

Microsoft.AI-102.vJan-2024.by.Osan.112q

Number: AI-102
Passing Score: 800
Time Limit: 120
File Version: 27.0

Exam Code: AI-102
Exam Name: Designing and Implementing a Microsoft Azure AI Solution



01 - Plan and Manage an Azure Cognitive Services Solution

QUESTION 1

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

Correct Answer: A, B, F

Section:

Explanation:



QUESTION 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.

Correct Answer: A, B, E

Section:

Explanation:

Reference:

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

QUESTION 3

HOTSPOT

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.

Hot Area:

Answer Area

▼ `https://management.azure.com/subscriptions/xxxxxxx-xxxx-`

PATCH
POST
PUT

`xxxx-xxxx-`
`xxxxxxxxxxxx/resourceGroups/RG1/providers/Microsoft.CognitiveServices/`
`accounts/CS1?api-version=2017-04-18`

```
{  
  "location": "West US",  
  "kind": "▼",  
    CognitiveServices  
    ComputerVision  
    TextAnalytics  
  "sku": {  
    "name": "S0"  
  },  
  "properties": {},  
  "identity": {  
    "type": "SystemAssigned"  
  }  
}
```



Answer Area:

Answer Area

▼ <https://management.azure.com/subscriptions/xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx/resourceGroups/RG1/providers/Microsoft.CognitiveServices/accounts/CS1?api-version=2017-04-18>

PATCH
POST
PUT

xxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx/resourceGroups/RG1/providers/Microsoft.CognitiveServices/accounts/CS1?api-version=2017-04-18

```
{
  "location": "West US",
  "kind": "CognitiveServices",
  "sku": {
    "name": "S0"
  },
  "properties": {},
  "identity": {
    "type": "SystemAssigned"
  }
}
```

CognitiveServices
ComputerVision
TextAnalytics



Section:

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>

Incorrect Answers:

PATCH is for updates.

Box 2: CognitiveServices

Microsoft Azure Cognitive Services provide us to use its pre-trained models for various Business Problems related to Machine Learning.

List of Different Services are:

Decision

Language (includes sentiment analysis)

Speech

Vision (includes OCR)

Web Search

Reference:

<https://docs.microsoft.com/en-us/rest/api/deviceupdate/resourcemanager/accounts/create>

<https://www.analyticsvidhya.com/blog/2020/12/microsoft-azure-cognitive-services-api-for-ai-development/>

QUESTION 4

DRAG DROP

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.

Select and Place:

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

Correct Answer:

Actions

-
-
-
-
- Build the image.
- Push the image to Docker Hub.

Answer Area

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

Section:

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>

QUESTION 5

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Answer Area:

```
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=xxxxxxxxxxxxxxxxxxxx
```

Section:

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}
```

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.azure.com> The endpoint for accessing the Text Analytics API.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-install-containers?tabs=sentiment>

QUESTION 6

HOTSPOT

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.

Hot Area:

Answer Area

```
var audioFormat = 
var speechConfig = SpeechConfig.FromSubscription("18c51a87-3a69-47a8-aedc-a54745f708a1", "westus");
var audioConfig = AudioConfig.FromStreamInput(pushStream, audioFormat);
using (var recognizer = new )
{
    var result = await recognizer.RecognizeOnceAsync();
    var text = result.Text;
}
```

The first dropdown menu shows the following options:

- AudioConfig.SetProperty
- AudioStreamFormat.GetCompressedFormat
- AudioStreamFormat.GetWaveFormatPCM
- PullAudioInputStream

The second dropdown menu shows the following options:

- KeywordRecognizer
- SpeakerRecognizer
- SpeechRecognizer
- SpeechSynthesizer

Answer Area:

Answer Area

```
var audioFormat = 
var speechConfig = SpeechConfig.FromSubscription("18c51a87-3a69-47a8-aedc-a54745f708a1", "westus");
var audioConfig = AudioConfig.FromStreamInput(pushStream, audioFormat);
using (var recognizer = new )
{
    var result = await recognizer.RecognizeOnceAsync();
    var text = result.Text;
}
```

Section:
Explanation:
Reference:
<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/how-to-use-codec-compressed-audio-input-streams?tabs=debian&pivots=programming-language-csharp>

QUESTION 7
HOTSPOT

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.

Hot Area:

Answer Area

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

	▼
Face	
Speech	
Text Analytics	

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

	▼
Face	
Speech	
Text Analytics	

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

	▼
Face	
Speech	
Text Analytics	



Answer Area:

Answer Area

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

 Face
 Speech
 Text Analytics

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

 Face
 Speech
 Text Analytics

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

 Face
 Speech
 Text Analytics

vdumps

Section:

Explanation:

Reference:

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

QUESTION 8

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 migrate to a Cognitive Search service that uses a higher tier.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

Section:

Explanation:

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

cost, which is why it is important to know the reason why throttling is occurring at all.

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-performance-analysis>

QUESTION 9

DRAG DROP

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.

Select and Place:

Services	Answer Area
Speaker Recognition	
Speech to Text	
Text Analytics	Detect the incoming language: <input type="text"/>
Text to Speech	Respond in the callers' own language: <input type="text"/>
Translator	

Correct Answer:

Services	Answer Area
Speaker Recognition	
Speech to Text	
	Detect the incoming language: <input type="text" value="Text Analytics"/>
Text to Speech	Respond in the callers' own language: <input type="text" value="Translator"/>



Section:

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.

Incorrect Answers:

Speaker Recognition which accurately verifies and identifies speakers by their unique voice characteristics.

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.

Reference:

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

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

QUESTION 10

You are developing a new sales system that will process the video and text from a public-facing website. You plan to notify users that their data has been processed by the sales system. Which responsible AI principle does this help meet?

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

Correct Answer: D

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/responsible-ai>

QUESTION 11

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



Correct Answer: B

Section:

Explanation:

The Azure Private Link should use a private endpoint, not a public endpoint.

Private Link service can be accessed from approved private endpoints in any public region.

Reference:

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

QUESTION 12

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 an IP firewall rule.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section:

Explanation:

Explanation/Reference:

Reference:

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

QUESTION 13

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

Correct Answer: C

Section:

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

Incorrect Answers:

A, B, D: Dispatch uses sample utterances for each of your bot's different tasks (LUIS, QnA Maker, or custom), and builds a model that can be used to properly route your user's request to the right task, even across multiple bots.

Reference:

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

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

QUESTION 14

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

Correct Answer: C

Section:

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>

QUESTION 15

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

Correct Answer: B, D

Section:

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>

QUESTION 16

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")

Correct Answer: B

Section:

Explanation:

Many of the Cognitive Services have a free tier you can use to try the service. To use the free tier, use F0 as the SKU for your resource. There are two tiers of keys for the Custom Vision service. You can sign up for a F0 (free) or S0 (standard) subscription through the Azure portal.

Incorrect Answers:

A: There is no free tier (F0) for ComputerVision.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/cognitive-services-apis-create-account-client-library?pivots=programming-language-csharp>

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/limits-and-quotas>

QUESTION 17

You successfully run the following HTTP request.

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

Correct Answer: B

Section:

Explanation:

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

QUESTION 18

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 servicel.

You need to ensure that app1 can connect directlyto servicel without routing traffic over the public internet. Solution: You deploy servicel and a public endpoint, and you configure a network security group (NSG) for vnet1.

Does this meet the goal?

- A. Yes
- B. No



Correct Answer: B

Section:

Explanation:

Explanation/Reference:

Reference:

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

QUESTION 19

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

Correct Answer: A

Section:

Explanation:

Explanation/Reference:

QUESTION 20

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

Correct Answer: D, E

Section:

Explanation:

Reference:

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

QUESTION 21

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 (IAM) page for the authoring resources in the Azure portal
- C. the Access control (IAM) page for the prediction resources in the Azure portal

Correct Answer: B

Section:

Explanation:

Reference:

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

QUESTION 22

DRAG DROP

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.

Select and Place:



Values

- AddPhraseListAsync
- Phraselist
- PhraselistCreateObject
- Phrases
- SavePhraselistAsync
- UploadPhraseListAsync

Answer Area

```
var phraselistId = await client.Features. [ ]
(appId, versionId, new [ ]
{
  EnabledForAllModels = false,
  IsExchangeable = true,
  Name = "PL1",
  Phrases = "item1,item2,item3,item4,item5"
});
```

Correct Answer:**Values**

- []
- Phraselist
- []
- Phrases
- SavePhraselistAsync
- UploadPhraseListAsync

Answer Area

```
var phraselistId = await client.Features. AddPhraseListAsync
(appId, versionId, new PhraselistCreateObject
{
  EnabledForAllModels = false,
  IsExchangeable = true,
  Name = "PL1",
  Phrases = "item1,item2,item3,item4,item5"
});
```

Section:**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

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/client-libraries-rest-api>

QUESTION 23

DRAG DROP

You plan to use a Language Understanding application named app1 that is deployed to a container.

App1 was developed by using a Language Understanding authoring resource named lu1.

App1 has the versions shown in the following table.

Version	Trained date	Published date
V1.2	None	None
V1.1	2020-10-01	None
V1.0	2020-09-01	2020-09-15

You need to create a container that uses the latest deployable version of app1.

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.)

Select and Place:

Actions

Run a container that has `version` set as an environment variable.

Export the model by using the Export as JSON option.

Select v1.1 of app1.

Run a container and mount the model file.

Select v1.0 of app1.

Export the model by using the Export for containers (GZIP) option.

Select v1.2 of app1.

Answer Area



Correct Answer:

Actions

Run a container that has `version` set as an environment variable.

Export the model by using the Export as JSON option.

Select v1.0 of app1.

Select v1.2 of app1.

Answer Area

Export the model by using the Export for containers (GZIP) option.

Select v1.1 of app1.

Run a container and mount the model file.



Section:

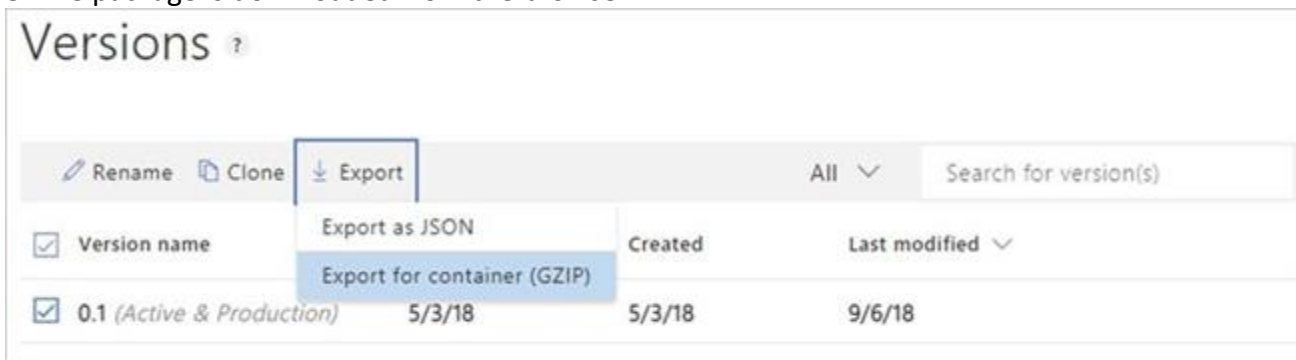
Explanation:

Step 1: Export the model using the Export for containers (GZIP) option.

Export versioned app's package from LUIS portal

The versioned app's package is available from the Versions list page.

1. Sign on to the LUIS portal.
2. Select the app in the list.
3. Select Manage in the app's navigation bar.
4. Select Versions in the left navigation bar.
5. Select the checkbox to the left of the version name in the list.
6. Select the Export item from the contextual toolbar above the list.
7. Select Export for container (GZIP).
8. The package is downloaded from the browser.



Step 2: Select v1.1 of app1.

A trained or published app packaged as a mounted input to the container with its associated App ID. Step 3: Run a contain and mount the model file.

Run the container, with the required input mount and billing settings.

Reference:

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

QUESTION 24

You have receipts that are accessible from a URL.

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

Which client and method should you use?

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

Correct Answer: D

Section:

Explanation:

To analyze receipts from a URL, use the StartRecognizeReceiptsFromUri method

Example code:

```
private static async Task AnalyzeReceipt(
    FormRecognizerClient recognizerClient, string receiptUri)
{
    RecognizedFormCollection receipts = await recognizerClient.StartRecognizeReceiptsFromUri(new Uri(receiptUri)).WaitForCompletionAsync();
}
```

Reference:

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

QUESTION 25

You have a collection of 50,000 scanned documents that contain text.

You plan to make the text available through Azure Cognitive Search.

You need to configure an enrichment pipeline to perform optical character recognition (OCR) and text analytics. The solution must minimize costs.

What should you attach to the skillset?

- A. a new Computer Vision resource
- B. a free (Limited enrichments) Cognitive Services resource
- C. an Azure Machine Learning pipeline
- D. a new Cognitive Services resource that uses the S0 pricing tier

Correct Answer: A

Section:

Explanation:

The Computer Vision API uses text recognition APIs to extract and recognize text information from images. Read uses the latest recognition models, and is optimized for large, text-heavy documents and noisy images.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/cognitive-search-with-skillsets>

QUESTION 26

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

Correct Answer: B

Section:

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>

QUESTION 27

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

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

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

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

You need to reduce the likelihood that search query requests are throttled.

Solution: You enable customer-managed key (CMK) encryption.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section:

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>

02 - Plan and Manage an Azure Cognitive Services Solution

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the

scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs.

When you are ready to answer a question, click the Question button to return to the question.

Overview

A company named Wide World Importers is developing an e-commerce platform.

You are working with a solutions architect to design and implement the features of the e-commerce platform. The platform will use microservices and a serverless environment built on Azure.

Wide World Importers has a customer base that includes English, Spanish, and Portuguese speakers.

Existing Environment

Applications

Wide World Importers has an App Service plan that contains the web apps shown in the following table.

Name	Description
Product Management	An app used by employees to create and manage products. The app and the expected inputs from the employees are in English.
Inventory Tracking	An app used by employees to manage inventory when dispatching orders, receiving refunds, and receiving consignments from suppliers.

Azure Resources

You have the following resources:

- An Azure Active Directory (Azure AD) tenant

An Azure Active Directory (Azure AD) tenant

- The tenant supports internal authentication.

- All employees belong to a group named AllUsers.

- Senior managers belong to a group named LeadershipTeam.

An Azure Functions resource

- A function app posts to Azure Event Grid when stock levels of a product change between OK, Low Stock, and Out of Stock. The function app uses the Azure Cosmos DB change feed. An Azure Cosmos DB account

- The account uses the Core (SQL) API.

- The account stores data for the Product Management app and the Inventory Tracking app.

- An Azure Storage account

An Azure Storage account

- The account contains blob containers for assets related to products.

- The assets include images, videos, and PDFs.

An Azure Cognitive Services resource named wwics

A Video Indexer resource named wwivi

Requirements

Business Goals

Wide World Importers wants to leverage AI technologies to differentiate itself from its competitors.

Planned Changes

Wide World Importers plans to start the following projects:

A product creation project: Help employees create accessible and multilingual product entries, while expediting product entry creation. A smart e-commerce project: Implement an Azure Cognitive Search solution to display products for customers to browse. A shopping on-the-go project: Build a chatbot that can be integrated into smart speakers to support customers.

Business Requirements

Wide World Importers identifies the following business requirements for all the projects:

Provide a multilingual customer experience that supports English, Spanish, and Portuguese.

Whenever possible, scale based on transaction volumes to ensure consistent performance.

Minimize costs.

Governance and Security Requirements

Wide World Importers identifies the following governance and security requirements:

Data storage and processing must occur in datacenters located in the United States.

Azure Cognitive Services must be inaccessible directly from the internet.

Accessibility Requirements

Wide World Importers identifies the following accessibility requirements:

All images must have relevant alt text.

All videos must have transcripts that are associated to the video and included in product descriptions. Product descriptions, transcripts, and alt text must be available in English, Spanish, and Portuguese.

Product Creation Requirements

Wide World Importers identifies the following requirements for improving the Product Management app:

Minimize how long it takes for employees to create products and add assets.

Remove the need for manual translations.

Smart E-Commerce Requirements

Wide World Importers identifies the following requirements for the smart e-commerce project:

Ensure that the Cognitive Search solution meets a Service Level Agreement (SLA) of 99.9% availability for searches and index writes. Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products. Support autocompletion and autosuggestion based on all product name variants.

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

Update the stock level field in the product index immediately upon changes.

Update the product index hourly.

Shopping On-the-Go Requirements

Wide World Importers identifies the following requirements for the shopping on-the-go chatbot:

Answer common questions.

Support interactions in English, Spanish, and Portuguese.

Replace an existing FAQ process so that all Q&A is managed from a central location.

Provide all employees with the ability to edit Q&As. Only senior managers must be able to publish updates. Support purchases by providing information about relevant products to customers. Product displays must include images and warnings when stock levels are low or out of stock.

Product JSON Sample

You have the following JSON sample for a product.

