic 3)

You are building an app that will use the Azure Video Indexer service. You plan to train a language model to recognize industry-specific terms. You need to upload a file that contains the industry-specific terms. Which file format should you use?

- A. PDF
- B. XML
- C. TXT
- D. XLS

**Answer:** C

**NEW QUESTION 47**

- (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 51**

- (Topic 3)

You have an Azure subscription that contains a multi-service Azure Cognitive Services Translator resource named Translator1.

You are building an app that will translate text and documents by using Translator1. You need to create the REST API request for the app. Which headers should you include in the request?

- A. the subscription key and the client trace ID
- B. the subscription key, the subscription region, and the content type
- C. the resource ID and the content language
- D. the access control request, the content type, and the content length

**Answer:** B

**NEW QUESTION 54**

HOTSPOT - (Topic 3)

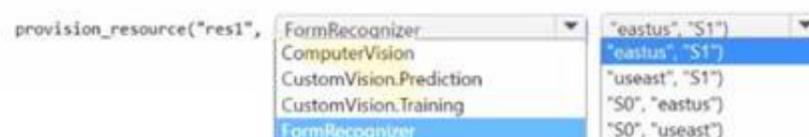
You plan to provision Azure Cognitive Services resources by using the following method.

```
{
  CognitiveServicesAccount parameters =
    new CognitiveServicesAccount(null, null, kind, location, name,
      new CognitiveServicesAccountProperties(), new Sku(tier));
  result = client.Accounts.Create(resource_group_name, tier, parameters);
}
```

You need to create a Standard tier resource that will convert scanned receipts into text. How should you call the method? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

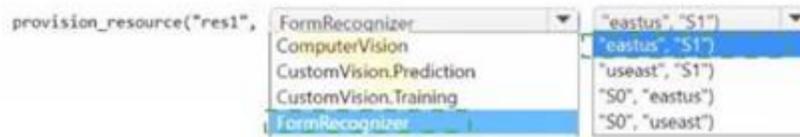


- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area



**NEW QUESTION 58**

- (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 60**

- (Topic 3)

You are building a natural language model. You need to enable active learning. What should you do?

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

**Answer:** C

**Explanation:**

Reference:

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

**NEW QUESTION 61**

DRAG DROP - (Topic 3)

You are developing an application that will recognize faults in components produced on a factory production line. The components are specific to your business. You need to use the Custom Vision API to help detect common faults.

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**

- Train the classifier model.
- Upload and tag images.
- Initialize the training dataset.
- Train the object detection model.
- Create a project.



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

- Step 1: Create a project Create a new project.
- Step 2: Upload and tag the images Choose training images. Then upload and tag the images.
- Step 3: Train the classifier model. Train the classifier

**NEW QUESTION 64**

- (Topic 3)

You have a Language Understanding resource named lu1. You build and deploy an Azure bot named bot1 that uses lu1. You need to ensure that bot1 adheres to the Microsoft responsible AI principle of inclusiveness.

How should you extend bot1?

- A. Implement authentication for bot1.
- B. Enable active learning for lu1.
- C. Host lu1 in a container.
- D. Add Direct Line Speech to bot1.

**Answer: D**

**Explanation:**

Inclusiveness: AI systems should empower everyone and engage people.

Direct Line Speech is a robust, end-to-end solution for creating a flexible, extensible voice assistant. It is powered by the Bot Framework and its Direct Line Speech channel, that is optimized for voice-in, voice-out interaction with bots.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/direct-line-speech>

**NEW QUESTION 65**

- (Topic 3)

You deploy a web app that is used as a management portal for indexing in Azure Cognitive Search. The app is configured to use the primary admin key. During a security review, you discover unauthorized changes to the search index. You suspect that the primary access key is compromised.

You need to prevent unauthorized access to the index management endpoint. The solution must minimize downtime.

What should you do next?

- A. Regenerate the primary admin key, change the app to use the secondary admin key, and then regenerate the secondary admin key.
- B. Change the app to use a query key, and then regenerate the primary admin key and the secondary admin key.
- C. Regenerate the secondary admin key, change the app to use the secondary admin key, and then regenerate the primary key.
- D. Add a new query key, change the app to use the new query key, and then delete all the unused query keys.

**Answer: C**

**NEW QUESTION 69**

- (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 72**

HOTSPOT - (Topic 3)

You have an Azure subscription that has the following configurations:

- Subscription ID: 8d3591aa-96b8-4737-ad09-00f9b1ed35ad
- Tenant ID: 3edfe572-cb54-3ced-ae12-c5c177f39a12

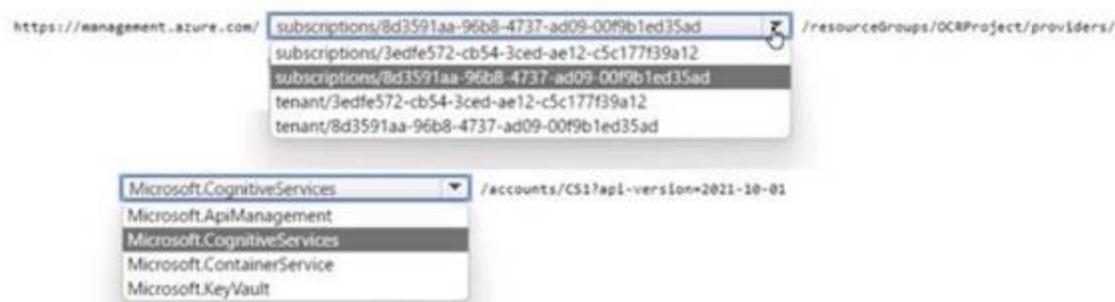
You plan to create a resource that will perform sentiment analysis and optical character recognition (OCR).

