Microsoft

Exam Questions AI-102

Designing and Implementing an Azure AI Solution

NEW QUESTION 1
- (Exam Topic 2)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You 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: B

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/private-link/private-link-overview

NEW QUESTION 2
- (Exam Topic 2)
You train a Custom Vision model used in a mobile app.
You receive 1,000 new images that do not have any associated data.
You need to use the images to retrain the model. The solution must minimize how long it takes to retrain the model.
Which three actions should you perform in the Custom Vision portal? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Text Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/getting-started-build-a-classifie

NEW QUESTION 3
- (Exam Topic 2)
You are developing a solution to generate a word cloud based on the reviews of a company's products. Which Text Analytics REST API endpoint should you use?

A. IceyPhrases
B. sentiment
C. languages
D. entities/recognition/general

Answer: A

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

NEW QUESTION 4
- (Exam Topic 2)
You are reviewing the design of a chatbot. The chatbot includes a language generation file that contains the following fragment.
Greet(user)

- ${Greeting()}, ${user.name}
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
`${user.name}` retrieves the user name by using a prompt.	○	○
`Greet()` is the name of the language generation template.	○	○
`${Greeting()}` is a reference to a template in the language generation file.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
Example: Greet a user whose name is stored in `user.name`
- ${ welcomeUser(user.name) }
Example: Greet a user whose name you don't know:
- ${ welcomeUser() }
Box 2: No
Greet(User) is a Send a response action.
Box 3: Yes
Reference:
https://docs.microsoft.com/en-us/composer/how-to-ask-for-user-input

NEW QUESTION 5
- (Exam Topic 2)
You are building a chatbot by using the Microsoft Bot Framework Composer. You have the dialog design shown in the following exhibit.