```
{
  "sku": "b1",
  "name": {
    "en": "Bicycle",
    "es": "Bicicleta",
    "pt": "Bicicleta"
  },
  "stocklevel": "Out of Stock",
  "description": {
    "en": "Bicycle",
    "es": "Bicicleta",
    "pt": "Bicicleta"
  },
  "image": {
    "uri": "https://upload.worldwideimporters.org/bicycle.jpg",
    "alttext": {
      "en": "Bicycle",
      "es": "Bicicleta",
      "pt": "Bicicleta"
    }
  },
  "createdUtc": "2020-02-14T06:08:39Z",
  "language": "en"
}
```



QUESTION 1

DRAG DROP

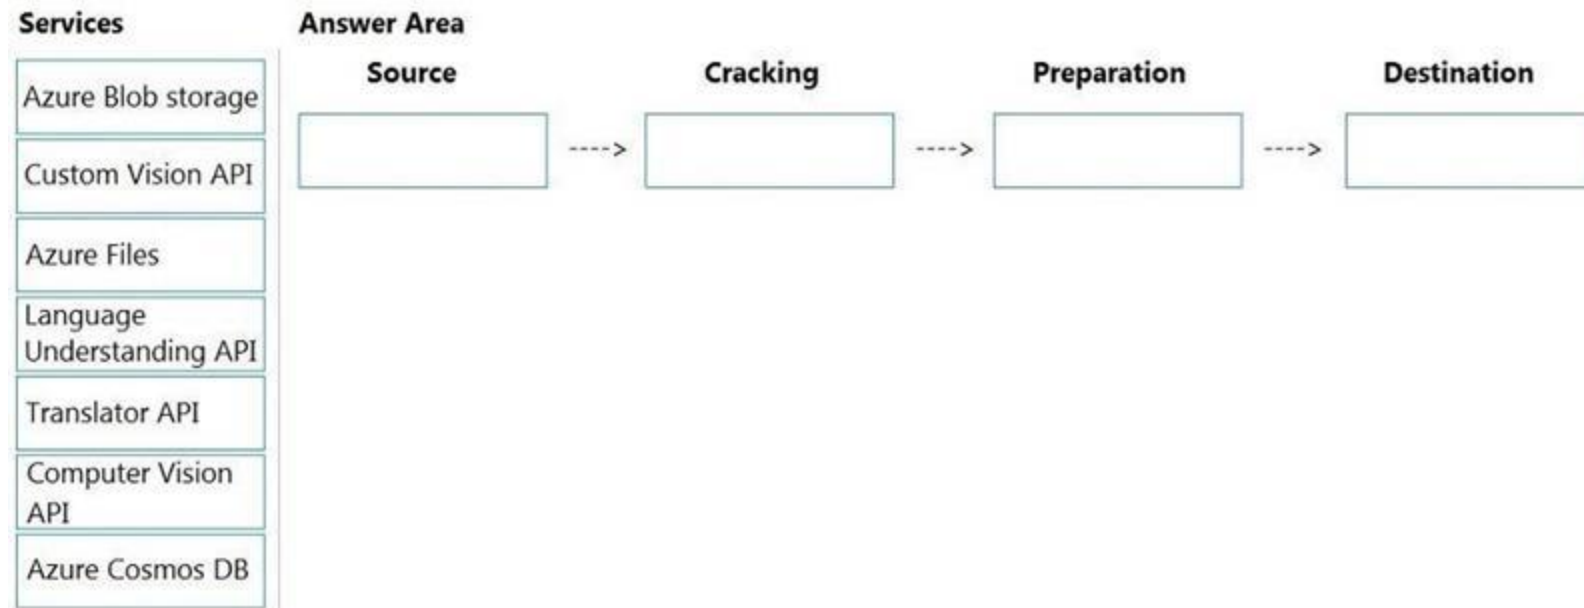
You are developing the smart e-commerce project.

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

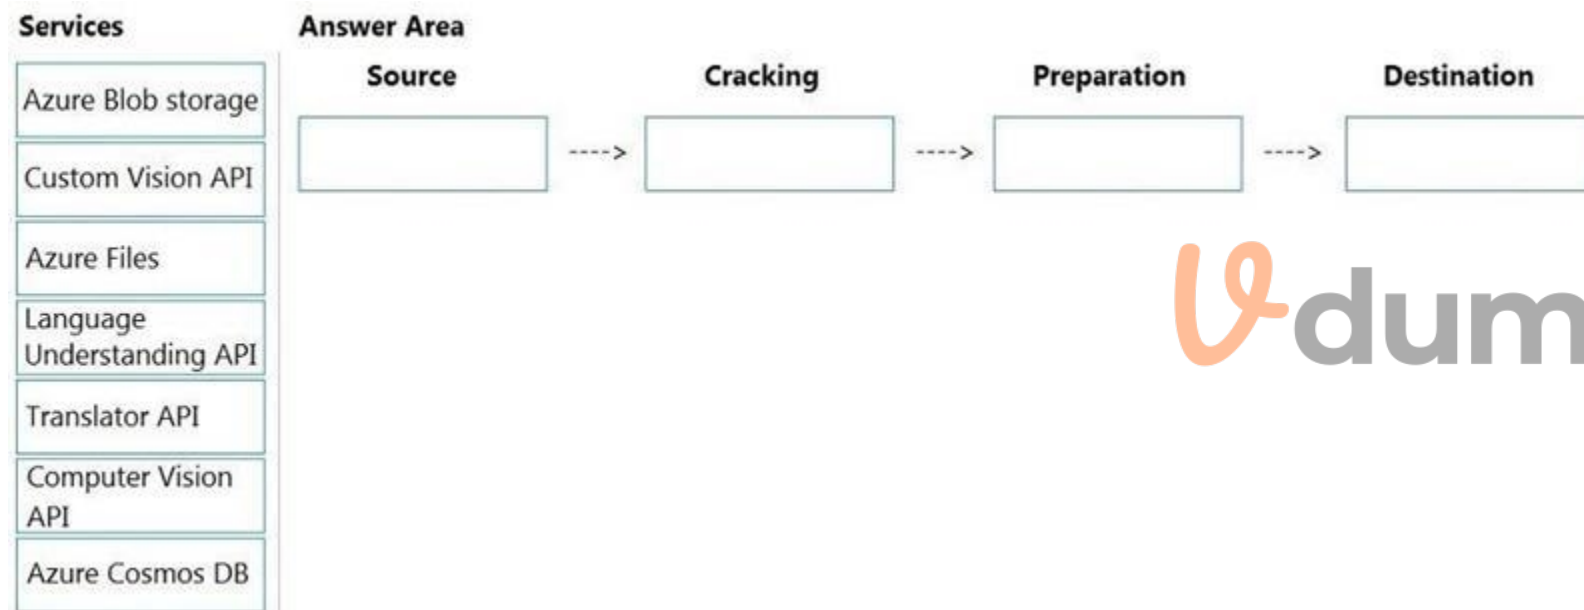
How should you complete the skillset design diagram? To answer, drag the appropriate services to the correct stages. Each service may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:



Correct Answer:



Section:

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 alt 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.

Incorrect Answers:

The custom vision API from Microsoft Azure learns to recognize specific content in imagery and becomes smarter with training and time.

Reference:

<https://docs.microsoft.com/en-us/azure/search/cognitive-search-concept-intro>

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/overview-ocr>

03 - Plan and Manage an Azure Cognitive Services Solution

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study. At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs.

When you are ready to answer a question, click the Question button to return to the question. Overview

General Overview

Contoso, Ltd. is an international accounting company that has offices in France, Portugal, and the United Kingdom. Contoso has a professional services department that contains the roles shown in the following table.

Name	Position	Office
Accountant	Manager	United Kingdom, France, Portugal
Accountant	Consultant	United Kingdom, France, Portugal
Customer Service	Manager	United Kingdom
Customer Service	Agent	United Kingdom
Bookkeeper	Manager	United Kingdom, France, Portugal
Bookkeeper	Consultant	United Kingdom, France, Portugal

Existing environment

Infrastructure

Contoso has the following subscriptions:

Azure

Microsoft 365

Microsoft Dynamics 365

Azure Active (Azure AD) Directory

Contoso has Azure Active Directory groups for securing role-based access. The company uses the following group naming conventions:

[Country]-[Level]-[Role]

[Level]-[Role]

Intellectual Property

Contoso has the intellectual property shown in the following table.

Content	Format	Language	Content store	Domain
Weekly webinars	Video	English	Azure Blob storage	Vid.contoso.com
Blogs	Text	English, French, Portuguese	WordPress	Pt-blog.contoso.com Blog.contoso.com Fr-blog.contoso.com
Wikis	Text	English, French, Portuguese	Azure Cosmos DB	Internal.contoso.com/wiki
Monthly conference recording	Video	English	SharePoint Online	Contoso.sharepoint.com
Frequently asked questions (FAQs)	Text	English	SharePoint Online	Contoso.sharepoint.com

Text-based content is provided only in one language and is not translated.

Requirements

Planned Projects

Contoso plans to develop the following:

A document processing workflow to extract information automatically from PDFs and images of financial documents

A customer-support chatbot that will answer questions by using FAQs

A searchable knowledgebase of all the intellectual property

Technical Requirements

Contoso identifies the following technical requirements:

All content must be approved before being published.

All planned projects must support English, French, and Portuguese.

All content must be secured by using role-based access control (RBAC).

RBAC role assignments must use the principle of least privilege.

RBAC roles must be assigned only to Azure Active Directory groups.

AI solution responses must have a confidence score that is equal to or greater than 70 percent.

When the response confidence score of an AI response is lower than 70 percent the response must be improved by human input.

Chatbot Requirements

Contoso identifies the following requirements for the chatbot:

Provide customers with answers to the FAQs.

Ensure that the customers can chat to a customer service agent.

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

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

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

Ensure that access to the customer service agents is managed by using Omnichannel for Customer Service. When the response confidence score is low, ensure that the chatbot can provide other response options to the customers.

Document Processing Requirements

Contoso identifies the following requirements for document processing:

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.

Members of a group named Management-Bookkeeper must define how to extract tables from the financial documents. Members of a group named Consultant-Bookkeeper must be able to process the financial documents.

Knowledgebase Requirements

Contoso identifies the following requirements for the knowledgebase:

- Supports searches for equivalent terms
- Can transcribe jargon with high accuracy
- Can search content in different formats, including video
- Provides relevant links to external resources for further research



QUESTION 1

DRAG DROP

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. Select and Place:

Select and Place:



Correct Answer:



Section:

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.

Reference: <https://docs.microsoft.com/en-us/azure/applied-ai-services/form-recognizer/label-tool>

QUESTION 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

Correct Answer: B

Section:

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).

Topic inference: Makes inference of main topics from transcripts. The 2nd-level IPTC taxonomy is included. Artifacts: Extracts rich set of "next level of details" artifacts for each of the models.

Sentiment analysis: Identifies positive, negative, and neutral sentiments from speech and visual text. Incorrect Answers:

C: Webinars Videos are in English.

Reference: <https://docs.microsoft.com/en-us/azure/azure-video-analyzer/video-analyzer-for-media-docs/video-indexer-overview>

01 - Implement Computer Vision Solutions

QUESTION 1

HOTSPOT

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.

Hot Area:

Vdumps

Answer Area

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 top-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"/>

Answer Area:

Answer Area

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 top-left corner of the rectangle that contains the brand logo of the displayed brands.	<input checked="" 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 checked="" type="checkbox"/>

Section:**Explanation:**

Box 1: Yes

Box 2: Yes

Coordinates of a rectangle in the API refer to the top left corner.

Box 3: No

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-brand-detection>

QUESTION 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 ReadUrl(ComputerVisionClient client, string urlFile)
{
    const int numberOfCharsInOperationId = 36;

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

    string opLocation = txtHeaders.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.

Correct Answer: B, D

Section:

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/ComputerVisionQuickstart.cs>

QUESTION 3

DRAG DROP

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.

Select and Place:

Values

"faceListId"
"LargeFaceListId"
"matchFace"
"matchPerson"

Answer Area

```
{
"faceId": "18c51a87-3a69-47a8-aedc-a54745f708a1",
: "employeefaces",
"maxNumOfCandidatesReturned": 1,
"mode":
}
```

Correct Answer:

Values

"faceListId"

"matchPerson"

Answer Area

```
{
"faceId": "18c51a87-3a69-47a8-aedc-a54745f708a1",
"LargeFaceListId": "employeefaces",
"maxNumOfCandidatesReturned": 1,
"mode": "matchFace"
}
```

Section:

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.

Incorrect Answers:

Not "faceListId": Add a face to a specified face list, up to 1,000 faces.

Box 2: matchFace

Find similar has two working modes, "matchPerson" and "matchFace". "matchPerson" is the default mode that it tries to find faces of the same person as possible by using internal same-person thresholds. It is useful to find a known person's other photos. Note that an empty list will be returned if no faces pass the internal thresholds. "matchFace" mode ignores same-person thresholds and returns ranked similar faces anyway, even the similarity is low. It can be used in the cases like searching celebrity-looking faces.

Reference:

<https://docs.microsoft.com/en-us/rest/api/faceapi/face/findsimilar>

QUESTION 4

DRAG DROP

You are developing a photo application that will find photos of a person based on a sample image by using the Face API.

You need to create a POST request to find the photos.

How should you complete the 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.

Select and Place:

Values

detect
findsimilar
group
identify
matchFace
matchPerson
verify

Answer Area

POST {Endpoint}/face/v1.0/
Request Body

```
{  
  "faceId": "c5c24a82-6845-4031-9d5d-978df9175426",  
  "largeFaceListId": "sample_list",  
  "maxNumOfCandidatesReturned": 10,  
  "mode": "  
}
```



Correct Answer:

Values

findsimilar
group
identify
matchFace

verify

Answer Area

POST {Endpoint}/face/v1.0/
Request Body

```
{  
  "faceId": "c5c24a82-6845-4031-9d5d-978df9175426",  
  "largeFaceListId": "sample_list",  
  "maxNumOfCandidatesReturned": 10,  
  "mode": "  
}
```

Section:

Explanation:

Box 1: detect

Face - Detect With Url: Detect human faces in an image, return face rectangles, and optionally with faceIds, landmarks, and attributes.

POST {Endpoint}/face/v1.0/detect

Box 2: matchPerson

Find similar has two working modes, "matchPerson" and "matchFace". "matchPerson" is the default mode that it tries to find faces of the same person as possible by using internal same-person thresholds. It is useful to find a known person's other photos. Note that an empty list will be returned if no faces pass the internal thresholds. "matchFace" mode ignores same-person thresholds and returns ranked similar faces anyway, even the similarity is low. It can be used in the cases like searching celebrity-looking faces.

Reference:

<https://docs.microsoft.com/en-us/rest/api/faceapi/face/detectwithurl>

<https://docs.microsoft.com/en-us/rest/api/faceapi/face/findsimilar>

QUESTION 5

DRAG DROP

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.

Select and Place:

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

Correct Answer:

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

Section:

Explanation:

Reference:
<https://docs.microsoft.com/en-us/azure/azure-video-analyzer/video-analyzer-for-media-docs/video-indexer-embed-widgets>

QUESTION 6

DRAG DROP

You train a Custom Vision model to identify a company's products by using the Retail domain.

You plan to deploy the model as part of an app for Android phones.

You need to prepare the model for deployment.

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.

Select and Place:

Actions

- Change the model domain.
- Retrain the model.
- Test the model.
- Export the model.

Answer Area



Correct Answer:

Actions

Four empty rectangular boxes stacked vertically, representing the source of actions for the correct answer.

Answer Area

Change the model domain.

Retrain the model.

Test the model.

Export the model.

The answer area shows the correct sequence of actions: 'Change the model domain.', 'Retrain the model.', 'Test the model.', and 'Export the model.'. To the left of the first three items are circular arrow icons pointing left, right, and right respectively. To the right of the first two items are circular arrow icons pointing up and down respectively.

Section:

Explanation:

Reference:
<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/export-your-model>

QUESTION 7

HOTSPOT

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.

Hot Area:

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"/>

Answer Area:

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

Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/use-persondirectory>

QUESTION 8

DRAG DROP

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.

Select and Place:

Actions

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

Answer Area



Correct Answer:

Actions

-
-
-
- Use the `ExportIteration` endpoint on acvdev.
- Use the `GetIterations` endpoint on acvdev.
- Use the `UpdateProject` endpoint on acvprod.

Answer Area



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



Section:

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/copy-move-projects>

QUESTION 9

DRAG DROP

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.

Select and Place:

Actions

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



Answer Area



Correct Answer:

Actions

- Train the classifier model.
-
- Initialize the training dataset.
-
-



Answer Area

- Create a project.
- Upload and tag images.
- Train the object detection model.



Section:

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

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/getting-started-build-a-classifier>

QUESTION 10

HOTSPOT

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.

Hot Area:

Answer Area

Project Types: ▼

Classification
Object Detection

Classification Types: ▼

Multiclass (Single tag per image)
Multilabel (Multiple tags per image)

Domains: ▼

Audit
Food
General
General (compact)
Landmarks
Landmarks (compact)
Retail
Retail (compact)



Answer Area:

Answer Area

Project Types: ▼

- Classification
- Object Detection

Classification Types: ▼

- Multiclass (Single tag per image)
- Multilabel (Multiple tags per image)

Domains: ▼

- Audit
- Food
- General
- General (compact)
- Landmarks
- Landmarks (compact)
- Retail
- Retail (compact)



Section:

Explanation:

Box 1: Classification

Incorrect Answers:

An object detection project is for detecting which objects, if any, from a set of candidates are present in an image.

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.

Incorrect Answers:

A multilabel classification project is similar, but each image can have multiple tags assigned to it.

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.

Reference:

<https://cran.r-project.org/web/packages/AzureVision/vignettes/customvision.html>

QUESTION 11

HOTSPOT

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.



Hot Area:

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);
        }
    }
});
```

Answer Area:

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);
        }
    }
});
```



Section:

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: CreateAsync

Create the persons for the PersonGroup. Persons are created concurrently.

Example:

```
await faceClient.PersonGroupPerson.CreateAsync(personGroupId, personName);
```

Reference:

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

QUESTION 12

HOTSPOT

You are developing an application that will use the Computer Vision client library. The application has the following code.


