Microsoft

Exam Questions AI-102

Designing and Implementing an Azure AI Solution

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

  All examinations will be up to date.

* 24/7 Quality Support

  We will provide service round the clock.

* 100% Pass Rate

  Our guarantee that you will pass the exam.

* Unique Gurantee

  If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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.

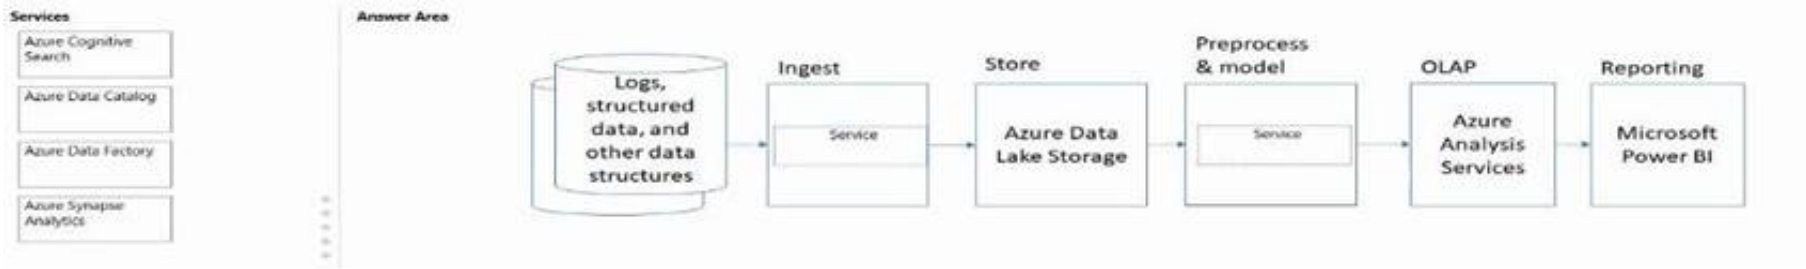

A. Mastered
B. Not Mastered

Answer: A

Explanation:

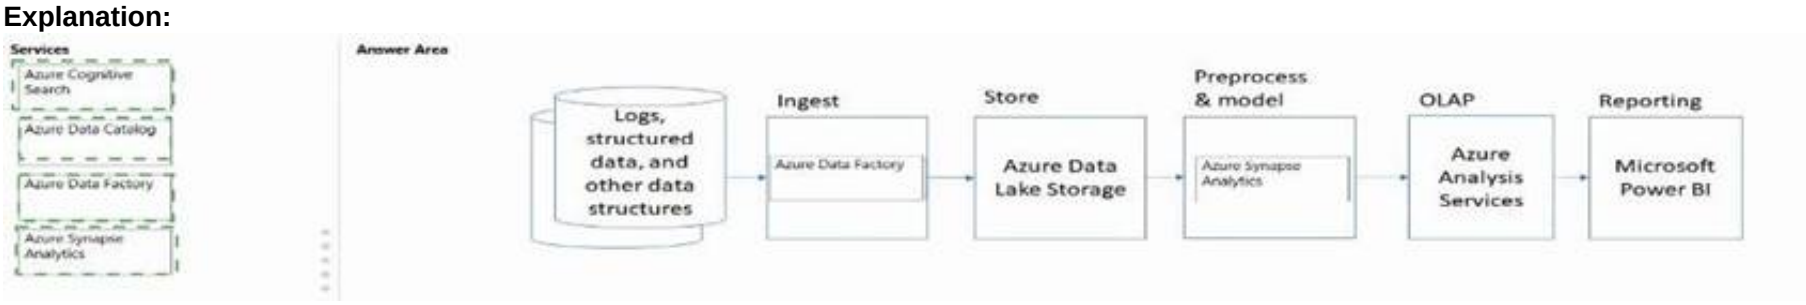

NEW QUESTION 2
HOTSPOT - (Topic 3)
You are building content for a video training solution.
You need to create narration to accompany the video content. The solution must use Custom Neural Voice.
What should you use to create a custom neural voice, and which service should you use to generate the narration? To answer, select the appropriate options in the answer area.
NOTE: Each correct answer is worth one point.

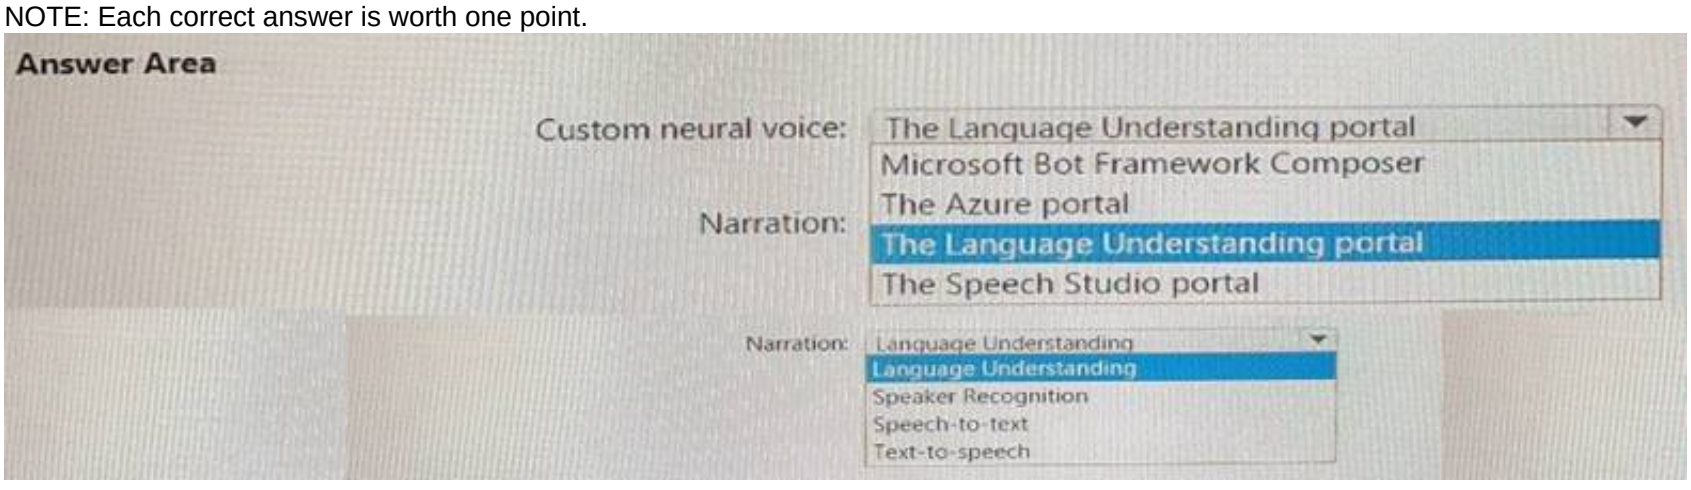

A. Mastered
B. Not Mastered

Answer: A

Explanation:

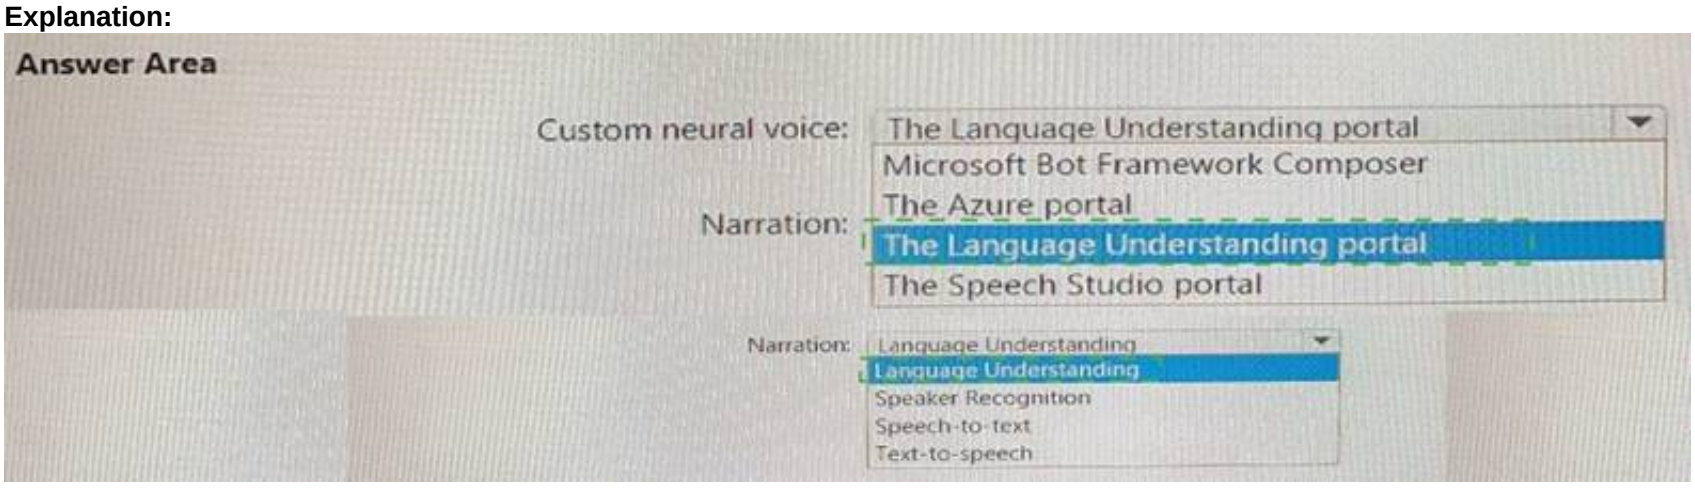

NEW QUESTION 3
HOTSPOT - (Topic 3)
Select the answer that correctly completes the sentence.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

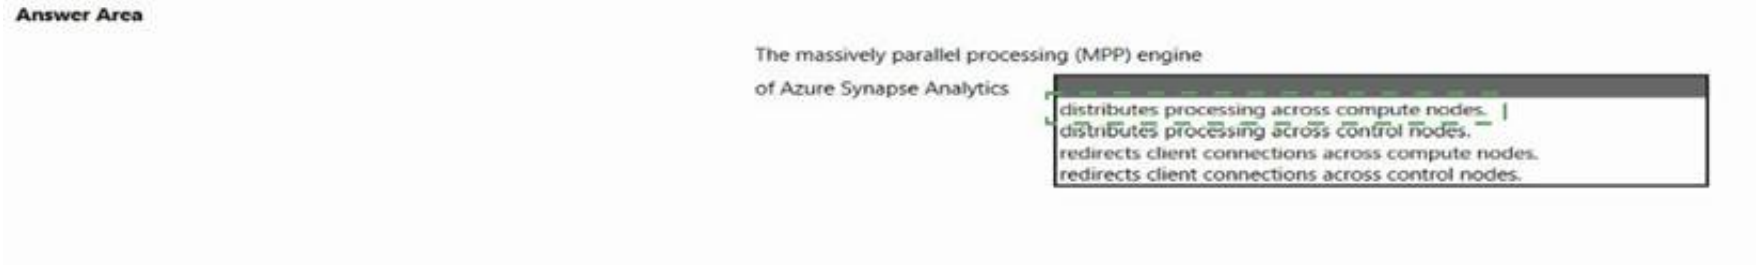

NEW QUESTION 4
HOTSPOT - (Topic 3)
You have an Azure Cognitive Search resource named Search 1 that is used by multiple apps You need to secure Search 1. The solution must meet the following requirements:
• Prevent access to Search1 from the internet.
• Limit the access of each app to specific queries
What should you do? 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
- (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 fram the model. Which type of custom entity should you use?

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

Answer: C

NEW QUESTION 6
- (Topic 3)
Your company needs to implement a relational database in Azure. The solution must minimize ongoing maintenance. Which Azure service should you use?

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

Answer: B

NEW QUESTION 7
HOTSPOT - (Topic 3)
For each of the following statements, select Yes if the statement is tine. 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.	○	○
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	○	○
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	○	○

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.	○	○
Batch processing must occur immediately and have latency in the order of seconds or milliseconds.	○	○
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.	○	○

NEW QUESTION 8
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

- Train the application.
- Publish the application.
- Add a new application.
- Add example utterances.
- Add the prebuilt domain ToDo.

Answer Area

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 9
- (Topic 3)
You build a conversational bot named bot1.
You need to configure the bot to use a QnA Maker application.

From the Azure Portal, where can you find the information required by bot1 to connect to the QnA Maker application?

A. Access control (IAM)
B. Properties
C. Keys and Endpoint
D. Identity

Answer: C

Explanation:
Obtain values to connect your bot to the knowledge base 1.In the QnA Maker site, select your knowledge base. 2. With your knowledge base open, select the SETTINGS tab. Record the value shown for service name. This value is useful for finding your knowledge base of interest when using the QnA Maker portal interface. It's not used to connect your bot app to this knowledge base.3. Scroll down to find Deployment details and record the following values from the Postman sample HTTP request:4. POST /knowledgebases/<knowledge-base-id>/generateAnswer 5.Host: <your-host-url>6.Authorization: EndpointKey <your-endpoint-key>
Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-qna

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 build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.
A conversational expert provides you with the following list of phrases to use for training.
Find contacts in London. Who do I know in Seattle? Search for contacts in Ukraine.
You need to implement the phrase list in Language Understanding. Solution: You create a new intent for location.
Does this meet the goal?

A. Yes
B. No

Answer: A

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

NEW QUESTION 10
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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

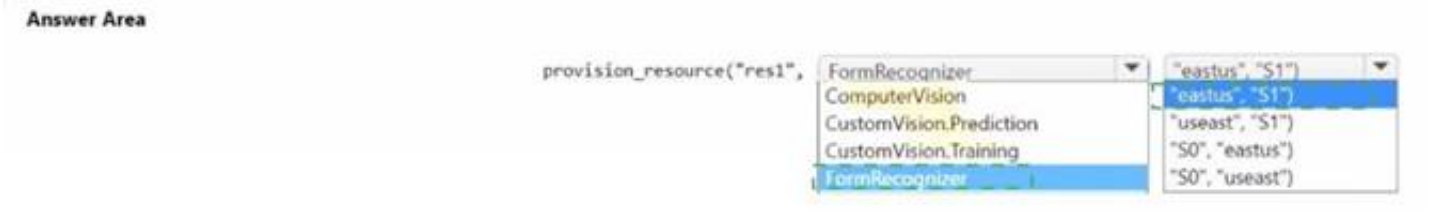

NEW QUESTION 15
- (Topic 3)
You are building an Al solution that will use Sentiment Analysis results from surveys to calculate bonuses for customer service staff. You need to ensure that the solution meets the Microsoft responsible Al principles. What should you do?

A. Add a human review and approval step before making decisions that affect the staffs financial situation
B. Include the Sentiment Analysis results when surveys return a low confidence score.
C. Use all the surveys, including surveys by customers who requested that their account be deleted and their data be removed.
D. Publish the raw survey data to a central location and provide the staff with access to the location.

Answer: A

NEW QUESTION 16
- (Topic 3)
You have an Azure subscription that contains a Language service resource named ta1 and a virtual network named vnet1. You need to ensure that only resources in vnet1 can access ta1. What should you configure?

A. a network security group (NSG) for vnet1
B. Azure Firewall for vnet1
C. the virtual network settings for ta 1
D. a Language service container for ta1

Answer: C

NEW QUESTION 19
- (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 24
- (Topic 3)
You need to build a chatbot that meets the following requirements:
? Supports chit-chat, knowledge base, and multilingual models
? Performs sentiment analysis on user messages
? Selects the best language model automatically
What should you integrate into the chatbot?

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

Answer: C

Explanation:
Language Understanding: An AI service that allows users to interact with your applications, bots, and IoT devices by using natural language.
QnA Maker is a cloud-based Natural Language Processing (NLP) service that allows you to create a natural conversational layer over your data. It is used to find the most appropriate answer for any input from your custom knowledge base (KB) of information.
Text Analytics: Mine insights in unstructured text using natural language processing (NLP)—no machine learning expertise required. Gain a deeper understanding of customer opinions with sentiment analysis. The Language Detection feature of the Azure Text Analytics REST API evaluates text input
Reference:
https://azure.microsoft.com/en-us/services/cognitive-services/text-analytics/ https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/overview/overview

NEW QUESTION 29
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.