You need to use an HTTP request to create the resource in the subscription. The solution must use a single key and endpoint.

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

NOTE: Each correct selection is worth one point.

Answer Area

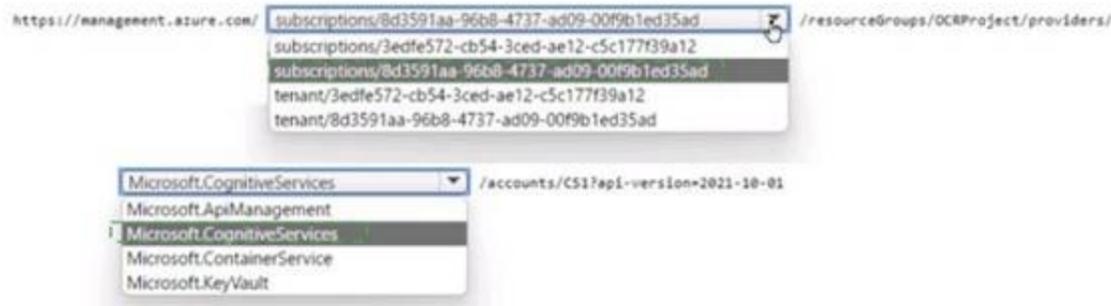


- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Answer Area



**NEW QUESTION 77**

- (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 82**

- (Topic 3)

You are building a Chatbot by using the Microsoft Bot Framework SDK. The bot will be used to accept food orders from customers and allow the customers to customize each food item. You need to configure the bot to ask the user for additional input based on the type of item ordered. The solution must minimize development effort. Which two types of dialogs should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

**Answer:** BC

**NEW QUESTION 85**

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	Answer Area
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.	⬆
	⬇

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Actions	Answer Area
Use the ExportProject endpoint on acvdev.	Use the GetProjects endpoint on acvdev.
Use the GetProjects endpoint on acvdev.	Use the ExportProject endpoint on acvdev.
Use the ImportProject endpoint on acvprod.	Use the ImportProject endpoint on acvprod.
Use the ExportIteration endpoint on acvdev.	
Use the GetIterations endpoint on acvdev.	
Use the UpdateProject endpoint on acvprod.	

**NEW QUESTION 88**

- (Topic 3)

You manage an application that stores data in a shared folder on a Windows server. You need to move the shared folder to Azure Storage. Which type of Azure Storage should you use?

- A. table
- B. queue
- C. file
- D. blob

**Answer: C**

**NEW QUESTION 93**

- (Topic 3)

You have the following data sources:

- ? Finance: On-premises Microsoft SQL Server database
- ? Sales: Azure Cosmos DB using the Core (SQL) API
- ? Logs: Azure Table storage
- ? HR: Azure SQL database

You need to ensure that you can search all the data by using the Azure Cognitive Search REST API. What should you do?

- A. Configure multiple read replicas for the data in Sales.
- B. Mirror Finance to an Azure SQL database.
- C. Migrate the data in Sales to the MongoDB API.
- D. Ingest the data in Logs into Azure Sentinel.

**Answer: B**

**Explanation:**

On-premises Microsoft SQL Server database cannot be used as an index data source. Note: Indexer in Azure Cognitive Search: : Automate aspects of an indexing operation by configuring a data source and an indexer that you can schedule or run on demand. This feature is supported for a limited number of data source types on Azure. Indexers crawl data stores on Azure.

- ? Azure Blob Storage
- ? Azure Data Lake Storage Gen2 (in preview)
- ? Azure Table Storage
- ? Azure Cosmos DB
- ? Azure SQL Database
- ? SQL Managed Instance
- ? SQL Server on Azure Virtual Machines

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-indexer-overview#supported-data-sources>

**NEW QUESTION 94**

- (Topic 3)

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

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

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

You need to add the new images to the classifier.

Solution: You add the new images and labels to the existing model. You retrain the model, and then publish the model.

Does this meet the goal?

- A. Yes
- B. No

**Answer: A**

**Explanation:**

The model needs to be extended and retrained.

**NEW QUESTION 98**

- (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 101**

- (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 105**

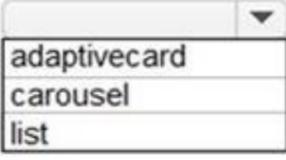
HOTSPOT - (Topic 3)

You are designing a conversation flow to be used in a chatbot.

You need to test the conversation flow by using the Microsoft Bot Framework Emulator. How should you complete the .chat file? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

user=User1
bot=watchbot
user: I want a new watch.
bot: [  ][Delay=3000]
bot: I can help you with that! Let me see what I can find.
bot: Here's what I found.
bot:
[AttachmentLayout=  ]
[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]
user: I like the first one.
bot: Sure, pulling up more information.
bot: [Attachment=cards\watchProfileCard.json  ]
user: That's nice! Thank you.
bot: Sure, you are most welcome!

```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

```

user=User1
bot=watchbot
user: I want a new watch.
bot: [ Attachment
        ConversationUpdate
        Typing ] [Delay=3000]

bot: I can help you with that! Let me see what I can find.
bot: Here's what I found.
bot:
[AttachmentLayout= adaptivecard
                   carousel
                   thumbnail ]

[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]
user: I like the first one.
bot: Sure, pulling up more information.
bot: [Attachment=cards\watchProfileCard.json ]
user: That's nice! Thank you.
bot: Sure, you are most welcome!
    