```

public async TaskAnalyzeImage(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.

Hot Area:

Answer Area			
	Statements	Yes	No
	The code will perform face recognition.	<input type="radio"/>	<input type="radio"/>
	The code will list tags and their associated confidence.	<input type="radio"/>	<input type="radio"/>
	The code will read a file from the local file system.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Statements	Yes	No
The code will perform face recognition.	<input type="radio"/>	<input checked="" type="radio"/>
The code will list tags and their associated confidence.	<input checked="" type="radio"/>	<input type="radio"/>
The code will read a file from the local file system.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

Box 1: No

Box 2: Yes

The ComputerVision.analyzeImageInStreamAsync operation extracts a rich set of visual features based on the image content.

Box 3: No

Images will be read from a stream.

Reference:

<https://docs.microsoft.com/en-us/java/api/com.microsoft.azure.cognitiveservices.vision.computervision.computervision.analyzeimageinstreamasync>

QUESTION 13

HOTSPOT

You have a Computer Vision resource named contoso1 that is hosted in the West US Azure region.

You need to use contoso1 to make a different size of a product photo by using the smart cropping feature.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

/vision/v3.1/

Answer Area:

```

Answer Area

curl -H "Ocp-Apim-Subscription-Key: xxx" /
-o "sample.png" -H "Content-Type: application/json" /
/vision/v3.1/
?width=100&height=100&smartCropping=true" /
"https://api.projectoxford.ai
"https://contoso1.cognitiveservices.azure.com
"https://westus.api.cognitive.microsoft.com
areaOfInterest
detect
generateThumbnail

-d "{ \"url\": \"https://upload.litwareinc.org/litware/bicycle.jpg\" }"

```

Section:

Explanation:

Reference:

<https://westus.dev.cognitive.microsoft.com/docs/services/computer-vision-v3-2/operations/56f91f2e778daf14a499f21b>

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-generating-thumbnails#examples>

QUESTION 14

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.



Correct Answer: A

Section:

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 newfacefor 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/whats-new/video-indexer/customize-person-model-with-api>

QUESTION 15

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)

Correct Answer: A, D

Section:

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/>

02 - Implement Computer Vision Solutions

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs.

When you are ready to answer a question, click the Question button to return to the question.

Overview

A company named Wide World Importers is developing an e-commerce platform.

You are working with a solutions architect to design and implement the features of the e-commerce platform. The platform will use microservices and a serverless environment built on Azure.

Wide World Importers has a customer base that includes English, Spanish, and Portuguese speakers.

Existing Environment

Applications

Wide World Importers has an App Service plan that contains the web apps shown in the following table.

Name	Description
Product Management	An app used by employees to create and manage products. The app and the expected inputs from the employees are in English.
Inventory Tracking	An app used by employees to manage inventory when dispatching orders, receiving refunds, and receiving consignments from suppliers.

Azure Resources

You have the following resources:

- An Azure Active Directory (Azure AD) tenant

An Azure Active Directory (Azure AD) tenant

- The tenant supports internal authentication.

- All employees belong to a group named AllUsers.

- Senior managers belong to a group named LeadershipTeam.

An Azure Functions resource

- A function app posts to Azure Event Grid when stock levels of a product change between OK, Low Stock, and Out of Stock. The function app uses the Azure Cosmos DB change feed. An Azure Cosmos DB account

- The account uses the Core (SQL) API.

- The account stores data for the Product Management app and the Inventory Tracking app.

- An Azure Storage account

An Azure Storage account

- The account contains blob containers for assets related to products.

- The assets include images, videos, and PDFs.

An Azure Cognitive Services resource named wwics

A Video Indexer resource named wwivi

Requirements

Business Goals

Wide World Importers wants to leverage AI technologies to differentiate itself from its competitors.

Planned Changes

Wide World Importers plans to start the following projects:

A product creation project: Help employees create accessible and multilingual product entries, while expediting product entry creation. A smart e-commerce project: Implement an Azure Cognitive Search solution to display products for customers to browse. A shopping on-the-go project: Build a chatbot that can be integrated into smart speakers to support customers.

Business Requirements

Wide World Importers identifies the following business requirements for all the projects:

Provide a multilingual customer experience that supports English, Spanish, and Portuguese.

Whenever possible, scale based on transaction volumes to ensure consistent performance.

Minimize costs.

Governance and Security Requirements

Wide World Importers identifies the following governance and security requirements:

Data storage and processing must occur in datacenters located in the United States.

Azure Cognitive Services must be inaccessible directly from the internet.

Accessibility Requirements

Wide World Importers identifies the following accessibility requirements:

All images must have relevant alt text.

All videos must have transcripts that are associated to the video and included in product descriptions. Product descriptions, transcripts, and alt text must be available in English, Spanish, and Portuguese.

Product Creation Requirements

Wide World Importers identifies the following requirements for improving the Product Management app:

Minimize how long it takes for employees to create products and add assets.

Remove the need for manual translations.

Smart E-Commerce Requirements

Wide World Importers identifies the following requirements for the smart e-commerce project:

Ensure that the Cognitive Search solution meets a Service Level Agreement (SLA) of 99.9% availability for searches and index writes. Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products. Support autocompletion and autosuggestion based on all product name variants.

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

Update the stock level field in the product index immediately upon changes.

Update the product index hourly.

Shopping On-the-Go Requirements

Wide World Importers identifies the following requirements for the shopping on-the-go chatbot:

Answer common questions.

Support interactions in English, Spanish, and Portuguese.

Replace an existing FAQ process so that all Q&A is managed from a central location.

Provide all employees with the ability to edit Q&As. Only senior managers must be able to publish updates. Support purchases by providing information about relevant products to customers. Product displays must include images and warnings when stock levels are low or out of stock.

Product JSON Sample

You have the following JSON sample for a product.



```
{
  "sku": "b1",
  "name": {
    "en": "Bicycle",
    "es": "Bicicleta",
    "pt": "Bicicleta"
  },
  "stocklevel": "Out of Stock",
  "description": {
    "en": "Bicycle",
    "es": "Bicicleta",
    "pt": "Bicicleta"
  },
  "image":
  {"uri": "https://upload.worldwideimporters.org/bicycle.jpg",
    "alttext": {
      "en": "Bicycle",
      "es": "Bicicleta",
      "pt": "Bicicleta"
    }
  },
  "createdUtc": "2020-02-14T06:08:39Z",
  "language": "en"
}
```

QUESTION 1

DRAG DROP

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.)

Select and Place:



Actions

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.

Answer Area

Correct Answer:

Actions

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.
Upload the video to file storage.

Answer Area

Upload the video to blob storage.
Index the video by using the Video Indexer API.
Extract the transcript from the Video Indexer API.
Translate the transcript by using the Translator API.



Section:

Explanation:

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

Reference:

<https://azure.microsoft.com/en-us/blog/get-video-insights-in-even-more-languages/>

<https://docs.microsoft.com/en-us/azure/media-services/video-indexer/video-indexer-output-json-v2>

QUESTION 2

HOTSPOT

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.

Hot Area:

Answer Area

```
public static async Task<string> SuggestAltText(ComputerVisionClient client,
{
    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);
}
```

image

Dictionary
stream
string

VisualFeatureTypes.Description
VisualFeatureTypes.ImageType
VisualFeatureTypes.Objects
VisualFeatureTypes.Tags

var c = results.Brands.DetectedBrands[0]
var c = results.Description.Captions[0]
var c = results.Metadata[0]
var c = results.Objects[0]

Answer Area:

Answer Area

```
public static async Task<string> SuggestAltText(ComputerVisionClient client,
{
    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);
}
```

image

Dictionary
stream
string

VisualFeatureTypes.Description
VisualFeatureTypes.ImageType
VisualFeatureTypes.Objects
VisualFeatureTypes.Tags

var c = results.Brands.DetectedBrands[0]
var c = results.Description.Captions[0]
var c = results.Metadata[0]
var c = results.Objects[0]

Section:

Explanation:

Reference:

 **vdumps**

<https://github.com/Azure-Samples/cognitive-services-dotnet-sdk-samples/blob/master/documentation-samples/quickstarts/ComputerVision/Program.cs>

03 - Implement Computer Vision Solutions

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study. At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs.

When you are ready to answer a question, click the Question button to return to the question. Overview

General Overview

Contoso, Ltd. is an international accounting company that has offices in France, Portugal, and the United Kingdom. Contoso has a professional services department that contains the roles shown in the following table.

Name	Position	Office
Accountant	Manager	United Kingdom, France, Portugal
Accountant	Consultant	United Kingdom, France, Portugal
Customer Service	Manager	United Kingdom
Customer Service	Agent	United Kingdom
Bookkeeper	Manager	United Kingdom, France, Portugal
Bookkeeper	Consultant	United Kingdom, France, Portugal

Existing environment

Infrastructure

Contoso has the following subscriptions:

Azure

Microsoft 365

Microsoft Dynamics 365

Azure Active (Azure AD) Directory

Contoso has Azure Active Directory groups for securing role-based access. The company uses the following group naming conventions:

[Country]-[Level]-[Role]

[Level]-[Role]

Intellectual Property

Contoso has the intellectual property shown in the following table.



Content	Format	Language	Content store	Domain
Weekly webinars	Video	English	Azure Blob storage	Vid.contoso.com
Blogs	Text	English, French, Portuguese	WordPress	Pt-blog.contoso.com Blog.contoso.com Fr-blog.contoso.com
Wikis	Text	English, French, Portuguese	Azure Cosmos DB	Internal.contoso.com/wiki
Monthly conference recording	Video	English	SharePoint Online	Contoso.sharepoint.com
Frequently asked questions (FAQs)	Text	English	SharePoint Online	Contoso.sharepoint.com

Text-based content is provided only in one language and is not translated.

Requirements

Planned Projects

Contoso plans to develop the following:

A document processing workflow to extract information automatically from PDFs and images of financial documents

A customer-support chatbot that will answer questions by using FAQs

A searchable knowledgebase of all the intellectual property

Technical Requirements

Contoso identifies the following technical requirements:

All content must be approved before being published.