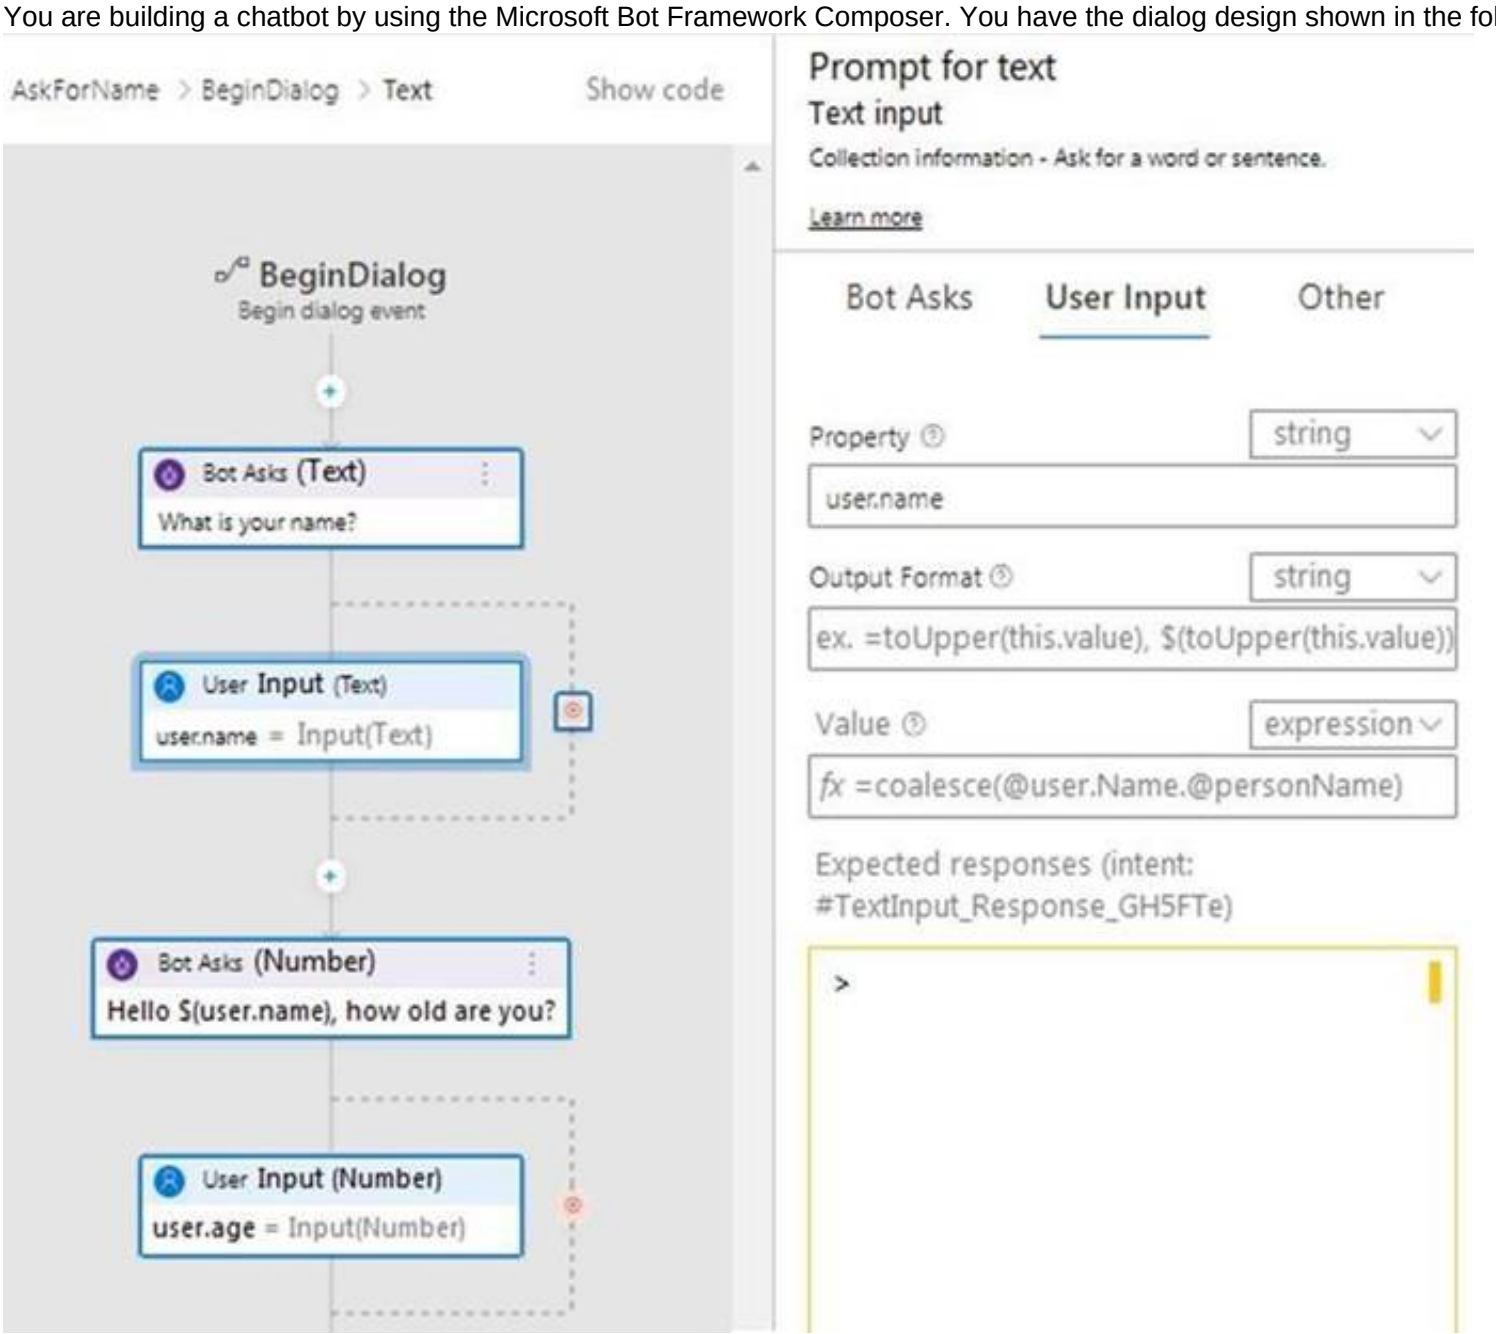

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
`user.name` is an entity.	○	○
The dialog asks for a user name and a user age and assigns appropriate values to the `user.name` and `user.age` properties.	○	○
The chatbot attempts to take the first non-null entity value for `userName` or `personName` and assigns the value to `user.name`.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
User.name is a property.
Box 2: Yes
Box 3: Yes
The coalesce() function evaluates a list of expressions and returns the first non-null (or non-empty for string) expression.
Reference:
https://docs.microsoft.com/en-us/composer/concept-language-generation https://docs.microsoft.com/en-us/azure/data-explorer/kusto/query/coalescefunction

NEW QUESTION 6
- (Exam Topic 2)
You build a custom Form Recognizer model.
You receive sample files to use for training the model as shown in the following table.