```

**NEW QUESTION 109**

HOTSPOT - (Topic 3)

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<input type="radio"/>	<input type="radio"/>
Normalizing a database reduces data redundancy.	<input type="radio"/>	<input type="radio"/>
Normalization improves data integrity.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	<input type="radio"/>	<input checked="" type="radio"/>
Normalizing a database reduces data redundancy.	<input checked="" type="radio"/>	<input type="radio"/>
Normalization improves data integrity.	<input type="radio"/>	<input type="radio"/>

**NEW QUESTION 112**

- (Topic 3)

You create a bot by using the Microsoft Bot Framework SDK.

You need to configure the bot to respond to events by using custom text responses. What should you use?

- A. an adaptive card
- B. an activity handler
- C. a dialog
- D. a skill

Answer: B

**NEW QUESTION 114**

- (Topic 3)

What is a characteristic of a non-relational database?

- A. full support for Transact-SQL
- B. a fixed schema
- C. self describing entities

Answer: C

**NEW QUESTION 116**

HOTSPOT - (Topic 3)

You are developing an application to recognize employees' faces by using the Face Recognition API. Images of the faces will be accessible from a URI endpoint. The application has the following code.

```
static async void AddFace(string subscription_key, string personGroupId, string personId, string imageURI)
{
    var client = new HttpClient();
    client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", subscription_key);
    var endpointURI = $"https://westus.api.cognitive.microsoft.com/face/v1.0/persongroups/{personGroupId}/persons/{personId}/persistedFaces";
    HttpResponseMessage response;
    var body = "{ \"url\": \"\" + imageURI + \"\"}";
    var content = new StringContent(body, Encoding.UTF8, "application/json");
    var response = await client.PutAsync(endpointURI, content);
}
```

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
The code will add a face image to a person object in a person group.	<input type="radio"/>	<input type="radio"/>
The code will work for a group of 10,000 people.	<input type="radio"/>	<input type="radio"/>
AddFace can be called multiple times to add multiple face images to a person object.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

B: see this example code from documentation that uses PersonGroup of size 10,000 : <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-add-faces>  
 the question wants to trick you into thinking you need to use a LargePersonGroup for a size of 10,000 - but the documentation for it doesn't include this limitation or criteria: <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-use-large-scale>

**NEW QUESTION 118**

HOTSPOT - (Topic 3)

You run the following command.

```
docker run --rm -it -p 5000:5000 --memory 10g --cpus 2 \
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment\
Eula=accept \
Billing={ENDPOINT_URI} \
ApiKey={API_KEY}
```

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

Statements	Yes	No
Going to <a href="http://localhost:5000/status">http://localhost:5000/status</a> will query the Azure endpoint to verify whether the API key used to start the container is valid.	<input type="radio"/>	<input type="radio"/>
The container logging provider will write log data.	<input type="radio"/>	<input type="radio"/>
Going to <a href="http://localhost:5000/swagger">http://localhost:5000/swagger</a> will provide the details to access the documentation for the available endpoints.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Yes  
<http://localhost:5000/status> : Also requested with GET, this verifies if the api-key used to start the container is valid without causing an endpoint query.  
 Box 2: Yes  
 The command saves container and LUIS logs to output mount at C:\output, located on container host  
 Box 3: Yes  
<http://localhost:5000/swagger> : The container provides a full set of documentation for the endpoints and a Try it out feature. With this feature, you can enter your settings into a web-based HTML form and make the query without having to write any code. After the query returns, an example CURL command is provided to demonstrate the

HTTP headers and body format that's required.

**NEW QUESTION 119**

- (Topic 3)

You are building a Conversational Language Understanding model.

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

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

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

- A. a location Intent
- B. a change setting entity
- C. a device intent
- D. a change setting intent
- E. a query setting intent
- F. a device entity

**Answer:** BCF

**NEW QUESTION 121**

HOTSPOT - (Topic 3)

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

You have the following code segment.

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

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 124**

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

**Answer Area**

A block of code that runs in a database is called

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

A block of code that runs in a database is called

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

**NEW QUESTION 127**

HOTSPOT - (Topic 3)

You are building an app by using the Speech SDK. The app will translate speech from French to German by using natural language processing. You need to define the source language and the output 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**

```
var speechTranslationConfig =
SpeechTranslationConfig.FromSubscription(speechKey, speechRegion);
speechTranslationConfig.
```

- SpeechRecognitionLanguage ("fr")
- AddTargetLanguage
- SpeechRecognitionLanguage
- SpeechSynthesisLanguage
- TargetLanguages
- VoiceName

```
speech_translation_config.
```

- SpeechSynthesisLanguage ("de")
- AddTargetLanguage
- SpeechRecognitionLanguage
- SpeechSynthesisLanguage
- TargetLanguages
- VoiceName

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

```
var speechTranslationConfig =
SpeechTranslationConfig.FromSubscription(speechKey, speechRegion);
speechTranslationConfig.
```

- SpeechRecognitionLanguage ("fr")
- AddTargetLanguage
- SpeechRecognitionLanguage
- SpeechSynthesisLanguage
- TargetLanguages
- VoiceName