All planned projects must support English, French, and Portuguese.

All content must be secured by using role-based access control (RBAC).

RBAC role assignments must use the principle of least privilege.

RBAC roles must be assigned only to Azure Active Directory groups.

AI solution responses must have a confidence score that is equal to or greater than 70 percent.

When the response confidence score of an AI response is lower than 70 percent the response must be improved by human input.

Contoso identifies the following requirements for the chatbot:

Provide customers with answers to the FAQs.

Ensure that the customers can chat to a customer service agent.

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

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

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

Ensure that access to the customer service agents is managed by using Omnichannel for Customer Service. When the response confidence score is low, ensure that the chatbot can provide other response options to the customers.

Document Processing Requirements

Contoso identifies the following requirements for document processing:

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.

Members of a group named Management-Bookkeeper must define how to extract tables from the financial documents. Members of a group named Consultant-Bookkeeper must be able to process the financial documents.

Knowledgebase Requirements

Contoso identifies the following requirements for the knowledgebase:

- Supports searches for equivalent terms
- Can transcribe jargon with high accuracy
- Can search content in different formats, including video
- Provides relevant links to external resources for further research



QUESTION 1

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/v3.1/read/analyzeResults
- B. /formrecognizer/v2.0/custom/models/{modelId}/analyze
- C. /formrecognizer/v2.0/prebuilt/receipt/analyze
- D. /vision/v3.1/describe
- E. /vision/v3.1/read/analyze

Correct Answer: C, E

Section:

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/v3.1/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-v2preview/operations/GetAnalyzeReceiptResult> <https://docs.microsoft.com/en-us/rest/api/computervision/3.1/read/read>

QUESTION 2

HOTSPOT


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.

Hot Area:



Answer Area

```
{
  "@odata.type": "#Microsoft.Skills.Text.EntityRecognitionSkill",
  "categories": [],
  "categories": [ "Email", "Persons", "Organizations" ],
  "categories": [ "Locations", "Persons", "Organizations" ],
  "defaultLanguageCode": "en",
  "includeTypelessEntities": true,
  "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" }
  ]
}
```



Answer Area:

Answer Area

```
{
  "@odata.type": "#Microsoft.Skills.Text.EntityRecognitionSkill",
  "categories": [],
  "categories": [ "Email", "Persons", "Organizations" ],
  "categories": [ "Locations", "Persons", "Organizations" ],
  "defaultLanguageCode": "en",
  "includeTypelessEntities": true,
  "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" }
  ]
}
```

Section:

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.

Reference: <https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-entity-recognition>

QUESTION 3

You are developing the knowledgebase by using Azure Cognitive Search.

You need to process wiki content to meet the technical requirements.

What should you include in the solution?

- A. an indexer for Azure Blob storage attached to a skillset that contains the language detection skill and the text translation skill
- B. an indexer for Azure Blob storage attached to a skillset that contains the language detection skill

- C. an indexer for Azure Cosmos DB attached to a skillset that contains the document extraction skill and the text translation skill
- D. an indexer for Azure Cosmos DB attached to a skillset that contains the language detection skill and the text translation skill

Correct Answer: C

Section:

Explanation:

The wiki contains text in English, French and Portuguese.

Scenario: All planned projects must support English, French, and Portuguese.

The Document Extraction skill extracts content from a file within the enrichment pipeline. This allows you to take advantage of the document extraction step that normally happens before the skillset execution with files that may be generated by other skills.

Note: The Translator Text API will be used to determine the from language. The Language detection skill is not required. Incorrect Answers:

Not A, not B: The wiki is stored in Azure Cosmos DB.

Reference:

<https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-document-extraction> <https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-text-translation>

QUESTION 4

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

Correct Answer: A

Section:

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>

Exam G

QUESTION 1

DRAG DROP

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

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

NOTE: Each correct match is worth one point.

Select and Place:



APIs

Cassandra API
Gremlin API
MongoDB API
Table API

Answer Area

API	Graph data
API	JSON documents
API	Key/value data

Correct Answer:

APIs

Cassandra API

Answer Area

Gremlin API	Graph data
MongoDB API	JSON documents
Table API	Key/value data

Section:

Explanation:



QUESTION 2

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

account
container
database
item

Answer Area:

Answer Area

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

account
container
database
item

Section:

Explanation:

QUESTION 3

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

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

Correct Answer: B

Section:

QUESTION 4

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.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section:

QUESTION 5

You develop a Conversational Language Understanding model by using Language Studio

During testing, users receive incorrect responses to requests that do NOT relate to the capabilities of the model. You need to ensure that the model identifies spurious requests.

What should you do?

- A. Enable active learning.
- B. Add examples to the custom intents.
- C. Add examples to the None intent
- D. Add entities.

Correct Answer: A

Section:

QUESTION 6

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

Correct Answer: C



Section:

QUESTION 7

HOTSPOT

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.

Hot Area:

Answer Area

Custom neural voice: The Language Understanding portal
Microsoft Bot Framework Composer
The Azure portal

Narration: The Language Understanding portal
The Speech Studio portal

Narration: Language Understanding
Language Understanding
Speaker Recognition
Speech-to-text
Text-to-speech

Answer Area:

Answer Area

Custom neural voice: The Language Understanding portal
Microsoft Bot Framework Composer
The Azure portal

Narration: The Language Understanding portal
The Speech Studio portal

Narration: Language Understanding
Language Understanding
Speaker Recognition
Speech-to-text
Text-to-speech

Section:

Explanation:

QUESTION 8

HOTSPOT

You have a bot that was built by using the Microsoft Bot Framework composer as shown in the following exhibit.



Use the drop-down menus 10 select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

If a user asks "what is the weather like in New York", the bot will [answer choice].

- Identify New York as a city entity
- change to a different dialog
- identify New York as a city entity
- identify New York as a state entity
- respond with the weather in Seattle

The GetWeather dialog uses a [answer choice] trigger.

- Language Understanding Intent recogn...
- Custom events
- Dialog events
- Language Understanding Intent recognized
- QnA Intent recognized

Answer Area:

Answer Area

If a user asks "what is the weather like in New York", the bot will [answer choice].

- Identify New York as a city entity
- change to a different dialog
- identify New York as a city entity
- identify New York as a state entity
- respond with the weather in Seattle

The GetWeather dialog uses a [answer choice] trigger.

- Language Understanding Intent recogn...
- Custom events
- Dialog events
- Language Understanding Intent recognized
- QnA intent recognized

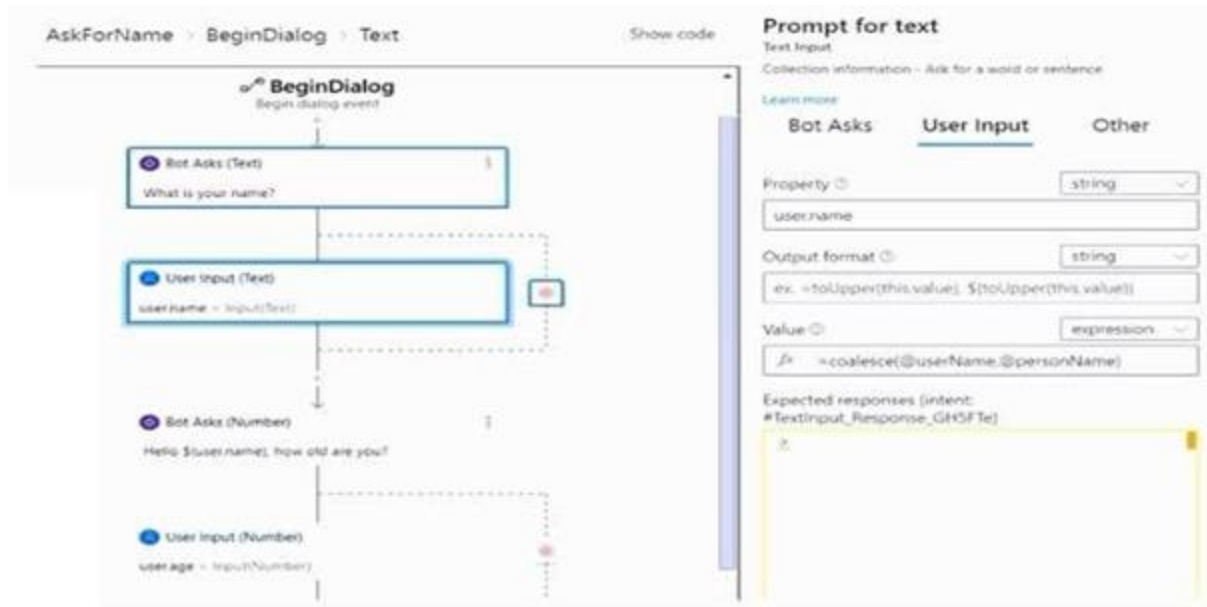
Section:

Explanation:

QUESTION 9

HOTSPOT

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.

Hot Area:

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"/>

Answer Area:

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"/>

Section:

Explanation:

QUESTION 10

DRAG DROP

You plan to build a chatbot to support task tracking.

You create a Conversational Language Understanding service named lu1.

You need to build a Conversational 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.

Select and Place:

Actions	Answer Area
Add the prebuilt domain ToDo.	
Add a new application.	
Add example utterances.	
Train the application.	
Publish the application.	

Correct Answer:

Actions	Answer Area
Add the prebuilt domain ToDo.	Add a new application.
	Add example utterances.
	Train the application.
	Publish the application.

Section:

Explanation:

QUESTION 11

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

Correct Answer: B

Section:

QUESTION 12

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 SQL query. What should you create?

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

Correct Answer: B

Section:

QUESTION 13

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

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

Correct Answer: A, D

Section:

QUESTION 14

HOTSPOT

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

Which RBAC role should you assign to each group? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Management-Accountants	<input type="text"/> Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor
Consultant-Accountants	<input type="text"/> Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor
Agent-CustomerServices	<input type="text"/> Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor

Answer Area:



Section:

Explanation:

Box 1: Cognitive Service User

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

Approve=publish.