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

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

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

Answer: ACF

Explanation:

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

NEW QUESTION 7
- (Exam Topic 2)
You are building a retail chatbot that will use a QnA Maker service.
You upload an internal support document to train the model. The document contains the following question: "What is your warranty period?"
Users report that the chatbot returns the default QnA Maker answer when they ask the following question: "How long is the warranty coverage?"
The chatbot returns the correct answer when the users ask the following question: 'What is your warranty period?"
Both questions should return the same answer.
You need to increase the accuracy of the chatbot responses.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. (Choose three.)

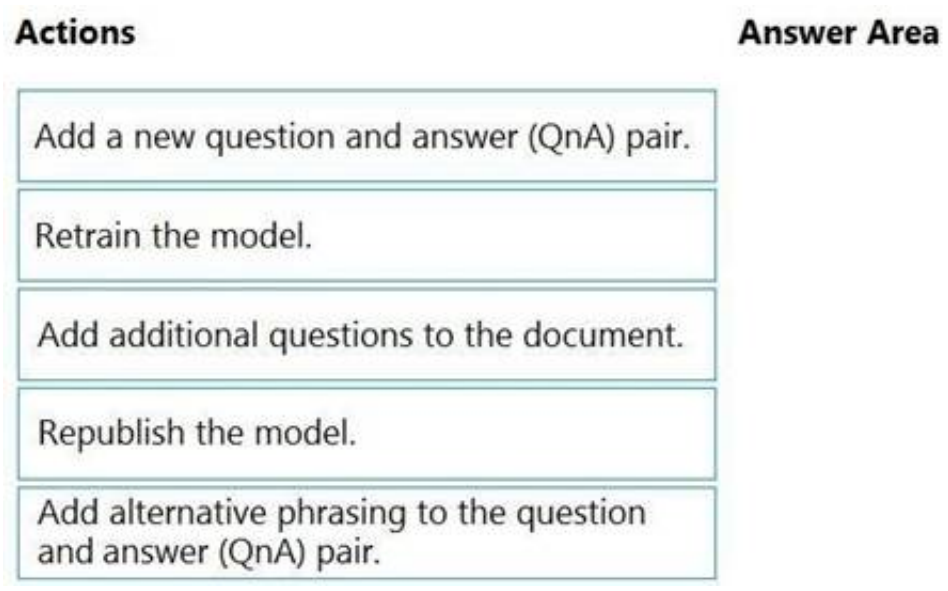

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Add alternative phrasing to the question and answer (QnA) pair.
Add alternate questions to an existing QnA pair to improve the likelihood of a match to a user query. Step 2: Retrain the model.
Periodically select Save and train after making edits to avoid losing changes. Step 3: Republish the model
Note: A knowledge base consists of question and answer (QnA) pairs. Each pair has one answer and a pair contains all the information associated with that answer.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/edit-knowledge-base

NEW QUESTION 8
- (Exam Topic 2)
You are creating an enrichment pipeline that will use Azure Cognitive Search. The knowledge store contains unstructured JSON data and scanned PDF documents that contain text.
Which projection type should you use for each data type? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