```
speech_translation_config.
```

- SpeechSynthesisLanguage ("de")
- AddTargetLanguage
- SpeechRecognitionLanguage
- SpeechSynthesisLanguage
- TargetLanguages
- VoiceName

**NEW QUESTION 130**

- (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 private endpoint to vnet1. Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.

The service could be an Azure service such as:

- ? Azure Storage
- ? Azure Cosmos DB
- ? Azure SQL Database
- ? Your own service using a Private Link Service.

Reference:

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

**NEW QUESTION 132**

DRAG DROP - (Topic 3)

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

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: LargeFaceListID

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

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

**NEW QUESTION 136**

DRAG DROP - (Topic 3)

You have 100 chatbots that each has its own Language Understanding model. Frequently, you must add the same phrases to each model.

You need to programmatically update the Language Understanding models to include the new phrases.

How should you complete the code? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

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

NOTE: Each correct selection is worth one point.

Values	Answer Area
AddPhraselistAsync	<pre>var phraselistId = await client.Features. [ ] (appId, versionId, new [ ] {   EnabledForAllModels = false,   IsExchangeable = true,   Name = "PL1",   Phrases = "item1,item2,item3,item4,item5" });</pre>
Phraselist	
PhraselistCreateObject	
Phrases	
SavePhraselistAsync	
UploadPhraselistAsync	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: AddPhraselistAsync Example: Add phraselist feature

```
var phraselistId = await client.Features.AddPhraselistAsync(appId, versionId, new PhraselistCreateObject
```

```
{
EnabledForAllModels = false, IsExchangeable = true, Name = "QuantityPhraselist", Phrases = "few,more,extra"
});
```

Box 2: PhraselistCreateObject

#### NEW QUESTION 140

- (Topic 3)

You are building a chatbot.

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

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

**Answer: B**

#### NEW QUESTION 143

- (Topic 3)

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

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

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

- A. create\_resource(client, "res1", "ComputerVision", "F0", "westus")
- B. create\_resource(client, "res1", "CustomVision.Prediction", "F0", "westus")
- C. create\_resource(client, "res1", "ComputerVision", "S0", "westus")
- D. create\_resource(client, "res1", "CustomVision.Prediction", "S0", "westus")

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 147

- (Topic 3)

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

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

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

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

You need to reduce the likelihood that search query requests are throttled. Solution: You add indexes.

Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

#### Explanation:

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>

#### NEW QUESTION 148

- (Topic 3)

You are developing a new sales system that will process the video and text from a public-facing website.

You plan to monitor the sales system to ensure that it provides equitable results regardless of the user's location or background.

Which two responsible AI principles provide guidance to meet the monitoring requirements? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

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

**Answer:** BC

**Explanation:**

<https://docs.microsoft.com/en-us/learn/modules/get-started-ai-fundamentals/8-understand-responsible-ai>

**NEW QUESTION 149**

DRAG DROP - (Topic 3)

You are developing a webpage that will use the Video Indexer service to display videos of internal company meetings.

You embed the Player widget and the Cognitive Insights widget into the page. You need to configure the widgets to meet the following requirements:

? Ensure that users can search for keywords.

? Display the names and faces of people in the video.

? Show captions in the video in English (United States).

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

NOTE: Each correct selection is worth one point.

Values	Answer Area
en-US	<b>Cognitive Insights Widget</b> https://www.videoindexer.ai/embed/insights/<accountId>/<videoId>/?widgets= [Value] controls= [Value]
false	
people,keywords	<b>Player Widget</b> https://www.videoindexer.ai/embed/player/<accountId>/<videoId>/? showcaptions= [Value] captions= [Value]
people,search	
search	
true	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 150**

- (Topic 3)

You are building an app that will share user images.

You need to configure the app to perform the following actions when a user uploads an image:

- Categorize the image as either a photograph or a drawing.
- Generate a caption for the image.

The solution must minimize development effort.

Which two services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. object detection in Computer Vision
- B. content tags in Computer Vision
- C. image descriptions in Computer Vision
- D. image type detection in Computer Vision
- E. image classification in Custom Vision

**Answer:** CD

**Explanation:**

? According to the Microsoft documentation, Computer Vision is a cloud-based service that provides developers with access to advanced algorithms for processing images and returning information. By uploading an image or specifying an image URL, Computer Vision algorithms can analyze visual content in different ways based on inputs and user choices.

? According to the Microsoft documentation, image type detection is one of the features of Computer Vision that can categorize an image as either a photograph or a drawing. You can use the image type detection feature by calling the Analyze Image API with the visualFeatures parameter set to ImageType. The API will return a JSON response with an imageType field that indicates whether the image is a photo or a clipart.

? According to the Microsoft documentation, image descriptions is another feature of Computer Vision that can generate a caption for an image. You can use the image descriptions feature by calling the Analyze Image API with the visualFeatures parameter set to Description. The API will return a JSON response with a description field that contains a list of captions for the image, each with a confidence score.

? Therefore, by using these two features of Computer Vision, you can achieve your app requirements with minimal development effort. You don't need to use any other services, such as object detection, content tags, or Custom Vision, which are designed for different purposes.

**NEW QUESTION 155**

DRAG DROP - (Topic 3)

Match the types of workloads to the appropriate scenarios.

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

NOTE: Each correct match is worth one point.

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 159**

DRAG DROP - (Topic 3)

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.

Entity Types	Answer Area
Email	
List	Paris: <input type="text"/>
Regex	email@domain.com: <input type="text"/>
GeographyV2	2 audit business: <input type="text"/>
Machine learned	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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.

**NEW QUESTION 161**

- (Topic 3)

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

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

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

You need to add the new images to the classifier.

Solution: You add the new images, and then use the Smart Labeler tool. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

The model need to be extended and retrained.

Note: Smart Labeler to generate suggested tags for images. This lets you label a large number of images more quickly when training a Custom Vision model.

**NEW QUESTION 162**

DRAG DROP - (Topic 3)

You are building a customer support chatbot.

You need to configure the bot to identify the following:

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

Which Azure Cognitive Service for Language feature should you use for each requirement? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content

NOTE: Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

**NEW QUESTION 163**

- (Topic 3)

Which Azure Storage service implements the key/value model?

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

Answer: C

**NEW QUESTION 164**

- (Topic 3)

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

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

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

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

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

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

Does this meet the goal?

- A. Yes
- B. No

Answer: B

**Explanation:**

Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance.

Define a set of intents that corresponds to actions users want to take in your application. Reference:

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

**NEW QUESTION 166**

HOTSPOT - (Topic 3)

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

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

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

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

NOTE: Each correct selection is worth one point.