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

1. Add new members to roles.
2. Create new resources.

Box 2: Cognitive Services QnA Maker Editor

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

QnA Maker Editor: API permissions:

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

Box 3: Cognitive Services QnA Maker Read

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

QnA Maker Read: API Permissions:

1. Download KB API
2. List KBs for user API

3. Get Knowledge base details

4. Download AlterationsGenerate Answer

Reference: <https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/concepts/role-based-access-control>

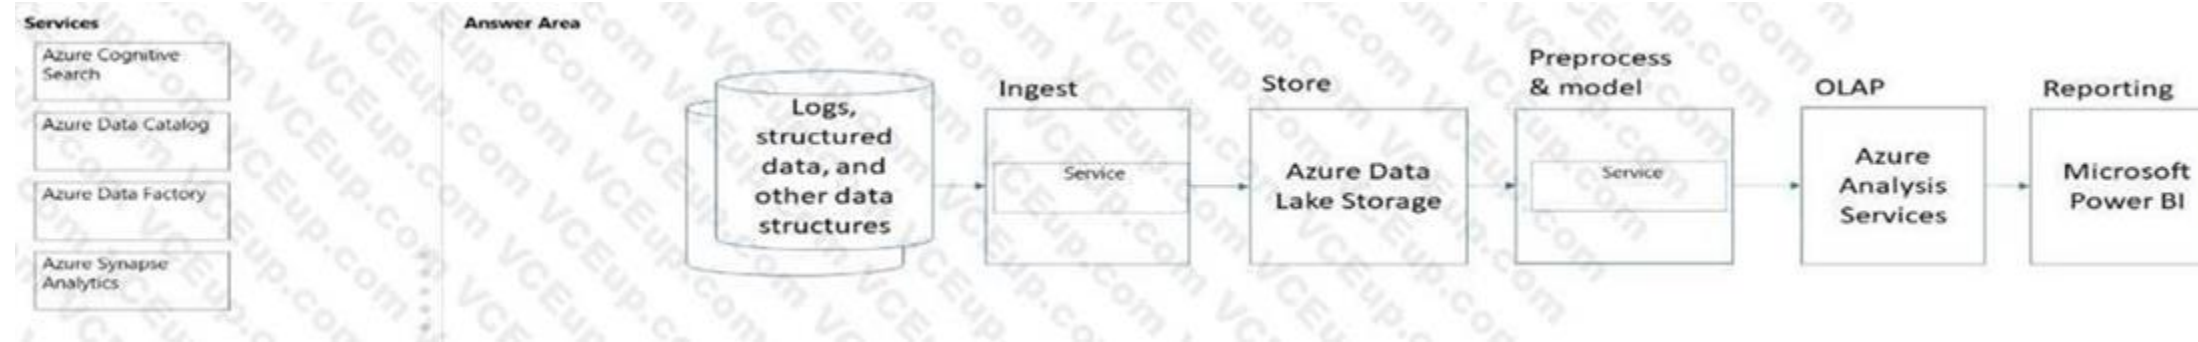
QUESTION 15

DRAG DROP

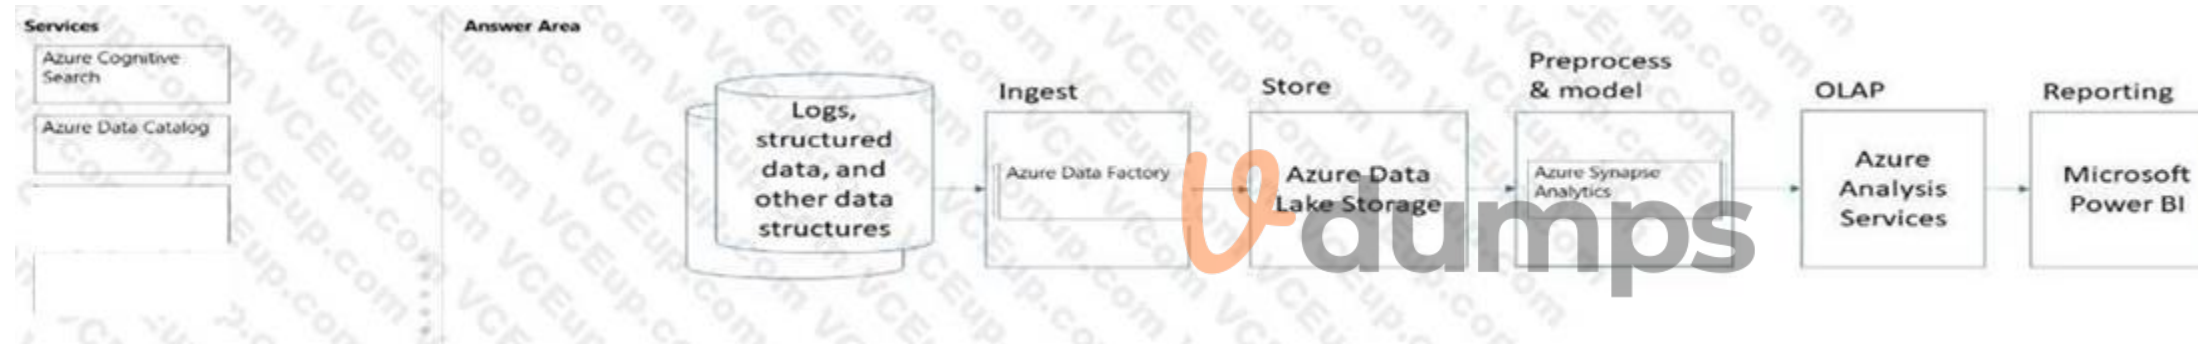
Match the Azure services to the appropriate locations in the architecture.

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.

Select and Place:



Correct Answer:



Section:

Explanation:

QUESTION 16

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	<input type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	<input type="radio"/>	<input type="radio"/>
All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	<input type="radio"/>	<input type="radio"/>

Answer Area:

Answer Area

Statements	Yes	No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	<input checked="" type="radio"/>	<input type="radio"/>
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	<input type="radio"/>	<input checked="" type="radio"/>
All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	<input type="radio"/>	<input checked="" type="radio"/>

Section:

Explanation:

QUESTION 17

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
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"/>

Answer Area:

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 checked="" type="radio"/>	<input type="radio"/>

Section:

Explanation:

QUESTION 18

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A block of code that runs in a database is called

a stored procedure.

a table.

a view.

an index.

Answer Area:

Answer Area

A block of code that runs in a database is called

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

Section:

Explanation:

QUESTION 19

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

The

- ALTER
- JOIN
- SET
- WHERE

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

Answer Area:

Answer Area

The

- ALTER
- JOIN
- SET
- WHERE

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

dumps

Section:

Explanation:

QUESTION 20

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

In an infrastructure as a service (IaaS) instance of Microsoft SQL Server on Azure, you manage the

- elastic pool
- MySQL server
- PostgreSQL server
- virtual machine

that hosts SQL Server.

Answer Area:

Answer Area

In an infrastructure as a service (IaaS) instance of Microsoft SQL Server on Azure, you manage the

- elastic pool
- MySQL server
- PostgreSQL server
- virtual machine

that hosts SQL Server.

Section:

Explanation:

QUESTION 21

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

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

- changes to relationships between entities.
- geographically distributed writes.
- transactional writes.
- writes that have varying data structures.

Answer Area:

Answer Area

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

- changes to relationships between entities.
- geographically distributed writes.
- transactional writes.
- writes that have varying data structures.

Section:

Explanation:

QUESTION 22

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

A JSON document is an example of

- graph data.
- relational data.
- semi-structured data.
- unstructured data.

Answer Area:

Answer Area

A JSON document is an example of

- graph data.
- relational data.
- semi-structured data.
- unstructured data.

Section:

Explanation:

QUESTION 23

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:



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.

Answer Area:

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.

Section:

Explanation:

QUESTION 24

HOTSPOT

Select the answer that correctly completes the sentence.

Hot Area:

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a

- graph.
- table.
- partition key.
- document.

Answer Area:

Answer Area

When using the Azure Cosmos DB Gremlin API, the container resource type is projected as a

- graph.
- table.
- partition key.
- document.

Section:

Explanation:

QUESTION 25

You need to store event log data that is semi-structured and received as the logs occur. What should you use?

- A. Azure Table storage
- B. Azure Queue storage
- C. Azure Files

Correct Answer: A

Section:

QUESTION 26

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

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

Correct Answer: D

Section:

QUESTION 27

What is a characteristic of a non-relational database?

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

Correct Answer: C

Section:

QUESTION 28

You have a factory that produces food products.

You need to build a monitoring solution for staff compliance with personal protective equipment (PPE) requirements. The solution must meet the following requirements:

- identify staff who have removed masks or safety glasses.
- Perform a compliance check every 15 minutes.
- Minimize development effort.
- Minimize costs.

Which service should you use?

- A. Face
- B. Computer Vision
- C. Azure Video Analyzer for Media (formerly Video indexer)

Correct Answer: A

Section:

QUESTION 29

HOTSPOT

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.

Hot Area:



Answer Area

The screenshot shows two dropdown menus in the Azure portal. The top dropdown is for selecting a subscription, with the URL `https://management.azure.com/` on the left and `/resourceGroups/OCRProject/providers/` on the right. The options are: `subscriptions/8d3591aa-96b8-4737-ad09-00f9b1ed35ad` (selected), `subscriptions/3edfe572-cb54-3ced-ae12-c5c177f39a12`, `subscriptions/8d3591aa-96b8-4737-ad09-00f9b1ed35ad`, `tenant/3edfe572-cb54-3ced-ae12-c5c177f39a12`, and `tenant/8d3591aa-96b8-4737-ad09-00f9b1ed35ad`. The bottom dropdown is for selecting a provider, with the URL `/accounts/CS1?api-version=2021-10-01` on the right. The options are: `Microsoft.CognitiveServices` (selected), `Microsoft.ApiManagement`, `Microsoft.CognitiveServices`, `Microsoft.ContainerService`, and `Microsoft.KeyVault`.

Answer Area:

Answer Area

This screenshot is identical to the one above, but with a large 'Vdumps' watermark overlaid in the center. The top dropdown shows the same subscription options, and the bottom dropdown shows the same provider options.

Section:

Explanation:

QUESTION 30

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

Correct Answer: C, D

Section:

QUESTION 31

You are building an AI 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 AI 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.

Correct Answer: A

Section:

QUESTION 32

You have a chatbot that was built by using Microsoft Bot Framework and deployed to Azure.

You need to configure the bot to support voice interactions. The solution must support multiple client apps. Which type of channel should you use?

- A. Cortana
- B. Microsoft Teams
- C. Direct Line Speech

Correct Answer: C

Section:



QUESTION 33

You are building a retail kiosk system that will use a custom neural voice. You acquire audio samples and consent from the voice talent. You need to create a voice talent profile. What should you upload to the profile?

- A. a five-minute wav or mp3 file of the voce talent describing the kiosk system
- B. a five-minute wav or mp3 file of the voce talent describing the kiosk system
- C. a five-minute .flac audio file and the associated transcript as a w file
- D. a .wav or mp3 file of the voice talent consenting to the creation of a synthetic version of their voice
- E. a.zip file that contains 10-second .wav files and the associated transcripts as .txt files

Correct Answer: D

Section:

QUESTION 34

You have an app named App1 that uses an Azure Cognitive Services model to identify anomalies in a time series data stream. You need to run App1 in a location that has limited connectivity. The solution must minimize costs. What should you use to host the model?

- A. Azure Kubernetes Services (AKS)
- B. a Kubernetes cluster hosted in an Azure Stack Hub integrated system
- C. Azure Container instances
- D. the Docker Engine

Correct Answer: B

Section:

QUESTION 35

You plan create an index for an Azure Cognitive Search service by using the Azure portal. The Cognitive Search service will connect to an Azure SQL database The Azure SQL database contains a table named UserMessages. Each row m User Messages has a field named MessageCopy that contains the text of social media messages sent by a user Users win perform full text searches against the MessageCopy field, and the values of the field will be shown to the users. You need to configure the properties of the index for the MessageCopy field to support the solution. Winch attributes should you enable for the field?

- A. Searchable arc Retrievable
- B. Sortable and Retrievable
- C. Searchable arc Facetable
- D. Filterable and Retrievable

Correct Answer: A

Section:

QUESTION 36

HOTSPOT

You are developing a text processing solution.

You have the function shown below.

```
static void GetKeywords(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.RecognizeEntities (text);
    Console.WriteLine("Key words:");

    foreach (CategorizedEntity entity in response.Value)
    {
        Console.WriteLine($"{entity.Text}");
    }
}
```



Hot Area:

Answer Area

Statements

- The output will include the following words: our and included.
- The output will include the following words: Paris, Eiffel, and Tower.
- The function will output all the key phrases from the input string to the console.