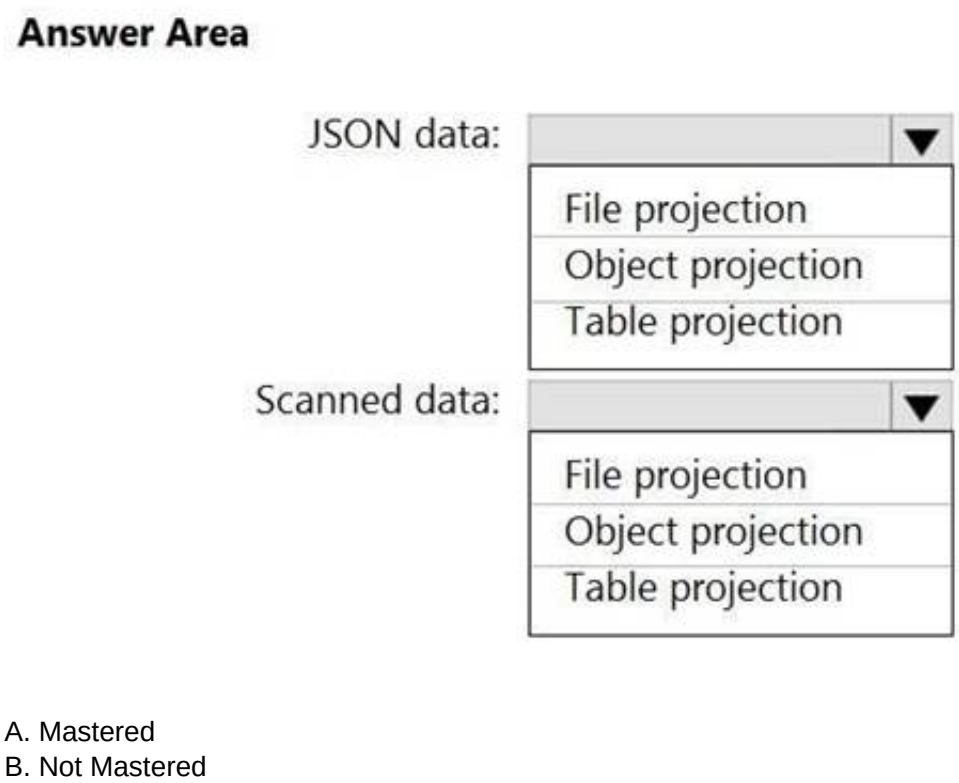

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Object projection
Object projections are JSON representations of the enrichment tree that can be sourced from any node. Box 2: File projection
File projections are similar to object projections and only act on the normalized_images collection. Reference:
https://docs.microsoft.com/en-us/azure/search/knowledge-store-projection-overview

NEW QUESTION 9
- (Exam Topic 2)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an application to identify species of flowers by training a Custom Vision model. You receive images of new flower species.
You need to add the new images to the classifier.
Solution: You create a new model, and then upload the new images and labels. Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:
The model needs to be extended and retrained.

NEW QUESTION 10
- (Exam Topic 2)
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 10
- (Exam Topic 2)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You develop an application to identify species of flowers by training a Custom Vision model. You receive images of new flower species.
You need to add the new images to the classifier.
Solution: You 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 12
- (Exam Topic 2)
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.