```
docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \
  http://contoso.blob.core.windows.net
  https://contoso.cognitiveservices.azure.com
  mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
  mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment
Eula=accept \
Billing=
  http://contoso.blob.core.windows.net
  https://contoso.cognitiveservices.azure.com
  mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
  mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment
ApiKey=xxxxxxxxxxxxxxxxxxxxxxxx
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 169**

**HOTSPOT - (Topic 3)**

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

**Answer Area**

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

prebuilt-document fileUri);

- prebuilt-document
- prebuilt-contract
- prebuilt-read

0.75

- 0.1
- 0.75
- 1.0

- A. Mastered
- B. Not Mastered

**Answer:** A

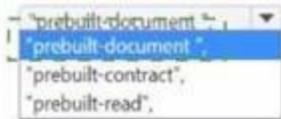
**Explanation:**

Answer Area

```

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

```



**NEW QUESTION 170**

- (Topic 3)

What is the primary purpose of a data warehouse?

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

Answer: C

**NEW QUESTION 175**

DRAG DROP - (Topic 3)

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

**Actions**

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

**Answer Area**

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

**Actions**

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

**Answer Area**

- Create a new classification model.
- Change the classification type.
- Export the model.

**NEW QUESTION 179**

DRAG DROP - (Topic 3)

You have an Azure subscription that contains an Azure OpenAI resource named AH. You plan to build an app named App1 that will write press releases by using AM. You need to deploy an Azure OpenAI model for App1. The solution must minimize development effort. Which three actions should you perform in sequence in Azure OpenAI Studio? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

**Actions**

- 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.

**Answer Area**

⬆  
⬇

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Actions**

- 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.

**Answer Area**

- 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.

⬆  
⬇

**NEW QUESTION 183**

- (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 187**

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 191**

- (Topic 3)

You successfully run the following HTTP request.

POST <https://management.azure.com/subscriptions/18c51a87-3a69-47a8-aedc-a54745f708a1/resourceGroups/RG1/providers/Microsoft.CognitiveServices/accounts/contoso/sol/regenerateKey?api-version=2017-04-18>

Body{"keyName": "Key2"}

Body{"keyName": "Key2"}

What is the result of the request?

- A. A key for Azure Cognitive Services was generated in Azure Key Vault.
- B. A new query key was generated.
- C. The primary subscription key and the secondary subscription key were rotated.
- D. The secondary subscription key was reset.

**Answer: B**

**Explanation:**

<https://docs.microsoft.com/en-us/rest/api/searchmanagement/2021-04-01-preview/query-keys/create>

**NEW QUESTION 195**

HOTSPOT - (Topic 3)

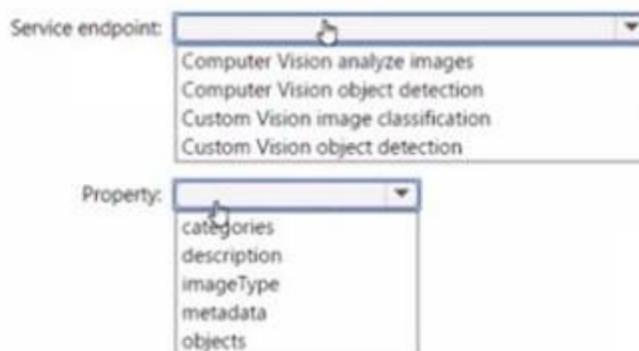
You have a library that contains thousands of images.

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

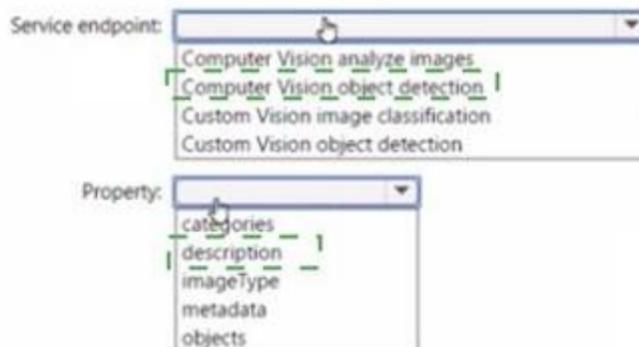


- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Answer Area



**NEW QUESTION 200**

DRAG DROP - (Topic 3)