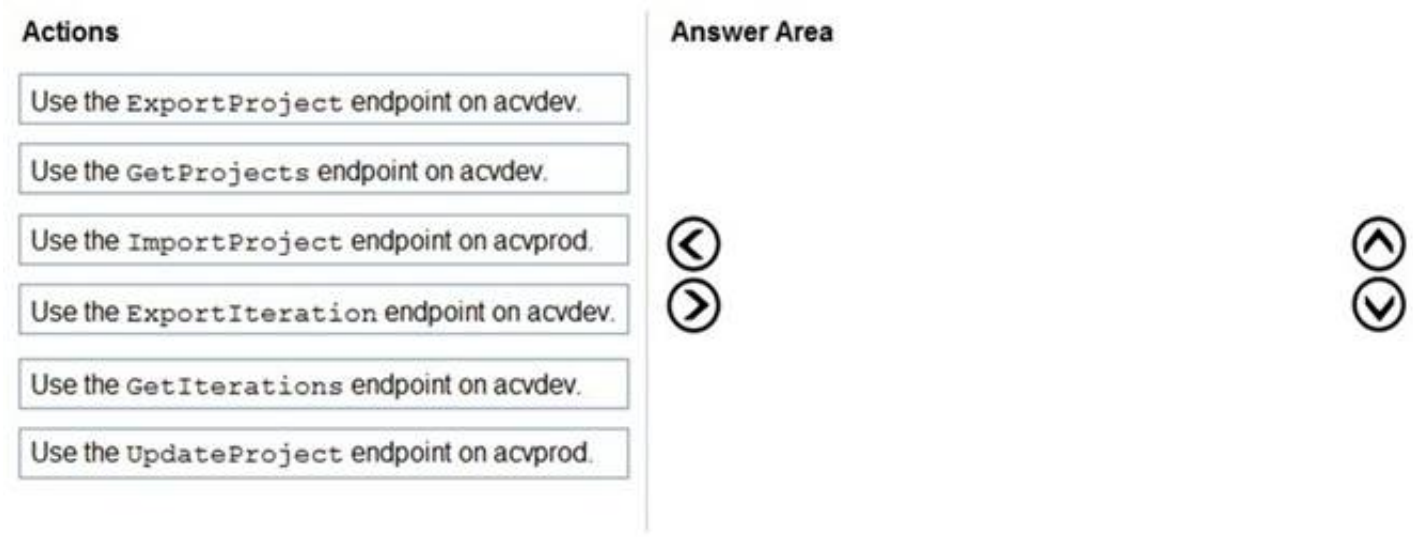

A. Mastered
B. Not Mastered

Answer: A

Explanation:

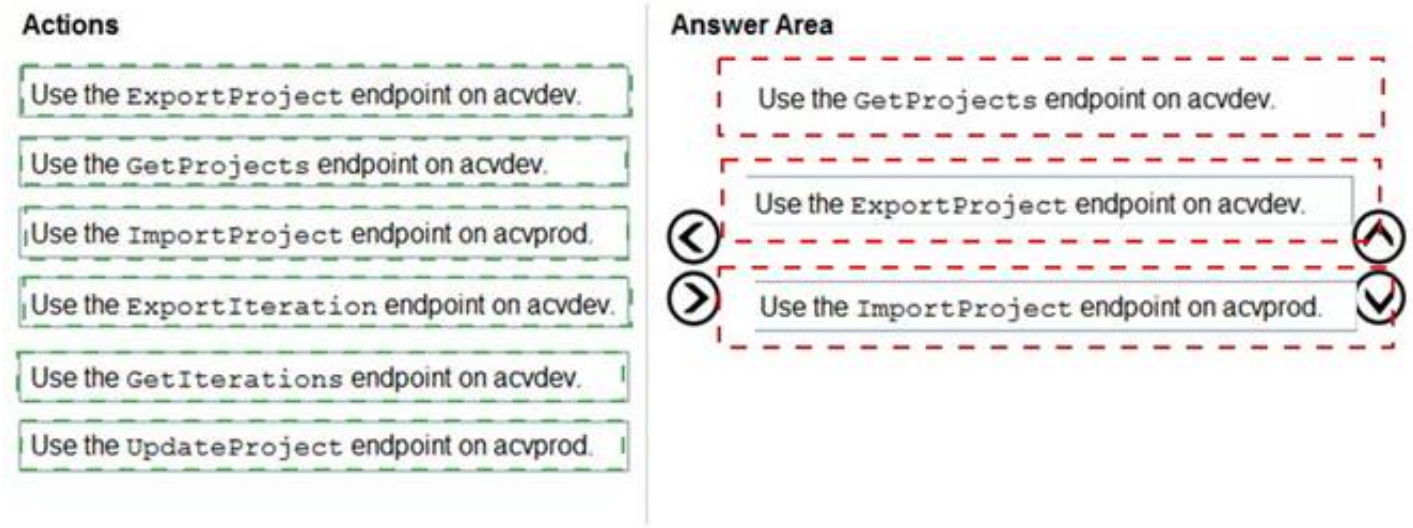

NEW QUESTION 30
- (Topic 3)
What should you use to build a Microsoft Power Bi paginated report?

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

Answer: A

NEW QUESTION 35
- (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 39
- (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 44
- (Topic 3)
Which property of a transactional workload guarantees that each transaction is treated as a single unit that either succeeds completely or tails completely?

A. isolation
B. atomicity
C. consistency
D. durability

Answer: B

NEW QUESTION 48
- (Topic 3)
You have an Azure Cognitive Search solution and a collection of handwritten letters stored as JPEG files.
You plan to index the collection. The solution must ensure that queries can be performed on the contents of the letters.
You need to create an indexer that has a skillset. Which skill should you include?

A. key phrase extraction
B. optical character recognition (OCR)
C. document extraction
D. image analysis

Answer: B

NEW QUESTION 52
- (Topic 3)
Your company wants to reduce how long it takes for employees to log receipts in expense reports. All the receipts are in English.
You need to extract top-level information from the receipts, such as the vendor and the transaction total. The solution must minimize development effort.
Which Azure Cognitive Services service should you use?

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

Answer: B

Explanation:
Azure Form Recognizer is a cognitive service that lets you build automated data processing software using machine learning technology. Identify and extract text, key/value pairs, selection marks, tables, and structure from your documents—the service outputs structured data that includes the relationships in the original file, bounding boxes, confidence and more.
Form Recognizer is composed of custom document processing models, prebuilt models for invoices, receipts, IDs and business cards, and the layout model.
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/form-recognizer

NEW QUESTION 56
HOTSPOT - (Topic 3)
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.

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

NEW QUESTION 61
- (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 64
- (Topic 3)
You have a text-based chatbot.
You need to enable content moderation by using the Text Moderation API of Content Moderator. Which two service responses should you use? Each correct answer presents part of the solution NOTE: Each correct selection is worth one point.

A. the adult classification score
B. optical character recognition (OCR)
C. personal data
D. text classification
E. the racy classification score

Answer: AD

NEW QUESTION 68
- (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 to a new virtual network, and you configure Azure Private Link.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/cognitive-services-virtual-networks?tabs=portal#use-private-endpoints

NEW QUESTION 69
HOTSPOT - (Topic 3)
Select the answer that correctly completes the sentence.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 74
DRAG DROP - (Topic 3)
You are building an app that will scan confidential documents and use the Language service to analyze the contents.
You provision an Azure Cognitive Services resource.
You need to ensure that the app can make requests to the Language service endpoint. The solution must ensure that confidential documents remain on-premises.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 75
- (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 spurious
questions
You need to ensure that the chatbot provides formal responses to spurious questions. Solution: From Language Studio, you change the chitchat source to qna_chitchat_professional.tsv. and then retrain and republish the model. Does this meet the goal?

A. Yes
B. No

Answer: A

NEW QUESTION 78
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.