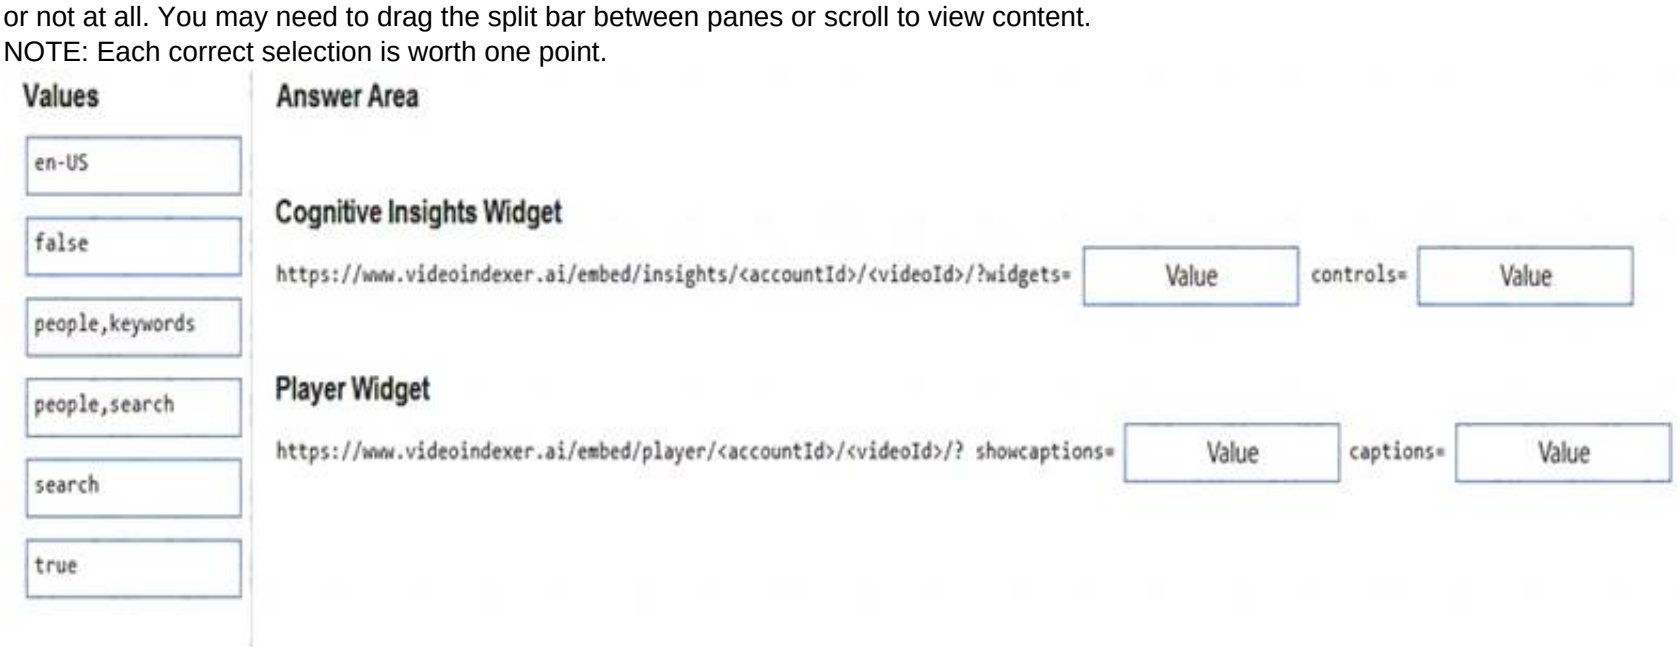

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 15
- (Exam Topic 2)
You are developing a new sales system that will process the video and text from a public-facing website. You plan to 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: BD

Explanation:
AI systems should treat all people fairly.
AI systems should perform reliably and safely. Reference:
https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/responsible-ai

NEW QUESTION 19
- (Exam Topic 2)
You are developing a method that uses the Computer Vision client library. The method will perform optical character recognition (OCR) in images. The method has the following code.

```
public static async Task ReadFileUrl(ComputerVisionClient client, string urlFile)
{
        const int numberOfCharsInOperationId = 36;

        var txtHeaders = await client.ReadAsync(urlFile, language: "en");

        string opLocation = textHeaders.OperationLocation;
        string operationId = opLocation.Substring(opLocation.Length -
numberOfCharsInOperationId);

        ReadOperationResult results;

        results = await client.GetReadResultAsync(Guid.Parse(operationId));

        var textUrlFileResults = results.AnalyzeResult.ReadResults;
        foreach (ReadResult page in textUrlFileResults)
        {
                foreach (Line line in page.Lines)
                {
                        Console.WriteLine(line.Text);
                }
        }
}
```

During testing, you discover that the call to the GetReadResultAsync method occurs before the read operation is complete.
You need to prevent the GetReadResultAsync method from proceeding until the read operation is complete. Which two actions should you perform? Each correct answer presents part of the solution. (Choose two.) NOTE: Each correct selection is worth one point.

A. Remove the Guid.Parse(operationId) parameter.
B. Add code to verify the results.Status value.
C. Add code to verify the status of the txtHeaders.Status value.
D. Wrap the call to GetReadResultAsync within a loop that contains a delay.

Answer: BD

Explanation:

Example code : do
{
results = await client.GetReadResultAsync(Guid.Parse(operationId));
}
while ((results.Status == OperationStatusCodes.Running || results.Status == OperationStatusCodes.NotStarted)); Reference:
https://github.com/Azure-Samples/cognitive-services-quickstart-code/blob/master/dotnet/ComputerVision/Comp

NEW QUESTION 23
- (Exam Topic 2)
You are developing a service that records lectures given in English (United Kingdom).
You have a method named AppendToTranscriptFile that takes translated text and a language identifier.
You need to develop code that will provide transcripts of the lectures to attendees in their respective language. The supported languages are English, French, Spanish, and German.
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

```
static async Task TranslateSpeechAsync()
{
    var config =SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");

    var lang = new List<string> [ ]
                                  ("en-GB")
                                  ("fr", "de", "es")
                                  ("French", "Spanish", "German")
                                  (languages)

    config.SpeechRecognitionLanguage = "en-GB";
    lang.ForEach(config.AddTargetLanguage);

    using var audioConfig = AudioConfig.FromDefaultMicrophoneInput();
    using var recognizer = new [ ] (config, audioConfig);
                                  IntentRecognizer
                                  SpeakerRecognizer
                                  SpeechSynthesizer
                                  TranslationRecognizer

    var result = await recognizer.RecognizeOnceAsync();
    if (result.Reason == ResultReason.TranslatedSpeech)
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: {"fr", "de", "es"}
A common task of speech translation is to specify target translation languages, at least one is required but multiples are supported. The following code snippet sets both French and German as translation language targets.
static async Task TranslateSpeechAsync()
{
var translationConfig =
SpeechTranslationConfig.FromSubscription(SPEECH SUBSCRIPTION KEY, SPEECH SERVICE REGION);
translationConfig.SpeechRecognitionLanguage = "it-IT";
// Translate to languages. See, https://aka.ms/speech/sttt-languages translationConfig.AddTargetLanguage("fr"); translationConfig.AddTargetLanguage("de");
}
Box 2: TranslationRecognizer
After you've created a SpeechTranslationConfig, the next step is to initialize a TranslationRecognizer. Example code:
static async Task TranslateSpeechAsync()
{
var translationConfig =
SpeechTranslationConfig.FromSubscription(SPEECH SUBSCRIPTION KEY, SPEECH SERVICE REGION);
var fromLanguage = "en-US";
var toLanguages = new List<string> { "it", "fr", "de" }; translationConfig.SpeechRecognitionLanguage = fromLanguage;
toLanguages.ForEach(translationConfig.AddTargetLanguage);
using var recognizer = new TranslationRecognizer(translationConfig);
}

NEW QUESTION 27
- (Exam Topic 2)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.
A conversational expert provides you with the following list of phrases to use for training. Find contacts in London.
Who do I know in Seattle? Search for contacts in Ukraine.
You need to implement the phrase list in Language Understanding. Solution: You create a new 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 32
- (Exam Topic 2)
You are building a bot and that will use Language Understanding. You have a LUDown file that contains the following content.

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

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

## Logout
- forget me
- log out

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

## SelectNone
- none

@ 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:
Graphical user interface, text, application, email Description automatically generated
Reference:
https://github.com/solliancenet/tech-immersion-data-ai/blob/master/ai-exp1/README.md

NEW QUESTION 34
- (Exam Topic 2)
You are developing an application that includes language translation.
The application will translate text retrieved by using a function named getTextToBeTranslated. The text can be in one of many languages. The content of the text must remain within the Americas Azure geography.
You need to develop code to translate the text to a single language.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
. . .
var endpoint = [ ▼ ] ;
```

- "https://api.cognitive.microsofttranslator.com/translate";
- "https://api.cognitive.microsofttranslator.com/transliterate";
- "https://api-apc.cognitive.microsofttranslator.com/detect";
- "https://api-nam.cognitive.microsofttranslator.com/detect";
- "https://api-nam.cognitive.microsofttranslator.com/translate";

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

- var uri = endpoint + "?from=en";
- var uri = endpoint + "?suggestedFrom=en";
- var uri = endpoint + "?to=en";

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

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 35
- (Exam Topic 2)
You plan to use containerized versions of the Anomaly Detector API on local devices for testing and in on-premises datacenters.
You need to ensure that the containerized deployments meet the following requirements:

- Prevent billing and API information from being stored in the command-line histories of the devices that run the container.
- Control access to the container images by using Azure role-based access control (Azure RBAC). Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. (Choose four.)

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

Actions

- Create a custom Dockerfile.
- Pull the Anomaly Detector container image.
- Distribute a `docker run` script.
- Push the image to an Azure container registry.
- Build the image.
- Push the image to Docker Hub.

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Pull the Anomaly Detector container image.
Step 2: Create a custom Dockerfile
Step 3: Push the image to an Azure container registry.
To push an image to an Azure Container registry, you must first have an image.
Step 4: Distribute the docker run script
Use the docker run command to run the containers. Reference:
https://docs.microsoft.com/en-us/azure/container-registry/container-registry-intro

NEW QUESTION 37

- (Exam Topic 2)
You are developing a streaming Speech to Text solution that will use the Speech SDK and MP3 encoding. You need to develop a method to convert speech to text for streaming MP3 data.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

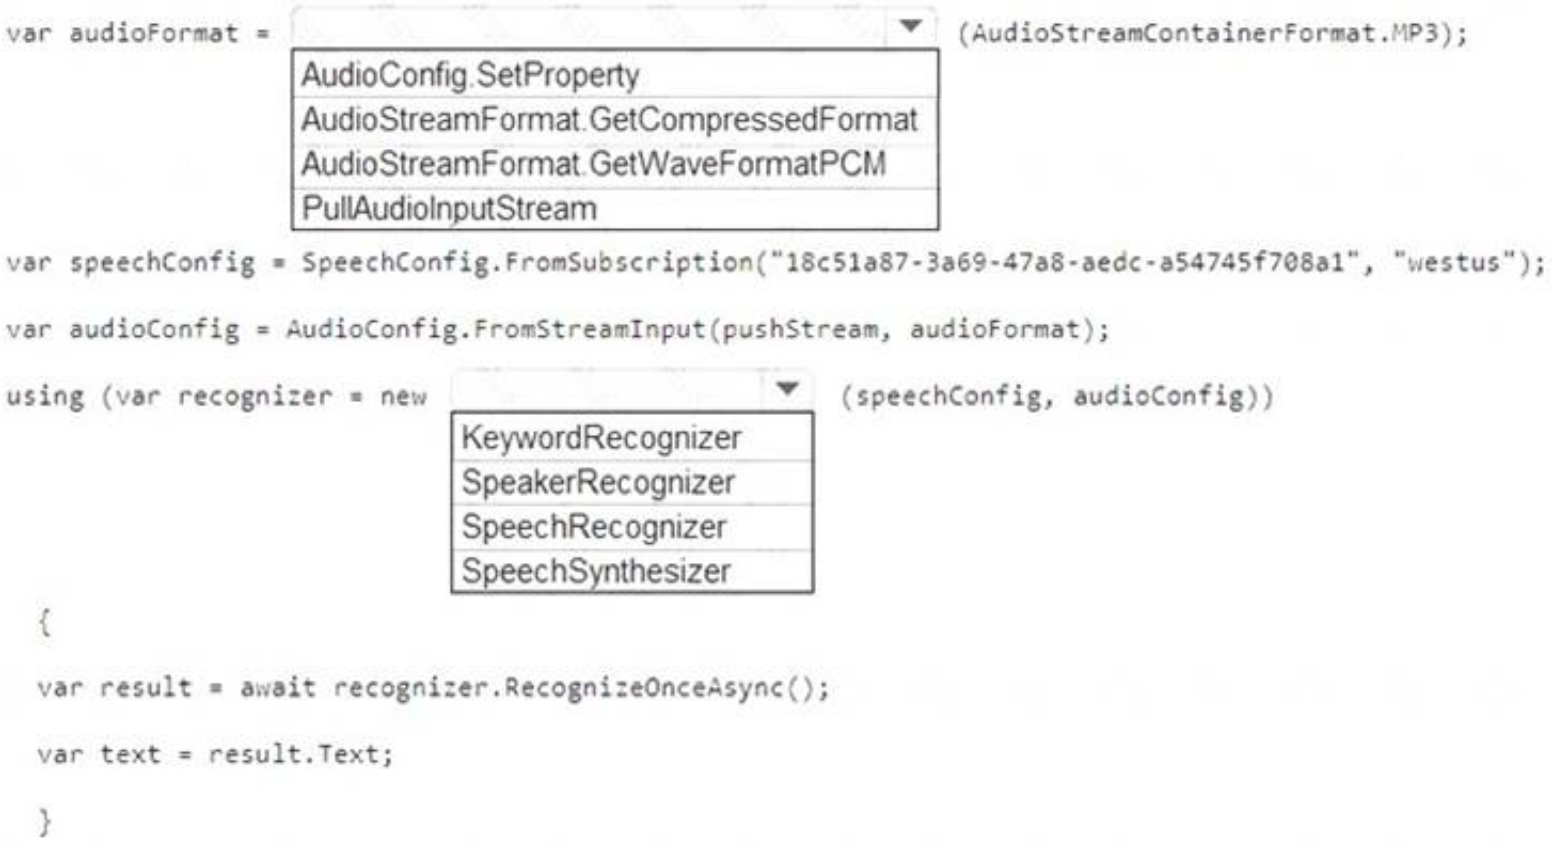

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Graphical user interface, text, application, email Description automatically generated
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/how-to-use-codec-compressed-audio-i

NEW QUESTION 42
- (Exam Topic 2)
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.

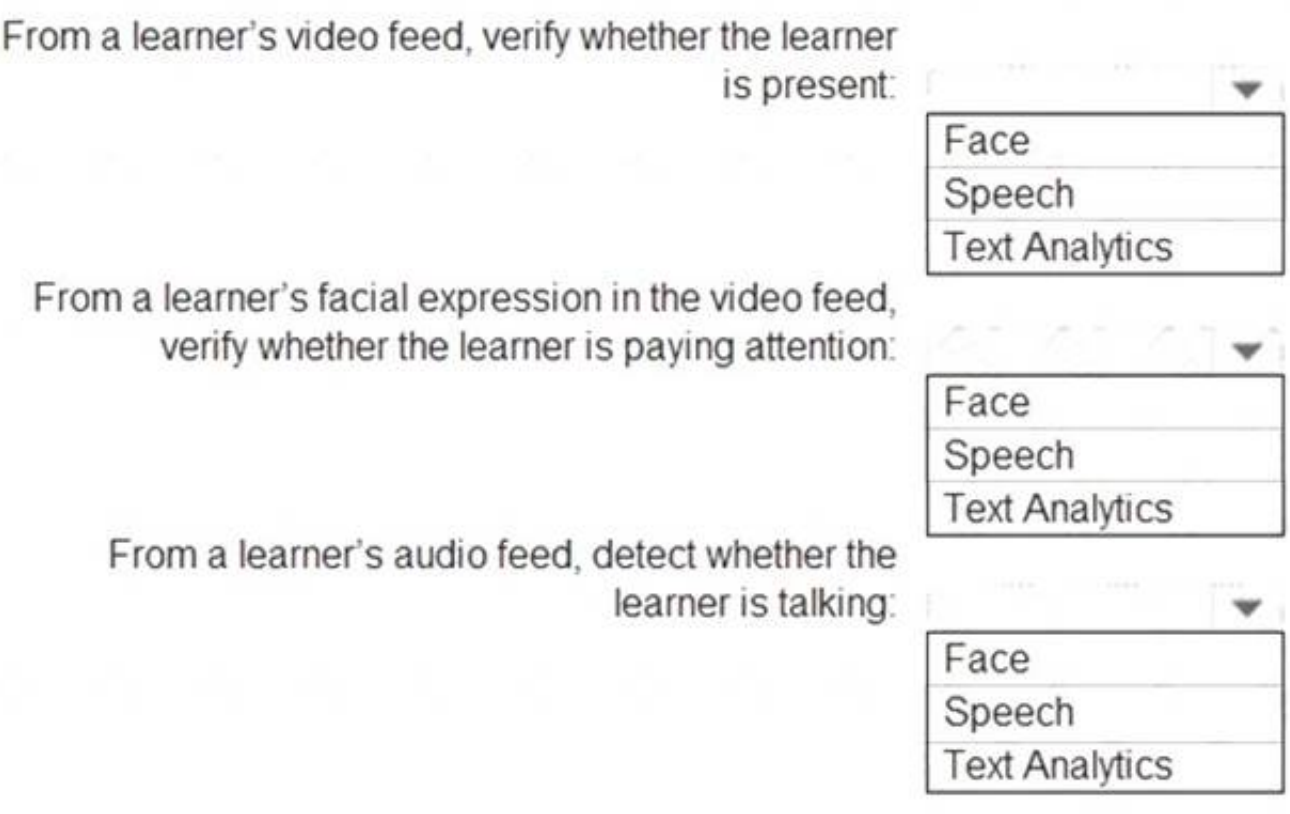

A. Mastered
B. Not Mastered

Answer: A

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

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/what-are-cognitive-services

NEW QUESTION 47
- (Exam Topic 2)
A customer uses Azure Cognitive Search.
The customer plans to enable a server-side encryption and use customer-managed keys (CMK) stored in Azure.
What are three implications of the planned change? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

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

Answer: ABE

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/search/search-security-overview

NEW QUESTION 52
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons 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 FirstTry

* AI-102 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year