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

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

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

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

Actions

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



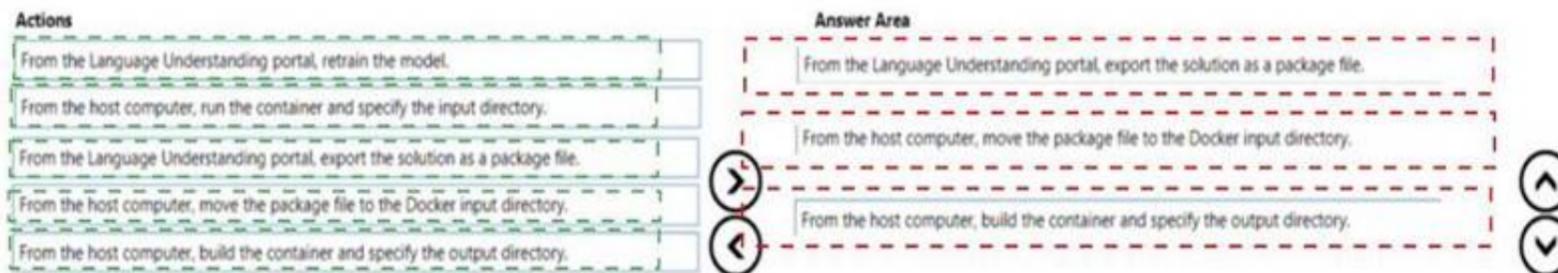
Answer Area



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**



**NEW QUESTION 203**

- (Topic 3)

You are building a social media extension that will convert text to speech. The solution must meet the following requirements:

- Support messages of up to 400 characters.
- Provide users with multiple voice options.
- Minimize costs.

You create an Azure Cognitive Services resource.

Which Speech API endpoint provides users with the available voice options?

- A. <https://uksouth.customvoice.apispeech.microsoft.com/api/texttospeech/v3.0/longaudiosynthesisis/voices>
- B. <https://uksouth.tts.speech.microsoft.com/cognitiveservices/voices/list>
- C. <https://uksouth.voice.speech.microsoft.com/cognitiveservices/v1?deploymentId={deploymentId}>
- D. <https://uksouth.api.cognitive.microsoft.com/speechtotext/v3.0/models/base>

**Answer: B**

**NEW QUESTION 205**

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 (
            t = File.OpenRead(imagePath))
        {
            await faceClient.PersonGroupPerson.
                (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 208**

- (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 in 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 209**

HOTSPOT - (Topic 3)

You are developing an internet-based training solution for remote learners.

Your company identifies that during the training, some learners leave their desk for long periods or become distracted.

You need to use a video and audio feed from each learner's computer to detect whether the learner is present and paying attention. The solution must minimize development effort and identify each learner.

Which Azure Cognitive Services service should you use for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

From a learner's video feed, verify whether the learner is present:

From a learner's facial expression in the video feed, verify whether the learner is paying attention:

From a learner's audio feed, detect whether the learner is talking:

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

From a learner's video feed, verify whether the learner is present:

From a learner's facial expression in the video feed, verify whether the learner is paying attention:

From a learner's audio feed, detect whether the learner is talking:

**NEW QUESTION 213**

- (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 214

- (Topic 3)

You need to implement a table projection to generate a physical expression of an Azure Cognitive Search index.

Which three properties should you specify in the skillset definition JSON configuration table node? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

- A. tableName
- B. generatedKeyName
- C. dataSource
- D. dataSourceConnection
- E. source

**Answer:** ABE

#### Explanation:

Defining a table projection.

Each table requires three properties:

tableName: The name of the table in Azure Storage.

generatedKeyName: The column name for the key that uniquely identifies this row. source: The node from the enrichment tree you are sourcing your enrichments from. This node is usually the output of a shaper, but could be the output of any of the skills.

Reference:

<https://docs.microsoft.com/en-us/azure/search/knowledge-store-projection-overview>

#### NEW QUESTION 216

- (Topic 3)

You are building a language model by using a Language Understanding service. You create a new Language Understanding resource.

You need to add more contributors. What should you use?

- A. a conditional access policy in Azure Active Directory (Azure AD)
- B. the Access control (1AM) page for the authoring resources in the Azure portal
- C. the Access control (1AM) page for the prediction resources in the Azure portal

**Answer:** B

#### Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-how-to-collaborate>

#### NEW QUESTION 217

- (Topic 3)

You are developing an application that will use Azure Cognitive Search for internal documents.

You need to implement document-level filtering for Azure Cognitive Search.

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

NOTE: Each correct selection is worth one point.

- A. Send Azure AD access tokens with the search request.
- B. Retrieve all the groups.
- C. Retrieve the group memberships of the user.
- D. Add allowed groups to each index entry.
- E. Create one index per group.
- F. Supply the groups as a filter for the search requests.

**Answer:** CDF

#### Explanation:

Your documents must include a field specifying which groups have access. This information becomes the filter criteria against which documents are selected or rejected from the result set returned to the issuer.

D: A query request targets the documents collection of a single index on a search service. CF: In order to trim documents based on group\_ids access, you should issue a search

query with a group\_ids/any(g:search.in(g, 'group\_id1, group\_id2,...')) filter, where 'group\_id1, group\_id2,...' are the groups to which the search request issuer belongs.

Reference:

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

#### NEW QUESTION 221

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.

Statements	Yes	No
Azure Databricks is an Apache Spark-based analytics platform.	<input type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input type="radio"/>	<input type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
Azure Databricks is an Apache Spark-based analytics platform.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input checked="" type="checkbox"/>	<input type="radio"/>

**NEW QUESTION 226**

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

xxx-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 230**

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

- A data analyst
- A data engineer
- A data scientist

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

- A data analyst
- A data engineer
- A data scientist

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

**NEW QUESTION 231**

- (Topic 3)

You need to create a visualization of running sales totals per quarter as shown in the following exhibit.

Sales by Year and Quarter

● Increase ● Decrease ● Total



What should you create in Power BI Desktop;1

- A. a waterfall chart
- B. a ribbon chart
- C. a bar chart
- D. a decomposition tree

Answer: C

**NEW QUESTION 234**

- (Topic 3)

You are building an Azure WebJob that will create knowledge bases from an array of URLs.

You instantiate a QnAMakerClient object that has the relevant API keys and assign the object to a variable named client.

You need to develop a method to create the knowledge bases.

Which two actions should you include in the method? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

Answer: AC

Explanation:

Reference:

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

**NEW QUESTION 239**

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.

Answer Area

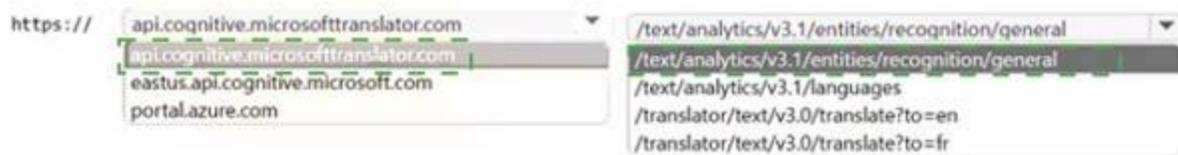


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 243

- (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 246

- (Topic 3)

You have an Azure Cognitive Services model named Model that identifies the intent of text input.

You develop an app in C# named App1. You need to configure App1 to use Model1. Which package should you add to App1?

- A. Azure.AI.Language.Conversations
- B. SpeechServicesToolkit
- C. Universal.Microsoft.CognitiveServices.Speech
- D. Xamarin.Cognitive.Speech

Answer: C

NEW QUESTION 251

- (Topic 3)

You have an Azure Cognitive Search instance that indexes purchase orders by using Form Recognizer

You need to analyze the extracted information by using Microsoft Power BI. The solution must minimize development effort. What should you add to the indexer?

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

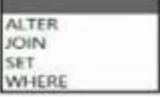
Answer: C

**NEW QUESTION 252**

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

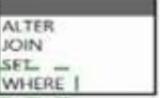
The  clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

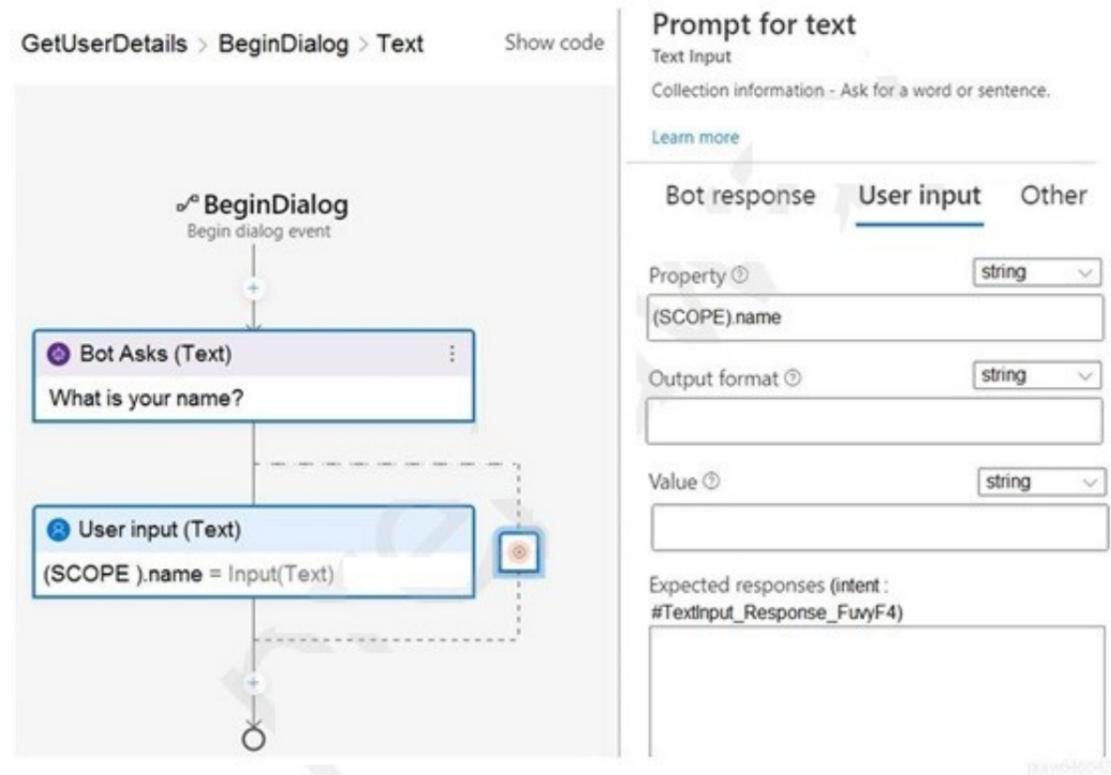
Answer Area

The  clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

**NEW QUESTION 256**

- (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 259**

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 261**

- (Topic 3)

What is used to define a query in a stream processing jobs in Azure Stream Analytics?

- A. SQL
- B. XML
- C. YAML
- D. KOL

**Answer:** A

**NEW QUESTION 263**

- (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 266**

- (Topic 3)

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

You plan to deploy the bot to Azure.

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

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

NOTE: Each correct selection is worth one point.

A. botId

B. tenantId

C. appId

D. objectId

E. appSecret

**Answer:** CE

**Explanation:**

Reference:

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

**NEW QUESTION 267**

- (Topic 2)

You are developing the document processing workflow.

You need to identify which API endpoints to use to extract text from the financial documents. The solution must meet the document processing requirements.