Yes

No

Answer Area:

Answer Area

Statements

- The output will include the following words: our and included.
- The output will include the following words: Paris, Eiffel, and Tower.
- The function will output all the key phrases from the input string to the console.

Yes	No
<input checked="" type="radio"/>	<input type="radio"/>
<input checked="" type="radio"/>	<input type="radio"/>
<input checked="" type="radio"/>	<input type="radio"/>

Section:
Explanation:

QUESTION 37

HOTSPOT

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.

Hot Area:

Answer Area

Service endpoint:

- Computer Vision analyze images
- Computer Vision object detection
- Custom Vision image classification
- Custom Vision object detection

Property:

- categories
- description
- imageType
- metadata
- objects

Answer Area:

Answer Area

Service endpoint:

- Computer Vision analyze images
- Computer Vision object detection
- Custom Vision image classification
- Custom Vision object detection

Property:

- categories
- description
- imageType
- metadata
- objects

Section:
Explanation:

QUESTION 38

You use the Microsoft Bot Framework Composer to build a chatbot that enables users to purchase items. You need to ensure that the users can cancel in-progress transactions. The solution must minimize development effort.

What should you add to the bot?

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

Correct Answer: D

Section:

QUESTION 39

You have an Azure subscription that contains an Anomaly Detector resource. You deploy a Docker host server named Server 1 to the on-premises network. You need to host an instance of the Anomaly Detector service on Server 1. Which parameter should you include in the docker run command?

- A. Fluentd
- B. Billing
- C. Http Proxy
- D. Mounts

Correct Answer: B

Section:

QUESTION 40

HOTSPOT

You have a chatbot.

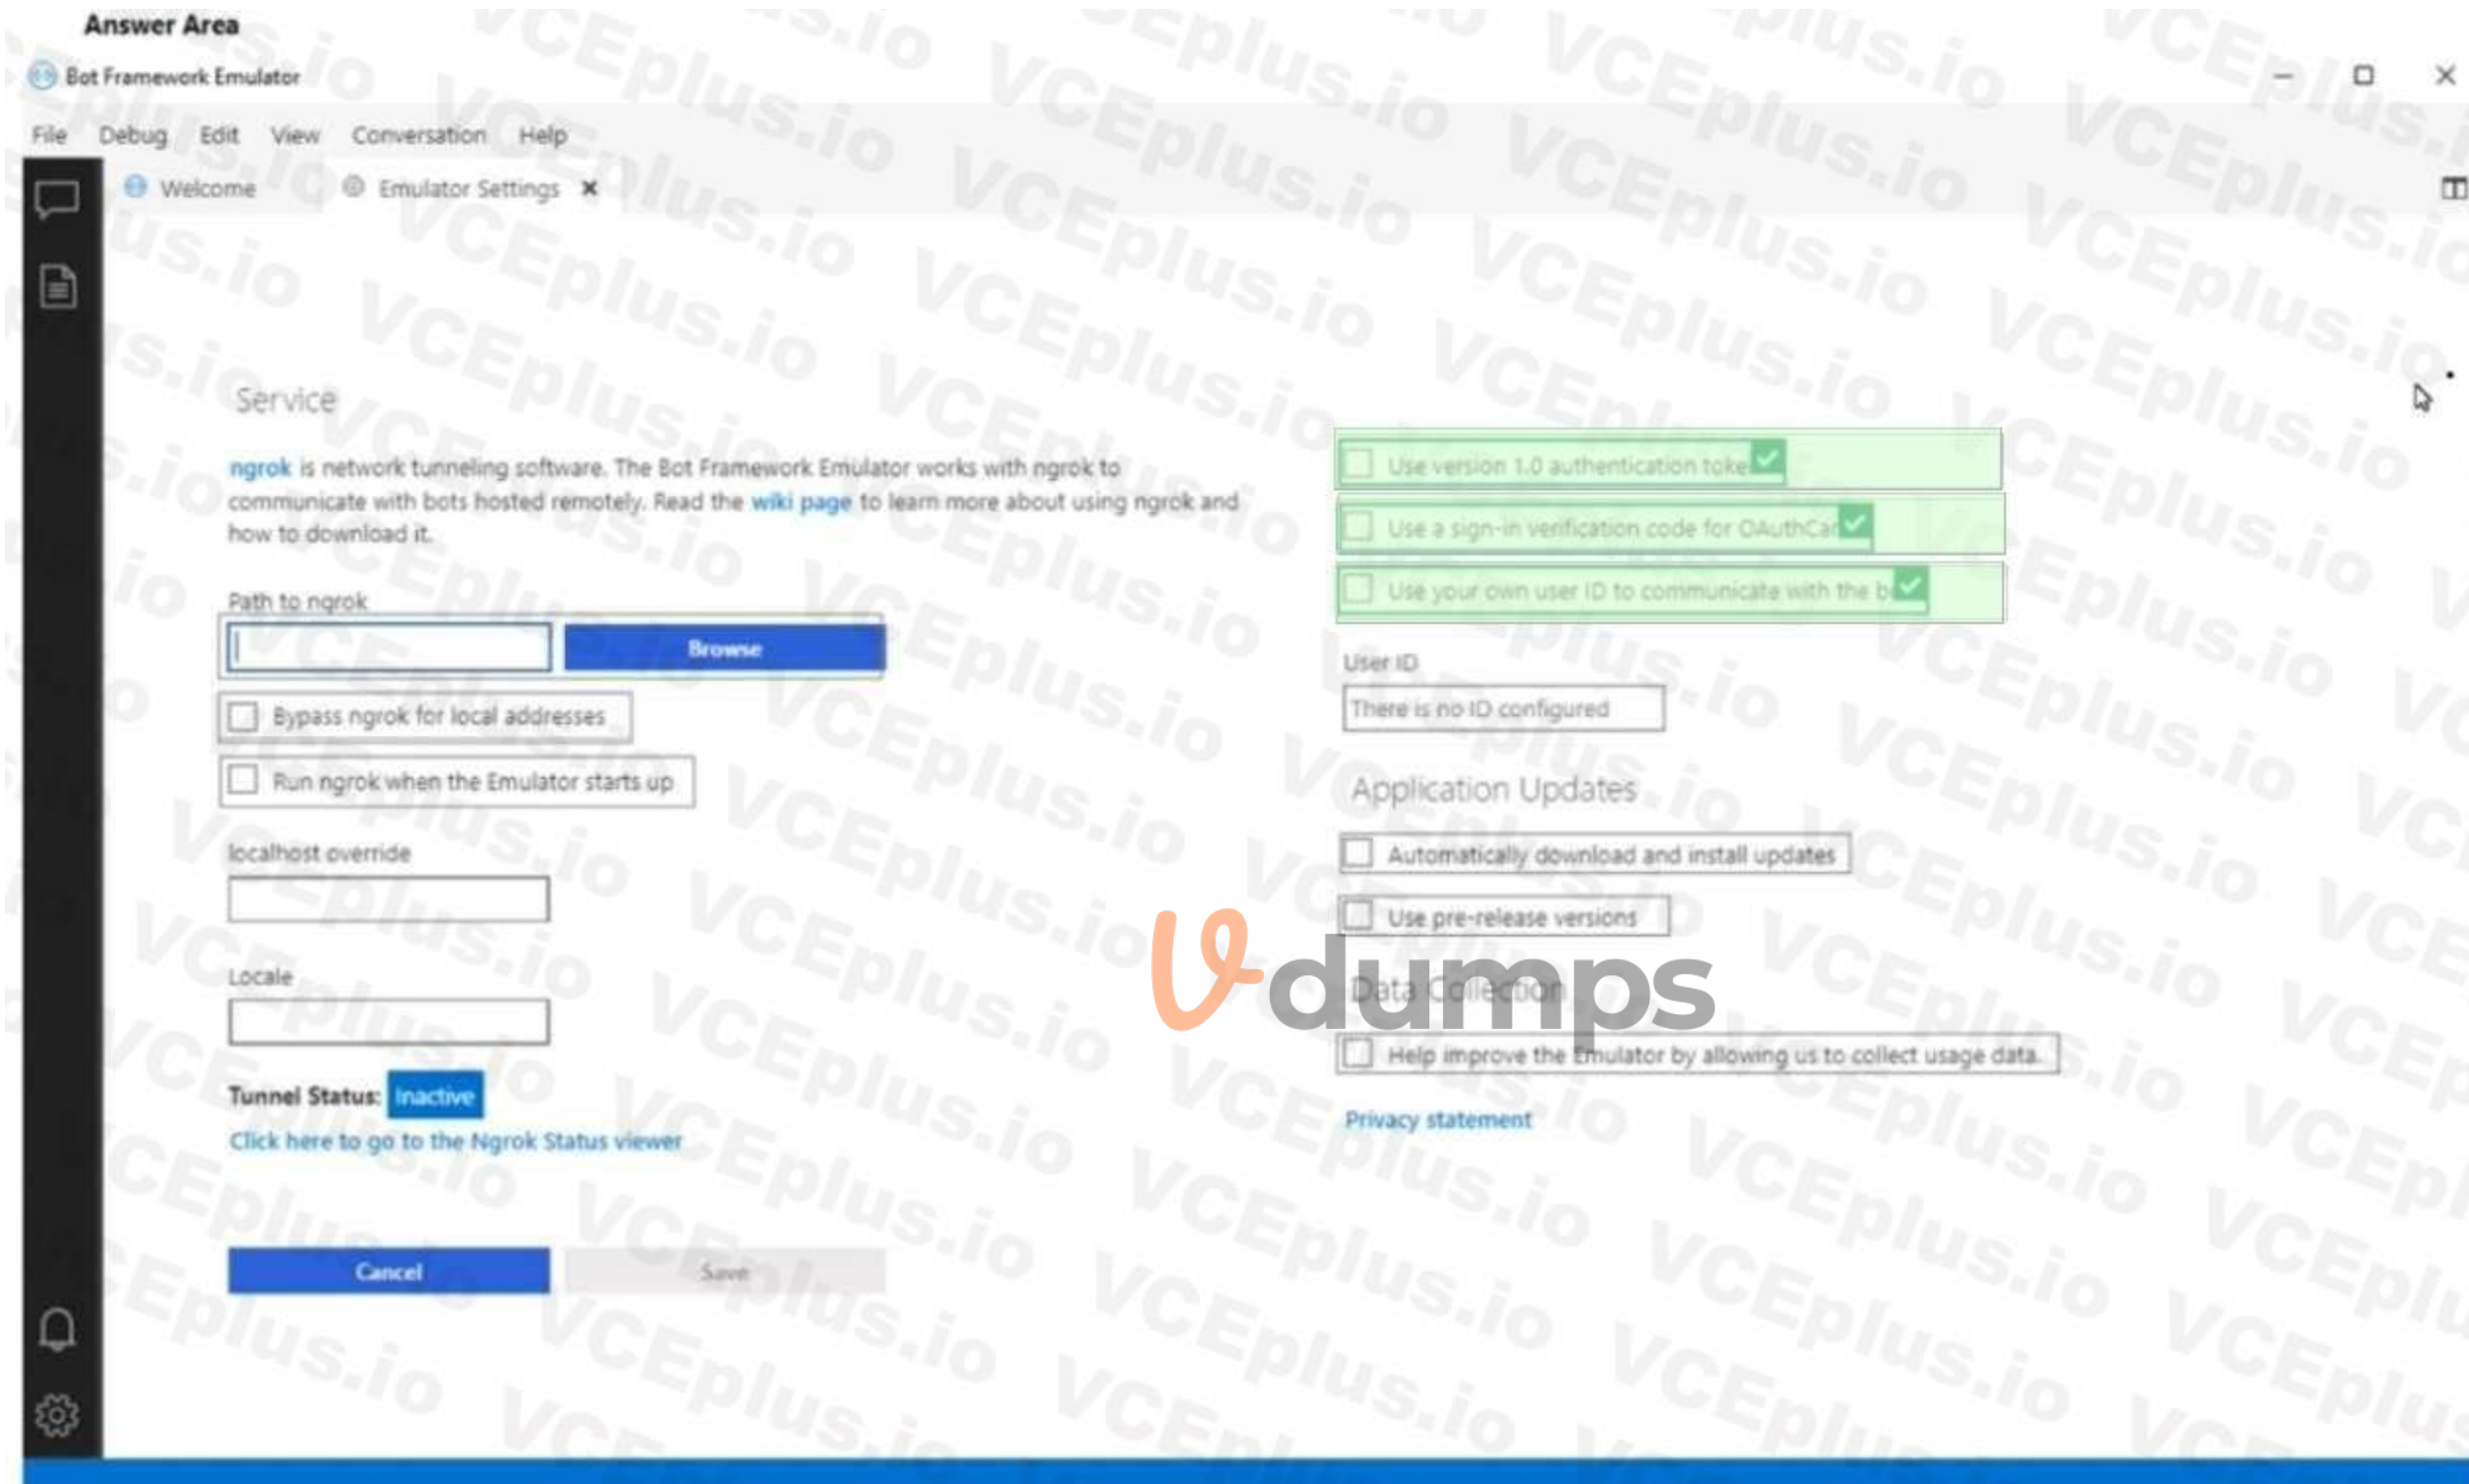
You need to test the bot by using the Bot Framework Emulator. The solution must ensure that you are prompted for credentials when you sign in to the bot. Which three settings should you configure? To answer, select the appropriate settings in the answer area. NOTE Each correct selection is worth one point.