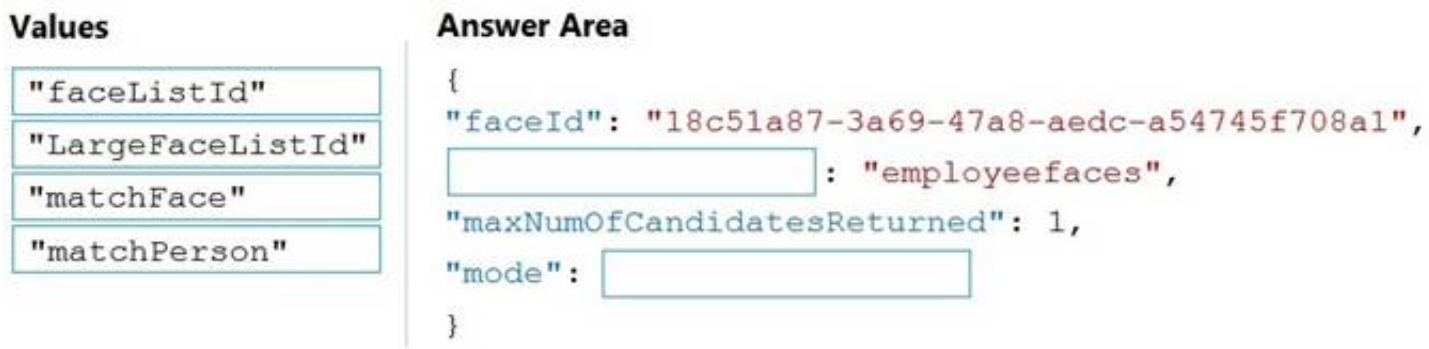

A. Mastered
B. Not Mastered

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 82
- (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 84
- (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 87
DRAG DROP - (Topic 3)
You develop an app in O named App1 that performs speech-to-speech translation. You need to configure App1 to translate English to German.
How should you complete the speechTransiationConf ig object? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

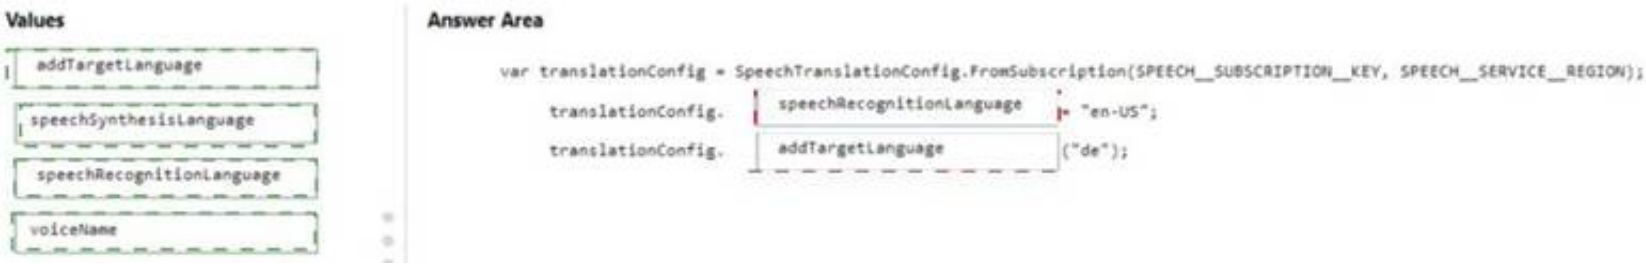

NEW QUESTION 88
DRAG DROP - (Topic 3)
You are developing a webpage that will use the Video Indexer service to display videos of internal company meetings.
You embed the Player widget and the Cognitive Insights widget into the page. You need to configure the widgets to meet the following requirements:
? 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.

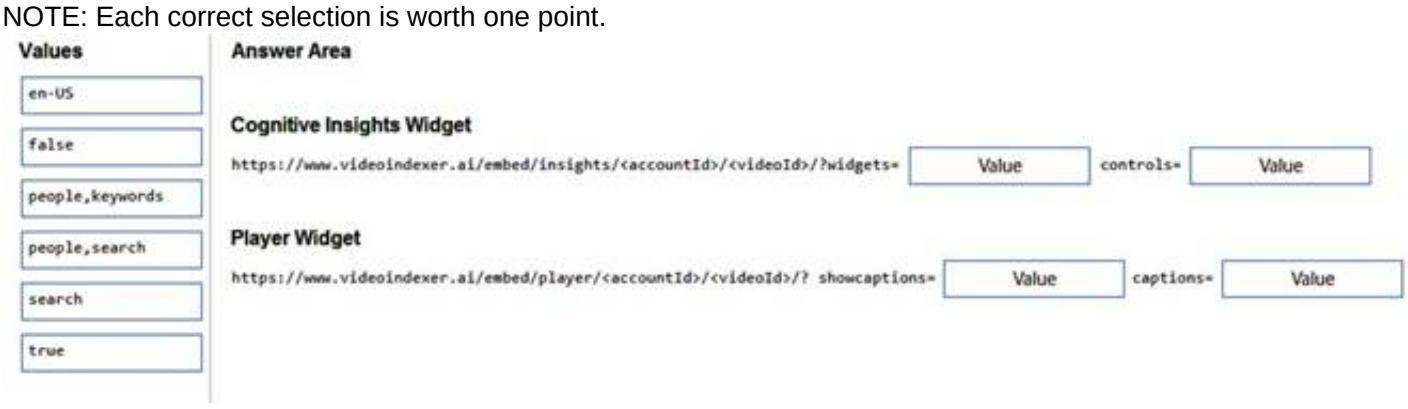

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 93
- (Topic 3)
You are building a chatbot that will use question answering in Azure Cognitive Service for Language.
You have a PDF named Docl.pdf that contains a product catalogue and a price list You upload Docl.pdf and train the model.
During testing, users report that the chatbot responds correctly to the following question: What is the price of < product>?
The chatbot fails to respond to the following question: How much does <product* cost? You need to ensure that the chatbot responds correctly to both questions.
Solution: From Language Studio, you create an entity for cost, and then retrain and republish the model.
Does this meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 96
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.

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 97
- (Topic 3)
You need to measure the public perception of your brand on social media messages. Which Azure Cognitive Services service should you use?

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

Answer: A

Explanation:
Text Analytics Cognitive Service could be used to quickly determine the public perception for a specific topic, event or brand.
Example: A NodeJS app which pulls Tweets from Twitter using the Twitter API based on a specified search term. Then pass these onto Text Analytics for sentiment scoring before storing the data and building a visualisation in PowerBI. The Architecture looked something like this:

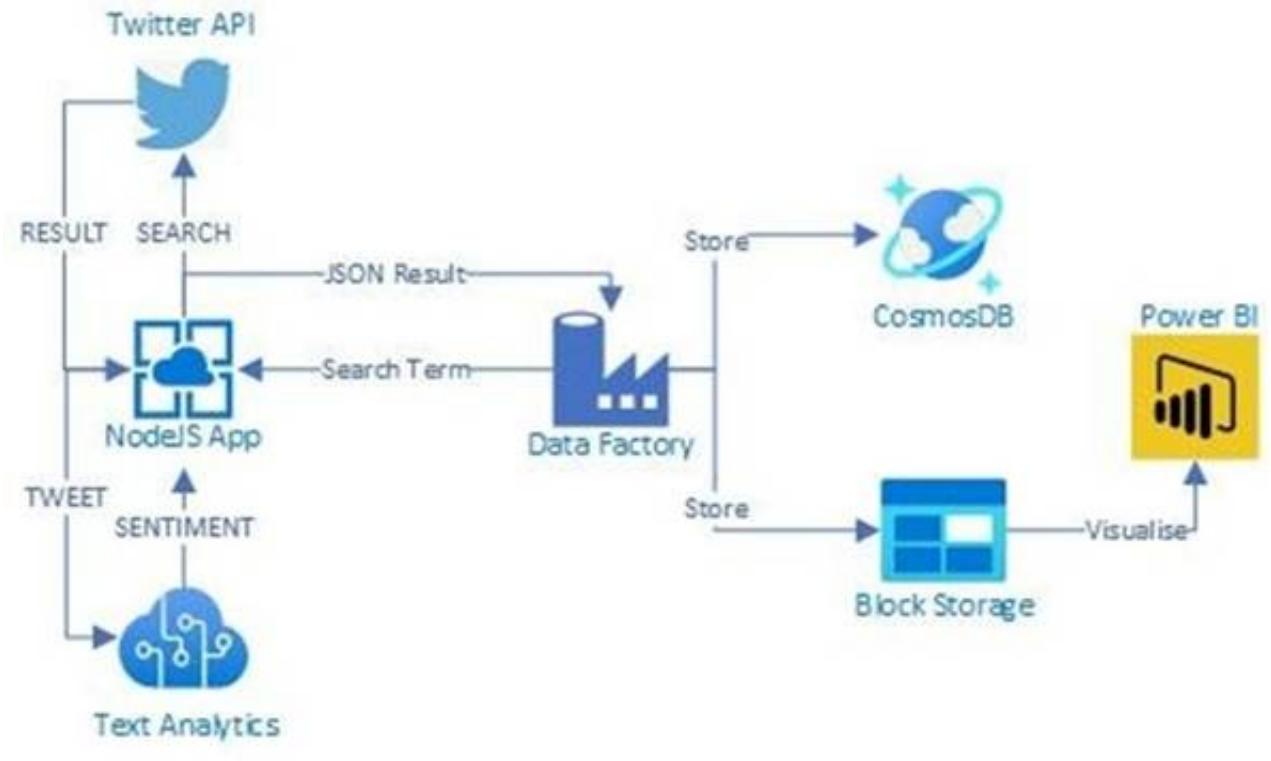

Reference:
https://www.linkedin.com/pulse/measuring-public-perception-azure-cognitive-services- steve-dalai

NEW QUESTION 100
DRAG DROP - (Topic 3)
You need to develop an automated call handling system that can respond to callers in their own language. The system will support only French and English.
Which Azure Cognitive Services service should you use to meet each requirement? To answer, drag the appropriate services to the correct requirements. Each service may be used once, more than once, or not at all. You may need to drag the split bat between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

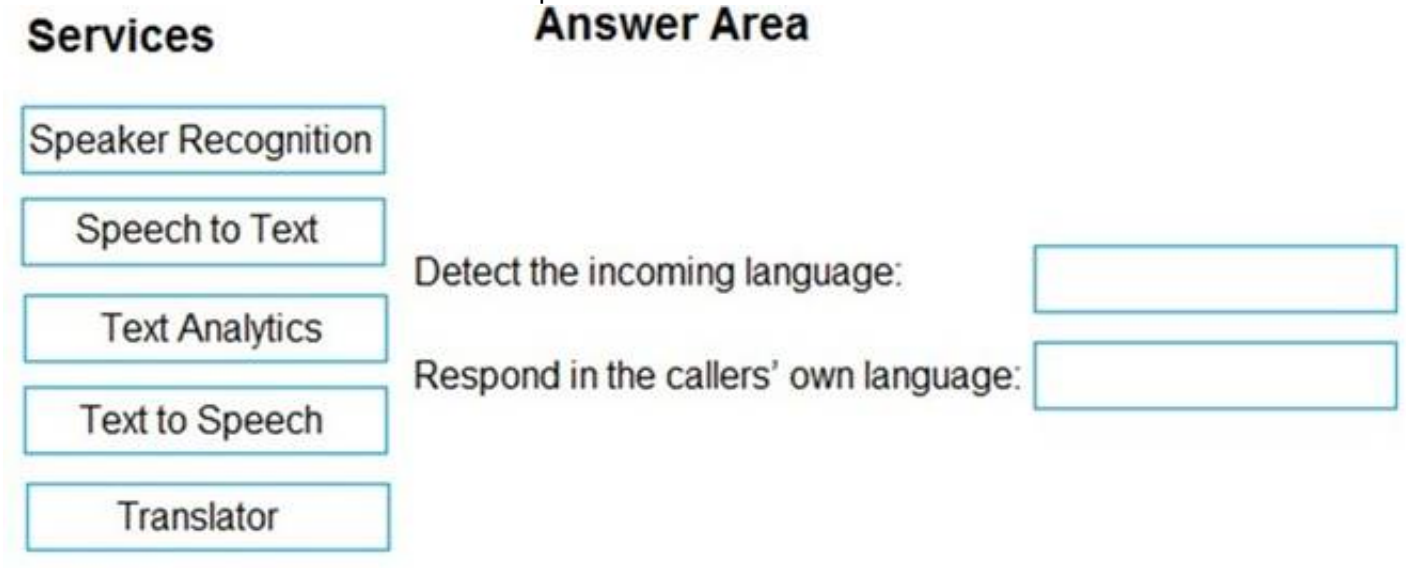

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Box 1: Text Analytics
The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.
Box 2: Translator
Translator is a cloud-based neural machine translation service that is part of the Azure Cognitive Services family of REST APIs. Translator can be used with any operating system and powers many Microsoft products and services used by thousands of businesses worldwide to perform language translation and other language-related operations.

NEW QUESTION 102
- (Topic 3)
You are building a bot by using Microsoft Bot Framework.
You need to configure the bot to respond to spoken requests. The solution must minimize development effort.
What should you do?

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

Answer: B

NEW QUESTION 105
- (Topic 3)
You have a SQL query that combines customer data and order data. The query includes calculated columns. You need to create a database object that would allow other users to rerun the same SOL query. What should you create?

A. an Index
B. a view
C. a scalar function
D. a table

Answer: B

NEW QUESTION 108
- (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 112
HOTSPOT - (Topic 3)
You are building a model to detect objects in images.
The performance of the model based on training data is shown in the following exhibit.

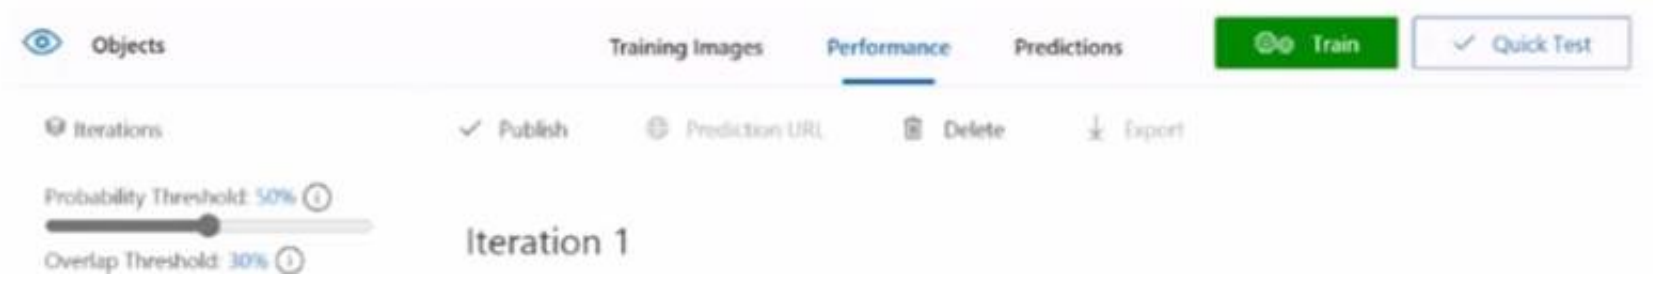

A. Mastered
B. Not Mastered

Answer: A

Explanation:

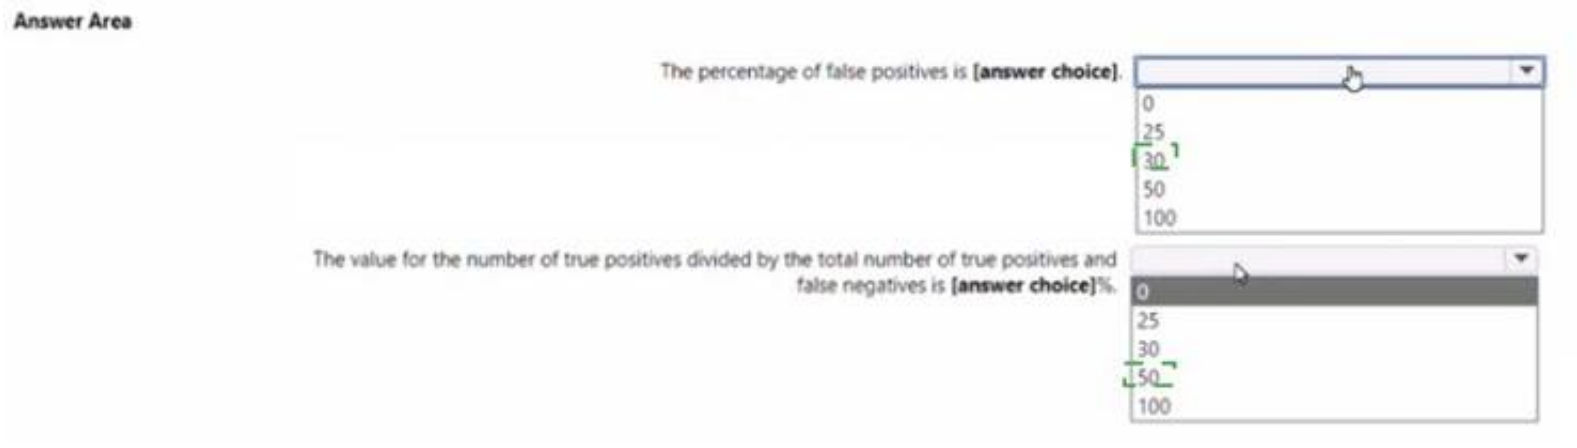

NEW QUESTION 113
DRAG DROP - (Topic 3)
You are building a transcription service for technical podcasts.
Testing reveals that the service fails to transcribe technical terms accurately. You need to improve the accuracy of the service.
Which five actions should you perform in sequence? To answer, move 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.

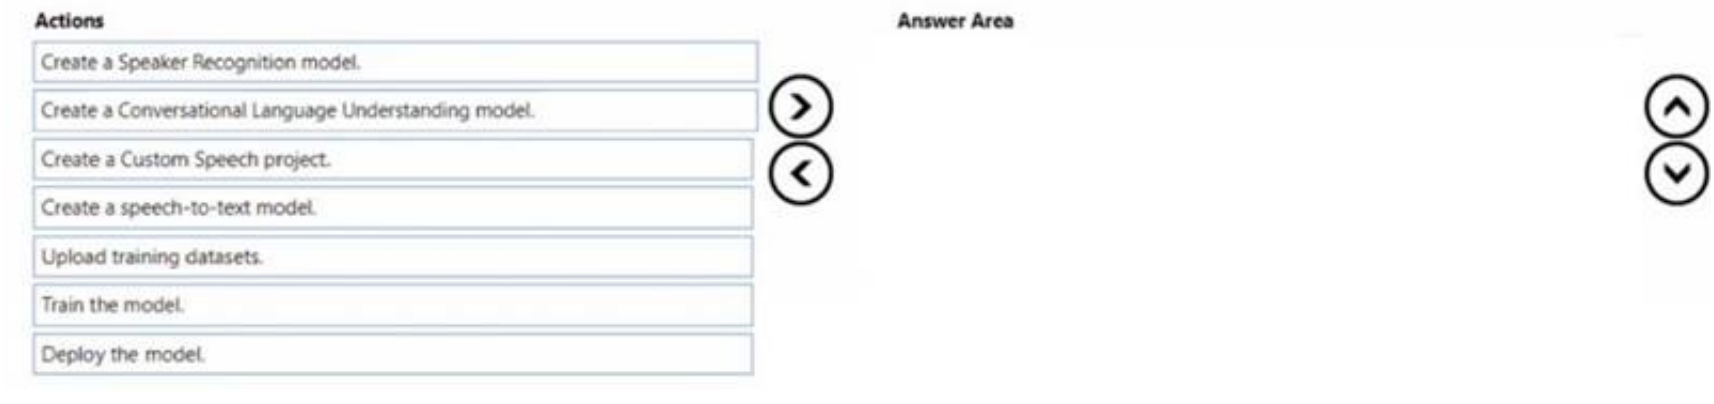

A. Mastered
B. Not Mastered

Answer: A

Explanation:
https://learn.microsoft.com/en-us/azure/cognitive-services/speech-service/custom-speech- overview#how-does-it-work
With Custom Speech, you can upload your own data, test and train a custom model, compare accuracy between models, and deploy a model to a custom endpoint.
- Create a project and choose a model. Use a Speech resource that you create in the Azure portal. If you will train a custom model with audio data, choose a Speech resource region with dedicated hardware for training audio data.
- Upload test data. Upload test data to evaluate the speech to text offering for your applications, tools, and products.
- Train a model. Provide written transcripts and related text, along with the corresponding audio data. Testing a model before and after training is optional but recommended.
- Deploy a model. Once you're satisfied with the test results, deploy the model to a custom endpoint. With the exception of batch transcription, you must deploy a custom endpoint to use a Custom Speech model.

NEW QUESTION 117
- (Topic 3)
You have an existing Azure Cognitive Search service.
You have an Azure Blob storage account that contains millions of scanned documents stored as images and PDFs.
You need to make the scanned documents available to search as quickly as possible. What should you do?

A. Split the data into multiple blob container
B. Create a Cognitive Search service for each containe

C. Within each indexer definition, schedule the same runtime execution pattern.
D. Split the data into multiple blob container
E. Create an indexer for each containe
F. Increase the search unit
G. Within each indexer definition, schedule a sequential execution pattern.
H. Create a Cognitive Search service for each type of document.
I. Split the data into multiple virtual folder
J. Create an indexer for each folde
K. Increase the search unit
L. Within each indexer definition, schedule the same runtime execution pattern.

Answer: D

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/search/search-howto-indexing-azure-blob-storage

NEW QUESTION 119
- (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 123
- (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 124
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
Azure Databricks is an Apache Spark-based analytics platform.	○	○
Azure Analysis Services is used for transactional workloads.	○	○
Azure Data Factory orchestrates data integration workflows.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 125
- (Topic 3)
You are building a multilingual chatbot.
You need to send a different answer for positive and negative messages.
Which two Text Analytics APIs should you use? Each correct answer presents part of the solution. (Choose two.)
NOTE: Each correct selection is worth one point.

A. Linked entities from a well-known knowledge base
B. Sentiment Analysis
C. Key Phrases
D. Detect Language
E. Named Entity Recognition

Answer: BD

Explanation:
B: The Text Analytics API's Sentiment Analysis feature provides two ways for detecting positive and negative sentiment. If you send a Sentiment Analysis request, the API will return sentiment labels (such as "negative", "neutral" and "positive") and confidence scores at the sentence and document-level.
D: The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.
This capability is useful for content stores that collect arbitrary text, where language is unknown. Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text- analytics-how-to- sentiment-analysis?tabs=version-3-1
https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text- analytics-how-to- language-detection

NEW QUESTION 127
HOTSPOT - (Topic 3)
You create a knowledge store for Azure Cognitive Search by using the following JSON.

```
"knowledgeStore" : {
        "storageConnectionString": "DefaultEndpointsProtocol=https;AccountName=<Acct Name>;AccountKey=<Acct Key>;",
        "projections": [
            {
                "tables": [
                    {
                        "tableName": "unrelatedDocument",
                        "generatedKeyName": "Documentid",
                        "source": "/document/pbiShape"
                    },
                    {
                        "tableName": "unrelatedKeyPhrases",
                        "generatedKeyName": "KeyPhraseid",
                        "source": "/document/pbiShape/keyPhrases"
                    }
                ],
                "objects": [

                    {
                        "storageContainer": "unrelatedocrlayout",
                        "source": null,
                        "sourceContext": "/document/normalized_images/*/layoutText",
                        "inputs": [
                            {
                                "name": "ocrLayoutText",
                                "source": "/document/normalized_images/*/layoutText"
                            }
                        ]
                    }
                ],
                "files": []
            }
        ]
}
```

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 128
HOTSPOT - (Topic 3)
You are developing an application that will use the Computer Vision client library. The application has the following code.

```
public async Task AnalyzeImage(ComputerVisionClient client, string localImage)
{
        List<VisualFeatureTypes> features = new List<VisualFeatureTypes>()
                {
                        VisualFeatureTypes.Description,
                        VisualFeatureTypes.Tags,
                };
        using (Stream imageStream = File.OpenRead(localImage))
        {
                try
                {
                  ImageAnalysis results = await client.AnalyzeImageInStreamAsync(imageStream, features);

                  foreach (var caption in results.Description.Captions)
                  {
                        Console.WriteLine($"{caption.Text} with confidence {caption.Confidence}");
                  }

                  foreach (var tag in results.Tags)
                  {
                        Console.WriteLine($"{tag.Name} {tag.Confidence}");
                  }
                }
                catch (Exception ex)
                {
                        Console.WriteLine(ex.Message);
                }
        }
}
```

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

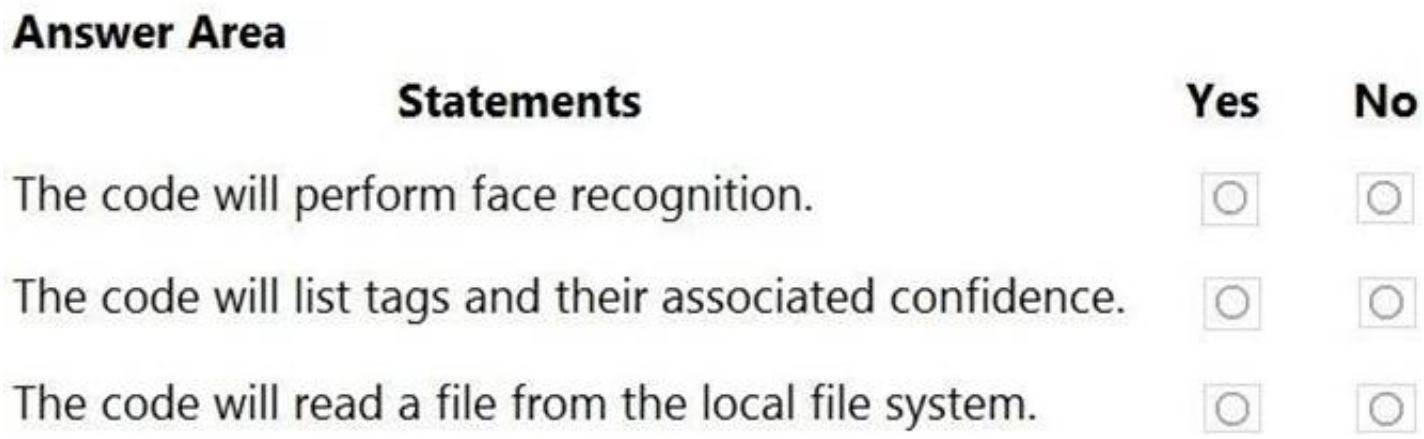

A. Mastered
B. Not Mastered

Answer: A

Explanation:

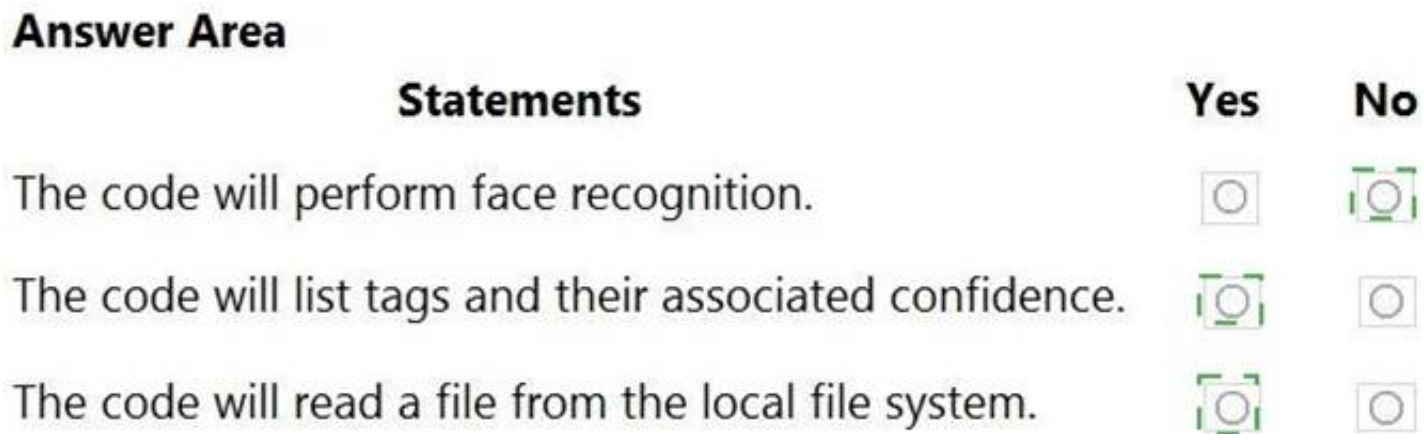

NEW QUESTION 132
- (Topic 3)
You need to measure the public perception of your brand on social media by using natural language processing. Which Azure service should you use?

A. Content Moderator
B. Form Recognizer

C. Computer Vision
D. Language service

Answer: D

NEW QUESTION 133
HOTSPOT - (Topic 3)
You are building a model that will be used in an iOS app.
You have images of cats and dogs. Each image contains either a cat or a dog.
You need to use the Custom Vision service to detect whether the images is of a cat or a dog.
How should you configure the project in the Custom Vision portal? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

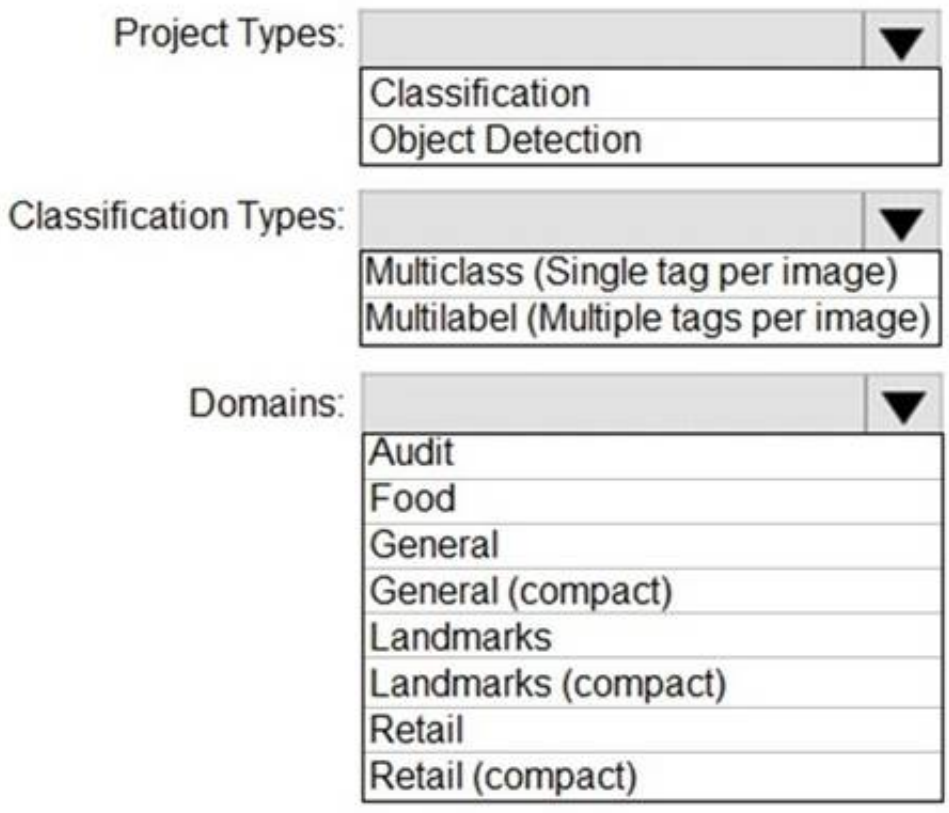

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Classification
Box 2: Multiclass
A multiclass classification project is for classifying images into a set of tags, or target labels. An image can be assigned to one tag only.
Box 3: General
General: Optimized for a broad range of image classification tasks. If none of the other specific domains are appropriate, or if you're unsure of which domain to choose, select one of the General domains.

NEW QUESTION 138
- (Topic 3)
Which database transaction property ensures that transactional changes to a database are preserved during unexpected operating system restarts?

A. durability
B. atomicity
C. consistency
D. isolation

Answer: A

NEW QUESTION 140
HOTSPOT - (Topic 3)
Select the answer that correctly completes the sentence.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 145
- (Topic 3)
You have an Azure subscription that contains an Azure Cognitive Service for Language resource. You need to identify the URL of the REST interface for the Language service. Which blade should you use in the Azure portal?

A. Identity
B. Keys and Endpoint
C. Properties
D. Networking

Answer: B

NEW QUESTION 149
- (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 150
HOTSPOT - (Topic 3)
Select the answer that correctly completes the sentence.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 151
- (Topic 3)
You build a Language Understanding model by using the Language Understanding portal. You export the model as a JSON file as shown in the following sample.

```
{
  "text": "average amount of rain by month at chicago last year",
  "intent": "Weather.CheckWeatherValue",
  "entities": [
    {
      "entity": "Weather.WeatherRange",
      "startPos": 0,
      "endPos": 6,
      "children": []
    },
    {
      "entity": "Weather.WeatherCondition",
      "startPos": 18,
      "endPos": 21,
      "children": []
    },
    {
      "entity": "Weather.Historic",
      "startPos": 23,
      "endPos": 30,
      "children": []
    }
  ]
}
```

To what does the Weather.Historic entity correspond in the utterance?

A. by month
B. chicago
C. rain
D. location

Answer: A

NEW QUESTION 153
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 157
- (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 162
- (Topic 2)
You need to develop an extract solution for the receipt images. The solution must meet the document processing requirements and the technical requirements.
You upload the receipt images to the From Recognizer API for analysis, and the API ret urns the following JSON.

```
"documentResults":[
      {
        "docType":"prebuilt:receipt",
        "pageRange":[
          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 163
- (Topic 2)
You are developing the knowledgebase by using Azure Cognitive Search.
You need to meet the knowledgebase requirements for searching equivalent terms. What should you include in the solution?

A. a synonym map
B. a suggester
C. a custom analyzer
D. a built-in key phrase extraction skill

Answer: A

Explanation:
Within a search service, synonym maps are a global resource that associate equivalent terms, expanding the scope of a query without the user having to actually provide the term. For example, assuming "dog", "canine", and "puppy" are mapped synonyms, a query on "canine" will match on a document containing "dog".
Create synonyms: A synonym map is an asset that can be created once and used by many indexes.
Reference:
https://docs.microsoft.com/en-us/azure/search/search-synonyms

NEW QUESTION 165

HOTSPOT - (Topic 1)
You need to develop code to upload images for the product creation project. The solution must meet the accessibility requirements.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
public static async Task<string> SuggestAltText(ComputerVisionClient client, [Dictionary | stream | string] image)
{
    List<VisualFeatureTypes?> features = new List<VisualFeatureTypes?>()
    {
        [VisualFeatureTypes.Description | VisualFeatureTypes.ImageType | VisualFeatureTypes.Objects | VisualFeatureTypes.Tags]
    };
ImageAnalysis results = await client.AnalyzeImageAsync(image, features);
[var c = results.Brands.DetectedBrands[0] | var c = results.Description.Captions[0] | var c = results.Metadata[0] | var c = results.Objects[0]]
if(c.Confidence>0.5) return(c.Text);
}
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:

```
public static async Task<string> SuggestAltText(ComputerVisionClient client, [Dictionary | stream | string] image)
{
    List<VisualFeatureTypes?> features = new List<VisualFeatureTypes?>()
    {
        [VisualFeatureTypes.Description | VisualFeatureTypes.ImageType | VisualFeatureTypes.Objects | VisualFeatureTypes.Tags]
    };
ImageAnalysis results = await client.AnalyzeImageAsync(image, features);
[var c = results.Brands.DetectedBrands[0] | var c = results.Description.Captions[0] | var c = results.Metadata[0] | var c = results.Objects[0]]
if(c.Confidence>0.5) return(c.Text);
}
```

NEW QUESTION 168
- (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 169
DRAG DROP - (Topic 1)
You are developing the smart e-commerce project.
You need to design the skillset to include the contents of PDFs in searches.
How should you complete the skillset design diagram? To answer, drag the appropriate services to the correct stages. Each service may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

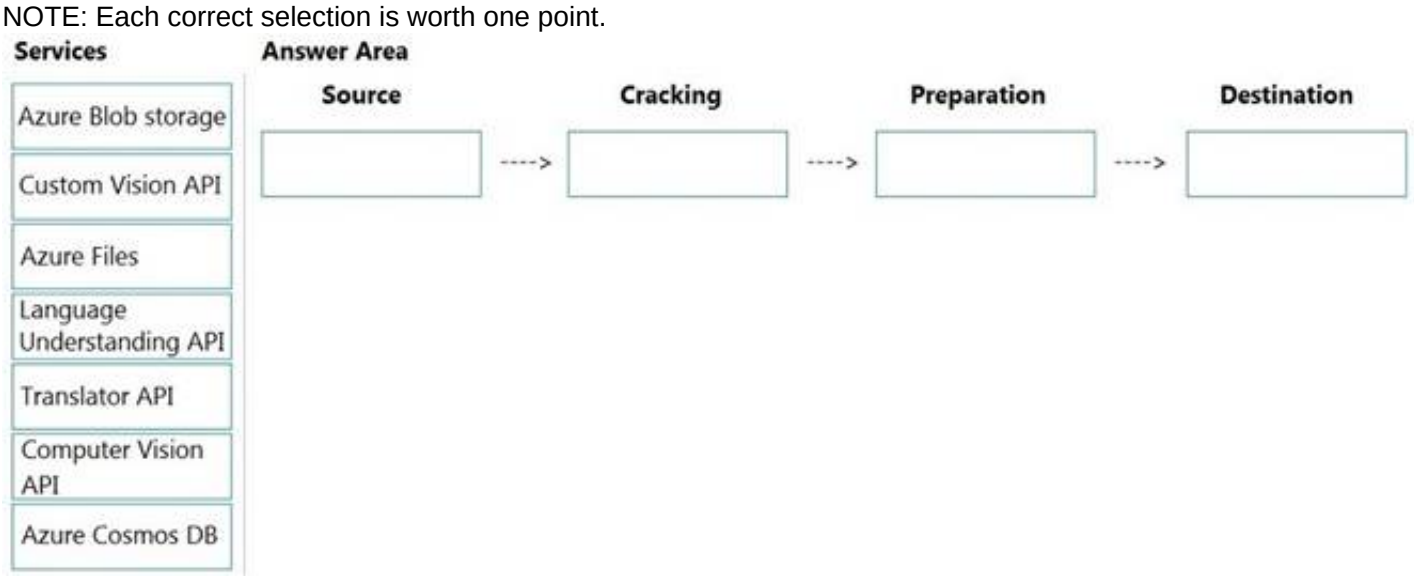

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Azure Blob storage
At the start of the pipeline, you have unstructured text or non-text content (such as images, scanned documents, or JPEG files). Data must exist in an Azure data storage service that can be accessed by an indexer.
Box 2: Computer Vision API
Scenario: Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products.
The Computer Vision Read API is Azure's latest OCR technology (learn what's new) that extracts printed text (in several languages), handwritten text (English only), digits, and currency symbols from images and multi-page PDF documents.
Box 3: Translator API
Scenario: Product descriptions, transcripts, and all text must be available in English, Spanish, and Portuguese.
Box 4: Azure Files
Scenario: Store all raw insight data that was generated, so the data can be processed later.

NEW QUESTION 171
.....

Relate Links

100% Pass Your AI-102 Exam with Exambible Prep Materials

https://www.exambible.com/AI-102-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/