Which two API endpoints should you identify? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. /vision/v3.2/read/analyzeResults

B. /formrecognizer/v2.0/prebuilt/receipt/analyze

C. /vision/v3.2/read/analyze

D. /vision/v3.2/describe

E. /formercognizer/v2.0/custom/models{modelId}/analyze

**Answer:** BC

**Explanation:**

C: Analyze Receipt - Get Analyze Receipt Result.

Query the status and retrieve the result of an Analyze Receipt operation. Request URL: <https://{endpoint}/formrecognizer/v2.0-preview/prebuilt/receipt/analyzeResults/{resultId}>

E: POST {Endpoint}/vision/v3.2/read/analyze

Use this interface to get the result of a Read operation, employing the state-of-the-art Optical Character Recognition (OCR) algorithms optimized for text-heavy documents.

Scenario: Contoso plans to develop a document processing workflow to extract information automatically from PDFs and images of financial documents

? The document processing solution must be able to process standardized financial

documents that have the following characteristics:

- Contain fewer than 20 pages.

- Be formatted as PDF or JPEG files.

- Have a distinct standard for each office.

? \*The document processing solution must be able to extract tables and text from the financial documents.

? The document processing solution must be able to extract information from receipt images.

Reference:

<https://westus2.dev.cognitive.microsoft.com/docs/services/form-recognizer-api-v2-preview/operations/GetAnalyzeReceiptResult>

<https://docs.microsoft.com/en-us/rest/api/computervision/3.1/read/read>

**NEW QUESTION 272**

- (Topic 2)

You are developing the chatbot.

You create the following components:

\* A QnA Maker resource

\* A chatbot by using the Azure Bot Framework SDK.

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

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

**Answer: D**

**Explanation:**

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

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

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

Reference:

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

**NEW QUESTION 274**

HOTSPOT - (Topic 2)

You are developing the knowledgebase by using Azure Cognitive Search. You need to build a skill that will be used by indexers.

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

NOTE: Each correct selection is worth one point.

**Answer Area**

```

{
  "@odata.type": "#Microsoft.Skills.Text.EntityRecognitionSkill",
  "categories": [],
  "categories": [ "Email", "Persons", "Organizations" ],
  "categories": [ "Locations", "Persons", "Organizations" ],
  "minimumPrecision": 0.7,
  "inputs": [
    { "name": "text",
      "source": "/document/content" }
  ],
  "outputs": [
    { "name": "persons", "targetName": "people" },
    { "name": "locations", "targetName": "locations" },
    { "name": "organizations", "targetName": "organizations" },
    { "name": "entities" },
    { "name": "categories" },
    { "name": "namedEntities" }
  ]
}

```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: "categories": ["Locations", "Persons", "Organizations"], Locations, Persons, Organizations are in the outputs.

Scenario: Contoso plans to develop a searchable knowledgebase of all the intellectual property

Note: The categories parameter is an array of categories that should be extracted. Possible category types: "Person", "Location", "Organization", "Quantity", "Datetime", "URL", "Email". If no category is provided, all types are returned.

Box 2: {"name": " entities"}

The include wikis, so should include entities in the outputs.

Note: entities is an array of complex types that contains rich information about the entities extracted from text, with the following fields

name (the actual entity name. This represents a "normalized" form) wikipediaId

wikipediaLanguage

wikipediaUrl (a link to Wikipedia page for the entity) etc.

**NEW QUESTION 279**

- (Topic 2)

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

**Answer: B**

**Explanation:**

Can search content in different formats, including video

Audio and video insights (multi-channels). When indexing by one channel, partial result for those models will be available.

Keywords extraction: Extracts keywords from speech and visual text.

Named entities extraction: Extracts brands, locations, and people from speech and visual text via natural language processing (NLP).

**NEW QUESTION 281**

- (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 Form 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 284**

DRAG DROP - (Topic 2)

You are developing a solution for the Management-Bookkeepers group to meet the document processing requirements. The solution must contain the following components:

- ? A form Recognizer resource
- ? An Azure web app that hosts the Form Recognizer sample labeling tool

The Management-Bookkeepers group needs to create a custom table extractor by using the sample labeling tool.

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

Actions	Answer Area
Train a custom model	
Label the sample documents	➤
Create a new project and load sample documents	➤
Create a composite model	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Step 1: Create a new project and load sample documents  
 Create a new project. Projects store your configurations and settings.  
 Step 2: Label the sample documents  
 When you create or open a project, the main tag editor window opens.  
 Step 3: Train a custom model. Finally, train a custom model.

**NEW QUESTION 286**

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 289**

- (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 292**

HOTSPOT - (Topic 1)

You are developing the shopping on-the-go project. You need to build the Adaptive Card for the chatbot.

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**

```

version": "1.3",
"body": [
  {
    "type": "TextBlock",
    "size": "Medium",
    "weight": "Bolder",
    "text": "${
  },
  {
    "type": "TextBlock",
    "color": Attention
  },
  {
    "type": "Image",
    "url": "${image.uri}",
    "size": "Medium",
    "altText": "${

```

if(language == 'en', 'en', name)  
 name  
 name.en  
 name[language]

"\$when": "\${stockLevel != 'OK'}"  
 "\$when": "\${stockLevel == 'OK'}"  
 "\$when": "\${stockLevel.OK}"

image.altText.en  
 image.altText.language  
 image.altText["language"]  
 image.altText[language]

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Description automatically generated  
 Box 1: name.en  
 Box 2: "\$when": "\${stockLevel != 'OK'}"  
 Product displays must include images and warnings when stock levels are low or out of stock.  
 Box 3: image.altText.en

**NEW QUESTION 297**

.....

## 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](#)