Hot Area:





Answer Area:



Section:

Explanation:

QUESTION 41

You have an Azure Cognitive Search solution and an enrichment pipeline that performs Sentiment Analysis on social media posts.

You need to define a knowledge store that will include the social media posts and the Sentiment Analysis results.

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

NOTE: Each correct selection is worth one point.

A. storageContainer

- B. tables
- C. storageConnectionString
- D. files
- E. objects

Correct Answer: D, E

Section:

QUESTION 42

HOTSPOT

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.

Hot Area:

Answer Area

.....

To capture messages:

- Speaker Recognition
- Speech-to-text**
- Text-to-speech
- Translator

To replay messages:

- Speech-to-text only
- Speech-to-text and Language
- Speaker Recognition and Language
- Text-to-speech and Language
- Text-to-speech and Translator**

Answer Area:

Answer Area

.....

To capture messages:

- Speaker Recognition
- Speech-to-text**
- Text-to-speech
- Translator

To replay messages:

- Speech-to-text only
- Speech-to-text and Language
- Speaker Recognition and Language
- Text-to-speech and Language
- Text-to-speech and Translator**

Section:

Explanation:

QUESTION 43

HOTSPOT

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.

Hot Area:

Answer Area

https://

api.cognitive.microsofttranslator.com	/text/analytics/v3.1/entities/recognition/general
api.cognitive.microsofttranslator.com	/text/analytics/v3.1/entities/recognition/general
eastus.api.cognitive.microsoft.com	/text/analytics/v3.1/languages
portal.azure.com	/translator/text/v3.0/translate?to=en
	/translator/text/v3.0/translate?to=fr

Answer Area:

Answer Area

https://

api.cognitive.microsofttranslator.com	/text/analytics/v3.1/entities/recognition/general
api.cognitive.microsofttranslator.com	/text/analytics/v3.1/entities/recognition/general
eastus.api.cognitive.microsoft.com	/text/analytics/v3.1/languages
portal.azure.com	/translator/text/v3.0/translate?to=en
	/translator/text/v3.0/translate?to=fr

Section:


Explanation:

QUESTION 44

HOTSPOT

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.



Hot Area:

Answer Area

The percentage of false positives is [answer choice].

0
25
30
50
100

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

0
25
30
50
100

Answer Area:

Answer Area

The percentage of false positives is [answer choice].

0
25
30
50
100

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

0
25
30
50
100



Section:

Explanation:

QUESTION 45

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

Correct Answer: D

Section:

QUESTION 46

You train a Conversational Language Understanding model to understand the natural language input of users. You need to evaluate the accuracy of the model before deploying it.

What are two methods you can use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Correct Answer: A, C

Section:

QUESTION 47

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

Correct Answer: C

Section:

QUESTION 48

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

Correct Answer: B, C

Section:

QUESTION 49

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

Correct Answer: B

Section:

QUESTION 50

DRAG DROP

You have a Docker host named Host1 that contains a container base image.

You have an Azure subscription that contains a custom speech-to-text model named model1.

You need to run model1 on Host1.

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.

Select and Place:

Actions

- Configure disk logging.
- Export model1 to Host1.
- Request approval to run the container.
- Retrain the model.
- Run the container.

Answer Area



Correct Answer:

Actions

-
-
- Request approval to run the container.
- Retrain the model.
-



Answer Area

- Export model1 to Host1.
- Run the container.
- Configure disk logging.



Section:

Explanation:

Export model1 to Host1

Run the container

Configure disk logging

QUESTION 51

DRAG DROP

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

You need to test the bot interactively on a local machine.

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.

Select and Place:

Actions

- Register the bot with the Azure Bot Service.
- Open the Bot Framework Composer.
- Build and run the bot.
- Open the Bot Framework Emulator.
- Connect to the bot endpoint.

Answer Area

Correct Answer:

Actions

- Register the bot with the Azure Bot Service.
-
- Open the Bot Framework Emulator.

Answer Area

- Build and run the bot.
- Open the Bot Framework Composer.
- Connect to the bot endpoint.

Section:

Explanation:

QUESTION 52

DRAG DROP

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 o' actions to the answer area and arrange them in the correct order.

Select and Place:

Actions

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

Answer Area

Correct Answer:

Actions

- Change Domains to **General (compact)**.
- Retrain the model.
-

Answer Area

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

Section:

Explanation:

QUESTION 53

DRAG DROP

You have a question answering project in Azure Cognitive Service for Language.

You need to move the project to a Language service instance in a different Azure region.

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.

Select and Place:

Actions	Answer Area
From the new Language service instance, import the project file.	
From the new Language service instance, enable custom text classification.	
From the new Language service instance, train and publish the project.	
From the original Language service instance, export the existing project.	
From the new Language service instance, regenerate the keys.	
From the original Language service instance, train and publish the model.	

Correct Answer:

Actions	Answer Area
From the new Language service instance, import the project file.	From the original Language service instance, export the existing project.
From the new Language service instance, enable custom text classification.	From the new Language service instance, regenerate the keys.
From the new Language service instance, train and publish the project.	From the original Language service instance, train and publish the model.

Section:

Explanation:

QUESTION 54

DRAG DROP

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.

Select and Place:

Actions	Answer Area
Pull an image from Docker Hub.	
Provision an on-premises Kubernetes cluster that has internet connectivity.	
Provision an Azure Kubernetes Service (AKS) resource.	
Run the container and specify an App ID and Client Secret.	
Provision an on-premises Kubernetes cluster that is isolated from the internet.	
Pull an image from the Microsoft Container Registry (MCR).	
Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.	

Correct Answer:

Actions	Answer Area
Pull an image from Docker Hub.	Provision an on-premises Kubernetes cluster that is isolated from the internet.
Provision an on-premises Kubernetes cluster that has internet connectivity.	Pull an image from the Microsoft Container Registry (MCR).
Provision an Azure Kubernetes Service (AKS) resource.	Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.
Run the container and specify an App ID and Client Secret.	

Section:

Explanation:

Provision an on-premises Kubernetes...

Pull an image from the Microsoft Container Registry...

Run the container and specify an API key...

QUESTION 55

You are developing a monitoring system that will analyze engine sensor data, such as rotation speed, angle, temperature, and pressure. The system must generate an alert in response to atypical values. What should you include in the solution?

- A. Application Insights in Azure Monitor
- B. metric alerts in Azure Monitor
- C. Multivariate Anomaly Detection
- D. Univariate Anomaly Detection

Correct Answer: C

Section:

QUESTION 56

DRAG DROP

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.

Select and Place:



Values

- addTargetLanguage
- speechSynthesisLanguage
- speechRecognitionLanguage
- voiceName

Answer Area

```
var translationConfig = SpeechTranslationConfig.FromSubscription(SPEECH_SUBSCRIPTION_KEY, SPEECH_SERVICE_REGION);
translationConfig. [ ] = "en-US";
translationConfig. [ ] ("de");
```

Correct Answer:

Values

- []
- speechSynthesisLanguage
- []
- voiceName

Answer Area

```
var translationConfig = SpeechTranslationConfig.FromSubscription(SPEECH_SUBSCRIPTION_KEY, SPEECH_SERVICE_REGION);
translationConfig. speechRecognitionLanguage = "en-US";
translationConfig. addTargetLanguage ("de");
```



Section:

Explanation:

QUESTION 57

DRAG DROP

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.

Select and Place:

Actions

- Create a Speaker Recognition model.
- Create a Conversational Language Understanding model.
- Create a Custom Speech project.
- Create a speech-to-text model.
- Upload training datasets.
- Train the model.
- Deploy the model.



Answer Area



Correct Answer:

Actions

- Create a Speaker Recognition model.
- Create a Conversational Language Understanding model.
-
-
-
-
-



Answer Area

- Create a Custom Speech project.
- Create a speech-to-text model.
- Upload training datasets.
- Train the model.
- Deploy the model.



Section:

Explanation:

- Create a Custom Speech project.
- Create a speech-to-text model.
- Upload training datasets.
- Train the model.
- Deploy the model.

QUESTION 58

You have a file share that contains 5,000 images of scanned invoices. You need to analyze the images. The solution must extract the following data:

- * Invoice items
- * Sales amounts
- * Customer details

What should you use?

- A. Custom Vision

- B. Computer Vision
- C. Immersive Reader
- D. Form Recognizer

Correct Answer: C

Section:

QUESTION 59

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

- * Include the category field in the search results.
- * Ensure that users can search for words in the category field.
- * Ensure that users can perform drill down filtering based on category.

Which index attributes should you configure for the category field?

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

Correct Answer: B

Section:

QUESTION 60

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 are building a chatbot that will use question answering in Azure Cognitive Service for Language.

You upload Doc1.pdf and train that contains a product catalogue and a price list.

During testing, users report that the chatbot responds correctly to the following question: What is the price of ?

The chatbot fails to respond to the following question: How much does cost?

You need to ensure that the chatbot responds correctly to both questions.

Solution: from Language Studio, you create an entity for price, and then retrain and republish the model.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section:

QUESTION 61

You are developing a system that will monitor temperature data from a data stream. The system must generate an alert in response to atypical values. The solution must minimize development effort.

What should you include in the solution?

- A. Univariate Anomaly Detection
- B. Azure Stream Analytics
- C. metric alerts in Azure Monitor
- D. Multivariate Anomaly Detection

Correct Answer: D

Section:

