ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

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.

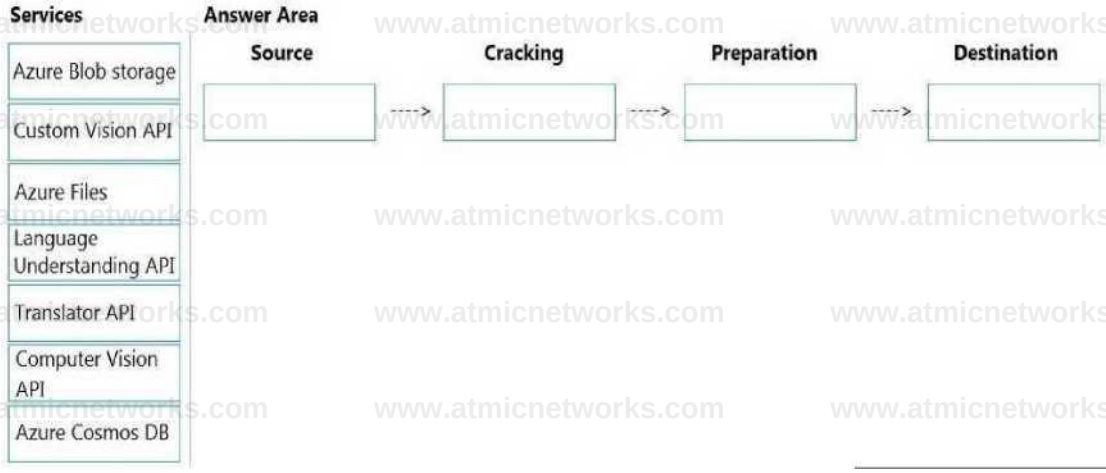

Answer:

Explanation:

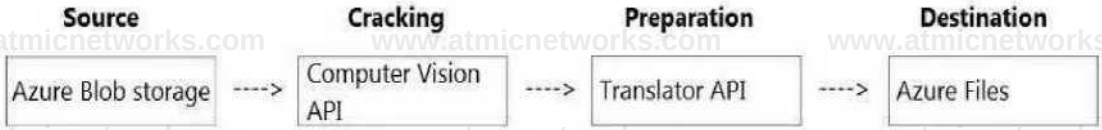

Box 1: Azure Blob storage

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

Box 2: Computer Vision API

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

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

Box 3: Translator API

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

Box 4: Azure Files

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

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

Question: 2

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.)

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

Answer:

Explanation:

Actions	Answer Area
Index the video by using the Video Indexer API.	Upload the video to blob storage.
Upload the video to blob storage.	Index the video by using the Video Indexer API.
Analyze the video by using the Computer Vision API.	Extract the transcript from the Video Indexer API.
Extract the transcript from Microsoft Stream.	Translate the transcript by using the Translator API.
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.	

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

Step 1: Upload the video to blob storage
Given a video or audio file, the file is first dropped into a Blob Storage. T

Step 2: Index the video by using the Video Indexer API.
When a video is indexed, Video Indexer produces the JSON content that contains details of the specified video insights. The insights include: transcripts, OCRs, faces, topics, blocks, etc.

Step 3: Extract the transcript from the Video Indexer API.

Step 4: Translate the transcript by using the Translator API.

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: 3

HOTSPOT

You are planning the product creation project.

You need to build the REST endpoint to create the multilingual product descriptions.

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

NOTE: Each correct selection is worth one point.

Answer Area

▼ ▼ ?api-version=3.Osto=esSto=pt

api.cognitive.microsofttranslator.com	/detect
api-nam.cognitive.microsofttranslator.com	/languages
westus.tts.speech.microsoft.com	/text-to-speech
wwics.cognitiveservices.azure.com/translator	/translate

Answer:

Explanation:

Box 1: api-nam.cognitive.microsofttranslator.com
https://docs.microsoft.com/en-us/azure/cognitive-services/translator/reference/v3-0-reference

Box 2: /translate

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/translator/reference/v3-0-translate

Question: 4

You are developing the smart e-commerce project.

You need to implement autocompletion as part of the Cognitive Search solution.

Which three actions should you perform? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

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

Answer: ABF

Explanation:

Scenario: Support autocompletion and autosuggestion based on all product name variants.

A: Call a suggester-enabled query, in the form of a Suggestion request or Autocomplete request, using an

API. API usage is illustrated in the following call to the Autocomplete REST API.

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

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

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

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

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

Question: 5

HOTSPOT

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

You are configuring access to the QnA Maker resources.

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

NOTE: Each correct selection is worth one point.

Answer Area

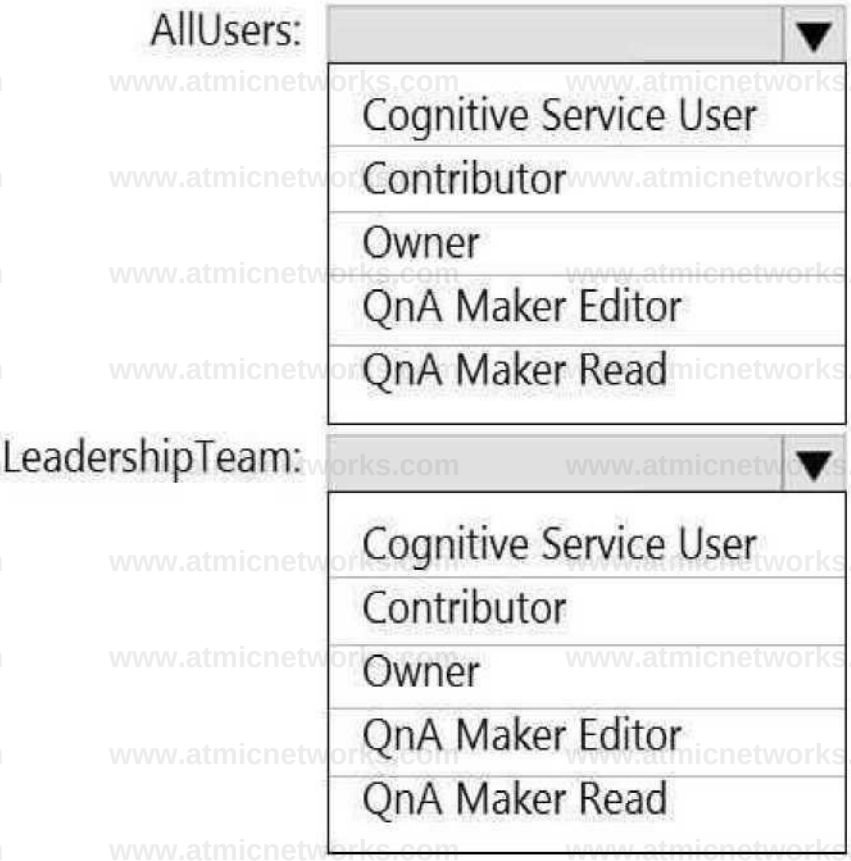

Answer:

Explanation:

Answer Area

Box 1: QnA Maker Editor
Scenario: Provide all employees with the ability to edit Q&As.

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

Box 2: Contributor
Scenario: Only senior managers must be able to publish updates.
Contributor permission: All except ability to add new members to roles

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

Question: 7

HOTSPOT

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

You need to build the Adaptive Card for the chatbot.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
version": "1.3 "body": [

type": "TextBlock", size": "Medium",
weight": "Bolder", text": [
```

(((language = = ‘en’, ‘en’, name) name name.en
name[language]

```
"type": "TextBlock",
```

"$when': “${stockLevel! = 'OK’}’ ■$when’:“${sto(*Level = = W
Swhen" S{stockLevel OK}"

```
color : Attention

"type                    "Image",
"url"  "S{image.uri}", "Medium",
": "${
alText
```

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

Answer:

Explanation:

```
version": "1.3", "body": [ {

"type":    "TextBlock",    "size":    "Medium",
"weight": "Bolder", "text": "${|     i
```

if(language = = en’, en’. name) name name.en
name[language] b

```
"type": "TextBlock",
```

'Swhen" "${stockLevel I = 'OK’}" Swhen" "${stockLevel = = OK }
“$when": “${stockLevel.OK}"

```
color : Attention
```

```
{
        "type": "Image",
        "url": "S{image.uri}", "size": "Medium",
        "alText": "${ "
```

- limage.altText.en
  image.altText language image altText[language] image altText[language]

Box 1: name.en

Box 2: "$when": "${stockLevel != 'OK'}"
Product displays must include images and warnings when stock levels are low or out of stock.

Box 3:image.altText.en

Topic 2, Contoso, Ltd.

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.

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

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

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:

- ICountryJ-[Levell-[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	1 nt er na 1 .contoso.com/wiki
Monthly conference recordings	Video	English	SharePoint Online	Contoso-sharepoint.com
Frequently asked questions (FAQs)	Text	English	SharePoint Online	Contoso.sliatepoint.com

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

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: 8

You need to develop an extract solution for the receipt images. The solution must meet the document processing requirements and the technical requirements.

You upload the receipt images to the From Recognizer API for analysis, and the API returns the following JSON.

```
"documentResults":[
        {
    "docType":"prebuiIt:receipt", "pageRange":[
      1,
      1
                    ],
                    "fields":{
  "ReceiptType":{
    "type":"string",
    "valuestring":" 11* mi . *1",
    "conf idence":0.672

  "MerchantName":{
    "type":"str ing",
    "valuestring":"Tailwind",
    "text":"Tailwind",
    "boundingBox":[],
    "page":1,
    "confidence":0.913,
    "elements":[
     "#/readResults/0/lines/0/words/0"
```

Which expression should you use to trigger a manual review of the extracted information by a member of the Consultant-Bookkeeper group?

A. documentResults.docType == "prebuilt:receipt"

B. documentResults.fields.".confidence < 0.7

C. documentResults.fields.ReceiptType.confidence > 0.7

D. documentResults.fields.MerchantName.confidence < 0.7

Answer: D

Explanation:

Need to specify the field name, and then use < 0.7 to handle trigger if confidence score is less than 70%.

Reference:
https://docs.microsoft.com/en-us/azure/applied-ai-services/form-recognizer/api-v2-0/reference- sdk-api-v2-0

Question: 9

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.

Answer Area

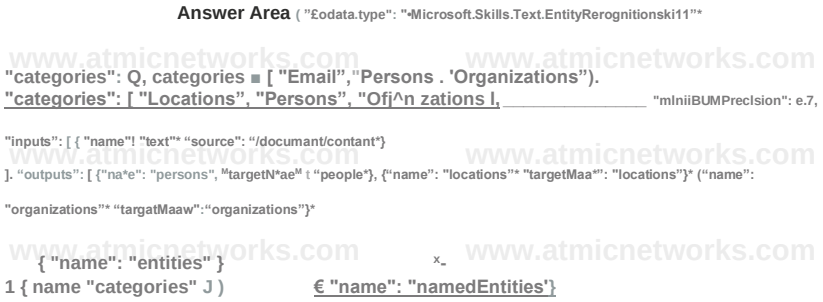

Answer:

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: 10

You are developing the chatbot.

You create the following components:

- A QnA Maker resource
- A chatbot by using the Azure Bot Framework SDK

You need to add an additional component to meet the technical requirements and the chatbot requirements. What should you add?

A. Dispatch
B. chatdown
C. Language Understanding
D. Microsoft Translator

Answer: A

Explanation:

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

If a bot uses multiple LUIS models and QnA Maker knowledge bases (knowledge bases), you can use the Dispatch tool to determine which LUIS model or QnA Maker knowledge base best matches the user input. The dispatch tool does this by creating a single LUIS app to route user input to the correct model.

Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-tutorial-dispatch

Question: 11

You are developing the document processing workflow.
You need to identify which API endpoints to use to extract text from the financial documents. The solution must meet the document processing requirements.
Which two API endpoints should you identify? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. /vision/v3.2/read/analyzeResults
B. /formrecognizer/v2.0/prebuilt/receipt/analyze
C. /vision/v3.2/read/analyze
D. /vision/v3.2/describe
E. /formercognizer/v2.0/custom/models{modelId}/ analyze

Answer: BC

Explanation:

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

E: POST {Endpoint}/vision/v3.2/read/analyze
Use this interface to get the result of a Read operation, employing the state-of-the-art Optical Character Recognition (OCR) algorithms optimized for text-heavy documents.

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

The document processing solution must be able to process standardized financial documents that have the following characteristics:

- Contain fewer than 20 pages.
- Be formatted as PDF or JPEG files.
- Have a distinct standard for each office.

*The document processing solution must be able to extract tables and text from the financial documents.
The document processing solution must be able to extract information from receipt images.

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

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

Question: 12

You are developing the chatbot.
You create the following components:
* A QnA Maker resource
* A chatbot by using the Azure Bot Framework SDK.
You need to integrate the components to meet the chatbot requirements.
Which property should you use?

A. QnADialogResponseOptions.CardNoMatchText
B. Qna MakerOptions-ScoreThreshold
C. Qna Maker Op t ions StrickFilters
D. QnaMakerOptions.RankerType

Answer: D

Explanation:

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

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

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

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

Question: 13

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.

Answer Area

Answer:

Explanation:

Box 1: Cognitive Service User

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

Approve=publish.

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

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

Box 2: Cognitive Services QnA Maker Editor

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

QnA Maker Editor: API permissions:

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

Box 3: Cognitive Services QnA Maker Read

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

QnA Maker Read: API Permissions:

1. Download KB API
2. List KBs for user API
3. Get Knowledge base details
4. Download Alterations

Generate Answer

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

Question: 14

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 From Recognizer resource
An Azure web app that hosts the Form Recognizer sample labeling tool

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

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

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

Answer:

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: 15

You are developing the knowledgebase.

You use Azure Video Analyzer for Media (previously Video indexer) to obtain transcripts of webinars.

You need to ensure that the solution meets the knowledgebase requirements.

What should you do?

A. Create a custom language model
B. Configure audio indexing for videos only
C. Enable multi-language detection for videos
D. Build a custom Person model for webinar presenters

Answer: B

Explanation:

Can search content in different formats, including video

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

Keywords extraction: Extracts keywords from speech and visual text.

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

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.

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

Question: 16

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

Answer: C

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.

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: 17

You are developing the knowledgebase by using Azure Cognitive Search.

You need to meet the knowledgebase requirements for searching equivalent terms.

What should you include in the solution?

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

Answer: A

Explanation:

Within a search service, synonym maps are a global resource that associate equivalent terms, expanding the scope of a query without the user having to actually provide the term. For example, assuming "dog", "canine", and "puppy" are mapped synonyms, a query on "canine" will match on a document containing "dog".

Create synonyms: A synonym map is an asset that can be created once and used by many indexes.

Reference:

https://docs.microsoft.com/en-us/azure/search/search-synonyms

Topic 3, Misc. Questions

Question: 18

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.

Values

```
|AddPhraseListAsync
|Phraselist
PhraselistCreateObject
Phrases
|SavePhrasellstAsync
UploadPhraseListAsync
```

Answer Area

```
var phraselistld = await client.Features.                    |
(appld, versionld, new                                       |

    EnabledForAllModels = false.
    IsExchangeable = true, Name = "PLl",
    Phrases = "iteml,item2,items, item4, items" });
```

Answer:

Explanation:

Values

```
lAdaPhraseListAsync
|Phraselist
|PhraselistCreateObject|
|Phrases
|SavePhrasellstAsync
|UploadPhraseListAsync |
```

Answer Area

```
var phraselistld - await client.Features. |AddPhraseLlstAsync
(appld, versionld, new |PhraselistCreateObject| {
    EnabledForAllModels = false,
    IsExphangeabla = true.
    Name - "PLl",
    Phrases = "Iteml,item2,item3,item4,items"

});
```

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: 19

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*
Vl.l	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.)

Actions

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

Export the model by using the Export as JSON option.

Select vl.l of appl.

Run a container and mount the model file.

Select vl.O of appl.

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

Select vl.2 of appl.

Answer Area

Answer:

Explanation:

Actions		Answer Area
Run a container that has version set as an environment variable.		Export the model by using the Export for containers (GZIP) option.
Export the model by using the Export as JSON option.		Select vl.l of appl.
Select vl.l of appl.		Run a container and mount the model file.

Run a container and mount the model file.

Select vl.O of appl.

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

Select vl.2 of appl.

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.

Sign on to the LUIS portal.

Select the app in the list.

Select Manage in the app's navigation bar.

Select Versions in the left navigation bar.

Select the checkbox to the left of the version name in the list.

Select the Export item from the contextual toolbar above the list.

Select Export for container (GZIP).

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: 20

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

Supports chit-chat, knowledge base, and multilingual models
Performs sentiment analysis on user messages
Selects the best language model automatically

What should you integrate into the chatbot?

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

Answer: C

Explanation:

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

QnA Maker is a cloud-based Natural Language Processing (NLP) service that allows you to create a natural conversational layer over your dat
a. 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: 21

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

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

Which Azure Cognitive Services service should you use?

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

Answer: C

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: 22

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.

Answer Area

```
V https://management.azure.com/subscriptions/xxxxxxxx-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": "SO"

  "properties": {},
  "identity": {
    "type": "SystemAssigned"
  }
```

Answer:

Explanation:

Answer Area

```
          ▼ https://management.azure.com/subscriptions/xxxxxxxx-xxxx-
    PATCH
    POST
    PUT
  xxxx-xxxx-
  xxxxxxxxxxxx/resourceGroups/RGl/providers/Microsoft.CognitiveServices/
  accounts/CSl?api-version=2017-04-18

    "location": "West US",
    "kind": "                    ▼
              CognitiveServices
              Computervision
              TextAnalytics
    "sku": { "name": "SO"

    "properties": { },
    "identity": {
      "type": "SystemAssigned"
    }
```

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: 23

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: BC

Explanation:

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

Question: 24

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.

Actions	Answer Area
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:

Explanation:

Step 1: Pull the Anomaly Detector container image.
Step 2: Create a custom Dockerfile
Step 3: Build the image
Step 4: Push the image to an Azure container registry.

https://docs.microsoft.com/en-us/azure/cognitive-services/containers/container-reuse-recipe

Question: 25

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.

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

[▼]

- http://contoso.blob.core.windows.net
- https://contoso.cognitiveservices.azure.cbm
- 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
- mor.Microsoft.com/azure-cognitive-services/textanalytics/keyphrase
- mcr.Microsoft.com/azure-cognitive-services/texfanalytics/sentiment

```
Ap i Ke y=xxx.xxxx.x.x.xxxxxxxxxx
```

Answer:

Explanation:

```
docker run —rm -it -p 5000:5000 —memory 8g. —cpus 1 \
I                                                    ▼ \
 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-acce.pt \
Billing—                                                    ▼ \
          http://contoso.blob.core.windows.net https://contoso.cognitiveservices.azure.com
          mcr.Microsoft.com/azure-cognitive-servrces/textanalytics/keyphrase
          mcr.Microsoft.com/azure-cognitive-services/textanalytics/sentiment

ApiKey=xxxxxxxxxxxxxxxxxx
```

Box 1: mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment

To run the Sentiment Analysis v3 container, execute the following docker run command. docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \ mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment \

Eula=accept \

Billing={ENDPOINT_URI} \

ApiKey={API_KEY} is the endpoint for accessing the Text Analytics API. https://<your-custom-subdomain>.cognitiveservices.azure.com

Box 2: https://contoso.cognitiveservices.azure.com

{ENDPOINT_URI} is the endpoint for accessing the Text Analytics API: https://<your-custom-

subdomain>.cognitiveservices.a The endpoint for accessing the Text Analytics API. zure.com

Reference:

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

Question: 26

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

```
static void createresource(CognitiveServicesManagementClient client, string
resourcename, string kind, string account_tier, string location)

    CognitiveServicesAccount parameters =
            new CognitiveServicesAccount(null, null, kind, location, resource_name,
new CognitiveServicesAccountProperties(), new Sku(account_tier));
    var result = client.Accounts.Create(resource_group_name, account_tier, parameters);
```

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

Which code should you use?

A. create_resource(client, "res1", "ComputerVision", "F0", "westus")
B. create_resource(client, "res1", "CustomVision.Prediction", "F0", "westus")
C. create_resource(client, "res1", "ComputerVision", "S0", "westus")
D. create_resource(client, "res1", "CustomVision.Prediction", "S0", "westus")

Answer: B

Explanation:

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

Question: 27

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.

Answer: B

Explanation:

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

Question: 28

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
Filel	PDF	20 MB
File2	MP4	100 MB
File3	JPG	20 MB
File4	PDF	100 MB
File5	GIF	1 MB
File6	JPG	40 MB

Which three files can you use to train the model? Each correct answer presents a complete solution. (Choose three.)

NOTE: Each correct selection is worth one point.

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

Answer: A, B, F

Explanation:

Input requirements
Form Recognizer works on input documents that meet these requirements:

Format must be JPG, PNG, PDF (text or scanned), or TIFF. Text-embedded PDFs are best because there's no possibility of error in character extraction and location.
File size must be less than 50 MB.

Reference:

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

Question: 29

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

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

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

Answer: A

Explanation:

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

Note: Video Indexer supports face detection and celebrity recognition for video content. The celebrity recognition feature covers about one million faces based on commonly requested data source such as IMDB, Wikipedia, and top LinkedIn influencers. Faces that aren't recognized by the celebrity recognition feature are

detected but left unnamed. Once you label a face with a name, the face and name get added to your account's Person model. Video Indexer will then recognize this face in your future videos and past videos.

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

Question: 30

You use the Custom Vision service to build a classifier.

After training is complete, you need to evaluate the classifier.

Which two metrics are available for review? Each correct answer presents a complete solution. (Choose two.)

NOTE: Each correct selection is worth one point.

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

Answer: AD

Explanation:

Custom Vision provides three metrics regarding the performance of your model: precision, recall, and AP.

Reference:
https://www.tallan.com/blog/2020/05/19/azure-custom-vision/

Question: 31

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.

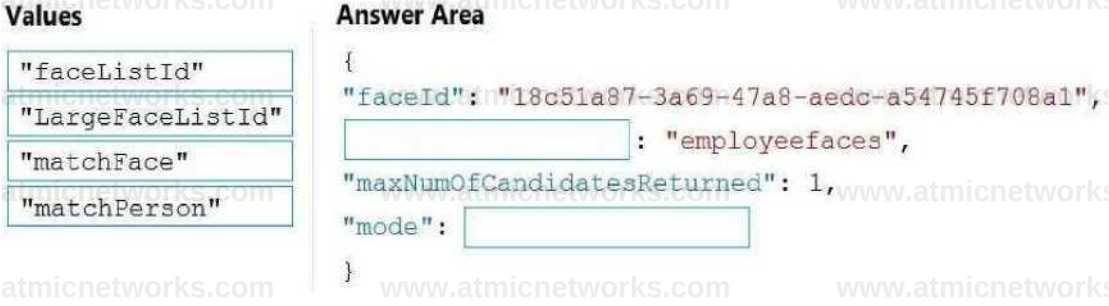

Answer:

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: 32

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.

findsintilars group identify matchFace matchPerson verify

Request Body

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

Answer:

Explanation:

Box 1: findsimilars

https://docs.microsoft.com/en-us/rest/api/faceapi/face/find-similar

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: 33

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

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
The code will return the name of each detected brand with a confidence equal to or higher than 75 percent.	O	O

The code will return coordinates for the bottom-left corner of the O O rectangle that contains the brand logo of the displayed brands.

The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.

0 0

Answer:

Explanation:

Box 1: Yes

Box 2: No

Box 3: No

Reference:

https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-detecting- faces

Question: 34

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.

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))
                File
                Stream
                Uri
                Uri

        await faceclient.PersonGroupPerson.                    Y
                                        AddFaceFromStreamAsync
                                        Add Fa cd Fr dmUr 1A s y n c CreateAsync
                                        GetAsync
(personGroupId, personId, t) ;
        }
    }
```

Answer:

Explanation:

Box 1: Stream

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

Example: Open the stream and read it back.

```
using (FileStream fs = File.OpenRead(path))
```

Box 2: AddFaceFromStreamAsync

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

Question: 35

HOTSPOT

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

```
public async TaskAnalyzelmage(ComputerVisionClient client, string localImage)

        List<VisualFeatureTypes> features = new Li3t<VisualFeatureTypes>()
                        VisualFeatureTypes.Description,
                        VisualFeatureTyces.Tags, 1;
        using (Stream imagestream = File.OpenRead(localImage)) { try

                ImageAnalysis results — await client.AnalyzeImageInStreamAsync(imagestream, features);

                foreach (var caption in results.Description.Captions) {
                      Console.Writeline($"{caption.Text) with confidence {caption.Confidence}"); )
                foreach (var tag in results.Tags) (
                      Console.Wxiteline($"(tag.Name} {tag.Confidence}");

            catch (Exception ex)

                  Console.Wxiteline(ex.Message);
        1 }
```

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The code will perform face recognition.		
The code will list tags and their associated confidence.		
The code will read a file from the local file system. Q o		

Answer:

Explanation:

Box 1: No

Box 2: Yes

Box 3: No

Question: 36

You are developing a method that uses the Computer Vision client library. The method will perform optical character recognition (OCR) in images. The method has the following code.

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

        var txtHeaders = await client.ReadAsync(urlFile, language: "en");
        string opLocation = textHeaders.OperationLocation;
        string operationId = opLocation.Substring(opLocation.Length -
numberOfCharsInOperationId);
        ReadOperationResult results;

        results = await client.GetReadResultAsync(Guid.Parse(operationId));
        var textUrlFileResults = results.AnalyzeResult.ReadResults;
        foreach (ReadResult page in textUrlFileResults)
                foreach (Line line in page.Lines)
                (
                        Console.WriteLine(line.Text);
                )}
```

During testing, you discover that the call to the GetReadResultAsync method occurs before the read operation is complete.

You need to prevent the GetReadResultAsync method from proceeding until the read operation is complete.

Which two actions should you perform? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

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

Answer: BD

Explanation:

Example code : do

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

```
}
while ((results.Status == OperationStatusCodes.Running || results.Status ==
    OperationStatusCodes.NotStarted));
```

Reference:
https://github.com/Azure-Samples/cognitive-services-quickstart-code/blob/master/dotnet/ComputerVision/ComputerVisionQuickstart.cs

Question: 37

You are building a Language Understanding model for an e-commerce platform. You need to construct an entity to capture billing addresses.
Which entity type should you use for the billing address?

A. machine learned
B. Regex
C. geographyV2
D. Pattern.any
E. list

Answer: A

Explanation:

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

Address: 4567 Main Street, NY, 98052, USA
Building Number: 4567
Street Name: Main Street
State: NY
Zip Code: 98052
Country: USA
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-entity-types

Question: 38

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

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

Answer: B

Explanation:

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

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

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

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

Question: 39

You are developing a method for an application that uses the Translator API.
The method will receive the content of a webpage, and then translate the content into Greek (el).
The result will also contain a transliteration that uses the Roman alphabet.

You need to create the URI for the call to the Translator API. You have the following URI.
https://api.cognitive.microsofttranslator.com/translate?api-version=3.0

Which three additional query parameters should you include in the URI? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

A. toScript=Cyrl
B. from=el
C. textType=html
D. to=el
E. textType=plain
F. toScript=Latn

Answer: A, D, F

Explanation:

C: textType is an optional parameter. It defines whether the text being translated is plain text or HTML text (used for web pages).

D: to is a required parameter. It specifies the language of the output text. The target language must be one of the supported languages included in the translation scope.

F: toScript is an optional parameter. It specifies the script of the translated text. We use Latin (Roman alphabet) script.

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/translator/reference/v3-0-translate

Question: 40

You have a chatbot that was built by using the Microsoft Bot Framework. You need to debug the chatbot endpoint remotely.

Which two tools should you install on a local computer? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

A. Fiddler
B. Bot Framework Composer
C. Bot Framework Emulator
D. Bot Framework CLI
E. ngrok
F. nginx

Answer: CE

Explanation:

Bot Framework Emulator is a desktop application that allows bot developers to test and debug bots, either locally or remotely.

ngrok is a cross-platform application that "allows you to expose a web server running on your local machine to the internet." Essentially, what we'll be doing is using ngrok to forward messages from external channels on the web directly to our local machine to allow debugging, as opposed to the

standard messaging endpoint configured in the Azure portal.

Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-service-debug-emulator

Question: 41

DRAG DROP

You are building a retail chatbot that will use a QnA Maker service.

You upload an internal support document to train the model. The document contains the following question "What is your warranty period?"

Users report that the chatbot returns the default QnA Maker answer when they ask the following question "How long is the warranty coverage?"

The chatbot returns the correct answer when the users ask the following question 'What is your warranty period?"

Both questions should return the same answer.

You need to increase the accuracy of the chatbot responses.

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

Actions

Add a new question and answer (QnA) pair.

Retrain the model.

Add additional questions to the document.

Republish the model.

Add alternative phrasing to the question and answer (QnA) pair.

Answer Area

Answer:

Explanation:

Add alternative phrasing to the question and answer (QnA) pair.

Retrain the model

Republish the model

Step 1: Add alternative phrasing to the question and answer (QnA) pair.

Add alternate questions to an existing QnA pair to improve the likelihood of a match to a user query.

Step 2: Retrain the model.

Periodically select Save and train after making edits to avoid losing changes.

Step 3: Republish the model

Note: A knowledge base consists of question and answer (QnA) pairs. Each pair has one answer and a pair contains all the information associated with that answer.

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/edit-knowledge-base

Question: 42

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

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

Answer: A

Explanation:

Text Analytics Cognitive Service could be used to quickly determine the public perception for a specific topic, event or brand.

Example: A NodeJS app which pulls Tweets from Twitter using the Twitter API based on a specified search term. Then pass these onto Text Analytics for sentiment scoring before storing the data and building a visualisation in PowerBI. The Architecture looked something like this:

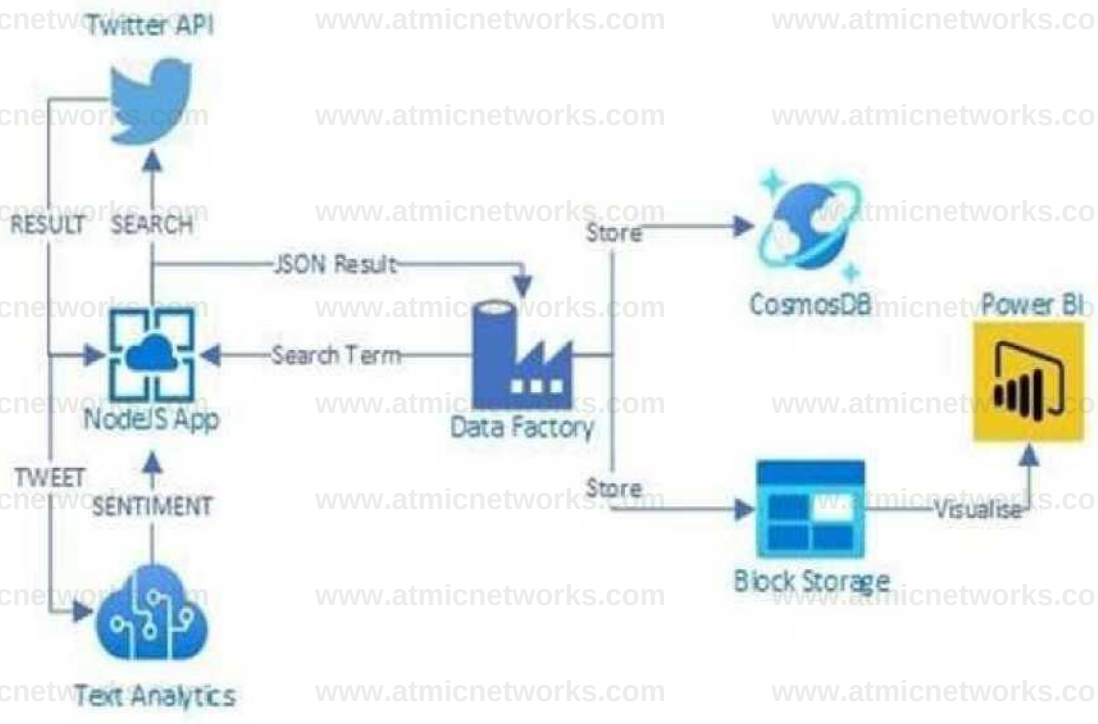

Reference:
https://www.linkedin.com/pulse/measuring-public-perception-azure-cognitive-services-steve-dalai

Question: 43

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

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

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.
A conversational expert provides you with the following list of phrases to use for training.

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

You need to implement the phrase list in Language Understanding.
Solution: You create a new intent for location.
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

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

Question: 44

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

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

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.
A conversational expert provides you with the following list of phrases to use for training.

Find contacts in London.

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

You need to implement the phrase list in Language Understanding.
Solution: You create a new entity for the domain.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance.
Define a set of intents that corresponds to actions users want to take in your application. Reference: https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-intent

Question: 45

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

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

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.
A conversational expert provides you with the following list of phrases to use for training.

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

You need to implement the phrase list in Language Understanding.
Solution: You create a new pattern in the FindContact intent.
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance.
Define a set of intents that corresponds to actions users want to take in your application. Reference: https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-intent

Question: 46

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

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

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

images of new flower species.
You need to add the new images to the classifier.

Solution: You add the new images, and then use the Smart Labeler tool.

Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

The model need to be extended and retrained.

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

Question: 47

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

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

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

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

Does this meet the goal?

A. Yes

B. No

Answer: A

Explanation:

The model needs to be extended and retrained.

Question: 48

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

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

will not appear in the review screen.

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

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

A. Yes
B. No

Answer: B

Explanation:

The model needs to be extended and retrained.

Question: 49

HOTSPOT

You are developing a service that records lectures given in English (United Kingdom).

You have a method named AppendToTranscriptFile that takes translated text and a language identifier.

You need to develop code that will provide transcripts of the lectures to attendees in their respective language. The supported languages are English, French, Spanish, and German.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
static async Task n.u   i _< u--Ji.. in ()

      var config =SpeechTran3lationConfig,FiomSubscription
("69cad5cc-Qab3-4704-bdff-afbf4aa07d85¹', "uksouth") ;

var lang = new List<string>      ▼

{"en-GB"}
{JM:r% "de ", "ea ″ ]    {"French", "Spanish", "German"}
{languages}

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

      using var audioConfig =
AudioConfig.EromDefaultMicrophonelnput();
      using var recognizer = new  ▼ (config, audioConfig);

IntentRecognizer
SpeakerRecognizer SpeechSynthesizer Translat ibnRe coghi z
er
      var result. = awit recognizer-RecognizeOnceAsync();
```

```
if (result.Reason = ResultReason.TranslatedSpeech)
```

Answer:

Explanation:

Box 1: {"fr", "de", "es"}

A common task of speech translation is to specify target translation languages, at least one is required but multiples are supported. The following code snippet sets both French and German as translation language targets.

```
static async Task TranslateSpeechAsync()
{
   var translationConfig =
      SpeechTranslationConfig.FromSubscription(SPEECH SUBSCRIPTION KEY,
SPEECH SERVICE REGION);

   translationConfig.SpeechRecognitionLanguage = "it-IT";

   // Translate to languages. See, https://aka.ms/speech/sttt-languages
   translationConfig.AddTargetLanguage("fr");
   translationConfig.AddTargetLanguage("de");
}
```

Box 2: TranslationRecognizer

After you've created a SpeechTranslationConfig, the next step is to initialize a TranslationRecognizer.

Example code:

```
static async Task TranslateSpeechAsync()
{
   var translationConfig =
      SpeechTranslationConfig.FromSubscription(SPEECH SUBSCRIPTION KEY,
SPEECH SERVICE REGION);

   var fromLanguage = "en-US";
   var toLanguages = new List<string> { "it", "fr", "de" };
   translationConfig.SpeechRecognitionLanguage = fromLanguage;
   toLanguages.ForEach(translationConfig.AddTargetLanguage);
```

```
    using var recognizer = new TranslationRecognizer(translationConfig);
}
```

Question: 50

HOTSPOT

You are developing a text processing solution.

You develop the following method.

```
static void GetKeyPhrases(TextAnalyticsClient textAnalyticsClient, string text) <
    var response = textAnalyticsClient.ExtractKeyPhrases(text);
    Console.WriteLine("Key phrases:");

    foreach (string keyphrase in response.Value) (
        Console.WriteLine($"\t(keyphrase)");
```

You call the method by using the following code.

```
GetKeyPhrases(textAnalyticsClient, "the cat sat on the mat");
```

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The call will output key phrases from the input string to the console.		
The output will contain the following words: the, cat, sat, on, and mat.		
The output will contain the confidence level for key phrases.		

Answer:

Explanation:

Answer Area

Statements	Yes	No
The call will output key phrases from the input string to the console.		
The output will contain the following words: the, cat, sat, on, and mat.		
The output will contain the confidence level for key phrases.		

Box 1: Yes

The Key Phrase Extraction API evaluates unstructured text, and for each JSON document, returns a list of key phrases.

Box 2: No
'the' is not a key phrase.

This capability is useful if you need to quickly identify the main points in a collection of documents. For example, given input text "The food was delicious and there were wonderful staff", the service returns the main talking points: "food" and "wonderful staff".

Box 3: No
Key phrase extraction does not have confidence levels.

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics- how-to-keyword-extraction

Question: 51

You deploy a web app that is used as a management portal for indexing in Azure Cognitive Search.
The app is configured to use the primary admin key.

During a security review, you discover unauthorized changes to the search index. You suspect that the primary access key is compromised.

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

What should you do next?

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

Answer: C

Explanation:

Question: 52

You have an existing Azure Cognitive Search service.

You have an Azure Blob storage account that contains millions of scanned documents stored as images and PDFs.

You need to make the scanned documents available to search as quickly as possible. What should you do?
A. Split the data into multiple blob containers. Create a Cognitive Search service for each container.

Within each indexer definition, schedule the same runtime execution pattern.

B. Split the data into multiple blob containers. Create an indexer for each container. Increase the search units. Within each indexer definition, schedule a sequential execution pattern.

C. Create a Cognitive Search service for each type of document.

D. Split the data into multiple virtual folders. Create an indexer for each folder. Increase the search units. Within each indexer definition, schedule the same runtime execution pattern.

Answer: D

Explanation:

Incorrect Answers:

A: Need more search units to process the data in parallel. B: Run them in parallel, not sequentially.

C: Need a blob indexer.

Note: A blob indexer is used for ingesting content from Azure Blob storage into a Cognitive Search index.

Index large datasets

Indexing blobs can be a time-consuming process. In cases where you have millions of blobs to index, you can speed up indexing by partitioning your data and using multiple indexers to process the data in parallel.

Here's how you can set this up:

Partition your data into multiple blob containers or virtual folders Set up several data sources, one per container or folder.

Create a corresponding indexer for each data source. All of the indexers should point to the same target search index.

One search unit in your service can run one indexer at any given time. Creating multiple indexers as described above is only useful if they actually run in parallel.

Reference:

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

Question: 53

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

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

NOTE: Each correct selection is worth one point.

A. tableName

B. generatedKeyName

C. dataSource

D. dataSourceConnection

E. source

Answer: ABE

Explanation:

Defining a table projection.

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

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

Question: 54

HOTSPOT

You are creating an enrichment pipeline that will use Azure Cognitive Search. The knowledge store contains unstructured JSON data and scanned PDF documents that contain text.

Which projection type should you use for each data type? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

JSON data: ▼

File projection
Object projection
Table projection

Scanned data: ▼

File projection
Object projection
Table projection

Answer:

Explanation:

Answer Area

JSON data: ▼

File projection

Object projection

Table projection

Scanned data: v

File projection Object projection

Table projection

Box 1: Object projection

Object projections are JSON representations of the enrichment tree that can be sourced from any node.

Box 2: File projection

File projections are similar to object projections and only act on the normalized_images collection.

Reference:

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

Question: 55

HOTSPOT

You are building an Azure Cognitive Search custom skill.

You have the following custom skill schema definition.

```
"@odata.type": "#Microsoft.Skills.Custom.WebApiSkill",
"description": "My custom skill description",
"uri": "https://contoso-webskill.azurewebsites.net/api/process"
"context": "/document/organizations/*",
"inputs": [
  {
    "name": "companyName",
    "source": "/document/organizations/*"

L
"outputs": [
  {
    "name": "companyDescription",
  }
```

i

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

NOTE: Each correct selection is worth one point.

Answer Area Statements Yes No

Company-Description is available for indexing.

The definition calls a web API as part of the enrichment process. |o| 1$

The enrichment step is called only for the first organization under "/document/organizations".

Answer:

Explanation:

Answer Area

Statements	Yes	No
Company-Description is available for indexing.	0	O
The definition calls a web API as part of the enrichment process.		
The enrichment step is called only for the first organization under "/document/organizations".	. .0	0

Box 1: Yes

Once you have defined a skillset, you must map the output fields of any skill that directly contributes values to a given field in your search index.

Box 2: Yes

The definition is a custom skill that calls a web API as part of the enrichment process.

Box 3: No

For each organization identified by entity recognition, this skill calls a web API to find the description of that organization.

Reference:
https://docs.microsoft.com/en-us/azure/search/cognitive-search-output-field-mapping

Question: 56

You have the following data sources:

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

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

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

Answer: B

Explanation:

On-premises Microsoft SQL Server database cannot be used as an index data source.

Note: Indexer in Azure Cognitive Search: : Automate aspects of an indexing operation by configuring a data source and an indexer that you can schedule or run on demand. This feature is supported for a limited number of data source types on Azure.

Indexers crawl data stores on Azure.

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

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

Question: 57

You are building a multilingual chatbot.

You need to send a different answer for positive and negative messages.

Which two Text Analytics APIs should you use? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

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

Answer: BD

Explanation:

8: The Text Analytics API's Sentiment Analysis feature provides two ways for detecting positive and negative sentiment. If you send a Sentiment Analysis request, the API will return sentiment labels (such as "negative", "neutral" and "positive") and confidence scores at the sentence and documentlevel.

D: The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis. This capability is useful for content stores that collect arbitrary text, where language is unknown.

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics- how-to-sentiment-analysis?tabs=version-3-1

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

Question: 58

DRAG DROP

You plan to build a chatbot to support task tracking.

You create a Language Understanding service named lu1.

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

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

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

Answer:

Explanation:

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

Question: 59

You are building a bot on a local computer by using the Microsoft Bot Framework. The bot will use an existing Language Understanding model.

You need to translate the Language Understanding model locally by using the Bot Framework CLI. What should you do first?

A. From the Language Understanding portal, clone the model.
B. Export the model as an .lu file.
C. Create a new Speech service.
D. Create a new Language Understanding service.

Answer: B

Explanation:

You might want to manage the translation and localization for the language understanding content for your bot independently.

Translate command in the @microsoft/bf-lu library takes advantage of the Microsoft text translation API to automatically machine translate .lu files to one or more than 60+ languages supported by the Microsoft text translation cognitive service.

What is translated?
An .lu file and optionally translate Comments in the lu file
LU reference link texts
List of .lu files under a specific path.

Reference:

https://github.com/microsoft/botframework-cli/blob/main/packages/luis/docs/translate- command.md

Question: 60

DRAG DROP

You are using a Language Understanding service to handle natural language input from the users of a web-based customer agent.

The users report that the agent frequently responds with the following generic response: "Sorry, I don't understand that."

You need to improve the ability of the agent to respond to requests.

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.)

Actions

Add prebuilt domain models as required.

Validate the utterances logged for review and modify the model.

Migrate authoring to an Azure resource authoring key.

Enable active learning.

Enable log collection by using Log Analytics.

Train and republish the Language Understanding model.

Answer Area

Answer:

Explanation:

- enable active learning
- validate the utterances
- train and republish

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

https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-prebuilt-model

Question: 61

You build a conversational bot named bot1.

You need to configure the bot to use a QnA Maker application.

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

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

Answer: C

Explanation:

Obtain values to connect your bot to the knowledge base

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

Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-qna

Question: 62

HOTSPOT

You are building a chatbot by using the Microsoft Bot Framework Composer.

You have the dialog design shown in the following exhibit.

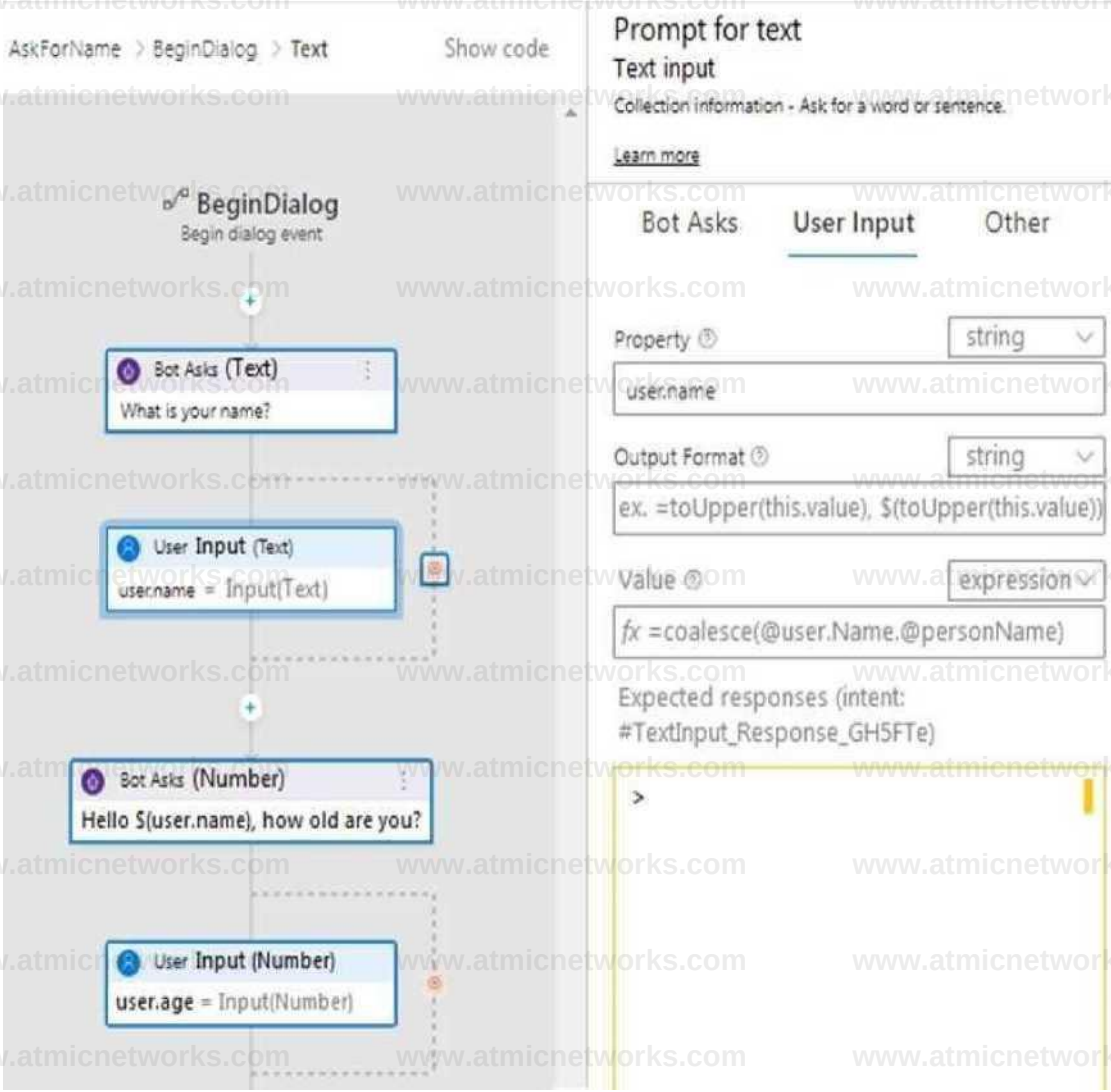

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area Statements Yes No

user.name is an entity.

The dialog asks for a user name and a user age and assigns appropriate values to the user.name and user.age properties.

The chatbot attempts to take the first non-null entity value for userName or personName and assigns the value to user .name.

Answer:

Explanation:

Answer Area Statements Yes No

`user.name` **is an entity.**

The dialog asks for a user name and a user age and assigns appropriate values to the user.name and user. age properties.

The chatbot attempts to take the first non-null entity value for `userName` or `personName` and assigns the value to `user . name`.

Box 1: No User.name is a property.

Box 2: Yes

Box 3: Yes

The coalesce() function evaluates a list of expressions and returns the first non-null (or non-empty for string) expression.

Reference:
https://docs.microsoft.com/en-us/composer/concept-language-generation

https://docs.microsoft.com/en-us/azure/data-explorer/kusto/query/coalescefunction

Question: 63

HOTSPOT

You are building a chatbot for a Microsoft Teams channel by using the Microsoft Bot Framework SDK. The chatbot will use the following code.

```
protected override async Task OnMembersAddedAsync(IList<ChannelAccount> membersAdded,
ITumContext<IConversationUpdateActivity> turnContext, CancellationToken cancellationToken) {
    foreach (var member in membersAdded)
        if (member. Id != turnContext. Activity. Recipient. Id)
            await turnContext.SendActivityAsync($"Hi there - (member.Name). {WelcomeMessage)",
cancellationToken: cancellationToken);
```

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

NOTE: Each correct selection is worth one point.

Answer Area Statements	Yes	No
`OnMembersAddedAsync` **will be triggered when a user joins the conversation, o**		
When a new user joins the conversation, the existing users in the conversation will see the chatbot greeting.		
`OnMembersAddedAsync` **will be initialized when a user sends a message.**		

Answer:

Explanation:

Answer Area Statements	Yes	No
`OnMembersAddedAsync` **will be triggered when a user joins the conversation, o**		
When a new user joins the conversation, the existing users in the conversation will see the chatbot greeting.		
`OnMembersAddedAsync` **will be initialized when a user sends a message.**		

Box 1: Yes
The ActivityHandler.OnMembersAddedAsync method overrides this in a derived class to provide logic for when members other than the bot join the conversation, such as your bot's welcome logic.

Box 2: Yes membersAdded is a list of all the members added to the conversation, as described by the conversation update activity.

Box 3: No

Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.bot.builder.activityhandler.onmembersaddedasync?view=botbuilder-dotnet-stable

Question: 64
HOTSPOT

You are building a chatbot by using the Microsoft Bot Framework SDK.

You use an object named UserProfile to store user profile information and an object named ConversationData to store information related to a conversation.

You create the following state accessors to store both objects in state.

```
var userStateAccessors = _userState.CreateProperty<UserProfile>(nameof(UserProfile));
```

```
var conversationStateAccessors =
_conversationState.CreateProperty<ConversationData>(nameof(ConversationData));
```

The state storage mechanism is set to Memory Storage.

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
The code will create and maintain the UserProf ile object in the underlying storage layer.	0	
The code will create and maintain the ConversationData object in the underlying storage layer.		
The UserProf ile and ConversationData objects will persist when the Bot Framework runtime terminates.	o	

Answer:

Explanation:

Statements	Yes	No
The code will create and maintain the UserProf ile object in the underlying storage layer.	O	
The code will create and maintain the ConversationData object in the underlying storage layer.	o	
The UserProfile and ConversationData objects will persist when the Bot Framework runtime terminates.		

Box 1: Yes

You create property accessors using the CreateProperty method that provides a handle to the BotState object. Each state property accessor allows you to get or set the value of the associated state property.

Box 2: Yes

Box 3: No

Before you exit the turn handler, you use the state management objects' SaveChangesAsync() method to write all state changes back to storage.

Reference:

https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-v4-state

Question: 65

HOTSPOT

You are building a chatbot that will provide information to users as shown in the following exhibit.

Passengers

Sarah Hum

Jeremy Goldberg

Evan Litvak

2 Stops

Tue, May 30, 2017 10:25 PM

San Francisco	San Francisco
Amsterdam	Amsterdam
SFO	SFO
AMS	AMS

Non-Stop

Fri, Jun 2, 2017 11:55 PM

San Francisco	San Francisco
Amsterdam	Amsterdam
SFO	SFO
AMS	AMS

Total	**$4,032.54**

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer Area

The chatbot is showing **[answer choice].**

an Adaptive Card a Hero Card
a Thumbnail Card

The card includes **[answer choice].**

an action set
an image
an image group media

Answer:

Explanation:

Box 1: Adaptive card

Box 2: an image

Reference:
https://docs.microsoft.com/en-us/microsoftteams/platform/task-modules-and-cards/cards/cards- reference

https://docs.microsoft.com/en-us/composer/how-to-send-cards?tabs=v1x

Question: 66

HOTSPOT

You are reviewing the design of a chatbot. The chatbot includes a language generation file that contains the following fragment.

- Greet(user)
- ${Greeting()}, ${user.name}

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
$ {user. name} retrieves the user name by using a prompt.		
Greet () is the name of the language generation template.		
$ {Greeting () } is a reference to a template in the language generation file.		

Answer:

Explanation: **StatementsYes No**

${user.name} retrieves the user name by using a prompt.

Greet () is the name of the language generation template.

$ {Greeting () } is a reference to a template in the language generation file.

Box 1: No
Example: Greet a user whose name is stored in `user.name`
- ${ welcomeUser(user.name) }

Example: Greet a user whose name you don't know:
- ${ welcomeUser() }

Box 2: No
Greet(User) is a Send a response action.

Box 3: Yes

Reference:
https://docs.microsoft.com/en-us/composer/how-to-ask-for-user-input

Question: 67

A customer uses Azure Cognitive Search.

The customer plans to enable a server-side encryption and use customer-managed keys (CMK) stored in Azure.

What are three implications of the planned change? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

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

Answer: ABF

Explanation:

"Customer-managed keys require an additional billable service, Azure Key Vault, which can be in a different region, but

under the same subscription, as Azure Cognitive Search. Enabling CMK encryption will increase index size and degrade query performance."
same document also lists Azure Key Vault as a requirement:

https://docs.microsoft.com/en-us/azure/search/search-security-overview#data-protection

Question: 68

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

You plan to notify users that their data has been processed by the sales system.

Which responsible AI principle does this help meet?

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

Answer: D

Explanation:

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

Question: 69

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

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

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

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

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

Solution: You deploy service1 and a public endpoint to a new virtual network, and you configure Azure Private Link.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

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

Question: 70

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

Answer: B

Explanation:

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

Question: 71

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

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

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

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

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

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

Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

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

Question: 72

You plan to perform predictive maintenance.

You collect IoT sensor data from 100 industrial machines for a year. Each machine has 50 different sensors that generate data at one-minute intervals. In total, you have 5,000 time series datasets.

You need to identify unusual values in each time series to help predict machinery failures.

Which Azure Cognitive Services service should you use?

A. Anomaly Detector
B. Cognitive Search
C. Form Recognizer
D. Custom Vision

Answer: A

Explanation:

Question: 73

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.

```
var audioFormat = I (AudioStreamContainerFormat.MP3);
                  AudioConfig SetProperty
                  AudioStreamFormat GetCompressedForm at
                  AudioStreamFormat GetWaveFormatPCM
                  PullAudioInputStream
var speechConfig • SpeechConfig.FromSubscription("18c51a87-3a69-47a8-aedc-a54745f708al", "westus");
var audioConfig = AudioConfig.FromStreamInput(pushStream, audioFormat);
using (var recogniter - new I (speechConfig, audioConfig)) KeywordRecognizer SpeakerRecognizer SpeechRecognizer SpeechSynthesizer
  {
   var result • await recognizer.RecognizeOnceAsync();
   var text * result.Text;
```

Answer:

Explanation:

```
var audioFormat * ^^                                    AudioStreamContainerFormat.MP3);
                  AudioConfig SetProperty
                  AudioStreamFormat GetCompressedFormat AudioStreamFormat GetWaveFormatPCM PullAudioInputStream
var speechconfig ■ Speechconfig.Froi»Subscription("18c51a87-3a69-47a8-aedc-a54745f708al", "westus")
var audioConfig * AudioConfig.FromStreaminput(pushstream, audioFormat);
using (var recognizer ■ new I(speechConfig, audioConfig))
                                  Keyword Recognizer SpeakerRecognizer SpeechRecognizer SpeechSynthesizer (
   var result ■ await recognizer.RecognizeOnceAsync();
   var text » result.Text;
```

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: 74

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.

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:

Explanation:

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

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

Question: 75

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

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

NOTE: Each correct selection is worth one point.

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

Answer: DE

Explanation:

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

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

Question: 76

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

You create a new Language Understanding resource.

You need to add more contributors.

What should you use?

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

Answer: B

Explanation:

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

Question: 77

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.

Answer:

Explanation:

westus.api.cognitive.microsoft.com

generateThumbnail

https://docs.microsoft.com/en-us/rest/api/computervision/3.1/generate-thumbnail/generate- thumbnail#examples

POST

https://westus.api.cognitive.microsoft.com/vision/v3.1/generateThumbnail?width=500&height=500

&smartCropping=True

Ocp-Apim-Subscription-Key: {API key}

Question: 78

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.

Answer:

Explanation:

Cognitive Insights Widget

```
https    ://www.videoindexer.    ai/embed/insights/caccount    IdVcvideoId^fidgets-    controls= search
peoplejkeywords
```

Player Widget

```
https://www.videoindexer.al/embed/player/<BCcountId>/<videoId>/? showcaptions- true    captions- en-U5
```

Question: 79

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.

Actions	Answer Area
Change the model domain	
Retrain the model.	
Test the model.	
Export the model.	

Answer:

Explanation:

In user want to change to deploy ofline model

1. Change model domain to compact model
2. Retrain compact model
3. Export model

Reference:

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

Question: 80

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("Oep-Apira-Subscription-Key", subscription_key);
    var endpointURI = J"https://westus.api.cognitive.Microsoft,coin/fa ce/vl. 6/per songroup s/{ per sonGroupId}/per sons/{personId}/persistedFaces";
    HttpResponseMessage response;
    var body = "{ \"url\": V* + imageURI + '\"}";
    var content ■ new StringContent(body, Encoding.UTFB, " app Heat ion/json");
    var response = await client.PutAsync(endpointURI, content);
}
```

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

NOTE: Each correct selection is worth one point.

Statements	Yes	No
The code will add a face image to a person object in a person group.	0	O
The code will work for a group of 10,000 people	0	o
AddFace can be called multiple times to add multiple face images to a person object	0	o

Answer:

Explanation:

A. True

B. True

C. True

B: see this example code from documentation that uses PersonGroup of size 10,000 :

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

the question wants to trick you into thinking you need to use a LargePersonGroup for a size of 10,000 - but the documentation for it doesn't include this limitation or criteria:

https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-use- large-scale

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

Question: 81

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.

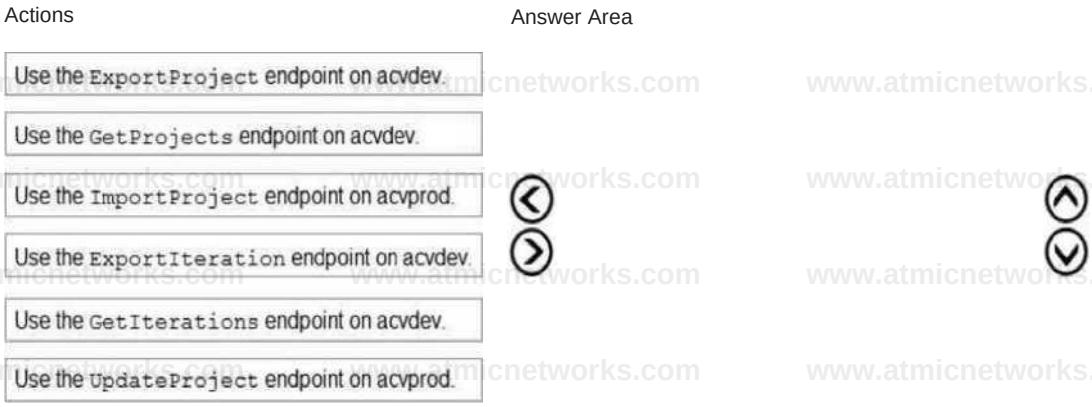

Answer:

Explanation:

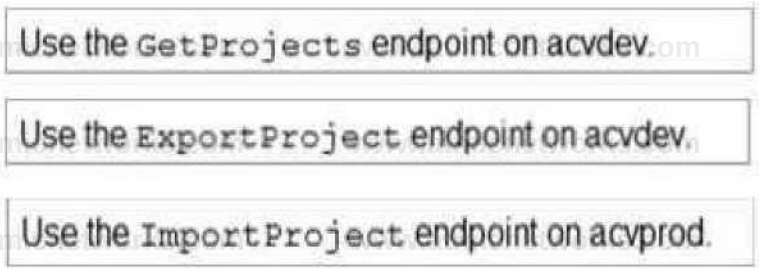

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-tutorial-pattern#what-did-this-tutorial-accomplish

Question: 82

DRAG DROP

You train a Custom Vision model used in a mobile app.

You receive 1,000 new images that do not have any associated data.

You need to use the images to retrain the model. The solution must minimize how long it takes to retrain the model.

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

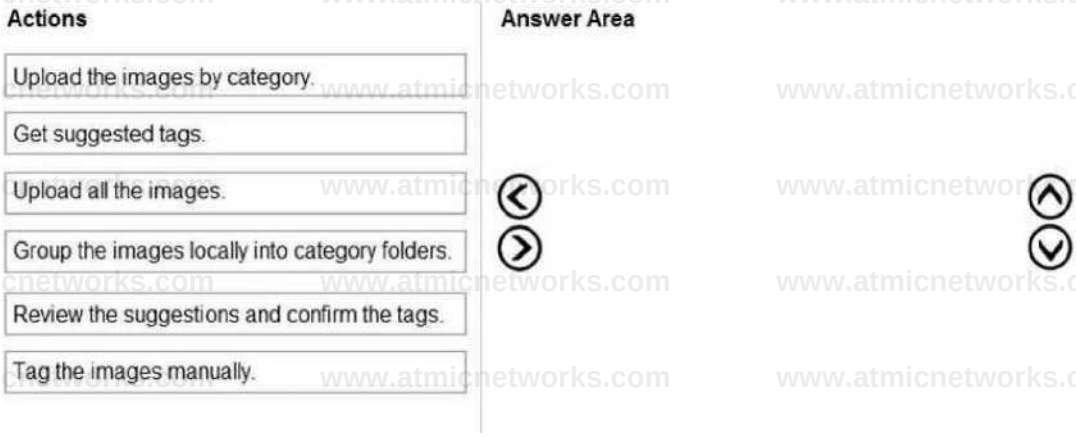

Answer:

Explanation:

1 .) upload all the images
2 .) Get suggested tags
3 .) Review the suggestions and confirm the tags

reference:

https://docs.microsoft.com/en-us/azure/cognitive-services/custom-vision-service/suggested-tags

Question: 83

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

You need to construct an entity to capture billing addresses.

Which entity type should you use?

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

Answer: B

Explanation:

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

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

Address: 4567 Main Street, NY, 98052, USA

Building Number: 4567 Street Name: Main Street State: NY
Zip Code: 98052
Country: USA

Question: 84

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

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

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

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

NOTE: Each correct selection is worth one point.

A. Create a list of FileDTO objects that represents data from the WebJob.
B. Call the client. Knowledgebase. CreateAsync method.
C. Create a list of QnADTO objects that represents data from the WebJob.
D. Create a CreaceKbDTO object.

Answer: AC

Explanation:

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

Question: 85

HOTSPOT

You are developing an application that includes language translation.

The application will translate text retrieved by using a function named getTextToBeTranslated. The text can be in one of many languages. The content of the text must remain within the Americas Azure geography.

You need to develop code to translate the text to a single language.

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

NOTE: Each correct selection is worth one point.

```
ver endpoint • _____
```

&quot https //api cognitive microsofttranslator com/translate&quot: &quot https://api cognitive microsofttranslator com/transliterate&quot &quot.https://api-apc cognitive microsofttranslator com/detect&quot,, &quot https api-nam cognitive

microsofttranslator com'detect&quot &quot, https //api-nam cognitive microsofttranslator com/translate&quot.

```
var apiKey - "FF956C68B83B21B38691A8D200A4C606";
var text = getTextToBeTranslated();
var body • '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add(“Ocp-Apim-Subscription-Key”, apiKey);
```

var un = endpoint + "?from=en";

var un = endpoint + &quot?suggestedFrom=en";
var un = endpoint + "?to=en";

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

Answer:

Explanation:

```
var endpoint « _____
```

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

```
var apiKey « "FF956C68883B21B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":*' + text + '"}]';
var client • new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscript ion-Key", apiKey);
```

var uri = endpoint + "?from=en";
var un = endpoint + &quot ?suggestedFrom=en&quot
var un = endpoint + "?to=en";

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

Question: 86

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

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

Answer: C

Explanation:

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

Question: 87

You are developing a solution to generate a word cloud based on the reviews of a company's products.

Which Text Analytics REST API endpoint should you use?

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

Answer: A

Explanation:

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

Question: 88

DRAG DROP

You have a web app that uses Azure Cognitive Search.

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

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

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

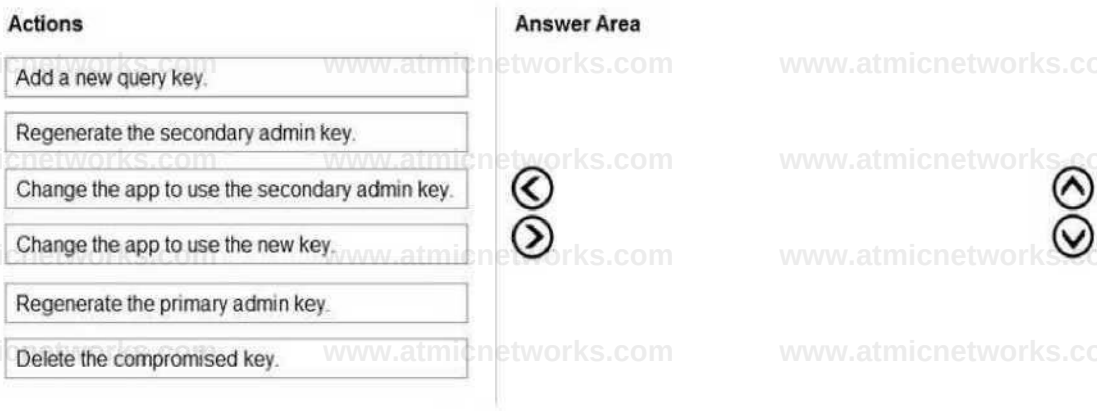

Answer:

Explanation:

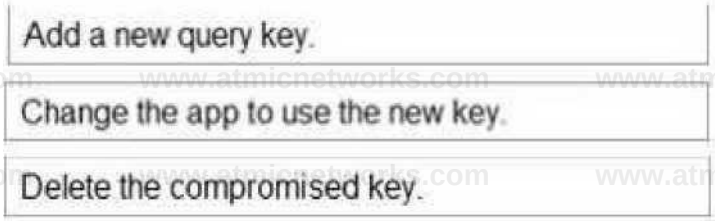

Reference:
https://docs.microsoft.com/en-us/azure/search/search-security-api-keys

Question: 89

HOTSPOT

You are building a bot and that will use Language Understanding.

You have a LUDown file that contains the following content.

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

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

- # Logout
- forget me
* log out

## Selectitem - choose last
- choose the {j0OirectionalReference=bottom left}
- choose {@OirectionalReference=top right}
- i like {@DirectionalReference=left} one

## SelectNone - none

0 ml DirectionalReference
0 prebuilt personName
```

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

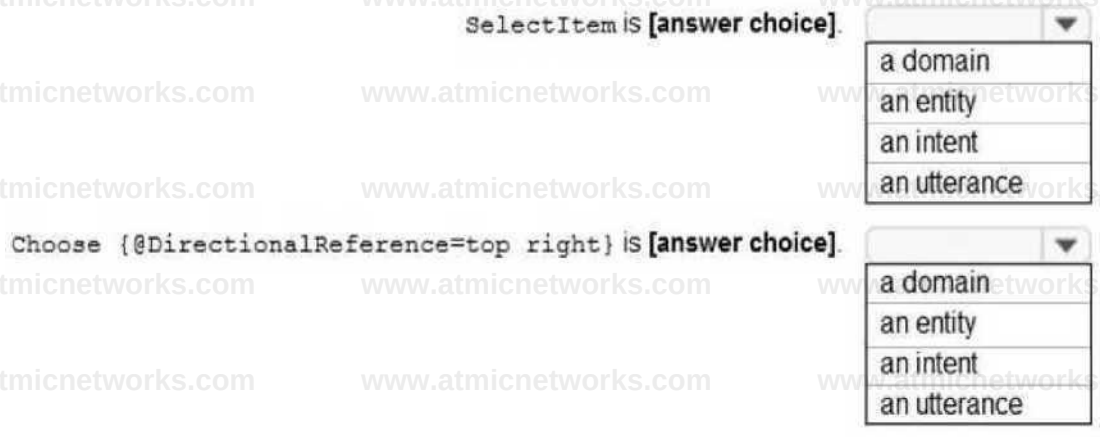

Answer:

Explanation:

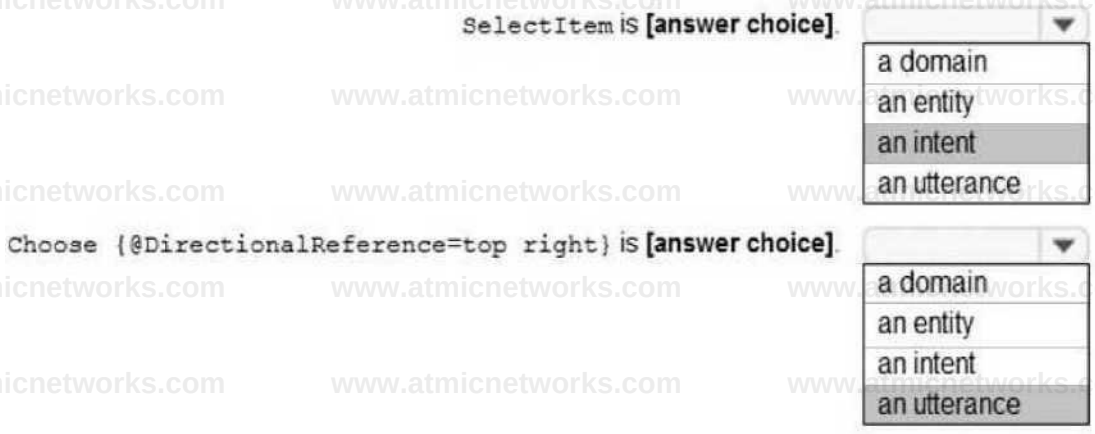

Reference:
https://github.com/solliancenet/tech-immersion-data-ai/blob/master/ai-exp1/README.md

Question: 90

HOTSPOT

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

You need to test the conversation flow by using the Microsoft Bot Framework Emulator.

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

NOTE: Each correct selection is worth one point.

u$er«Userl bot=watchbot user: I want a new watch.

bot: [˅][Delay=3C00] Attachment
ConversationUpdate Typing

bot: I can help you with that! Let me see what I can find, bot: Here's what I found.
bot:

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

adaptivecard
carousel list

Answer:

Explanation:

```
user'Userl
bot=watchbot
user: I want a new watch.
bot: [ v ] [Delay>3000]
```

Attachment
ConversationUpdate Typing

```
bot: I can help you with that! Let me see what I can find.
bot: Here's what I found.
bot:
[AttachmentLayout* v J
```

adaptivecard
carousel
thumbnail

```
[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.netZwatch02.jpg] user: I like the
first one.
bot: Sure, pulling up more information.
bot: [Attachment=cards\watchProfileCard.json ]
user: That's nice! Thank you.
bot: Sure, you are most welcome!
```

adaptivecard
carousel
list

Reference:
https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-add-media-attachments?view=azure-bot-service-4.0&tabs=csharp

Question: 91

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

You plan to deploy the bot to Azure.

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

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

NOTE: Each correct selection is worth one point.

A. botId
B. tenancld
C. appld
D. objeccld
E. appSecrec

Answer: CE

Explanation:

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

Question: 92

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

Answer: A

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: 93

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.

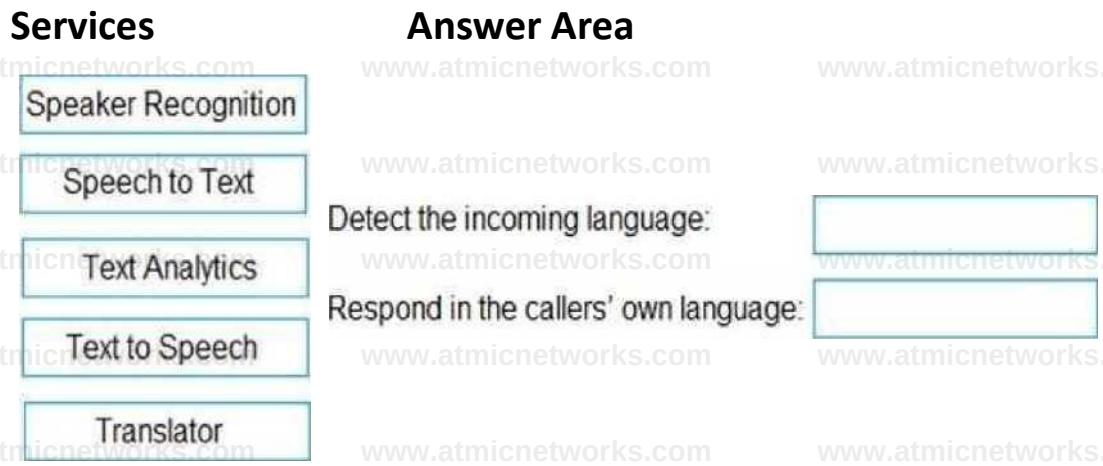

Answer:

Explanation:

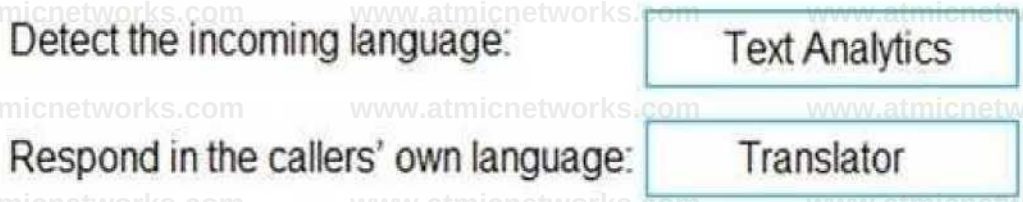

Box 1: Text Analytics

The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.

Box 2: Translator

Translator is a cloud-based neural machine translation service that is part of the Azure Cognitive Services family of REST APIs. Translator can be used with any operating system and powers many Microsoft products and services used by thousands of businesses worldwide to perform language translation and other language-related operations.

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: 94

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 FormRecognizerClienc client and the ScarcRecognizeConcencFromUri method
B. the FormTrainingClienc client and the ScarcRecognizeContentFromUri method
C. the FormRecognizerClienc client and the ScarcRecognizeReceipcsFromUri method
D. the FormTrainingClient client and the ScarcRecognizeReceipcsFromUri method

Answer: D

Explanation:

To analyze receipts from a URL, use the StartRecognizeReceiptsFromUri method

Example code:

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

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

Question: 95

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 SO pricing tier

Answer: A

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: 96

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

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

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

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

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

Solution: You add indexes.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Instead, you could migrate to a Cognitive Search service that uses a higher tier.

Note: A simple fix to most throttling issues is to throw more resources at the search service (typically replicas for query-based throttling, or partitions for indexing-based throttling). However, increasing replicas or partitions adds cost, which is why it is important to know the reason why throttling is occurring at all.

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

Question: 97

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

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

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

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

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

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

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Customer-managed key (CMK) encryption does not affect throttling.

Instead, you could migrate to a Cognitive Search service that uses a higher tier.

Note: A simple fix to most throttling issues is to throw more resources at the search service (typically replicas for query-based throttling, or partitions for indexing-based throttling). However, increasing replicas or partitions adds cost, which is why it is important to know the reason why throttling is occurring at all.

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

Question: 98

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.

Actions

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

Answer Area

Answer:

Explanation:

Create a project

Upload and tag images.

Train the classifier model.

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: 99

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.

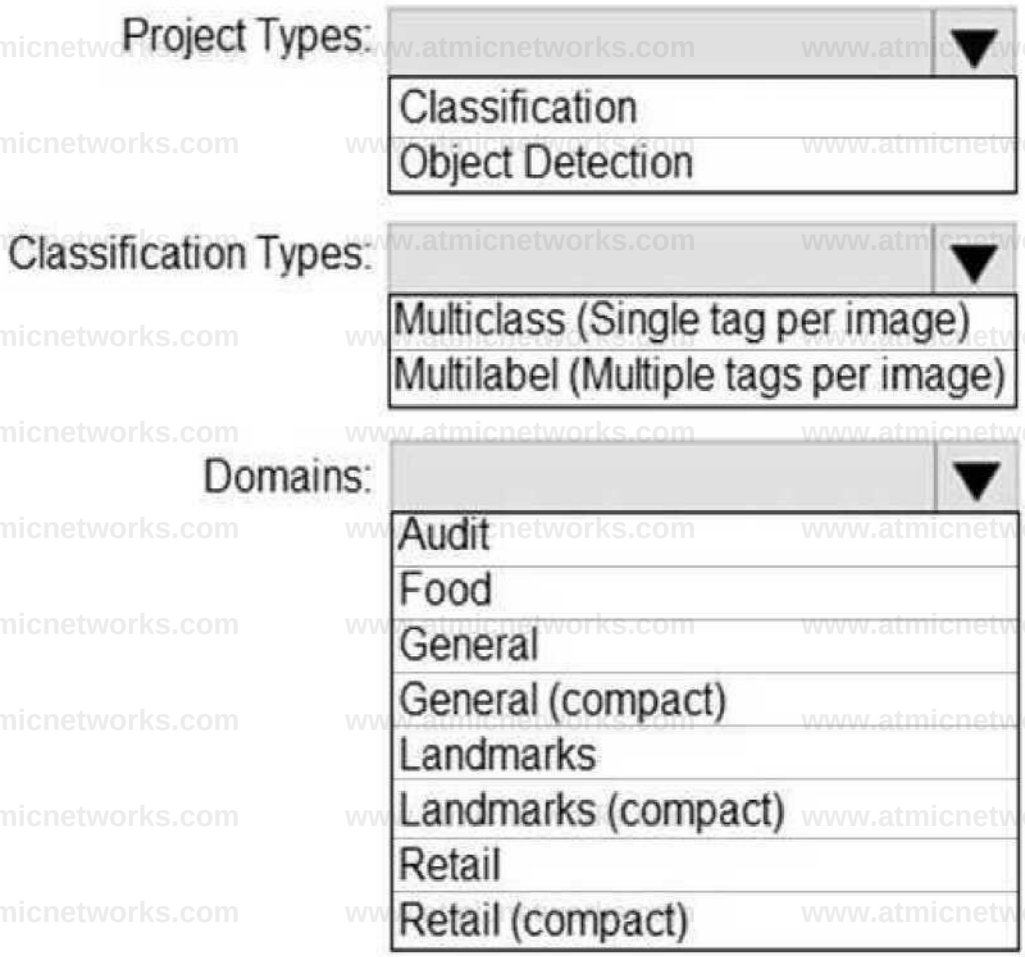

Answer:

Explanation:

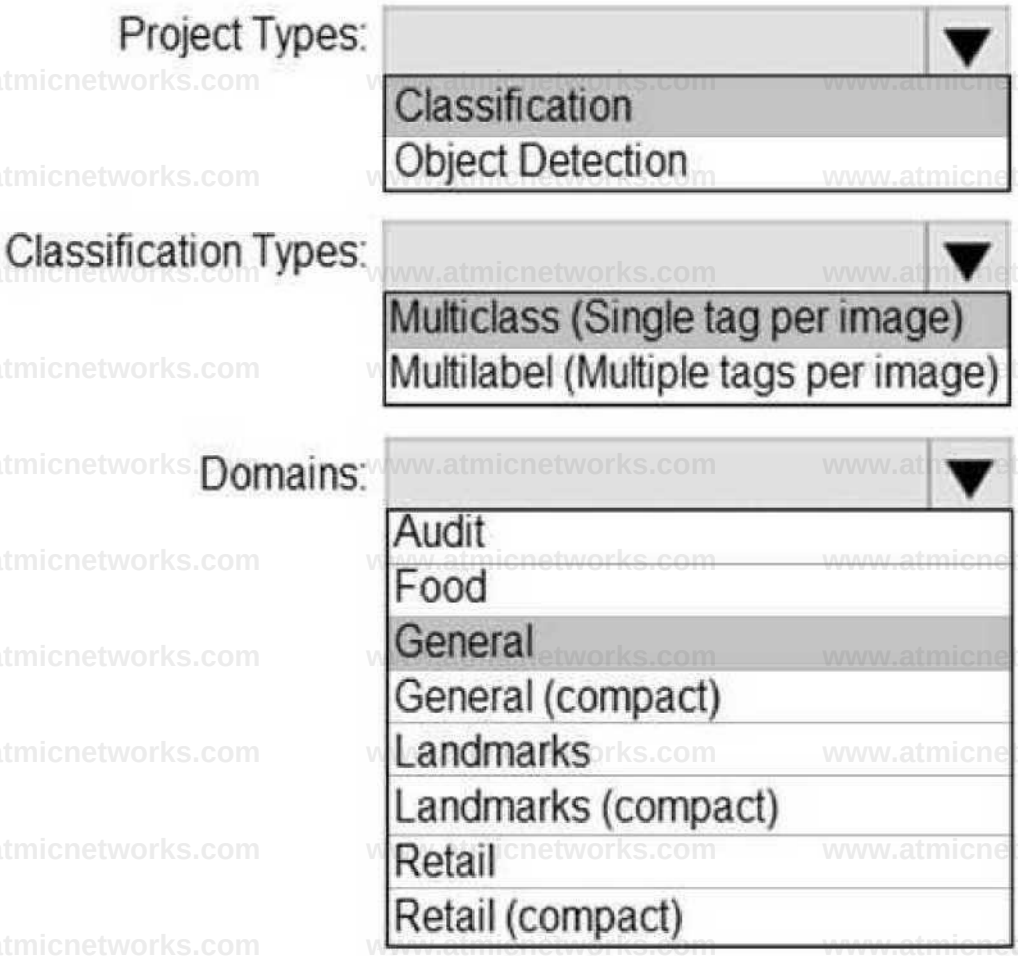

Box 1: Classification

Box 2: Multiclass

A multiclass classification project is for classifying images into a set of tags, or target labels. An image can be assigned to one tag only.

Box 3: General

General: Optimized for a broad range of image classification tasks. If none of the other specific domains are appropriate, or if you're unsure of which domain to choose, select one of the General domains.

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

Question: 100

HOTSPOT

You run the following command.

```
docker run –rm -it -p 5000:5000 –memory 10g –cpus 2 \
```

```
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment\
Eula=accept \
Billing={ENDPOINT_URI} \
ApiKey={API_KEY}
```

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

Statements	Yes	No
Going to http://localhost: 5000'status will query the Azure endpoint to verify whether the API key used to start the container is valid		
The container logging provider will write log data.	O	O
Going to http7/localhost'5000/swagger will provide the details to access the documentation for the available endpoints.	O	O

Answer:

Explanation:

Statements	Yes	No
Going to http //localhost 5000/status will query the Azure endpoint to venfy whether the API key used to start the container is valid		
The container logging provider will write log data	O	O
Going to http://localhost5000/swagger will provide the details to access the documentation for the available endpoints		

Box 1: Yes

http://localhost:5000/status : Also requested with GET, this verifies if the api-key used to start the container is valid without causing an endpoint query.

Box 2: Yes

The command saves container and LUIS logs to output mount at C:\output, located on container host

Box 3: Yes

http://localhost:5000/swagger : The container provides a full set of documentation for the endpoints and a Try it out feature. With this feature, you can enter your settings into a web-based HTML form and make the query without having to write any code. After the query returns, an example CURL command is provided to demonstrate the HTTP headers and body format that's required.

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

Question: 101

You are training a Language Understanding model for a user support system.

You create the first intent named GetContactDetails and add 200 examples.

You need to decrease the likelihood of a false positive.

What should you do?

A. Enable active learning.
B. Add a machine learned entity.
C. Add additional examples to the GetContactDetails intent.
D. Add examples to the None intent.

Answer: A

Explanation:

Active learning is a technique of machine learning in which the machine learned model is used to identify informative new examples to label. In LUIS, active learning refers to adding utterances from the endpoint traffic whose current predictions are unclear to improve your model.

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-glossary

Question: 102

DRAG DROP

You are building a Language Understanding model for purchasing tickets.

You have the following utterance for an intent named PurchaseAndSendTickets.

Purchase [2 audit business] tickets to [Paris] [next Monday] and send tickets to [email@domain.com]

You need to select the entity types. The solution must use built-in entity types to minimize training data whenever possible.

Which entity type should you use for each label? To answer, drag the appropriate entity types to the correct labels. Each entity type may be used once, more than once, or not at all.

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

Answer:

Explanation:

Box 1: GeographyV2

The prebuilt geographyV2 entity detects places. Because this entity is already trained, you do not need to add example utterances containing GeographyV2 to the application intents.

Box 2: Email

Email prebuilt entity for a LUIS app: Email extraction includes the entire email address from an utterance. Because this entity is already trained, you do not need to add example utterances containing email to the application intents.

Box 3: Machine learned

The machine-learning entity is the preferred entity for building LUIS applications.

Reference:

https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-geographyv2

https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-email

https://docs.microsoft.com/en-us/azure/cognitive-services/luis/reference-entity-machine-learned- entity

Question: 103

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: CDF

Explanation:

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

D. A query request targets the documents collection of a single index on a search service.

CF: In order to trim documents based on group_ids access, you should issue a search query with a group_ids/any(g:search.in(g, 'group_id1, group_id2,...')) filter, where 'group_id1, group_id2,...' are the groups to which the search request issuer belongs.

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

Question: 104

You are building a chatbot by using the Microsoft Bot Framework Composer as shown in the exhibit. (Click the Exhibit tab.)

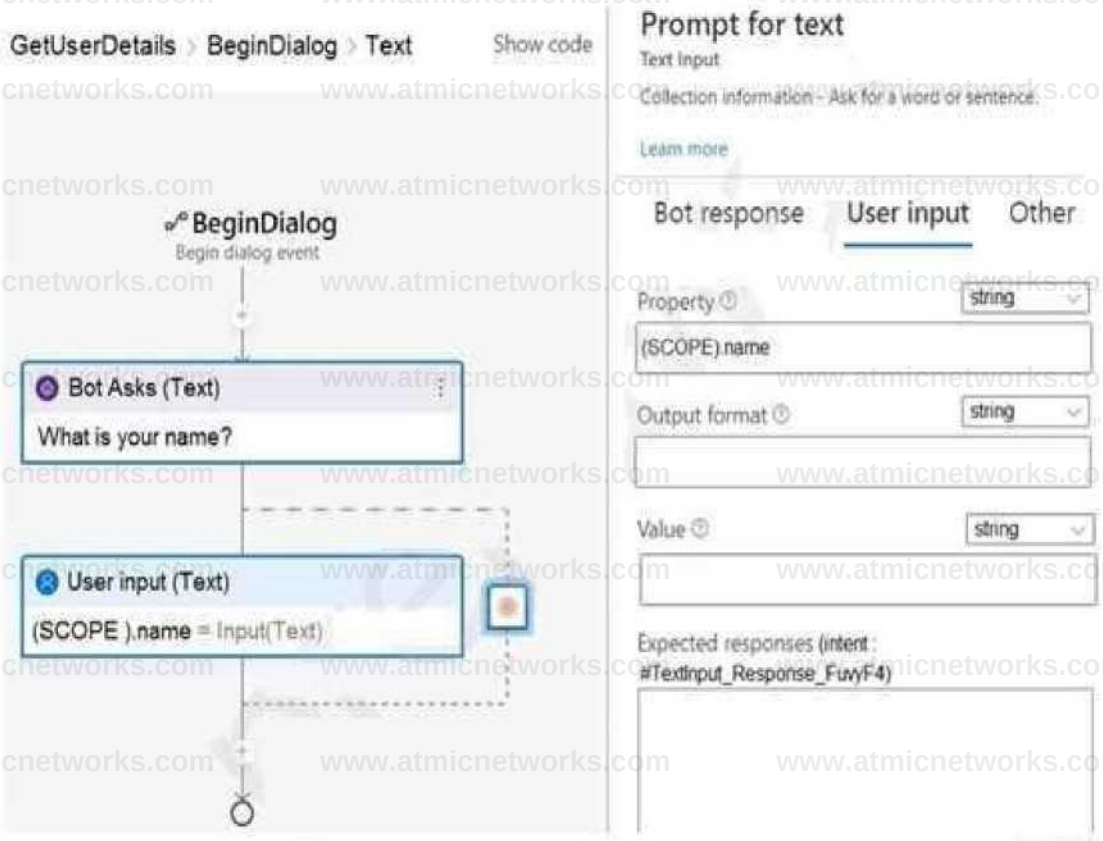

The chatbot contains a dialog named GetUserDetails. GetUserDetails contains a TextInput control that prompts users for their name.

The user input will be stored in a property named name.

You need to ensure that you can dispose of the property when the last active dialog ends.

Which scope should you assign to name?

A. dialog
B. user
C. curn
D. conversation

Answer: A

Explanation:

The dialog scope associates properties with the active dialog. Properties in the dialog scope are retained until the dialog ends.

Incorrect Answers:

A: The conversation scope associates properties with the current conversation. Properties in the conversation scope have a lifetime of the conversation itself. These properties are in scope while the bot is processing an activity associated with the conversation (for example, multiple users together in a Microsoft

Teams channel).

B: The user scope associates properties with the current user. Properties in the user scope do not expire. These properties are in scope while the bot is processing an activity associated with the user.

C: The turn scope associates properties with the current turn. Properties in the turn scope expire at the end of the turn.

Reference:
https://docs.microsoft.com/en-us/composer/concept-memory?tabs=v2x

Question: 105

DRAG DROP

You have a chatbot that uses a QnA Maker application.

You enable active learning for the knowledge base used by the QnA Maker application.

You need to integrate user input into 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.

Actions	Answer Area
Add a task to the Azure resource	
Approve and reject suggestions	
Publish the knowledge base	
Modify the automation task logic app to run an Azure Resource Manager template that creates the Azure Cognitive Services resource. ((v	
For the knowledge base, select Show active learning suggestions.	
Save and train the knowledge base	
Select the properties of the Azure Cognitive Services resource.	

Answer:

Explanation:

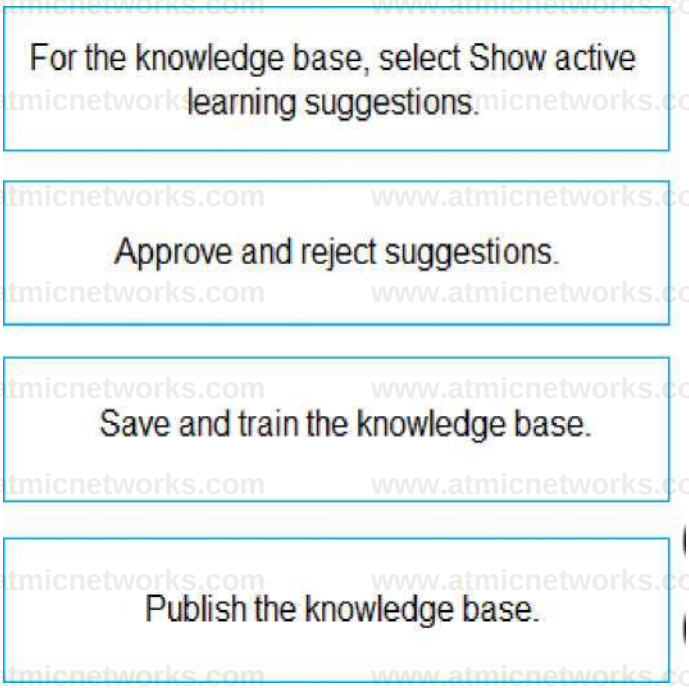

Step 1: For the knowledge base, select Show active learning suggestions.
In order to see the suggested questions, on the Edit knowledge base page, select View Options, then select Show active learning suggestions.

Step 2: Approve and reject suggestions.

Each QnA pair suggests the new question alternatives with a check mark, , to accept the question or an x to reject the suggestions. Select the check mark to add the question.

Step 3: Save and train the knowledge base.
Select Save and Train to save the changes to the knowledge base.

Step 4: Publish the knowledge base.

Select Publish to allow the changes to be available from the GenerateAnswer API.

When 5 or more similar queries are clustered, every 30 minutes, QnA Maker suggests the alternate questions for you to accept or reject.

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/how-to/improve-knowledge- base

Question: 106

You need to enable speech capabilities for a chatbot.

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

NOTE: Each correct selection is worth one point.

A. Enable WebSockets for the chatbot app.
B. Create a Speech service.
C. Register a Direct Line Speech channel.
D. Register a Cortana channel.
E. Enable CORS for the chatbot app.
F. Create a Language Understanding service.

Answer: ABC

Explanation:

You can use the Speech service to voice-enable a chat bot.

The Direct Line Speech channel uses the text-to-speech service, which has neural and standard voices.

You'll need to make a small configuration change so that your bot can communicate with the Direct Line Speech channel using web sockets.

Reference:
https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/tutorial-voice-enable- your-bot-speech-sdk

Question: 107

You are examining the Text Analytics output of an application.
The text analyzed is: "Our tour guide took us up the Space Needle during our trip to Seattle last week."
The response contains the data shown in the following table.

Text	Category	ConfidenceScore
Tour guide	Person Type	0.45
Space Needle	Location	0.38
Trip	Event	0.78
Seattle	Location	0.78
Last week	DateTjme	0.80 •

Which Text Analytics API is used to analyze the text?

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

Answer: B

Explanation:

Question: 108

HOTSPOT

You are developing an application that includes language translation.

The application will translate text retrieved by using a function named get_text_to_be_translated.

The text can be in one of many languages. The content of the text must remain within the Americas Azure geography.

You need to develop code to translate the text to a single language.

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

NOTE: Each correct selection is worth one point.

```
api_key  =  ■FF95  6C63B8  36  21.E336&lABt)2WA4C606"  text  •
g«_t^xt_to_bt_trMislMed() headers ■ {
        'Content -Type *: 'application/j son',„ "©cp-Apim-Subscription-
Key' : apikey } body — £
```

```
conn
response = conn.get response^)
response = conn.getrespond
```

```
Capi.cognirrve.mkrosofrtranslarorxom'J
Capt-apcxognmve.microsohtransfaiorxom")
Capi-nam.cognrttve niicrQ5offlranslator.com")

/translale?from=efi'
'/transit te?sugg»£tedFrom=wT
7Transiaie?ro=en“ /d elect ?to=en '
"/detect?trom=en *
```

Answer:

Explanation:

Question: 109

HOTSPOT

You create a knowledge store for Azure Cognitive Search by using the following JSON.

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

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

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

Answer Area

There will be [answer choice]

- no projection groups one
- projection group two
- projection groups four
- projection groups

Images will [answer choice!

- be saved to a blob container, be
- saved to file storage.
- be saved to an Azure Data lake.

Answer:

Explanation:

There will be (answer choice).

- two projection groups

Images will [answer choice)

- not be saved.

Question: 110

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

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

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

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

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

Solution: You deploy service1 and a private endpoint to vnet1.

Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:

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

The service could be an Azure service such as:

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

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

Question: 111

You have a Language Understanding resource named lu1.

You build and deploy an Azure bot named bot1 that uses lu1.

You need to ensure that bot1 adheres to the Microsoft responsible AI principle of inclusiveness. How should you extend bot1?

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

Answer: D

Explanation:

Inclusiveness: AI systems should empower everyone and engage people.

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

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

Question: 112

Your company uses an Azure Cognitive Services solution to detect faces in uploaded images. The method to detect the faces uses the following code.

```
itjr ' .ivnc Task ■ tortHerten(atrHnj ImaqoFHoPathl

    HltpClient client * new Httpcilent!);
    petaultPoquosUloadera .Add(w> •;.-ij in-'";t:: ::<■ , :.-? ,", subset ipt iCftKayl#
    t-tlti) request Parameter ■ *drt'-ci p-(Mr-l- I-ih 11 rt . n iH: riti-irnt-r>-1<1-1; ur >.t< ! unil-.irr i.-iridi.ir ki-l . l-r-*r
    nrrlnq url • endpoint ♦ */far<i/v|,o/<lntoct?* • request Par-ureters:
    HttpKoaponseHessaqe response;
    byte 11 byteJUt* - GetlnegesAaHyteAttey (leageFlleltith);
    1-   r 1 (BytcArrayContont content • n»» ByteArrayContunt (bytcliuta))

        Headers.Cent*r«tTypo - nee HedlaTypeliMkliuValtMl'applicAtlon/octet'Stroae'*);
        response - ■«-. 1 poatAnyiiclutk, content):
        it r mi; content$ttln<) - aw.ilt Content.HeaHAaStrlnttAnyncQ;
        ProcMsPetoelion (emt nntStt ing);
```

You discover that the solution frequently fails to detect faces in blurred images and in images that contain sideways faces.

A. Use a different version of the Face API.
B. Use the Computer Vision service instead of the Face service.
C. Use the Identify method instead of the Detect method.

D. Change the detection model.

You need to increase the likelihood that the solution can detect faces in blurred images and images that contain sideways faces.

What should you do?

Answer: D

Explanation:

Evaluate different models.

The best way to compare the performances of the detection models is to use them on a sample dataset. We recommend calling the Face - Detect API on a variety of images, especially images of many faces or of faces that are difficult to see, using each detection model. Pay attention to the number of faces that each model returns.

Reference:

https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/specify- detection-model

Question: 113

You have the following C# method.

You need to deploy an Azure resource to the East US Azure region. The resource will be used to perform sentiment analysis.

How should you call the method?

A. create_resource("res1", "ContentModerator", "S0", "eastus")
B. create_resource("res1", "TextAnalytics", "S0", "eastus")
C. create_resource("res1", "ContentModerator", "Standard", "East US")
D. create_resource("res1", "TextAnalytics", "Standard", "East US")

Answer: B

Explanation:

To perform sentiment analysis, we specify TextAnalytics, not ContentModerator.

Possible SKU names include: 'F0','F1','S0','S1','S2','S3','S4','S5','S6','S7','S8'

Possible location names include: westus, eastus

Reference:
https://docs.microsoft.com/en-us/powershell/module/az.cognitiveservices/new-azcognitiveservicesaccount

Question: 114

You build a Language Understanding model by using the Language Understanding portal.

You export the model as a JSON file as shown in the following sample.

```
"text": "average amount ol rain by month at Chicago last year", "intent":
"Weather.CheckWeatherValue", "entities": |

    "entity": "Weather.WeatherRange",
    "startPos": 0,
    "endPos": 6,
    "children": []
I,

    "entity": "Weather.Weathercondition",
    "startPos": 18,
    "endPos": 21, "children": |)

    "entity": "Weather.Historic" "startPos": 23, "endPos": 30,
    "children": J J
J
1
```

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

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

Answer: A

Explanation:

Question: 115

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.

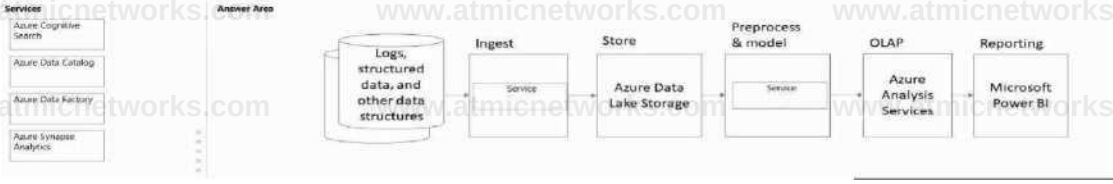

Answer:

Explanation:

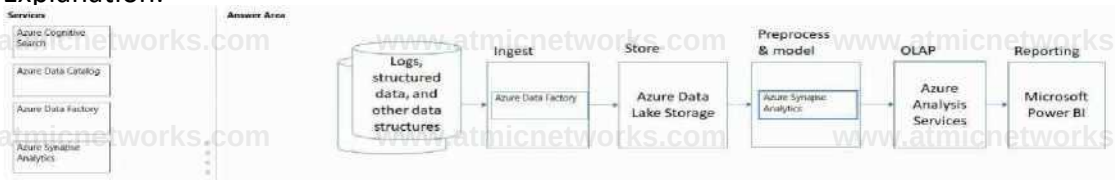

Question: 116

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.

Answer Areo

Statement*	Yes No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration 1 effort than infrastructure as a service (laaS) database offerings.	
Platform as a service (PaaS) database offerings In Azure provide end users with the ability to control and update the operating system version-	•
All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	Q

Answer:

Explanation:

Answer Area

Statements	Yes No
Platform as a service (PaaS) database offerings in Azure require less setup and configuration * effort tnan infrastructure as a service [laaS) database offerings,	
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	•
All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	@3

Question: 117

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

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

Answer: B

Explanation:

Question: 118

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.

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	O	O
Normalizing a database reduces data redundancy.	O	O
Normalization improves data integrity.	O	O

Answer:

Explanation:

Answer Area

Statements	Yes	No
Normalization involves eliminating relationships between database tables.	j	*
Normalizing a database reduces data redundancy.	®	O
Normalization improves data integrity.	®	O

Question: 119

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

A block of code that runs in a database is called

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

Answer:

Explanation:

Answer Area

A block of code that runs in a database is called a stored procedure

Question: 120

HOTSPOT

Select the answer that correctly completes the sentence.

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

Explanation

Answer Area

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

Question: 121

Your company needs to implement a relational database in Azure. The solution must minimize ongoing

maintenance. Which Azure service should you use?

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

Answer: B

Explanation:

Question: 122

You have a SQL query that combines customer data and order dat

a. The query includes calculated columns. You need to create a database object that would allow other users to rerun the same SOL query. What should you create?

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

Answer: B

Explanation:

Question: 123

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

Answer: A, D

Explanation:

Question: 124

You have data saved in the following format.

```
FirstName,LastName,Age,LeisureHobby,SportsHobby
John.baith.2J.Heading.Basketball
ben.^mith.21.buitar.Curling
```

Which format was used?

A. CSV
B. JSON
C. HTML
D. YAML

Answer: A

Explanation:

Question: 125

What is a primary characteristic of a relational database?

A. data is queried and manipulated by using a variant of the SOL language
B. a lack of dependencies between tables
C. a flexible data structure
D. a large amount of duplicate data

Answer: C

Explanation:

Question: 126

HOTSPOT
Select the answer that correctly completes the sentence.

Answer Area

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

A data analyst
A data engineer

Answer

Explanation

Answer Area

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

Question: 127

HOTSPOT
Select the answer that correctly completes the sentence.

Answer Area

is responsible for creating visuals and charts that help a company make informed decisions.

A data analyst
A data engineer
A data scientist
A database administrate!

Answer:

Explanation:

Question: 128

DRAG DROP

Match the types of workloads to the appropriate scenarios.

To answer, drag the appropriate workload type from the column on the left to its scenario on the right.

Each workload type may he used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Answer:

Explanation:

Question: 129

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

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

Answer: A, B

Explanation:

Question: 130

What should you use to build a Microsoft Power Bi paginated report?

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

Answer: A

Explanation:

Question: 131

HOTSPOT

Select the answer that correctly completes the sentence.

The massively parallel processing (MPP) engine of Azure Synapse Analytics

distributes processing across compute nodes.
distributes processing across control nodes, redirects client connections across compute nodes, redirects client connections across control nodes.

Answer:

Explanation:

Answer Area

The massively parallel processing (MPP) engine
Of Azure Synapse Analytics distributes processing across compute nodes.

Question: 132

Which scenario is an example of a streaming workload?

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

Answer: D

Explanation:

Question: 133

What is the primary purpose of a data warehouse?

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

Answer: C

Explanation:

Question: 134

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.

Statement!	V«	No
Azure Darabricks kan Apache Spark-based analytics platform.		
Azure Analysis Services is used for transactional workloads,		
Azure Data Factory orchestrates data integration workflows.	.:	.'

Answer:

Explanation:

Answer Area

Statements	Ve* No
Azure Databricks is an Apache Spark-based analytics platform. (>	
Azure Analysts Services is used for transactional workloads. ®	
Azure Data Factory orchestrates data integration workflows. ®	

Question: 135

You need to develop a solution to provide data to executives. The solution must provide an interactive graphical interface, depict various key performance indicators, and support data exploration by using drill down. What should you use in Microsoft Power Bl?

A. a report
B. Microsoft Power Apps
C. a view
D. a dataflow

Answer: C

Explanation:

Question: 136

Your company has a repotting solution that has paginated reports. The reports query a dimensional model in a data warehouse. Which type of processing does the reporting solution use?

A. Online Transaction Processing (OLTP)
B. Online Analytical Processing (OLAP)
C. batch processing
D. stream processing

Answer: B

Explanation:

Question: 137

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

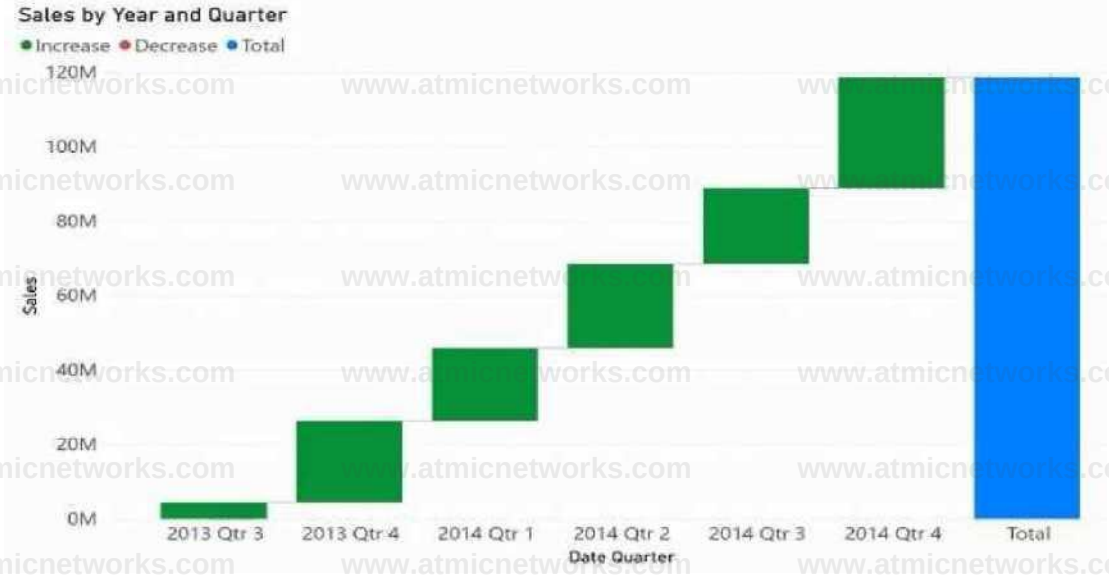

What should you create in Cower Bl Desktop;1

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

Answer: C

Explanation:

Question: 138

HOTSPOT

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

Answer Area

Statements	Ves	No
Stream processing has access to the most recent data received or data within a rolling time window.		O
Batch processing must occur immediately and have latency in the order of seconds or GO milliseconds.		
Stream processing is used for simple response functions, aggregates, or calculations such as O O rolling averages.		

Answer:

Explanation:

Answer Area

Statements	Ves	No
Stream processing has access to the most recent data received or data within a rolling time • window.		
Batch processing must occur immediately and have latency in the order of seconds or • milliseconds.		
Stream processing is used for simple response functions, aggregates, or calculations such as rolling averages.		•

Question: 139

Which Azure Storage service implements the key/value model?

A. Azure Files

B. Azure Blob
C. Azure Table
D. Azure Queue

Answer: C

Explanation:

Question: 140

HOTSPOT
Select the answer that correctly completes the sentence.

Answer Area

Varying fields for each entity in a J SON document is an example of

relational data, semi-Struclured data, structured data, unstructured data.

Answer

Explanation

Answer Area

Varying fields for each entity in a JSON document is an example of structured data.

Question: 141

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

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

Answer: A

Explanation:

Question: 142

HOTSPOT
Select the answer that correctly completes the sentence.

Answer Aren

tn 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

Explanation

Answer Area

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

Question: 143

HOTSPOT

Select the answer that correctly completes the sentence.

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:

Explanation:

Question: 144

Which property of a transactional workload guarantees that each transaction is treated as a single unit that either succeeds completely or tails completely?

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

Answer: B

Explanation:

Question: 145

Which database transaction property ensures that transactional changes to a database are preserved during unexpected operating system restarts?

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

Answer: A

Explanation:

Question: 146

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

A JSON document is an example of

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

Answer:

Explanation:

Answer Area

A JSON document is an example of relational data.

Question: 147

HOTSPOT

Select the answer that correctly completes the sentence.

Relational data is stored tn

a file system as unstructured data.
j hierarchal folder structure.
a tabular form of rows and co* lnns, comma-separated value (CSV) files.

Answer:

Explanation:

Answer Area

Relational data is stored tn a hierarchal folder structure.

Question: 148

Which database transaction property ensures that individual transactions are executed only once and either succeed in their entirety or roll back?

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

Answer: A

Explanation:

Question: 149

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

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

Answer: C

Explanation:

Question: 150

You need to recommend a non-relational data store that is optimized for storing and retrieving text

files, videos, audio streams, and virtual disk images. The data store must store data, some metadata, and a unique ID for each file. Which type of data store should you recommend?

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

Answer: D

Explanation:

Question: 151

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

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

graph table,
partition key.
document-

Answer:

Explanation:

Answer Area

When using the Azure Cosmos DB Gremlin API. the container resource type is projected as a graph-

Question: 152

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.

Answer Area

Cassandra API

Gremlin API

MongoDB API

Table API

Graph data

JSON documents

Key/value data

Answer:

Explanation:

Answer Area

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

Question: 153

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

Answer: A

Explanation:

Question: 154

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

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

Answer: D

Explanation:

Question: 155

What is a characteristic of a non-relational database?

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

Answer: C

Explanation:

Question: 156

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

When provisioning an Azure Cosmos DB [account / container / database / item], you need to specify which type of API you will use.

Answer:

Explanation:

Answer Area

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

Question: 157

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 .flac audio file and the associated transcript as a w file
C. a .wav or mp3 file of the voice talent consenting to the creation of a synthetic version of their voice
D. a .zip file that contains 10-second .wav files and the associated transcripts as .txt files

Answer: D

Explanation:

Question: 158

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

Answer: B

Explanation:

Question: 159

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

Answer: A

Explanation:

Question: 160

HOTSPOT

You are developing a text processing solution.

You have the function shown below.

```
static void 6etkeyw©rds.Te«t«nalyticsCliert textAnalyticsClient, string text)

    var response • te*MnalyticsClient.aacognl:e(ntitles (text);
    Console.WrittllneOey words:');

    foreach (Categoricedtntity entity in response.Value)

        Console.witallna. J* identity. Text}*);
```

Answer Area statements	Ves	No
	The output will include the fol towing words: our and included.	
	The output will include the foltowing words Pais Ethel, and Tower.	
	The function will output all the key phrases from the input string to the console.	

Answer:

Explanation:

Answer Arc*

Statements	Yes	No
The output will include the following words our and included	•	
The output mil include the following words Pans, bhel and Tower	•	
The tunaion will output ail the key phrases from the input string to the console	!■[1]	

Question: 161

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.

Answer Area

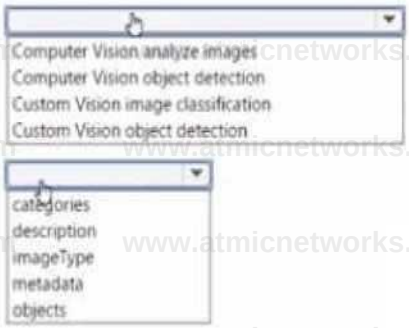

Answer:

Explanation:

Answer Area

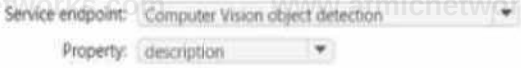

Question: 162

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.

Answer Area

The percentage of false positives is [answer choice].

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

Answer

Explanation:

Answer Arce

'he percentage of Wse positives c [answer choice] 30

The value lor the number of true pontnes divided by the total number of true positives and 50 fahe negate K [answer choree]

Question: 163

You have an Azure loT hub that receives series data from machinery. You need to build an app that

will perform the following actions:

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

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

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

Answer: D

Explanation:

Question: 164

You build a language model by using Conversational Language Understanding. The language model is used to search for information on a contact list by using an intent named Findcontact. A conversational expert provides you with the following list of phrases to use for training

- Find contacts in London.
- Who do I know in Seattle?
- Search for contacts m Ukraine.

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

A. Yes
B. No

Answer: B

Explanation:

Question: 165

You develop a Conversational Language Undemanding 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.

Answer: A

Explanation:

Question: 166

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 Bl. The solution must minimize development effort.

What should you add to the indexer?

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

Answer: C

Explanation:

Question: 167

You develop a custom question answering project in Azure Cognitive Service for Language. The project will be used by a chatbot. You need to configure the project to engage in multi-turn conversations. What should you do?

A. Add follow-up prompts.
B. Enable active learning.
C. Add alternate questions.
D. Enable chit-chat.

Answer: A

Explanation:

Question: 168

HOTSPOT

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

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

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

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

NOTE: Each correct selection is worth one point.

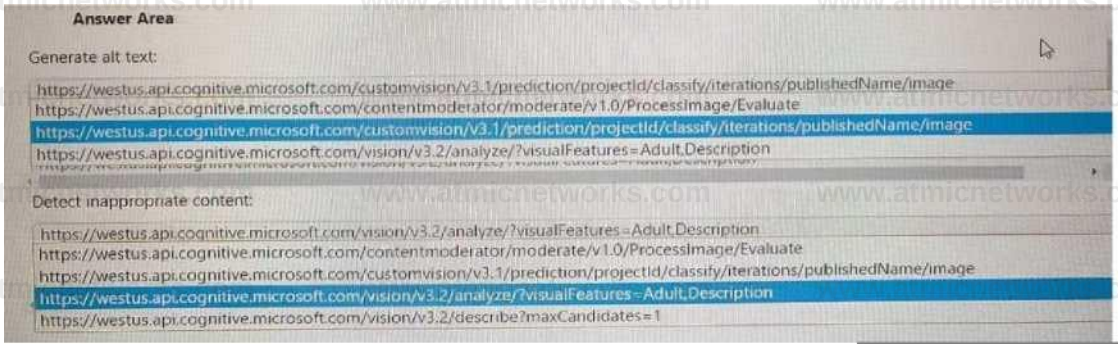

Answer:

Explanation:

Question: 169

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.

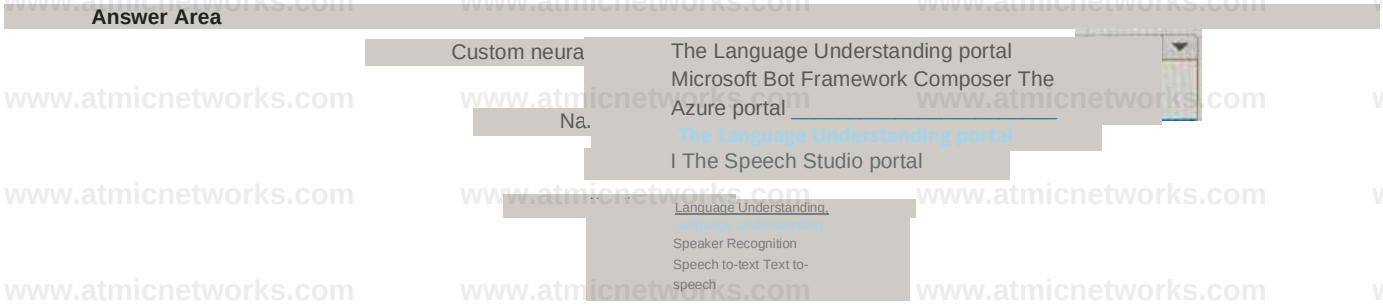

Answer:

Explanation:

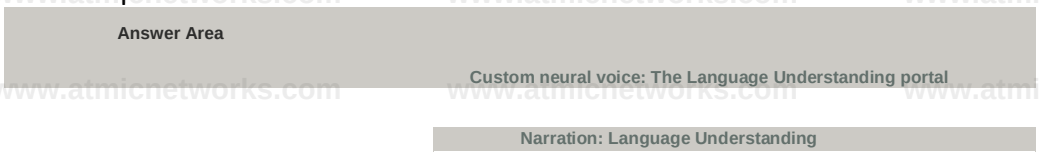

Question: 170

You need to measure the public perception of your brand on social media by using natural language processing. Which Azure service should you use?

A. Content Moderator
B. Form Recognizer
C. Computer Vision
D. Language service

Answer: D

Explanation:

Question: 171

You need to measure the public perception of your brand on social media by using natural language processing. Which Azure service should you use?

A. Content Moderator
B. Form Recognizer
C. Computer Vision
D. Language service

Answer: D

Explanation:

Question: 172

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.

Answer Area

if a user asks "what is the weather like in New York the bat win (answer choice) | identify N$v York at a . ity entity ^j
change to a different dialog
identify New York as a city entity
| identify New York as a state entity
respond with the weather nt Seattle

The GetWeather dialog uses a **[answer choice)** trigger I-^\Jii^^^i^J^H^itJJL^
Custom events
Dialog events
Language Understanding Intent recognized
OnA Intent recognized

Answer:

Explanation:

Answer Area

if a user asks what is the weather like m New York". the bot will **(answer choice]** identify New Yjik as a city entity *

The GetWeather dialog uses a **(answer choice)** trigger. |_dngw2£1^2£2i^!!!21!!l£^

Question: 173

HOTSPOT

You are building a chatbot by using the Microsoft Bot Framework Composer. You have the dialog

design shown in the following exhibit.

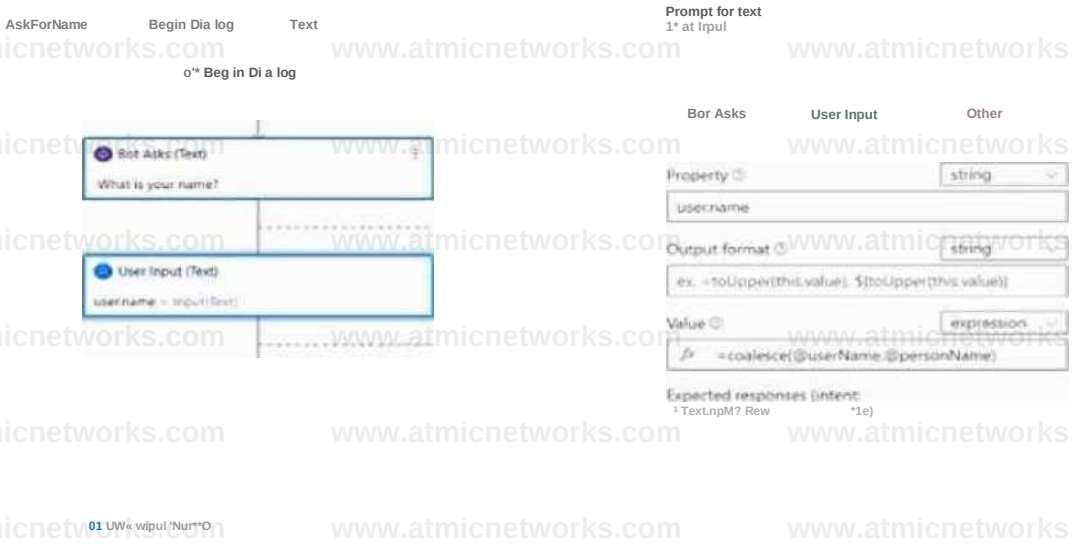

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection Is worth one point.

Answer Area

Statements	Yes	No
user, omv is an entity,	O	*Q
The dialog asks for a user name and a user age and assigns appropriate values to the user.name and user.age properties		
The chatbot attempts to take the first non*null entity value for userue*e Of personNaw and assigns the value to user. name.		1

Answer:

Explanation:

Answer Area

Statements	Yes	No
user.name is an entity.	O	W
The dialog asks for a user name and a user age and assigns appropriate values to the * user .name and user. age properties		
The chatbot attempts to take the first non-nuH entity value for username or personNaae * and assigns the value to user.name.		

Question: 174

You are examining the Language service output of an application.

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

The response contains the data shown in the following table.

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

Which Language service API is used 10 analyze the Text?

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

Answer: B

Explanation:

Question: 175

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.

Answer:

Explanation:

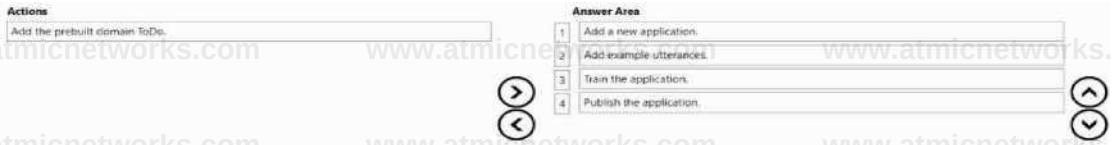

Question: 176

You are building a Conversational Language Understanding model.

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

- Set all the lights to on.
- Turn off the lights in the living room.
- What is the current thermostat temperature?
- Lower the temperature of the thermostat by five degrees.

Which three elements should you add to the model?

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

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

Answer: B, C, F

Explanation:

Question: 177

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.

Remitter the bot with th* Azure Bor Service

Open the Bot Framework Composer.

Build and run the hot.

| Open me Bor Framework Emulator

Connect to the bot endpoint.

Answer:

Explanation:

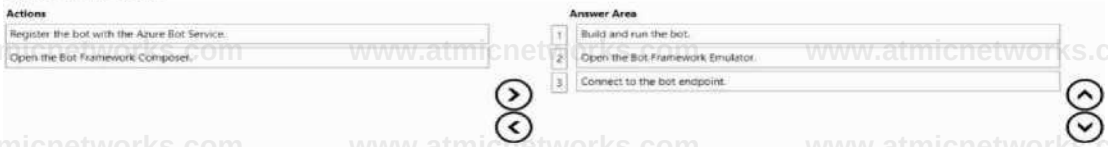

Question: 178

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

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

You create an Azure Cognitive Services resource.

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

A. https://uksouth.customvoice.apispeech.microsoft.eom/api/texttospeech/v3.0/longaudiosynthesis/v oices

B. https://uksouth.tts.speech.microsoft.conT/cognitiveservices/voices/list

C. https://uksouth.voice.speech.microsoft.com/cognitiveservices/v1?deploymentId = {deployment>

D. https://uksouth.api.cognitive.microsoft.eom/speechtotext/v3.0/models/base

Answer: C

Explanation:

Question: 179

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.

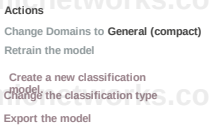

Answer

Explanation:

Question: 180

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.

Answer

Explanation:

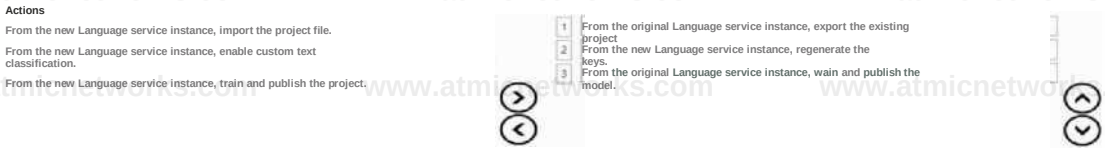

Question: 181

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

You have a PDF named Docl.pdf that contains a product catalogue and a price list

You upload Docl.pdf and train the model.

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

The chatbot fails to respond to the following question How much does <product* cost?

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

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

Does this meet the goal?

A. Yes

B. No

Answer: B

Question: 182

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.

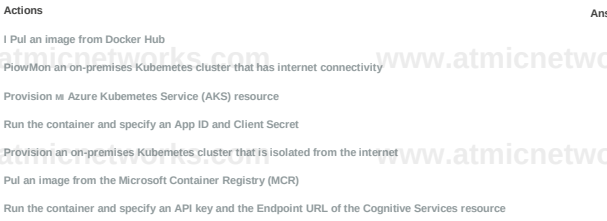

Answer:

Explanation:

Question: 183

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)

Answer: A

Explanation:

Question: 184

HOTSPOT

You have an Azure subscription that has the following configurations:

- Subscription ID: 8d3591aa-96b8-4737-ad09-OOf9b1ed35ad
- Tenant ID: 3edfe572-cbS4-3ced-ae12-c5c177f39a12

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

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

How should you complete the request? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

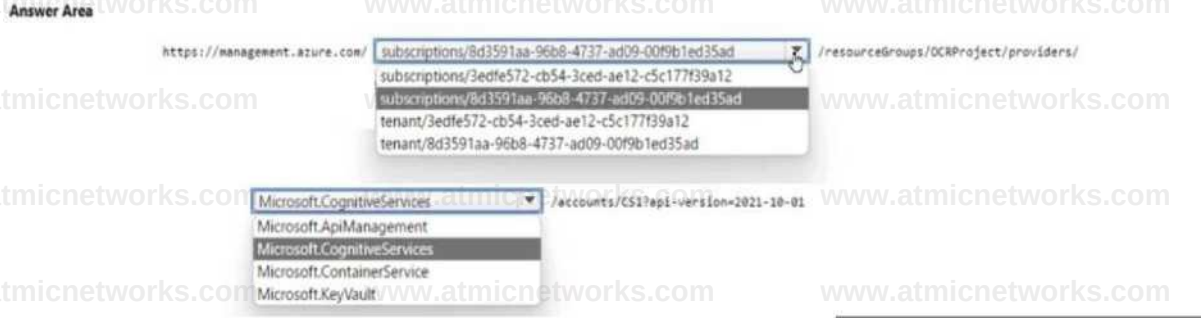

Answer:

Explanation:

Question: 185

You have a text-based chatbot.

You need to enable content moderation by using the Text Moderation API of Content Moderator. Which two service responses should you use? Each correct answer presents part of the solution NOTE: Each correct selection is worth one point.

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

Answer: C, D

Explanation:

Question: 186

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

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

Answer: A

Explanation:

Question: 187

DRAG DROP

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

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

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

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

Actions

| From the language Understanding portal, retrain the model

| From the host computer, run the container and specify the input directory.

From the Language Understanding portal export the solution as a package file

From the host computer, move the package file to the Docker input directory

From the host computer, build the container and specify the output directory

Answer Area

Answer:

Explanation:

Actions

| From the Language Understanding portal, retrain the model

I From the host computer, run the container and specify the input directory.

Answer Area

1 | From the Language Understanding portal, export the solution as a package file

2 From the host computer move the package file to the Dorter input directory.

3 From the host computer, build the container and specify the output directory.

Question: 188

DRAG DROP

You are building a customer support chatbot.

You need to configure the bot to identify the following:

- Code names for internal product development
- Messages that include credit card

The solution must minimize development effort.

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

NOTE: Each correct selection is worth one point.

Features

I Custom named entity recognition (NERI

Key phrase extraction

language detection

Named Entity Recognition (NER)

| Personally Identifiable Information (Psi) detection

Sentiment analysis

Answer Area

identify code names for internal product development

Identify messages that include credit card numbers

Answer:

Explanation:

Question: 189

You are building a chatbot.
You need to configure the bot to guide users through a product setup process.
Which type of dialog should you use?

A. component
B. waterfall
C. adaptive
D. action

Answer: B

Explanation:

Question: 190

You have an Azure subscription that contains an Azure Cognitive Service for Language resource. You need to identify the URL of the REST interface for the Language service. Which blade should you use in the Azure portal?

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

Answer: B

Explanation:

Question: 191

You are building a chatbot by using Microsoft Bot Framework Composer.
You need to configure the chatbot to present a list of available options. The solution must ensure that an image is provided for each option.
Which two features should you use? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. an Azure function

B. an adaptive card
C. an entity
D. a dialog
E. an utterance

Answer: B, D

Explanation:

Question: 192

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

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

Answer: B

Explanation:

Question: 193

DRAG DROP
You need to analyze video content to identify any mentions of specific company names.
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.

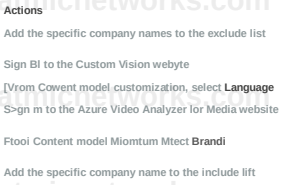

Answer:

Explanation:

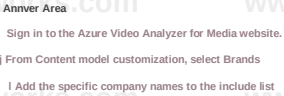

Question: 194

You plan to build an app that will generate a list of tags for uploaded images. The app must meet the

following requirements:

- Generate tags in a user's preferred language.
- Support English, French, and Spanish.
- Minimize development effort

You need to build a function that will generate the tags for the app. Which Azure service endpoint should you use?

A. Custom Vision image classification
B. Content Moderator Image Moderation
C. Custom Translator
D. Computer Vision Image Analysis

Answer: A

Explanation:

Question: 195

HOTSPOT

You have an Azure Cognitive Search resource named Search 1 that is used by multiple apps You need to secure Search 1. The solution must meet the following requirements:

- Prevent access to Search1 from the internet.
- Limit the access of each app to specific queries

What should you do? To answer, select the appropriate options in the answer area NOTE Each correct answer is worth one point.

Answer Area

To prevent access from the internet: [Create a private endpoint ˅^
Configure an IP firewall
Create a private endpoint.
Use Azure roles

To limit access to queries: use Azure roles.
Create a private endpoint.
Use Azure roles.
Use key authentication.

Answer:

Explanation:

Answer Area

To prevent access from the internet Create a private endpoint

To limit access to queues: Use Azure roles.

Question: 196

You have a mobile app that manages printed forms.

You need the app to send images of the forms directly to Forms Recognizer to extract relevant information.

For compliance reasons, the image files must not be stored in the cloud.

In which format should you send the images to the Form Recognizer API endpoint?

A. raw image binary
B. form URL encoded
D. JSON

Answer: A

Explanation:

Question: 197

DRAG DROP

You have a factory that produces cardboard packaging for food products. The factory has intermittent internet connectivity.

The packages are required to include four samples of each product.

You need to build a Custom Vision model that will identify defects in packaging and provide the location of the defects to an operator. The model must ensure that each package contains the four products.

Which project type and domain should you use? To answer, drag the appropriate options to the correct targets. Each option 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.

Options

^^

General

General (compact)

image classification

i-ogo]

Object detection

Answer Area

Project type

Domain:

Answer:

Explanation:

Options

2"-^

General

I1

Generat (compact)

image classification

| Logo

Object detection

Answer Area

Project type Object detection

Domain: I Food

Question: 198

You have an app that analyzes images by using the Computer Vision API.

You need to configure the app to provide an output for users who are vision impaired. The solution must provide the output in complete sentences.

Which API call should you perform?

A. readInStreamAsync
B. describeImageInStreamAsync
C. toglfiageInStreaaAsync

D. analyreImageByDomainInStreamAsync

Answer: D

Explanation:

Question: 199

You are designing a conversational interface for an app that will be used to make vacation requests. The interface must gather the following data:

- The start date of a vacation
- The end date of a vacation
- The amount of required paid time off

The solution must minimize dialog complexity. Which type of dialog should you use?

A. Skill
B. waterfall
C. adaptive
D. component

Answer: D

Explanation:

Question: 200

You have a Language service resource that performs the following:

- Sentiment analysis
- Named Entity Recognition (NER)
- Personally Identifiable Information (Pll) identification

You need to prevent the resource from persisting input data once the data is analyzed. Which query parameter in the Language service API should you configure?

A. loggingOptOut
B. piiCategories
C. showStats
D. Model-version

Answer: A

Explanation:

Question: 201

You have an Azure subscription that contains a Language service resource named ta1 and a virtual network named vnet1. You need to ensure that only resources in vnet1 can access ta1. What should you configure?

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

Answer: B

Explanation:

Question: 202

HOTSPOT

You have a collection of press releases stored as PDF files.

You need to extract text from the files and perform sentiment analysis.

Which service should you use for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

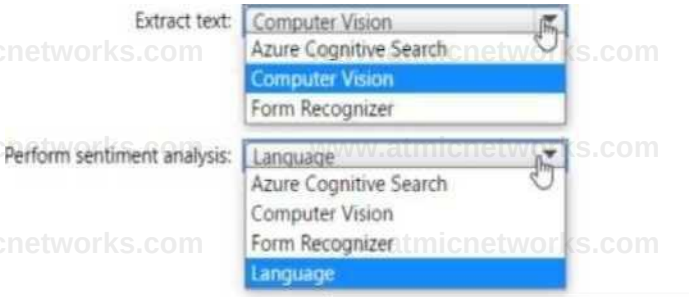

Answer:

Explanation:

Answer Area

Extract text Computer Vision

Perform sentiment analysis: Language

Question: 203

HOTSPOT

You are building a chatbot.

You need to use the Content Moderator service to identify messages that contain sexually explicit language.

Which section in the response from the service will contain the category score, and which category will be assigned to the message? To answer, select the appropriate options in the answer area, NOTE: Each correct selection is worth one point.

Answer Area

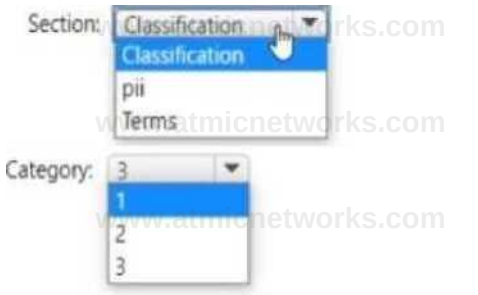

Answer:

Explanation:

Answer Area

Section: Classification

Category 3

Question: 204

HOTSPOT

You are building a solution that students will use to find references for essays. You use the following code to start building the solution.

```
using Azure;
using Sy5tea;
using Azure. AI.TextAnalytics;

private static readonly AzureKeyCredential credentials * new A:ureKeyCredential(*<key>*);
private static readonly Uri endpoint • new Uri(*<endpoint>");

static void Entitylinker(TextAnalyticsClient client)
{
    var response * client.UecognizeLinkedEntities(
        'Our tour guide took us up the Space Needle during our trip to Seattle last week.*)}
```

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

Answer Area

Statements	Yes	No
The code will detect the language of documents.	O	0
The wrl attribute returned for each linked entity will be a Bing search link.		0
The matches attribute returned for each linked entity will provide the location in a document [1] ■ where the entity is referenced.		O

Answer:

Explanation:

Answer Area

Statements	Yes	No
The code will detect the language of documents.	Q	||
The uri attribute returned for each linked entity will be a Bing search link.	•	
The Mtches attribute returned for each linked entity will provide the location in a document where the entity is referenced.	•	O

Question: 205

You are building a bot by using Microsoft Bot Framework.

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

What should you do?

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

Answer: B

Explanation:

Question: 206

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

Answer: C

Explanation:

Question: 207

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

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

Answer: C

Explanation:

Question: 208

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.

Answer:

Explanation:

Question: 209

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: A, C

Explanation:

Question: 210

You are building a Language Understanding solution.

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

You need to minimize the number of utterances used to fram the model.

Which type of custom entity should you use?

A. Pattera.any
B. machine-learning
C. list

D. regular expression

Answer: C

Explanation:

Question: 211

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

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

Answer: D

Explanation:

Question: 212

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

Answer: B

Explanation:

Question: 213

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.

Answer:

Explanation:

Question: 214

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. storaeeConnectionString
D. files
E. objects

Answer: D, E

Explanation:

Question: 215

HOTSPOT

You are building a call handling system that will receive calls from French-speaking and Germanspeaking 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 are

a. NOTE: Each correct selection is worth one point.

Answer Area

Answer:

Explanation:

Answer Area

To capture messages; Speech to text

To replay messages; Text-to-speech and Translator

Question: 216

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.

Answer

Explanation:

Question: 217

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

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

Answer: B, C

Explanation:

Question: 218

You are building an app that will include one million scanned magazine articles. Each article will be stored as an image file. You need to configure the app to extract text from the images. The solution must minimize development effort. What should you include in the solution?

A. Computer Vision Image Analysis
B. the Read API in Computer Vision
C. Form Recognizer
D. Azure Cognitive Service for Language

Answer: B

Explanation:

Question: 219

You have an Azure subscription that contains an AI enrichment pipeline in Azure Cognitive Search and an Azure Storage account that has 10 G8 of scanned documents and images. You need to index the documents and images in the storage account. The solution must minimize how long it takes to build the index.

What should you do?

A. From the Azure portal, configure parallel indexing.
B. Create a text-based indexer by using the REST API.
C. From the Azure portal, configure scheduled indexing.
D. Configure field mappings by using the REST API.

Answer: A

Explanation:

Question: 220

You are building a flight booking bot by using the Microsoft Bot Framework SDK.

The bot will ask users for their departure date. The bot must repeat the question until a valid date is given, or the users cancel the transaction.

Which type of dialog should you use?

A. prompt
B. action
C. waterfall
D. adaptive

Answer: A

Explanation:

Question: 221

You create five bots by using Microsoft Bot Framework Composer.

You need to make a single bot available to users that combines the bots. The solution must support dynamic routing to the bots based on user input.

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

NOTE: Each correct selection is worth one point.

A. Create an Orchestrator model.
B. Change the Recognizer/Dispatch type.
C. Create a composer extension.
D. Enable WebSockets.
E. Create a custom recognizer JSON file.
F. Install the Orchestrator package.

Answer: A, B, F

Explanation:

Question: 222

HOTSPOT

You are building an app by using the Speech SDK. The app will translate speech from French to German by using natural language processing.

You need to define the source language and the output language.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
var speechTr»nsi»tionConflg •

SpeechTranslationConfig.FromSubscriptionfspeechKeyj speechRegion); speech?ranslationCcnfig. i^eethfevOqnrtKnbincju^^

                                    AddTargetlanguage
                                    Speech Recognition Language
                                    SpeechSynthesislanguage
                                    Targetlanguages
                                    VoiceName

sp«ech_tr»nii»tion_conHg. | SpeechSynthesislanguage AddTargetlanguage
                                    SpeechRecognitionlanguage
                                    SpeechSynthesislanguage
                                    TargetLanguages
                                    VoiceName
```

Answer:

Explanation:

Answer Area

```
var speechTranslationConfig -

Speech?r anslationConfig. From5ubscri.pt ion(speechKey, speechRegion); speechTransiationConftg.

                                    SpeechRecognitionlanguage          ' «"fr*

speech_tr»n»i»tion_eonf£g.          SpeechSynthesislanguage            T ("de")
```

Question: 223

DRAG DROP

You have a collection of Microsoft Word documents and PowerPoint presentations in German.

You need to create a solution to translate the file to French. The solution must meet the following requirements:

* Preserve the original formatting of the files.
* Support the use of s custom glossary.

You create a blob container for German files and a blob container for French files. You upload the original files to the container for German files.

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

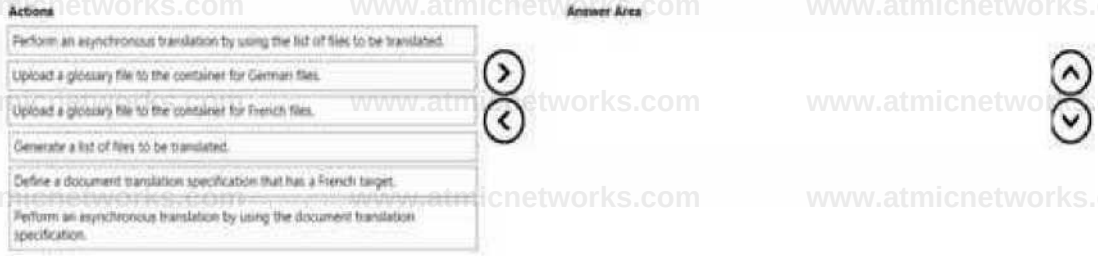

Answer:

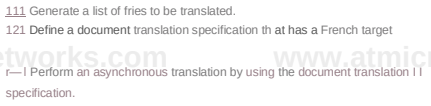

Explanation:

Question: 224

DRAG DROP

You have an app that manages feedback.

You need to ensure that the app can detect negative comments by using the Sentiment Analysis API in Azure Cognitive Service for Language. The solution must ensure that the managed feedback remains on your company's internal network.

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

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

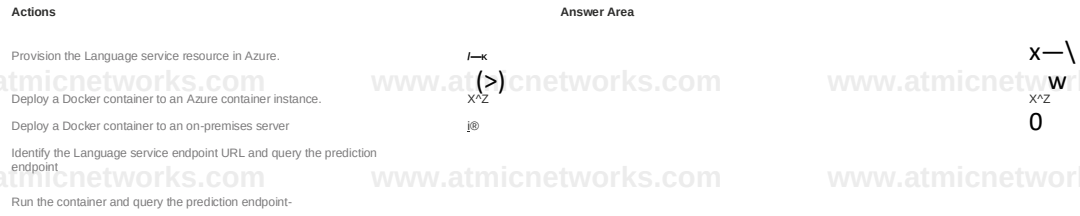

Answer:

Explanation:

Provision the Language service resource in Azure.

Deploy a Docker container to an on-premises server. Run the container and query the prediction endpoint.

According to the Microsoft documentation, the Language service is a cloud-based service that provides various natural language processing features, such as sentiment analysis, key phrase extraction, named entity recognition, etc. You can provision the Language service resource in Azure by following the steps in Create a Language resource. You will need to provide a name, a subscription, a resource group, a region, and a pricing tier for your resource. You will also get a key and an endpoint for your resource, which you will use to authenticate your requests to the Language service API.

According to the Microsoft documentation, you can also use the Language service as a container on your own premises or in another cloud. This option gives you more control over your data and network, and allows you to use the Language service without an internet connection. You can deploy a Docker container to an on-premises

server by following the steps in Deploy Language containers. You will need to have Docker installed on your server, pull the container image from the Microsoft Container Registry, and run the container with the appropriate parameters. You will also need to activate your container with your key and endpoint from your Azure resource.

According to the Microsoft documentation, once you have deployed and activated your container, you can run it and query the prediction endpoint to get sentiment analysis results. The prediction endpoint is a local URL that follows this format: http://<container IP address>:<port>/text/analytics/v3.1-preview.4/sentiment. You can send HTTP POST requests to this endpoint with your text input in JSON format, and receive JSON responses with sentiment labels and scores for each document and sentence in your input.

Question: 225

You are building an app that will share user images.

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

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

The solution must minimize development effort.

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

NOTE: Each correct selection is worth one point.

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

Answer: CD

Explanation:

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

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

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

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

Question: 226

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

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

You have a chatbot that uses question answering in Azure Cognitive Service for Language

Users report that the responses of the chatbot lack formality when answering spurious questions

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

Solution: From Language Studio, you change the chitchat source to qna_chitchat_professional.tsv. and then retrain and republish the model.

Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Question: 227

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 lark formality when answering spurious questions

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

Solution: From Language Studio, you change the chitchat source to qna_chitchit_friindly.tsv. and then retrain and republish the model.

Does this meet the goal?

A. Yes

B. No

Answer: B

Explanation:

Question: 228

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 lark formality when answering spurious questions You need to ensure that the chatbot provides formal responses to spurious questions.

Solution: From Language Studio, you remove all the chit-chat question and answer pairs, and then retrain and republish the model

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Removing all the chit-chat question and answer pairs from the project will not ensure that the chatbot provides formal responses to spurious questions. It will only make the chatbot unable to handle any chit-chat scenarios, which may result in a poor user experience and a loss of engagement. Instead, you should choose a chit-chat personality that matches the tone and style of your chatbot, such as Professional or Caring. You can also edit the chit-chat questions and answers to suit your specific needs, or add new ones that are not in the predefined data set12. This way, you can ensure that the chatbot responds appropriately to spurious questions, while still maintaining a conversational and engaging interaction with the user.

Question: 229

HOTSPOT

You are building an app that will share user images.

You need to configure the app to meet the following requirement's

- uploaded images must be scanned and any text must be extracted from the images.
- Extracted text must be analyzed for the presence of profane language.
- The solution must minimize development effort.

what should you use for each requirement? To answer, select the apocopate opinions in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

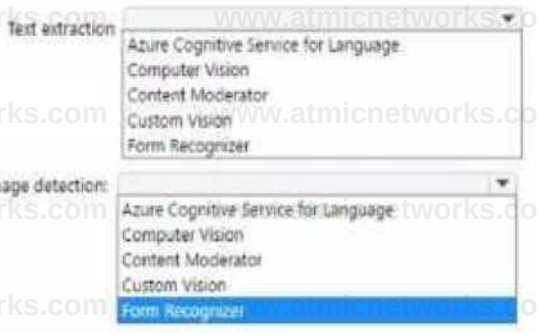

Answer:

Explanation:

Text Extraction: b) Computer Vision

Computer Vision is a service that can analyze images and extract text from them using optical character recognition (OCR) or read API1. OCR can detect and extract printed or handwritten text from images in various languages and formats, such as PDF, TIFF, or JPEG2. Read API can perform asynchronous batch processing of

text from multiple images, and can also handle text in tables, mixed languages, and rotated text3.

Profane Language Detection: c) Content Moderator

Content Moderator is a service that can detect and filter out potentially offensive or unwanted content in text, images, and videos. For text, Content Moderator can identify profane or abusive language, personal data, and custom terms that you define. You can also use Content Moderator to review and moderate the content manually or automatically using workflows.

Question: 230

You are building a solution that will detect anomalies in sensor data from the previous 24 hours.
You need to ensure that the solution scans the entire dataset, at the same time, for anomalies.
Which type of detection should you use?

A. batch
B. streaming
C. change point

Answer: A

Explanation:

Batch anomaly detection is a type of anomaly detection that scans the entire dataset at once for outliers and unusual patterns. Batch anomaly detection is suitable for ofline analysis of historical data, such as sensor data from the previous 24 hours. Batch anomaly detection can use various techniques, such as statistical methods, machine learning methods, or hybrid methods, to identify anomalies in the data123.

Question: 231

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

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

Answer: D

Explanation:

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

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

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

The best option is to upload file1.avi to a cloud storage service such as Microsoft OneDrive, and then generate a shareable link for the file. You can then paste the link in the Azure Video Indexer website and upload the video from the URL. This way, you can avoid the file size limitation and ensure that the file is accessible and

valid2.

Question: 232

DRAG DROP

You Build a bot in JavaScript.

From the Azure Command-Line interface (CLI), you run the following command.

az bot prepare-deploy

You need to deploy the bot to Azure.

Which three Azure CLI commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the client order.

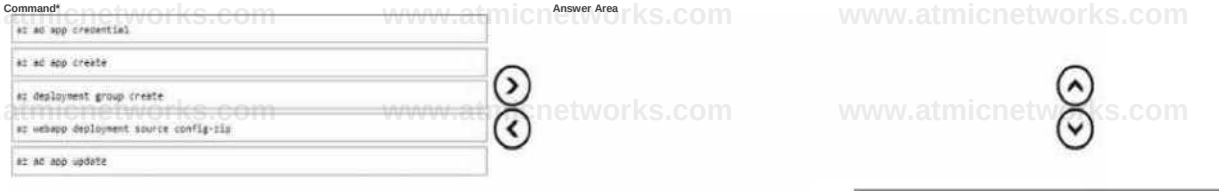

Answer:

Explanation:

To deploy the bot to Azure, you should run the following three Azure CLI commands in sequence: az deployment group create This command will create the Azure resources for your bot using an ARM template and a parameters file. You need to specify the resource group name, the template file path, and the parameters file path. For example:

az deployment group create --resource-group myResourceGroup --template-file "deploymentTemplates\template-with-preexisting-rg.json" --parameters "deploymentTemplates\parameters-for-template-BotApp-with-rg.json"

This command will also output the app ID and password of your bot, which you will need for the next command1.

az webapp deployment source config-zip This command will deploy your bot code to the app service that you created in the previous step. You need to specify the resource group name, the app service name, and the zip file path of your bot code. For example:

az webapp deployment source config-zip --resource-group myResourceGroup --name myBotAppService --src "code.zip"

This command will also output the URL of your bot endpoint, which you will need for the next command2.

az ad app update This command will update your bot registration with the endpoint URL of your bot.

You need to specify the app ID of your bot and the endpoint URL. For example:

az ad app update --id myBotAppId --set replyUrls="https://myBotAppService.azurewebsites.net/api/messages"

This command will complete the deployment process and make your bot ready to be tested3.

Question: 233

You are building a chatbot.

You need to ensure that the bot will recognize the names of your company's products and codenames. The solution must minimize development effort.

Which Azure Cognitive Service for Language service should you include in the solution?

A. custom text classification
B. entity linking
C. custom Named Entity Recognition (NER)

D. key phrase extraction

Answer: C

Explanation:

Question: 234

You are developing an app that will use the Decision and Language APIs.

You need to provision resources for the app. The solution must ensure that each service is accessed by using a single endpoint and credential.

Which type of resource should you create?

A. Language
B. Speech
C. Azure Cognitive Services
D. Content Moderator

Answer: C

Explanation:

Question: 235

You have an Azure subscription that contains an Azure App Service app named Appl.

You provision a multi-service Azure Cognitive Services resource named CSAccount1.

You need to configure App1 to access CSAccount1. The solution must minimize administrative effort.

What should you use to configure App1?

A. a system-assigned managed identity and an X.509 certificate
B. the endpoint URI and an OAuth token
C. the endpoint URI and a shared access signature (SAS) token
D. the endpoint URI and subscription key

Answer: C

Explanation:

Question: 236

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

You develop an app in C# named App1.

You need to configure App1 to use Model1.

Which package should you add to App1?

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

Answer: C

Explanation:

Question: 237

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

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

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

Answer: B

Explanation:

Question: 238

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

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

You have a chatbot that uses question answering in Azure Cognitive Service for Language.

Users report that the responses of the chatbot lack formality when answering random questions that are outside the scope of the knowledge base.

You need to ensure that the chatbot provides formal responses to these spurious questions. Solution: From Language Studio, you modify the question and answer pairs for the custom intents, and then retrain and republish the model.

Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:

Question: 239

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.

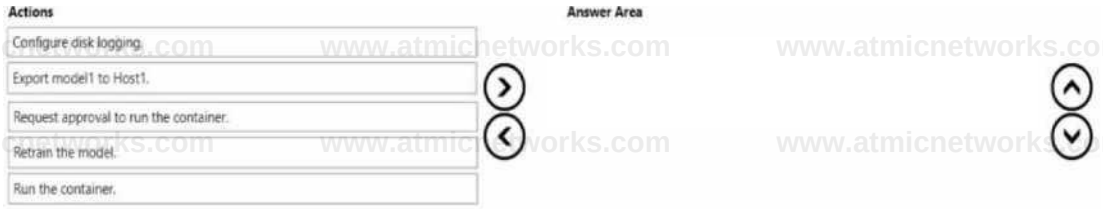

Answer:

Explanation:

According to the course AI-102T00: Designing and Implementing a Microsoft Azure AI Solution1, the correct sequence of actions to run model1 on Host1 is:

Export model1 to Host1

Run the container

Configure disk logging

The explanation and references are as follows:

Export model1 to Host1: This step is required to deploy the custom speech-to-text model to the Docker host. You can use the Azure portal or the Azure CLI to export the model as a container image2.

Run the container: This step is required to start the container and run the model on the Docker

host. You can use the Docker CLI to run the container image3.

Configure disk logging: This step is optional but recommended to monitor the performance and health of the container. You can use the Docker CLI to configure disk logging for the container4.

Question: 240

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.

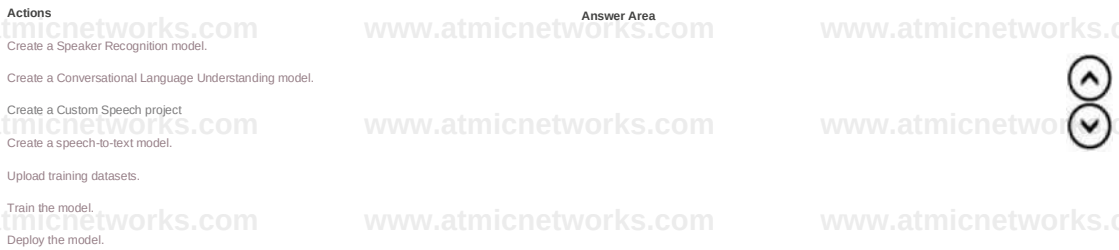

Answer:

Explanation:

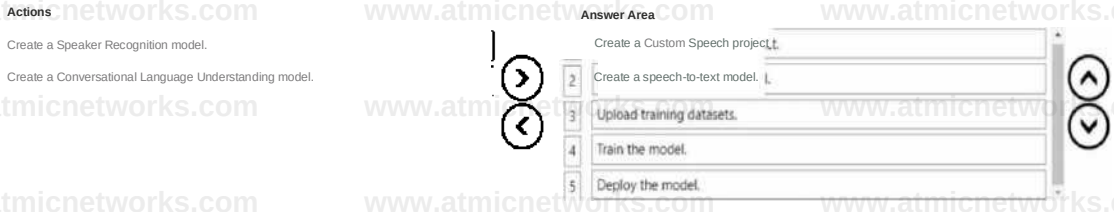

https://learn.microsoft.com/en-us/azure/cognitive-services/speech-service/custom-speech- overview#how-does-it-work

With Custom Speech, you can upload your own data, test and train a custom model, compare accuracy between models, and deploy a model to a custom endpoint.

- Create a project and choose a model. Use a Speech resource that you create in the Azure portal. If you will train a custom model with audio data, choose a Speech resource region with dedicated hardware for training audio data.

- Upload test dat

a. Upload test data to evaluate the speech to text offering for your applications, tools, and products.

- Train a model. Provide written transcripts and related text, along with the corresponding audio data. Testing a model before and after training is optional but recommended.

- Deploy a model. Once you're satisfied with the test results, deploy the model to a custom endpoint. With the exception of batch transcription, you must deploy a custom endpoint to use a Custom Speech model.

Question: 241

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

Answer: C

Explanation:

Question: 242

You have an Azure Cognitive Search solution and a collection of blog posts that include a category

field. You need to index the posts. The solution must meet the following requirements:

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

Which index attributes should you configure for the category field?

A. searchable, facetable, and retrievable

B. retrievable, filterable, and sortable

C. retrievable, facetable, and key

D. searchable, sortable, and retrievable

Answer: B

Explanation:

Question: 243

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

You are building an app that will translate text and documents by using Translator1.

You need to create the REST API request for the app.

Which headers should you include in the request?

A. the subscription key and the client trace ID

B. the subscription key, the subscription region, and the content type

C. the resource ID and the content language

D. the access control request, the content type, and the content length

Answer: B

Explanation:

Question: 244

You are building an app that will use the Azure Video Indexer service.

You plan to train a language model to recognize industry-specific terms.

You need to upload a file that contains the industry-specific terms.

Which file format should you use?

A. PDF

B. XML

C. TXT

D. XLS

Answer: C

Explanation:

Question: 245

HOTSPOT

You are building an Azure web app named App1 that will translate text from English to Spanish.

You need to use the Text Translation REST API to perform the translation. The solution must ensure that you have data sovereignty in the United States.

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

NOTE: Each correct selection is worth one point.

Answer Area

https://	api-nam.cognitive.micro5ofttranslator.com api.cognitive.mkrosofttranslator.com apt-na rn.cognitrveservices.azure.com eastus.api.cognitive.mrcrosoft.com	translate detect languages test-to-speech	?ipi-versionsJ.04to*fs

Answer

Explanation:

Answer Area

https://	api-nam.cognitive.mtcrosofttransldtOF.com	/ translate	?api-version»3.0Sto-es

1. api-nam.cognitive.microsofttranslator.com
2. translate

https://learn.microsoft.com/en-us/azure/cognitive-services/Translator/reference/v3-0- reference#base-urls

Requests to Translator are, in most cases, handled by the datacenter that is closest to where the request originated. If there's a datacenter failure when using the global endpoint, the request may be routed outside of the geography.

To force the request to be handled within a specific geography, use the desired geographical endpoint. All requests are processed among the datacenters within the geography.

- United States

api-nam.cognitive.microsofttranslator.com

https://learn.microsoft.com/en-us/azure/cognitive-services/translator/reference/rest-api-guide

- translate

Translate specified source language text into the target language text.

Question: 246

HOTSPOT

You have an app named App1 that uses Azure Al Document Intelligence to analyze medical You have an app named App1 that uses Azure Al Document Intelligence for patients. You send a request to App1 and receive

the following response.

“status": "succeeded", "createdOateTime": “2023-O9-14T21:01:02Z", "lastUpdatedDateTime":
"2023-09-14T21:01:03Z“, "analyzeResult": {
"apiVersion": "2023-07-31",
"modelId": "prebuilt-healthInsuranceCard.us", "stringIndexType": "utf16CodeUnit“, "content":
"Blood Pressure 118/72", "pages": (
(
"words": (
{ "content": "Blood", "polygon": [...], “confidence": 0.766, “span": (...)
L { "content": "Pressure", "polygon": [...], “confidence": 0.716, "span": { ... }), {
"content": "118/72",
"polygon": [...], “confidence": 0.761, "span": (...)

b

"documents": (
{ "docType": "healthInsuranceCard.us" "boundingRegions": (...)

) b "fields": {), "confidence": 1, "spans": [_]
}

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

NOTE: Each correct selection is worth point.

Answer Area Statements	Ves	No
The chosen mode! is suitable for the intended use case.	O	Q
The text content was recognized with greater than 70 percent confidence.		
The form elements were recognized with greater than 70 percent confidence.	0	o

Answer:

Explanation:

Answer Area

Statements	Yes	No
The chosen model is suitable for the intended use case.	O	®
The text content was recognized with greater than 70 percent confidence.	®	
The form elements were recognized with greater than 70 percent confidence.	0	®

Question: 247

HOTSPOT

You plan to deploy an Azure OpenAI resource by using an Azure Resource Manager (ARM) template.

You need to ensure that the resource can respond to 600 requests per minute.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
{
  "type": "Hicrosoft.CcgnitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample■resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId("Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource")]" 1, 'sku": {
    "name": "Standard",
                    600
                    1
                    60
                    100
                    600
  },
```

```
"properties": {
  "model": {
    “format": "OpenAI",
```

Answer:

Explanation:

Question: 248

You have an Azure subscription. The subscription contains an Azure OpenAI resource that hosts a GPT-4 model named Modell and an app named Appl. App1 uses Model!

You need to ensure that App1 will NOT return answers that include hate speech.

What should you configure for Model1?

A. the Frequency penalty parameter

B. abuse monitoring

C. a content filter

D. the Temperature parameter

Answer: B

Explanation:

Question: 249

You have an Azure subscription. The subscription contains an Azure OpenAI resource that hosts a GPT-3.5 Turbo model named Model1.

You configure Model1 to use the following system message: "You are an Al assistant that helps people solve mathematical puzzles. Explain your answers as if the request is by a 4-year-old."

Which type of prompt engineering technique is this an example of?

A. few-shot learning

B. affordance

C. chain of thought

D. priming

Answer: D

Explanation:

Question: 250

HOTSPOT

You build a chatbot by using Azure OpenAI Studio.

You need to ensure that the responses are more deterministic and less creative.

Which two parameters should you configure? To answer, select the appropriate parameters in the answer area.

NOTE: Each correct answer is worth one point.

Answer Area

Answer

Explanation:

Answer Area

Question: 251

You are building a chatbot for a travel agent. The chatbot will use the Azure OpenAI GPT 3.5 model and will be used to make travel reservations.

You need to maximize the accuracy of the responses from the chatbot.

What should you do?

A. Configure the model to include data from the travel agent's database.

B. Set the Top P parameter for the model to 0.

C. Set the Temperature parameter for the model to 0.

D. Modify the system message used by the model to specify that the answers must be accurate.

Answer: A

Explanation:

Question: 252

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($"\t{entity.Text}");
    }
}
```

For the second argument, you call the function and specify the following string.

Our tour of Paris included a visit to the Eiffel Tower.

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

NOTE: Each correct selection is worth one point.

Answer Area Statements	Yes	No
The output will include the following words; our and included.	Q	
The output will include the following words: Paris. Eiffel, and Tower.	Q	
The function will output all the key phrases from the input string to the console.		

Answer:

Explanation:

Answer Area Statements	Yes	No
The output will include the following words: our and included.	®	O
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.		

Question: 253

HOTSPOT

You have an Azure subscription that contains an Azure Al Document Intelligence resource named DM.

You build an app named App1 that analyzes PDF files for handwritten content by using DM.

You need to ensure that App1 will recognize the handwritten content.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
Uri fileUri - new Urif'cf ileUri>");
AnalyzeOocumentOperation operation = await client.AnalyzeDocumentFromUriAsyncCWaitUntil.Completed, AnalyzeResult result =  . piebuilt~d<?CUnien.t..\- T fileUri);
operation.Value;
foreach (Documentstyle style in resuit.Styles)
bool isHandwritten - style.IsHandwritten.HasValue 88 style.IsHandwritten ~ true; if (isHandwritten &S
style.Confidence >
{
Console.WriteLine($"Handwritten content found:");
foreach (Documentspan span in style.Spans)
{
Console.WriteLine($" Content: {result.Content.Substring(span.Index, span.Length)}");
}
```

"prebuilt-document"
"prebuilt-contracf.
"prebuilbread",

0.75
0.1
0.75
1.0

Answer:

Explanation:

Answer Area

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

Question: 254

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

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

What are two ways to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

A. Fine-tune the model.

B. Provide grounding content.

C. Add sample request/response pairs.

D. Retrain the language model by using your own data.

E. Train a custom large language model (LLM).

Answer: B, C

Explanation:

Question: 255

HOTSPOT

You have an Azure subscription that contains an Azure OpenAI resource named AH.

You build a chatbot that will use AI1 to provide generative answers to specific questions.

You need to ensure that the responses are more creative and less deterministic.

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

NOTE: Each correct selection is worth one point.

Answer Area

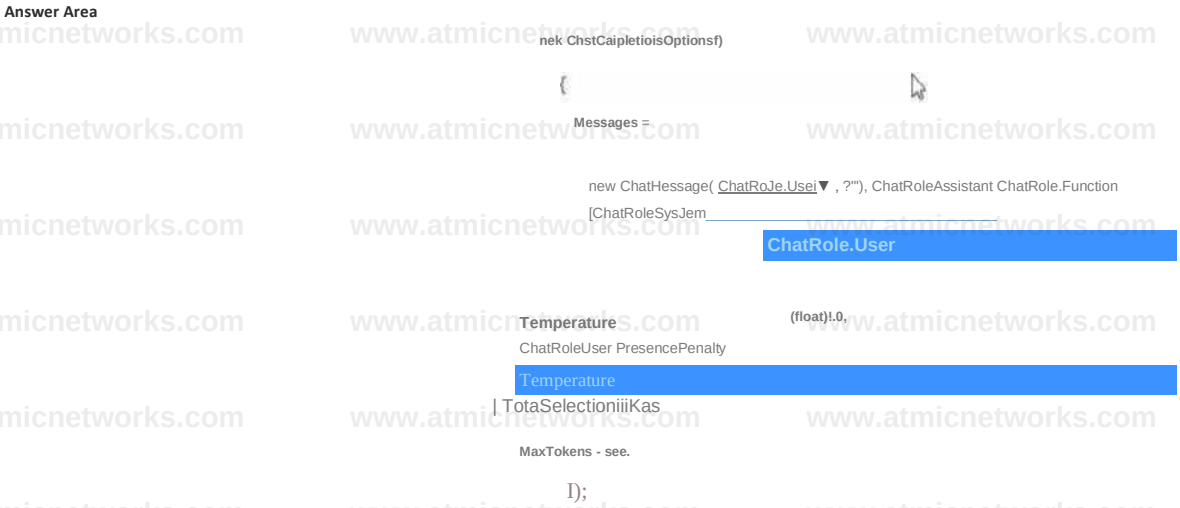

Answer:

Explanation:

Answer Area

```
new ChatCanpletionsOptions()

Messages =
{
    new ChatMessagef ChatfcileAJse^                    ᵀ . g""},
},
  Temperature          ' = (float)l.e,
MaxTokens - 808,
```

Question: 256

DRAG DROP

You have an Azure subscription that contains an Azure OpenA1 resource named AH.

You plan to build an app named App1 that will write press releases by using AM.

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

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

Actions

Create a deployment that uses the texL-enibeddmyada-OM model.

Apply the Default system message template

Create a deployment that uses the GPT-35 Turbo model

Apply the Marketing Writing Assistant system message template

Deploy the solution to a new web app.

Answer Area

Answer:

Explanation:

Actions

Create a deployment that uses the text-embedding-ada-002 model.

Apply the Default system message template.

Answer Area

1 Create a deployment that uses the GPT-35 Turbo model.

12 Apply the Marketing Writing Assistant system message template.

the solution to a new web app.

Question: 257

HOTSPOT

A1 You have an Azure OpenA1 resource named AH that hosts three deployments of the GPT 3.5 model. Each deployment is optimized for a unique workload.

You plan to deploy three apps. Each app will access AM by using the REST API and will use the deployment that was optimized for the apps intended workload.

You need to provide each app with access to AH and the appropriate deployment. The solution must ensure that only the apps can access AM.

What should you use to provide access to AM. and what should each app use to connect to its appropriate deployment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Answer:

Explanation:

Answer Area

Provide access to AI1 by using: An API key

Connect to the deployment by using: A deployment endpoint

Question: 258

HOTSPOT

You are building a chatbot.

You need to use the Content Moderator API to identify aggressive and sexually explicit language.

Which three settings should you configure? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

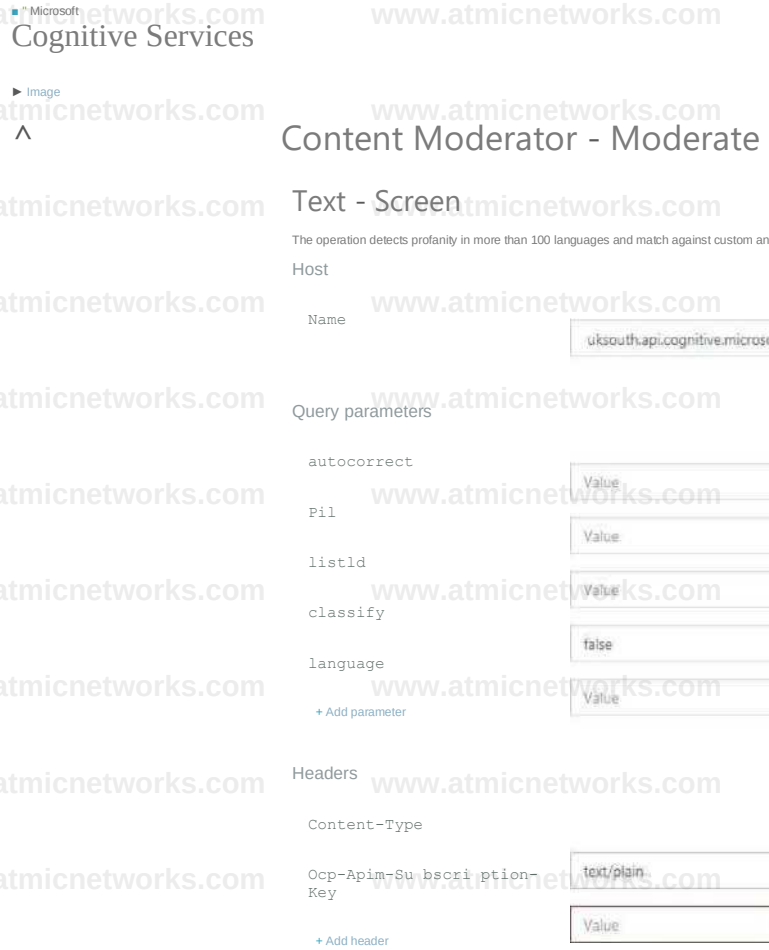

Explanation:

Answer:

Answer Area

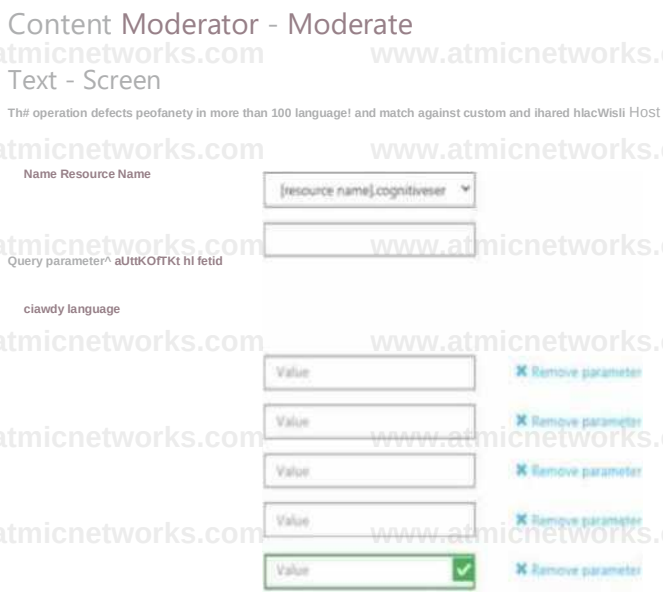

Question: 259

DRAG DROP

You have a monitoring solution that uses the Azure Al Anomaly Detector service.

You provision a server named Server! that has intermittent internet access.

You need to deploy the Azure Al Anomaly Detector to Server 1.

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.

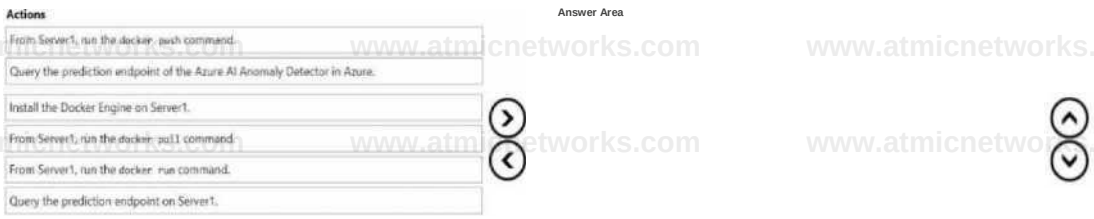

Answer:

Explanation:

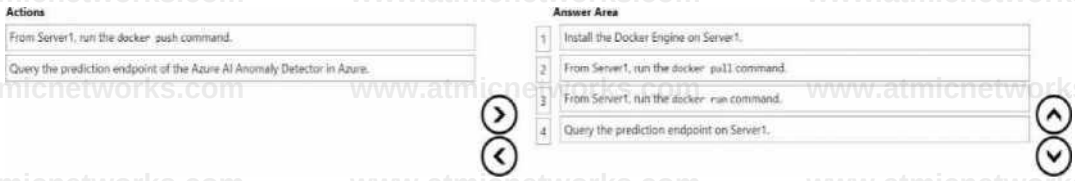

Question: 260

You have an Azure Cognitive Search solution and a collection of handwritten letters stored as JPEG files.

You plan to index the collection. The solution must ensure that queries can be performed on the contents of the letters.

You need to create an indexer that has a skillset.

Which skill should you include?

A. key phrase extraction

B. optical character recognition (OCR)

C. document extraction

D. image analysis

Answer: B

Explanation:

Question: 261

HOTSPOT

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

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

You need to create a Standard tier resource that will convert scanned receipts into text.

How should you call the method? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Computervision
CustomVision.Prediction
CustomVision,Training

"edstus", "STT
useast', "SV)
'SO", "eastus")
'SO", "useast")

Answer

Explanation:

Amwer Area

Question: 262

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

Answer: D

Question: 263

HOTSPOT
You have an Azure subscription that contains an Azure OpenA1 resource. You configure a model that has the following settings:

- Temperature: 1
- Top probabilities: 0.5
- Max response tokens: 100

You ask the model a question and receive the following response.

```
"choices": (
    "finish/ewofl": "stop",
    "index": 0,
    "message"; '
        "content": "The founders of Microsoft are Bill Gates and Paul Allen. They co-founded the company in 1975 "role": "assistant"
```

b

```
"created": 1679914554,
"id": "chatcapl-6u5fn2yyjki»ESe364jaQR6bDSc01M,
"model": "gp--3.5-turbo-0301',
"object": "chat.completion**,
 "usage": {
    "completion^tokens"; 86,
    "proapt-totaens": 37,
    "total-tokens": 123
```

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

Answer Area

Statements	Yes	No
The subscription will be charged 86 tokens for the execution of the session.		
The text completion was truncated because the Max response tokens value was exceeded.	O	
The prompt-tokens value will be included in the calculation of the Max response tokens value.	O	0

Answer:

Answer Area Statements	Yes	No
The subscription will be charged 86 tokens for the execution of the session.	9	
The text completion was truncated because the Max response tokens value was exceeded.	O	®
The prompt-tokens value will be included in the calculation of the Max response tokens value.	O	•

Question: 264

HOTSPOT

You have an Azure subscription that contains an Azure OpenA1 resource named All.

You plan to develop a console app that will answer user questions.

You need to call All and output the results to the console.

How should you complete the code? To answer, select the appropriate options in the answer area NOTE: Each correct selection is worth one point.

Answer Area

```
OcenAIClient client -

    new OptnAKlientfnew Uri(endpoint), new AzureKeyCredentUl(key));

Response<Co*nplerions> response •

    client. I GetCompletions                    ˅ I C^^ploy^ntMane, "What is Microsoft Azure?");
                Ge (Completions
            i GetEmbeddings
            | GetimageGenerations

    Console,writeline      [ (response. Value.ChoicesIOl,Text);
                           (response.Value.ChorcesIOl.Text);
                           (response.Value Id);
                           (response.Value PromptFilterResults)
```

Answer:

Answer Area

```
OpenAIClient client •

        new OpenAICMert(ne>* Uri(endpoint), new AzureKeyCredentIel(key));
```

```
Response<Co«pletions> response *
    client. GetComptetions                    ^ (deploymentNaae, "what is Microsoft Azure?");
Console, write Line 'response Value.Cho<ces|0].Text); ^
```

Question: 265

HOTSPOT

You have an Azure subscription.

You need to deploy an Azure Al Document Intelligence resource.

How should you complete the Azure Resource Manager (ARM) template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

"Sscheu": "https://schead.mandgeoent.a2ure.coa/scheus/2019-04-01/deploynentTefflplate.json#',
'contentversion": "1.0.0.0",
"parameters": {},
"variables": {},
"resources": (

"type": [Miaosoft.CognitiveServices "j /accounts",
' MicrosoftCognittveSearch |
Microsoft Cognitive Ser vices
| Microsoft.MachineLearning
Microsoft.MachineLearningSeivices
'•plVersiM'i '2023-05-01-,
"name": "Oocu«ntintelligenceOero",
"location': "westeurope",

-F0-

FormRecognizer" T
AiBuilder', CognitiveSeardT
Form Recognizer".
| OpenAT, |

Answer:

Answer Area

'5s enema": "https.'//schema,management.azure.com/schenas/2019-04-01/deploynentTemplate* jsane",
"contentversion": "1.0.0.0",
"parameters": {},
"variables": {},
"resources": [

■type': Microsoft,CognitiveServices ' /accounts",
"apiVorsion": '2023-05-01",
"name"; "DocumentIntelllgenceDemo",
"location"; "westeurope",
"sku": {
"name": "F0"

"kind": FormRecognizer", '

J
J

Question: 266

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 have a PDF named Doc1.pdf that contains a product catalogue and a price list

You upload Doc1.pdf and train the model.

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

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

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

Solution: From Language Studio, you add alternative phrasing to the question and answer pair, and then retrain and republish the model.

Does this meet the goal?

A. Yes

B. No

Answer: A

Question: 267

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 <product>?

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

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

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

Does this meet the goal?

A. Yes

B. No

Answer: B

Question: 268

You ace building an internet-based training solution. The solution requires that a user's camera and microphone remain enabled.

You need to monitor a video stream ol the user and detect when the user asks an instructor a question.

The solution must minimize development effort.

What should you include in the solution?

A. object detection in Azure AI Custom Vision

B. the Face service in Azure AI Vision

C. language detection in Azure AI Language Service

D. speech-to-text in the Azure AI Speech service

Answer: D

Explanation:

Question: 269

HOTSPOT

You are building an app that will answer customer calls about the status of an order. The app will query a database for the order details and provide the customers with a spoken response.

You need to identify which Azure AI service APIs to use. The solution must minimize development effort.

Which object should you use for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Convert customer calk into text queries Translation Rec oqnizer

SpeechRecognizer

VoKeProfileOmtit

Provide customers with the order details | SpeechSynthesaer :h ize<

TranslationRecognize i

Answer:

Explanation:

Answer Area

Convert customer calls into tert queries _ TranslationRecognizer

Provide customers with the order details Speech$ynthes<zer

Question: 270

HOTSPOT

You are building an app that will automatically translate speech from English to French, German, and Spanish by using Azure Al service.

You need to define the output languages and configure the Azure Al Speech service.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
static async Task Translat*$p»«chAsync()
{
var config . Speech!ranslationConflf.froaSufr*crlpti*i(SMKH_5UKatlPnON_«Y, $PeECH_SUVICe_l«e»);
var languages ■ new LMt<string> Ctf/dtVeS'I                                    * 1
```

Cen GB'|

["en","fr","de","es"]
["fr","de","es"]

('French'. "Ge man". "Spantsh'j

```
1anguages.ForEach(config.AddTargetlanguage),
usingvar recognizer * new
```

Tr*Bl«tionRKoaniHf
Intent Recognizer
SpeakerRecogmzer
SpeechSyntheuzer

```
}
```

Answer:

Explanation:

Answer Area

```
static async Task TranslateSpeechAsyncO

ver config = ^prechTranUatlonConf ig .Fra«SutoscriptIDn(S*EKH_SU8SC#IPTI0W_cE'r, $HKH_SHIV1CE.JIEAX») van languages « new LIst<strIng> nr'de'/es'I T ,

1anguage s.Fort acH conf ig.AddTarget Language);

us Ingvar recognizer * new TransbtronRecoqmzof                    " ;
```

Question: 271

You have a Microsoft OneDrive folder that contains a 20-GB video file named FileVavi. You need to index File1.avi by using the Azure Video Indexer website. What should you do?

A. Upload File1.avi to the www.youtube.com webpage. and then copy the URL of the video to the Azure Al Video Indexer website

B. From OneDrive. create a download link, and then copy the link to the Azure Al Video Indexer website.

C. From OneDrive, create a sharing link for File1.avi and then copy the link to the Azure Al Video Indexer website.

D. Download File1 avi to a local computer, and then upload the file to the Azure Al Video Indexer website.

Answer: D

Explanation:

Question: 272

You ace developing an app that will use the Speech and language APIs.

You need to provision resources for the app. The solution must ensure that each service is accessed by using a single endpoint and credential

Which type of resource should you create?

A. Azure Al Content Safety

B. Azure Al service

C. Azure Al Speech

D. Azure AI Language

Answer: B

Explanation:

Question: 273

DRAG DROP

You have an app that uses Azure AI and a custom trained classifier to identity products in images. You need to add new products to the classifier. The solution must meet the following requirements:

- Minimize how long it takes to add the products
- Minimize development effort.

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.

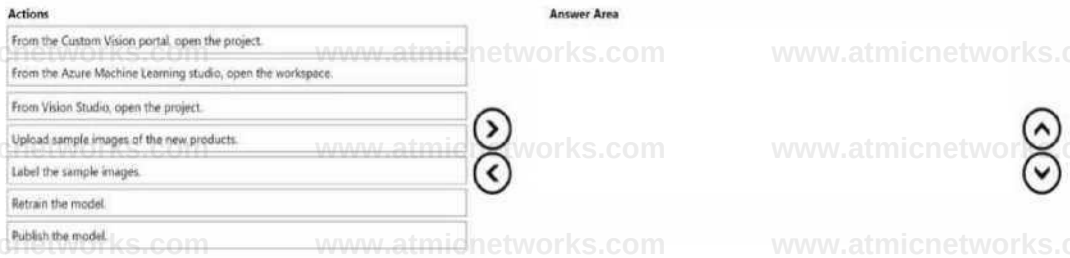

Answer:

Explanation:

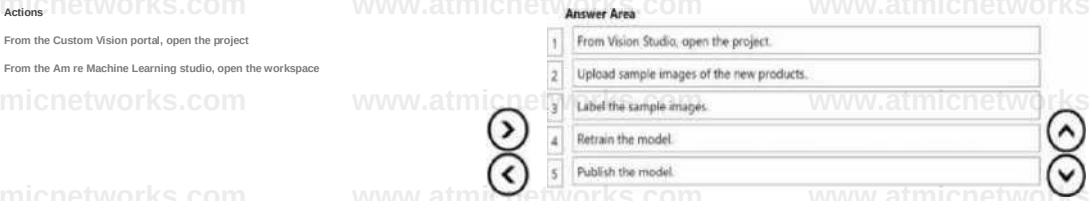

Question: 274

You have an Azure subscription that contains an Azure Al service resource named CSAccount1 and a virtual network named VNet1 CSAaccount1 is connected to VNet1

You need to ensure that only specific resources can access CSAccount1. The solution must meet the following requirements:

- Prevent external access to CSAccount1
- Minimize administrative effort

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point.

A. In VNet1, modify the virtual network settings.

B. In VNet1. enable a service endpoint for CSAccount1

C. In CSAccount1, configure the Access control (1AM) settings.

D. In VNet1, create a virtual subnet.

E. In CSAccount1, modify the virtual network settings.

Answer: B, D

Explanation:

Question: 275

HOTSPOT

You have an Azure subscription

You need to create a new resource that will generate fictional stories in response to user prompts. The solution must ensure that the resource uses a customer-managed key to protect data.

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

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Question: 276

HOTSPOT

You have an Azure subscription that contains an Azure Al Content Safety resource named CS1. You need to use the SDK to call CS1 to identify requests that contain harmful content. How should you complete the code? lo answer, select the appropriate options m the answer area.

NOTE: Each correct selection is worth one point

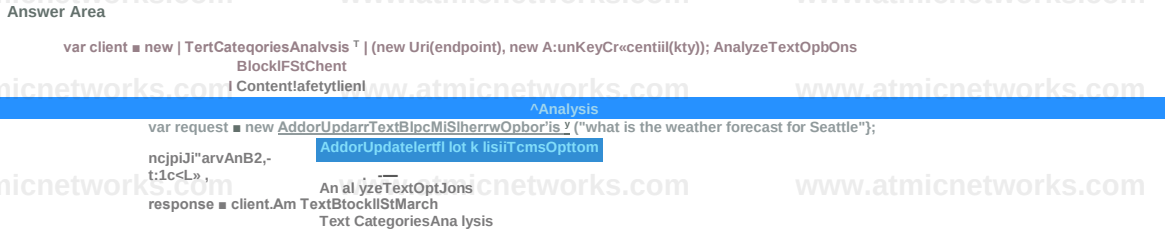

Answer:

Explanation:

Answer Area

```
var client * naw TextCateqonesAnalysii ˅ (naw Uri(endpoint), naw A:urexeyCredenti*l(key));
var request * new AddorUpdateTextBlocklistItermOptions * ("what is the weather forecast for Seattle");
Response*AnalyzeTentResuit> response;
response ■ client.AnalyzeText(request);
```

Question: 277

HOTSPOT

You have an Azure subscription that contains an Azure AI Content Safety resource named CS1.

You need to call CS1 to identify whether a user request contains hateful language.

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

NOTE: Each correct selection is worth one point

Answer Area

```
curl "location "request POST ' 'https://csl.cognitiveservices.asure.com/
--header ocp-Apim-subtcription-fey •your_subscription_xey>' \
--header Content-Type: application/jspn' \
—data-raw '{
'text". "What is the weather forecast for Seattle",
"categories': ("Hate"J
    "blcckliitMames": [
        "String'
    la
```

language/
completions/
contentsafety/
healthinaghts/
language/

1 embeddings^ *] ?api-verslwi«it2J-lB-er completions

embeddings
text analyze
text/block lists

Answer:

Explanation:

Answer Area

```
curl -location "request POST ' 'https://«I.co(nitiv«ervice».asure.com/ language/ ❝ embeddings ᵀ ?api'ver»ion«2O23-l®-01' \
--header Ocp-Apsm-Substriotion-Key <your„tubscription_key>' i
■■header Content-Type: application/json' \
--date-raw '{
    "tort": "What is the weather forecast for Seattle",
    "categories'. ["Hate"] "blocklistNames": { "string"

    "haltOnBiockiistHit" true,
    "outputType"■ 'FourSeverityLevels"
```

Question: 278

You have an Azure subscription that contains an Azure AI Document Intelligence resource named DM. DI1 uses the Standard SO pricing tier

You have the files shown in the following table.

Name	Size	Description
File1.pdf	800 MB	Contains scanned images
File2.jpg	1 KB	An image that has 25 x 25 pixels
File3.tiff	5 MB	An image that has 5000 x 5000 pixels

Which files can you analyze by using DI1?

A. File1.pdf only

B. File2.jpg only

C. File3 tiff only

D. File2.jpg and File3.tiff only

E. File1.pdf,File2,jpg, and File3,tiff

Answer: D

Explanation:

Question: 279

You have an Azure OpenAI model named All.

You are building a web app named App1 by using the Azure OpenAI SDK You need to configure App1 to connect to All What information must you provide?

A. the endpoint, key, and model name

B. the deployment name, endpoint. and key

C. the endpoint, key, and model type

D. the deployment name, key, and model name

Answer: B

Explanation:

Question: 280

You have an Azure subscription that contains an Azure Al Content Safety resource named CS1. You plan to build an app that will analyze user-gene rated documents and identify obscure offensive terms. You need to create a dictionary that will contain the offensive terms. The solution must minimize development effort. What should you use?

A. a text classifier

B. text moderation

C. language detection

D. a blacklist

Answer: D

Explanation:

Question: 281

You have an Azure subscription that contains an Azure OpenAI resource named All and a user named User1 You need to ensure that User1 can perform the following actions in Azure OpenAI Studio;

- Identify resource endpoints.
- View models that are available for deployment.
- Generate text and images by using the deployed models

The solution must follow the principle of least privilege Which role should you assign to User1?

A. Cognitive Services OpenAI User

B. Cognitive Services Contributor

C. Contributor

D. Cognitive Services OpenAI Contributor

Answer: A

Explanation:

Question: 282

You have an Azure subscription that contains an Azure OpenAI resource named AH and an Azure AI Content Safety resource named CS1.

You build a chatbot that uses All to provide generative answers to specific questions and CS1 to check input and output for objectionable content.

You need to optimize the content filter configurations by running tests on sample questions.

Solution: From Content Safety Studio, you use the Protected material detection feature to run the tests.

Does this meet the requirement?

A. Yes

B. No

Answer: B

Explanation:

Question: 283

You have an Azure subscription that contains an Azure OpenAI resource named All and an Azure AI Content Safety resource named CS1.

You build a chatbot that uses All to provide generative answers to specific questions and CS1 to check input and output for objectionable content.

You need to optimize the content filter configurations by running tests on sample questions.

Solution: From Content Safety Studio, you use the Safety metaprompt feature to run the tests

Does this meet the requirement?

A. Yes

B. No

Answer: B

Explanation:

Question: 284

You have an Azure subscription that contains an Azure OpenAI resource named All and an Azure AI Content Safety resource named CS1.

You build a chatbot that uses All to provide generative answers to specific questions and CSI to check input and output for objectionable content.

You need to optimize the content filter configurations by running tests on sample questions.

Solution: From Content Safety Studio, you use the Monitor online activity feature to run the tests

Does this meet the requirement?

A. Yes

B. No

Answer: B

Explanation:

Question: 285

You have an Azure subscription that contains an Azure Al Document Intelligence resource named Aldoc1 in the SO tier.

You have the lies shown in the following table.

Name	Format	Password-leaked	Size (MB)
Friel	JPG	No	400
File2	PDF	No	250
File3	PNG	Yes	180
File4	XLSX	No	900
F^'	PDF	Yes	160

You need to train a custom extraction model by using AIdoc1.

Which files can you upload to Document Intelligence Studio?

A. File1, and File2 only

B. File2, File4, and File5 only

C. File1, File2, and File4 only

D. File1, and File5 only

E. File1, File2. File3, File4, and File5

Answer: C

Explanation:

Question: 286

You have a custom Azure OpenAI model.

You have the files shown in the following table.

Name	Size
File1.tsv	80 MB
File2.xml	25 MB
File3.pdf	50 MB
File4.xlsx	200 MB

You need to prepare training data for the model by using the OpenAI CLI data preparation tool. Which files can you upload to the tool?

A. Filel.tsvonly

B. File2.xml only

C. File3.pdf only

D. File4.xlsx only

E. File1.tsv and File4.xslx only

F. File1.tsv.File2.xml and File4.xslx only

G. File1.tsv, File2.xml, Fil3.pdf and File4.xslx

Answer: A

Explanation:

Question: 287

You have an Azure subscription and 10,000 ASCII files.

You need to identify files that contain specific phrases. The solution must use cosine similarity.

Which Azure OpenAI model should you use?

A. text-embedding-ada-002

B. GPT-4

C. GPT-35 Turbo

D. GPT-4-32k

Answer: A

Explanation:

Question: 288

HOTSPOT

You have an Azure subscription that contains an Azure Al Content Safety resource.

You are building a social media app that will enable users to share images.

You need to configure the app to moderate inappropriate content uploaded by the users.

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

NOTE: Each correct selection is worth one point.

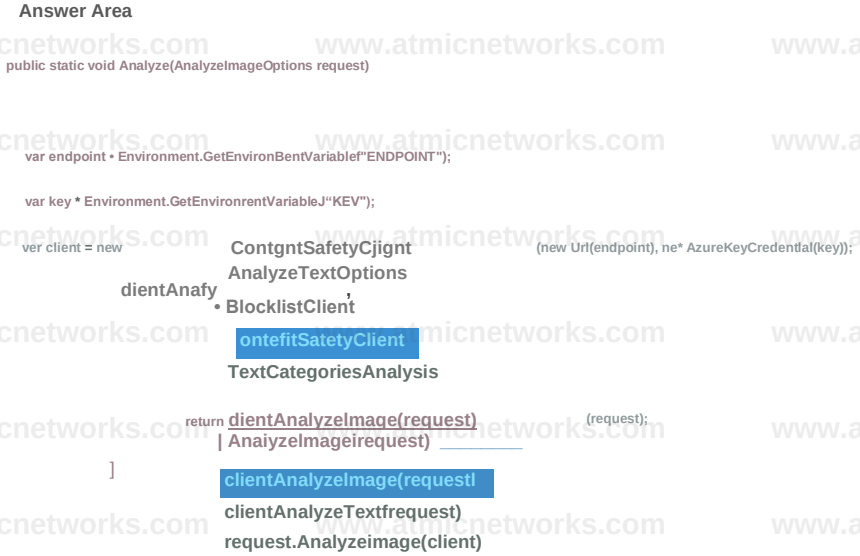

Answer:

Explanation:

Answer Area

```
public static void Analyxe(Analyzel«ageOptions request)
{
    var endpoint - Environment.GetEnvironnentVariable("ENDPOINT");
    var key = Environment.GetEnvironnentVariable "»'.Ev" ;
    ver client » new ContentSafetyCBent ^ (new Url(endpoInt), new AzureKeyCredentIal(key));
    return      dient.Analyzelmage. request)          ' (request) ;
```

Question: 289

HOTSPOT

You have a chatbot that uses Azure OpenAI to generate responses.

You need to upload company data by using Chat playground. The solution must ensure that the chatbot uses the data to answer user questions.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
var options = new
{
    Messages =
    [
        new ChatMessage(ChatRole.User, "What are the differences between Azure Machine Learning and Azure AI services?"),
    ],
    Extensions =
    [
        new
        {
            Search
            SearchKey = new AzureKeyCredential(searchKey),
            IndexName = searchIndex,
```

ChatCompletionsOptions()
CompletionsOptions()
StreamingChatCompletions()

AzureChatExtensionConfiguration
AzureChatExtensionsOptions
AzureCognitiveSearchChatExtensionConfiguration

Answer:

Explanation:

Question: 290

You have an Azure Subscription that contains an Azure OpenAI resource named AI1 and a user named User1.

You need to ensure that User1 can add custom data sources to AI1. The solution must follow the principle of least privilege.

Which role should you assign to User1?

A. Search Service Contributor
B. Cognitive Services OpenAI Contributor
C. Cognitive Services Contributor
D. Search index Data Contributor

Answer: C

Explanation:

Question: 291

You have an Azure DevOps pipeline named Pipeline1 that is used to deploy an app. Pipeline1 includes a step that will create an Azure AI services account.

You need to add a step to Pipeline1 that will identify the created Azure AI services account. The solution must minimize development effort.

Which Azure Command-Line interface (CLI) command should you run?

A. Az resource link

B. Az account list

C. Az cognitivesservices account network-rule

D. As cognitiveservices account show

Answer: D

Explanation:

Question: 292

You are building an app that will use the Azure AI Speech service.

You need to ensure that the app can authenticate to the service by using a Microsoft Entra ID token.

Which two action should you perform? Each answer part of the solution.

NOTE: Each correct selection is worth one point.

A. Create a Conditional Access

B. Create a private endpoint

C. Request an X.509 certificate

D. Certificate a custom subdomain.

E. Enable a virtual network service endpoint.

Answer: B, C

Explanation:

Question: 293

You are designing a content management system.

You need to ensure that the reading experience is optimized for users who have reduced comprehensive

and learning differences, such as dyslexia.

Which Azure service should you include in the solution?

A. Azure AI Translator

B. Azure AI Document Intelligence

C. Azure AI Immersive Reader

D. Azure AI Language

Answer: C

Explanation:

Question: 294

You are building an image sharing app that will use Azure AI to prevent users from sharing sexually explicit images.

You need to ensure that inappropriate images are identified correctly. The solution must minimize development effort.

What should you use?

A. Visual Studio

B. Vision Studio in Azure AI Vision

C. Azure AI Content Safety Studio

D. Azure AI Studio

Answer: B

Explanation:

Question: 295

You are building a solution in Azure that will use Azure AI Language service to process sensitive customer data.

You need to ensure that only specific Azure processes can access the Language service. The solution must minimize administrative.

What should you include in the solution?

A. Azure Application Gateway

B. IPsec rules

C. A virtual network gateway

D. Virtual network rules

Answer: D

Explanation:

Question: 296

HOTSPOT

You plan to provision Azure AI service resources by using the following method.

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

Anwx Aim

provision_resource("resi"

ComputerVison
CustomVuion Prediction
CurtorVuion Tianin;
EESEMBB

<•85^5' "Sn
use«t 'ST'
SO' "useasf i

Answer:

Explanation:

Answer Area

provision_resource("resi", FormRecognizer, "S0", "eastus")

Question: 297

HOTSPOT

You have an app that uses the AI Language custom question answering service.

You need to ad alternatives for the word testing by using the Authoring API.

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

NOTE: Each correct selection is worth one point.

Answer Area

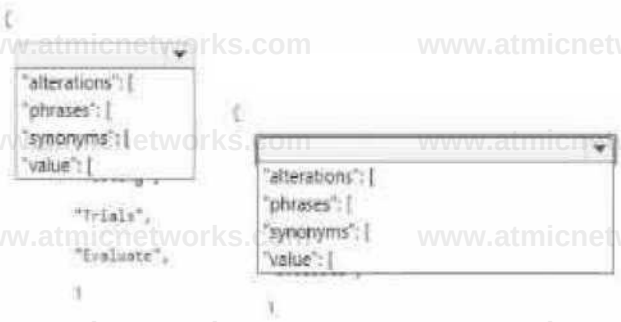

Answer:

Explanation:

Answer Area

Question: 298

You have the following C# function.

```
static void MyFunction(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.ExtractKeyPhrases(text);
    Console.WriteLine("Key phrases:");
    foreach (string keyphrase in response.Value)
    {
        Console.WriteLine($"{keyphrase}");
    }
}
```

You call the function by using the following code.

```
MyFunction(textAnalyticsClient, "the quick brown fox jumps over the lazy dog");
```

Following 'key phrases' what output will you receive?

A. Jumps over the

B. The quick brown fox jumps over the lazy dog

C. Quick brown fox lazy dog

D. The quick

Answer: C

Explanation:

Question: 299

DRAG DROP

You have a web app that uses Azure AI search.

When reviewing activity, you see greater than expected search query volumes. You suspect that the query key is compromised.

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

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

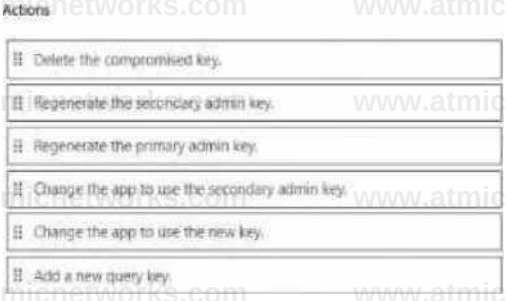

Answer:

Explanation:

Comprehensive Detailed Explanation

You are dealing with Azure AI Search (formerly Azure Cognitive Search). The issue is that the query key is

compromised, and you need to ensure minimal downtime while keeping users with read-only access.

Azure AI Search uses two types of keys:

- Admin keys (full control – manage indexes, data sources, etc.)
- Query keys (read-only access for client applications)

If a query key is compromised, the remediation must be safe and fast, without breaking the application unnecessarily.

Step 1 – Add a new query key

Instead of immediately deleting the compromised key, first generate a new query key in the Azure portal or via API. This ensures you have a replacement ready before cutting off the old key.

Step 2 – Change the app to use the new key

Update your web app configuration (for example, connection strings or environment variables) so that it authenticates against Azure AI Search using the newly created query key. This step ensures a smooth transition with minimal downtime.

Step 3 – Delete the compromised key

Finally, once your app is verified to work with the new query key, remove the old compromised key to prevent any further unauthorized access.

Correct Order:

1. Add a new query key.
2. Change the app to use the new key.
3. Delete the compromised key.

Microsoft Reference

- Manage admin and query keys in Azure AI Search
- Best practices for key management

Question: 300

HOTSPOT

You have an Azure subscription that contains an Azure AI Document intelligence resource named D1.

You create a PDF document named test.pdf that contain tabular data.

You need to analyze Test.pdf by using DI1.

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

NOTE: Each correct selection is worth one point.

Answer

Explanation:

Step 1: The Authorization Header

When calling Azure AI Document Intelligence REST API, the correct header to authenticate is:

-H "Ocp-Apim-Subscription-Key: {yourKey}"

- Key1 and Secret are not valid header names.
- Subscription-Key is also not the standard header.
- Correct choice: Ocp-Apim-Subscription-Key

Step 2: Selecting the Prebuilt Model

- Since the file contains tabular data (tables, structured document), the right prebuilt model is:

o prebuilt-document → extracts text, tables, structure, key-value pairs.

- prebuilt-contract → specialized for contracts.
- prebuilt-layout → extracts layout, lines, and words but not tables in structured form.
- prebuilt-read → simple OCR only.

So the correct choice is: prebuilt-document

Final Completed Command

curl -i -X POST "{endpoint}/formrecognizer/documentModels/prebuilt-document:analyze?api-version=2023-07-31" \

-H "Content-Type: application/json" \

-H "Ocp-Apim-Subscription-Key: {yourKey}" \

-d "{ 'urlSource': 'https://<your-storage>/test.pdf' }"

Verified Answer:

- Left box: Ocp-Apim-Subscription-Key
- Right box: prebuilt-document

Microsoft Reference:

- Azure AI Document Intelligence REST API
- Prebuilt Document model

Question: 301

You have an Azure subscription that contain an Azure OpenAI resource named AI1.

You build a chatbot that uses AI1 to provide generation answers to specific questions.

You need to ensure that the chatbot checks all input output for objectionable content.

Which types of resource should you create first?

A. Azure Machine Learning

B. Log Analytics

C. Azure AI Content Safety

D. Microsoft Defender Threat intelligence (Defender TI)

Answer: C

Explanation:

Question: 302

HOTSPOT

You are developing an app that will use the Azure AI vision API to analyze an image.

You need configure the request that will be used by the app to identity whether an image is clipart or a line drawing.

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

NOTE: Each correct select is worth one point.

Answer:

Explanation:

- HTTP Method: POST
- Visual Feature: imageType

Comprehensive Detailed Explanation

The question requires configuring an Azure AI Vision API request that determines if an image is clipart or a line drawing.

Step 1 – Determine the HTTP Method

The Analyze Image endpoint (/analyze) of Azure AI Vision expects an HTTP POST request.

- GET is not valid for analysis with an image body.

- PATCH is irrelevant here.
- POST is correct because you are submitting image content or a URL for analysis.

Reference: Azure AI Vision Analyze Image API

Step 2 – Select the correct Visual Feature

The visualFeatures query parameter determines what the API extracts.

- description → Generates natural language description of the image.
- tags → Identifies tags related to the image content.
- objects → Detects objects and their locations.
- imageType → Determines whether the image is a clipart, line drawing, or photo.

Since the requirement is to identify clipart vs line drawing, the correct choice is imageType.

Correct Request Completion:

- Method: POST
- Visual Feature: imageType

Microsoft Reference

- Analyze Image (REST API) – Visual Features
- Computer Vision API v3.2 Reference

Question: 303

DRAG DROP

You have an Azure subscription that contains an Azure OpenAI resource named AH.

You need to analyze an image to obtain a text descnption.

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

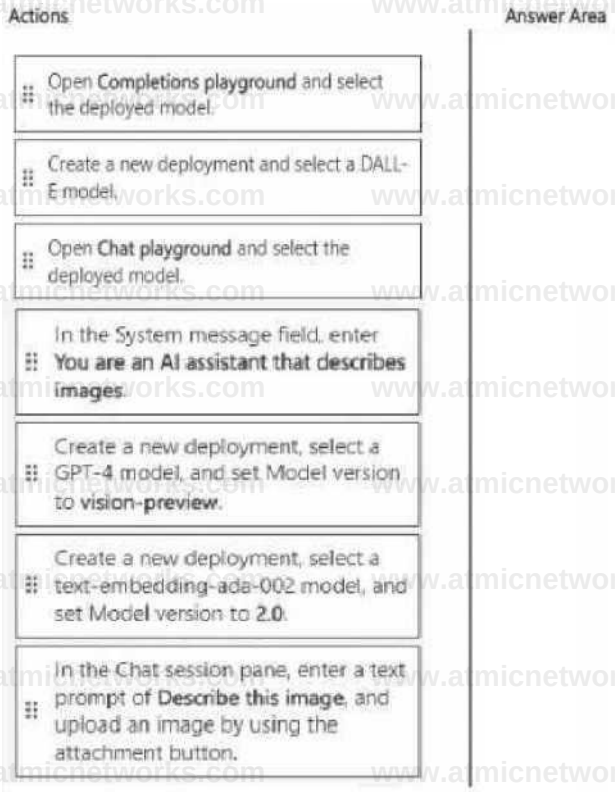

Answer

Explanation:

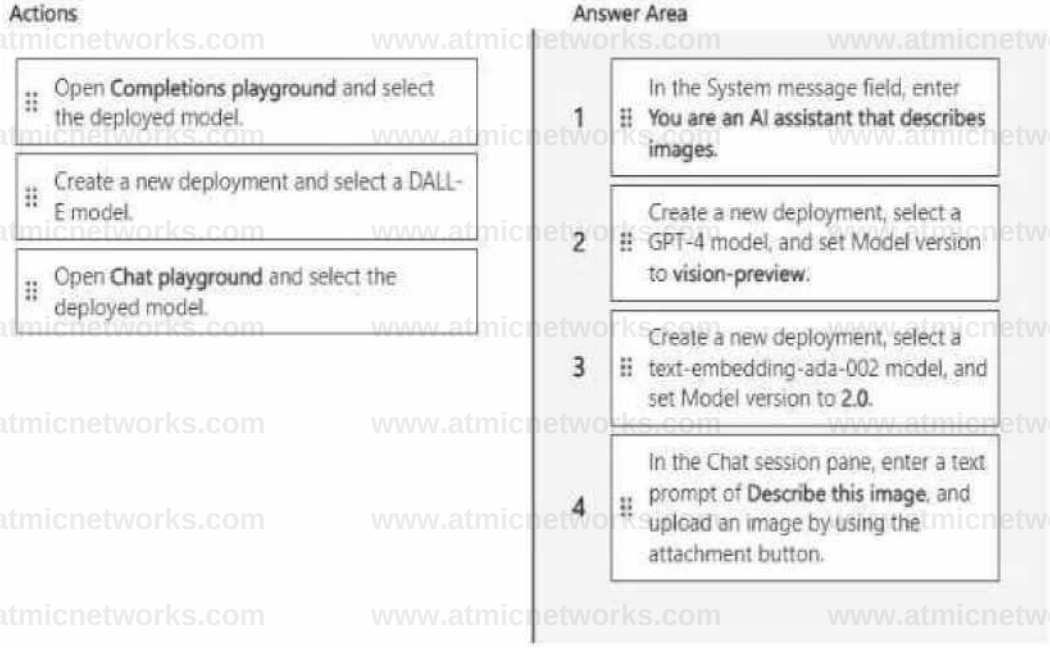

Key Knowledge:

- To analyze images with Azure OpenAI, you must use a GPT-4 model with vision support (model version vision-preview).

- This requires deploying a GPT-4 model for vision.
- Then you use the Chat playground (not completions, not DALL-E) since you want analysis, not generation.
- You need to give the assistant instructions in the system message field (e.g., “You are an AI assistant that describes images”).
- Finally, you upload an image and prompt the model to describe it.

Correct Sequence:

1. Create a new deployment, select a GPT-4 model, and set Model version to vision-preview.
   - Because GPT-4 with vision is required.
2. Open Chat playground and select the deployed model.
   - Vision capabilities are accessed through the Chat playground.
3. In the System message field, enter: “You are an AI assistant that describes images.”
   - This sets the assistant’s role.
4. In the Chat session pane, enter a text prompt of “Describe this image,” and upload an image by using the attachment button.
   - This triggers the model to analyze and describe the image.

Verified Answer (Correct Order):

1. Create a new deployment, select a GPT-4 model, and set Model version to vision-preview.
2. Open Chat playground and select the deployed model.
3. In the System message field, enter: You are an AI assistant that describes images.
4. In the Chat session pane, enter: Describe this image, and upload an image by using the attachment button.

Microsoft Reference:

Azure OpenAI vision capabilities

Azure OpenAI Studio overview

Question: 304

You have a local folder that contains the files shown in the following table.

Name	Format	Length (mins)	Size (MB)
Riel	WMV	34	400
Riel	AVI	90	1.200
FileS	MOV	300	980
Rle4	MP4	80	1.800

You need to analyze the files by using Azure Ai Video Indexer. Which files can you upload to the Video Indexer website?

A. Filel.FileZ and File4 only

B. File1, and File2 only

C. File1, File2, and File3 only

D. File1, File2. File3 and Fi1e4

E. File1, and File3 only

Answer: E

Explanation:

You have files with the following formats:

File1: WMV (34 mins, 400 MB)

File2: AVI (90 mins, 1,200 MB)

File3: MOV (300 mins, 980 MB)

File4: MP4 (80 mins, 1,800 MB)

You must determine which files can be uploaded to Azure AI Video Indexer.

Key Rules (Video Indexer upload limitations):

Supported formats: MP4, MOV, WMV, AVI, M2TS.

Maximum file length: 4 hours (240 minutes).

Maximum file size: 2 GB (2048 MB) for website uploads.

Question: 305

You have an app named App1 that uses a custom Azure AI Document Intelligence model to recognize contract documents. You need to ensure that the model supports an additional contract format. The solution must minimize development effort. What should you do?

A. Lower the confidence score threshold of App1.

B. Lower the accuracy threshold of App1.

C. Add the additional contract format to the existing training set. Retrain the model.

D. Create a new training set and add the additional contract format to the new training set.

E. Create and train a new custom model.

Answer: C

Explanation:

Analysis of Options:

- A. Lower the confidence score threshold of App1
  - Would just accept more low-confidence predictions, but won't add support for a new format. X
- B. Lower the accuracy threshold of App1
  - Similar reasoning: affects prediction acceptance, not model capabilities. X
- C. Add the additional contract format to the existing training set and retrain the model
  - Correct approach. You extend the training data to include the new format, retrain, and reuse the same model. Q
- D. Create a new training set and add the additional contract format
  - This means starting from scratch, which requires more effort. Not minimal. X
- E. Create and train a new custom model
  - Same as D, higher effort than necessary. X

Correct Answer:

C. Add the additional contract format to the existing training set. Retrain the model.

Microsoft Reference:

- Azure AI Document Intelligence – Custom models

Question: 306

HOTSPOT

You have an Azure subscription.

You plan to build a solution That will analyze scanned documents and export relevant fields to a database.

You need to recommend which Azure AI service to deploy for the following types of documents:

- Internal expenditure request authorization forms
- Supplier invoices

The solution must minimize development effort.

What should you recommend for each document type? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point

Answer Area

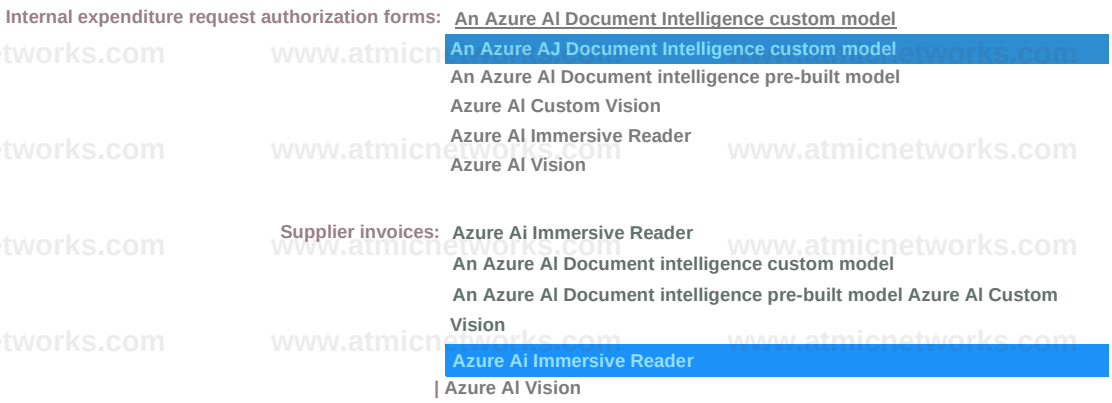

Answer

Explanation:

Answer Area

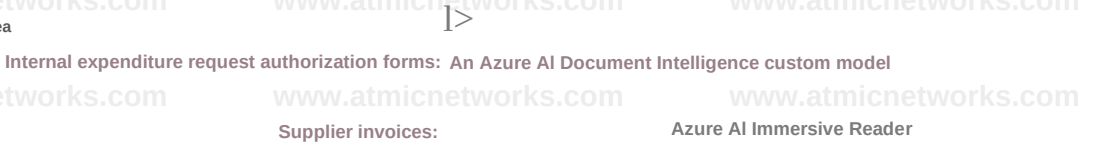

Azure AI Document Intelligence offers pre-built models for common document types like invoices, which extract key fields (e.g., vendor, amount, date) with minimal setup. For custom forms like internal authorization requests, the pre-built general document model handles layout and text extraction, but to minimize development effort, use the pre-built invoice model where applicable and a custom model only if pre-built options fall short. However, the standard recommendation prioritizes pre-built for both to avoid training.

Question: 307

You have an Azure subscription that contains an Azure OpenAI resource named AII and ari Azure AI Content Safety resource named CS1.

You build a chatbot that uses AII to provide generative answers to specific questions and CS1 to check input and output for objectionable content

You need to optimize the content filter configurations by running tests on sample questions.

Solution: From Content Safety Studio, you use the Moderate text content feature to run the tests.

Does this meet the requirement?

A. Yes

B. No

Answer: B

Explanation:

Comprehensive Detailed Explanation

- You have a chatbot that uses Azure OpenAI for generative responses and Azure AI Content Safety to filter objectionable content.
- The requirement is to optimize the content filter configurations by running tests on sample questions.
- The proposed solution: Using Content Safety Studio → Moderate text content feature.

This solution is correct because:

- Content Safety Studio provides tools to test and fine-tune content filters.
- The Moderate text content feature allows you to run sample prompts and responses against filters for categories such as sexual, violence, hate, and self-harm.
- This is exactly how you optimize and validate configurations before applying them to production workloads.

Correct Answer: Yes

Microsoft Reference:

- Content Safety Studio overview
- Azure AI Content Safety

Question: 308

You have an Azure subscription that contains an Azure OpenAI resource named OpenAI1 and a user named User1.

You need to ensure that User1 can upload datasets to OpenA11 and finetune the existing models.

The solution must follow the principle of least privilege.

Which role should you assign to User!?

A. Cognitive Services Contributor

B. Contributor

C. Cognitive Services OpenAI User

D. Cognitive Services OpenAI Contributor

Answer: C

Explanation:

• A user User1 must be able to upload datasets and fine-tune existing models in Azure OpenAI.

1 The role must follow principle of least privilege.

Role analysis:

2 A. Cognitive Services Contributor: Too broad; allows management of all Cognitive Services resources, not least privilege.

3 B. Contributor: Overly permissive; allows full management rights on the entire resource, not just OpenAI.

4 C. Cognitive Services OpenAI User: Grants permissions only to use models (query endpoints). Does not allow fine-tuning or dataset upload.

5 D. Cognitive Services OpenAI Contributor: Correct role. This provides the ability to manage OpenAI resources, including uploading training data and fine-tuning models.

Correct Answer: D. Cognitive Services OpenAI Contributor

Microsoft Reference:

6 Azure RBAC roles for Azure OpenAI

7 Built-in roles: Cognitive Services OpenAI Contributor

Question: 309

DRAG DROP

You have an Azure subscription.

You are building a chatbot that will use an Azure OpenAI model.

You need to deploy the model.

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.

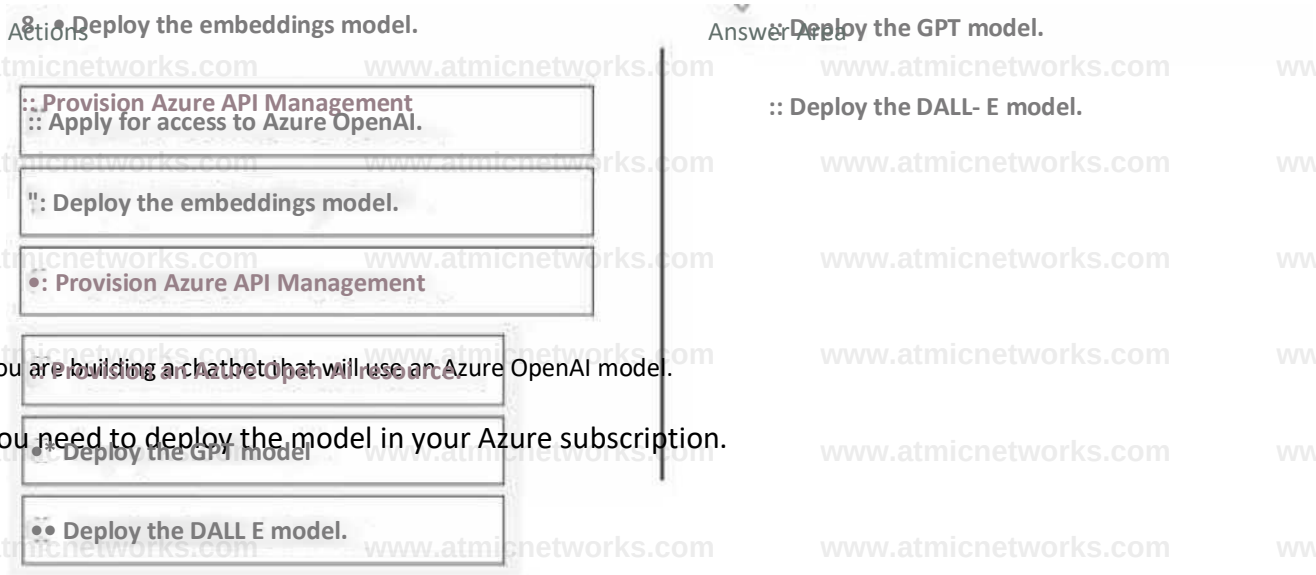

You are building a chatbot that will use an Azure OpenAI model.

You need to deploy the model in your Azure subscription.

Answer:

Explanation:

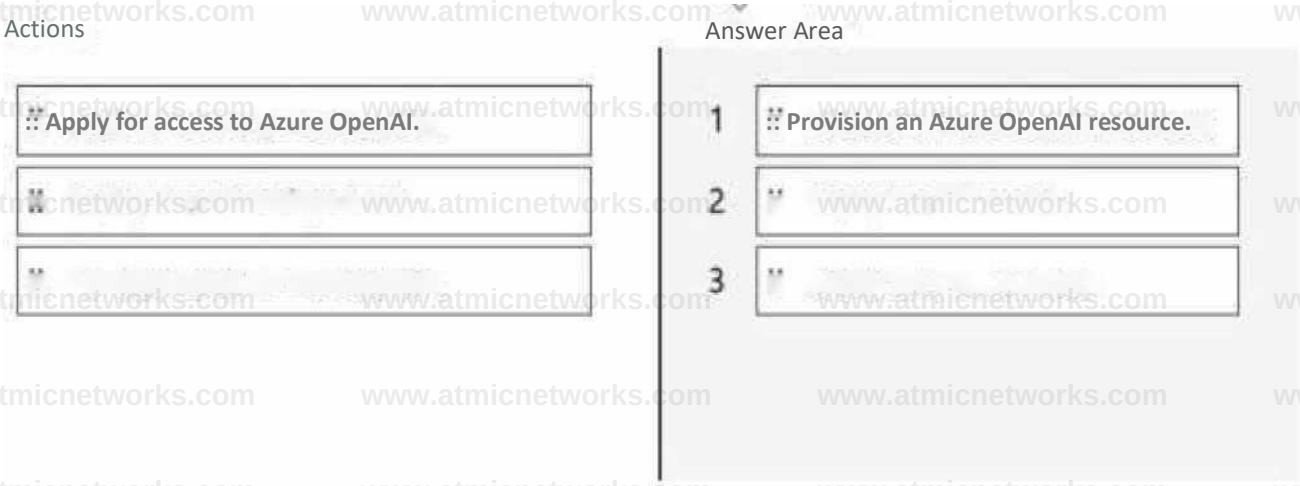

Key Steps to Deploy an Azure OpenAI Model for a Chatbot:

1. Apply for access to Azure OpenAI

o Azure OpenAI is a gated service and requires approval before you can provision resources.

2. Provision an Azure OpenAI resource

o After approval, you must provision an Azure OpenAI resource in your subscription.

3. Deploy the GPT model

o Since the chatbot requires natural language understanding and conversation, you deploy a GPT model (e.g., GPT-4, GPT-3.5).

Actions That Are Not Required for this

Question: 310

HOTSPOT

You have 1,000 scanned images of hand-written survey responses. The surveys do NOT have a consistent layout.

You have an Azure subscription that contains an Azure AI Document Intelligence resource named AIdoc1.

You open Document Intelligence Studio and create a new project.

You need to extract data from the survey responses. The solution must minimize development effort.

To where should you upload the images, and which type of model should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

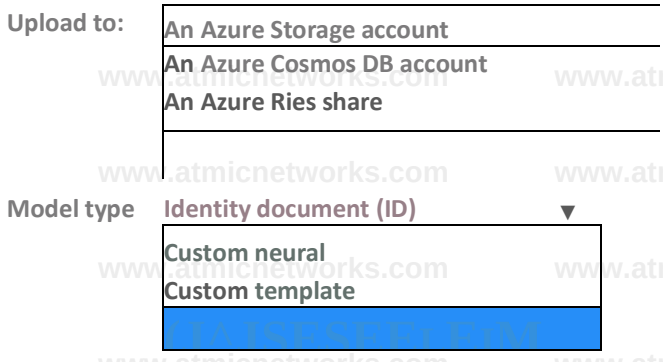

Answer:

Explanation:

Upload to: An Azure Storage account

- Model type: Custom neural

You have 1,000 scanned hand-written survey responses with no consistent layout. The goal is to extract data while minimizing development effort.

Step 1: Upload location

- For Azure AI Document Intelligence (formerly Form Recognizer), training data (images, PDFs) must be placed in a container in an Azure Storage account.
- Options like Cosmos DB or Azure Files share are not valid training data sources for Document Intelligence.
- Therefore, you upload the survey images to Azure Storage account.

Step 2: Model type

- The surveys are handwritten and do not have a consistent layout.
- Custom template models work only when the layout is fixed and consistent (e.g., invoices with the same fields in the same place).
- Identity document (ID) model is a prebuilt model for passports, driver's licenses, etc., not for surveys.
- Custom neural model is designed to handle unstructured or variable-layout forms, including handwritten content. This best fits the survey scenario.

Thus, the correct configuration is:

- Upload to Azure Storage account
- Use Custom neural model

Correct Answer:

- Upload to: An Azure Storage account
- Model type: Custom neural

Microsoft Reference

- Azure AI Document Intelligence – Custom models
- Custom neural vs custom template models
- Train a custom model with Azure Storage input

Question: 311

You are developing an app that will use the text-to-speech capability of the Azure AI Speech service. The app will be used in motor vehicles.

You need to optimize the quality of the synthesized voice output.

Which Speech Synthesis Markup Language (SSML) attribute should you configure?

A. the style attribute of the mstts: express -as element

B. the level attribute of the emphasis element

C. the pitch attribute of the prosody element

D. the effect attribute of the voice element

Answer: D

Explanation:

Question: 312

HOTSPOT

You are building a text-to-speech app that will use a custom neural voice.

You need to create an SSML file for the app. The solution must ensure that the voice profile meets the following requirements:

- Expresses a calm tone

- Imitates the voice of a young adult female

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

a. NOTE: Each correct selection is worth one point.

Answer Area

```
...
<mstts:express-as role ="YoungAdultFemale" style ="gentle">
    How can I assist you?
</mstts:express-as>
...
```

role: role, style, styledegree, type, voice

style: role, style, styledegree, type, voice

Answer:

Explanation:

Answer Area

<sistt5: express-as role * =“YoungAdultFeiwle“ Style * ="gentle“>

How can I assist you?

■ /instts:express-as>

Question: 313

You develop a custom question answering project in Azure Al Language. The project will be used by a chatbot.

You need to configure the project to engage in multi-turn conversations. What should you do?

A. Add alternate questions.

B. Enable chit-chat.

C. Add follow-up prompts.

D. Enable active learning.

Answer: C

Explanation:

Question: 314

HOTSPOT

You are building a language learning solution.

You need to recommend which Azure services can be used to perform the following tasks:

- Analyze lesson plans submitted by teachers and extract key fields, such as lesson times and required texts.
- Analyze learning content and provide students with pictures that represent commonly used words or phrases in the text

The solution must minimize development effort.

Which Azure service should you recommend for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

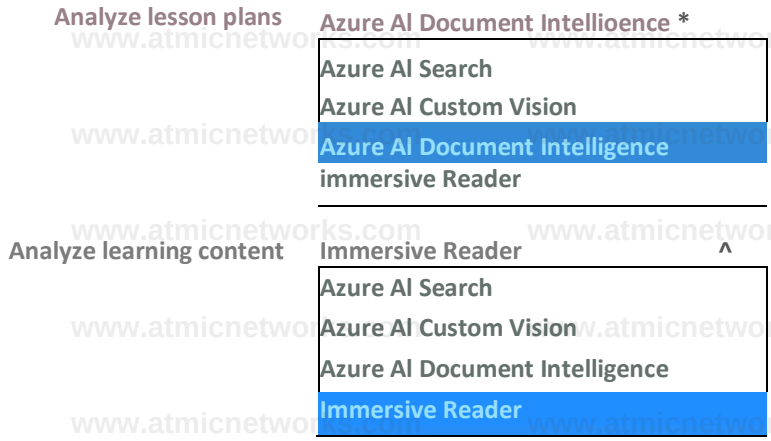

Answer:

Explanation:

Answer Area

Analyze lesson plans: Azure AI Document Intelligence

Analyze learning content Immersive Reader

Question: 315

You have a product knowledgebase that contains multiple PDF documents.

You need to build a chatbot that will provide responses based on data in the knowledgebase. The solution must minimize development effort and costs.

What should you include in the solution?

A. Azure Al Language conversational language understanding (CLU)

B. Azure Al language detection

C. Azure Al Language custom question answering

D. Azure OpenAI

Answer: C

Explanation:

Question: 316

You need to ensure that the chatbot can classify user input into separate categories. The categones must be dynamic and defined at the time of inference.

Which service should you use to classify the input?

A. Azure OpenAI text summarization

B. Azure OpenAI text classification

C. Azure Al Language custom named entity recognition (NER)

D. Azure Al Language custom text classification

Answer: D

Explanation:

Question: 317

In Azure Al Studio, you use Completions playground with the GPT-35 Turbo model.

You have a prompt that contains the following code.

```
•function F(n)

    var f = [0, 1];
    for (var 1=2; i < n; i++) f[i) = f[i-l] + f[l-2];
    return f;
```

You need the model to create an explanation of the code. The solution must minimize costs. What should you do?

A. Change the model to GPT-4-32lc

B. Add// what does function F do? to the prompt.

C. Add function F(explanation) to the prompt.

D. Set the temperature parameter to 1.

Answer: B

Explanation:

Question: 318

You have an Azure subscription that contains an Azure OpenAI resource named AM.

You build a chatbot that uses All to provide generative answers to specific questions.

You need to ensure that questions intended to circumvent built-in safety features are blocked.

Which Azure Al Content Safety feature should you implement?

A. Protected material text detection

B. Jailbreak risk detection

C. Monitor online activity

D. Moderate text content

Answer: B

Explanation:

Question: 319

HOTSPOT

You have an Azure subscription that contains an Azure OpenAl resource. Multiple different models are deployed to the resource.

You are building a chatbot by using Chat playground in Azure Al Studio.

You need to ensure that the chatbot generates text in concise formal business language. The solution must meet the following requirements:

- Reduce the cost of running the language model.
- Maintain the size of the chatbot history window.

Which two settings should you configure? To answer, select the appropriate settings in the answer are

a. NOTE: Each correct selection is worth one point.

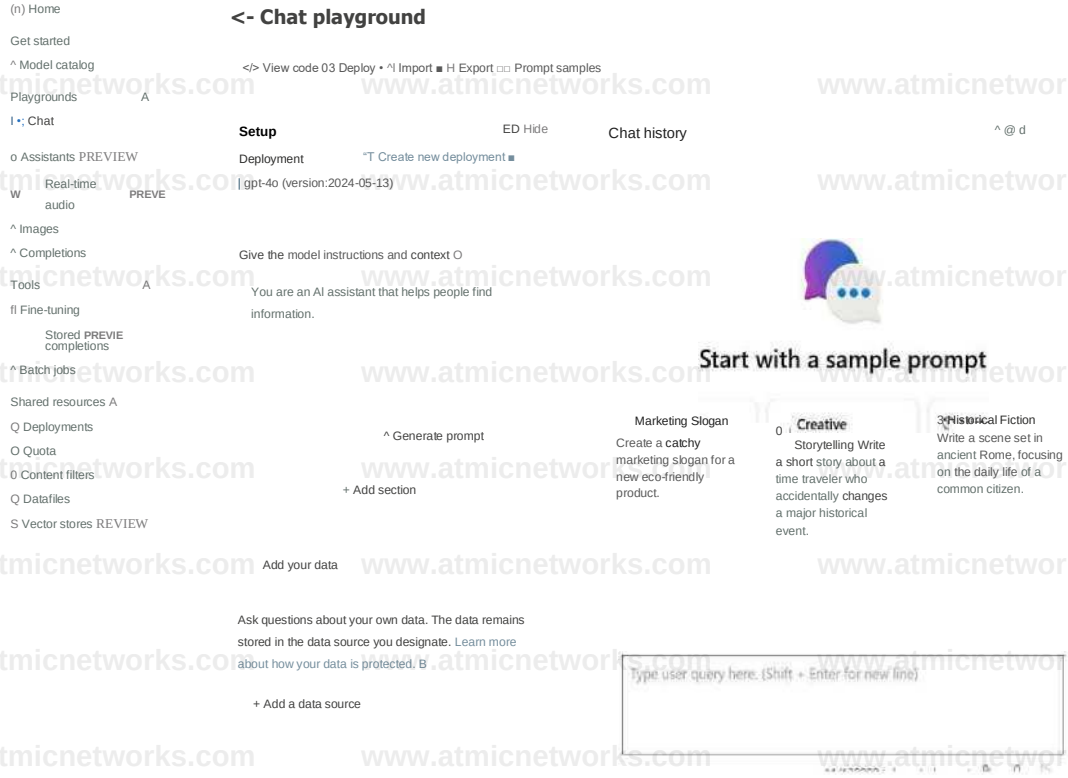

Answer:

Explanation:

Answer Area

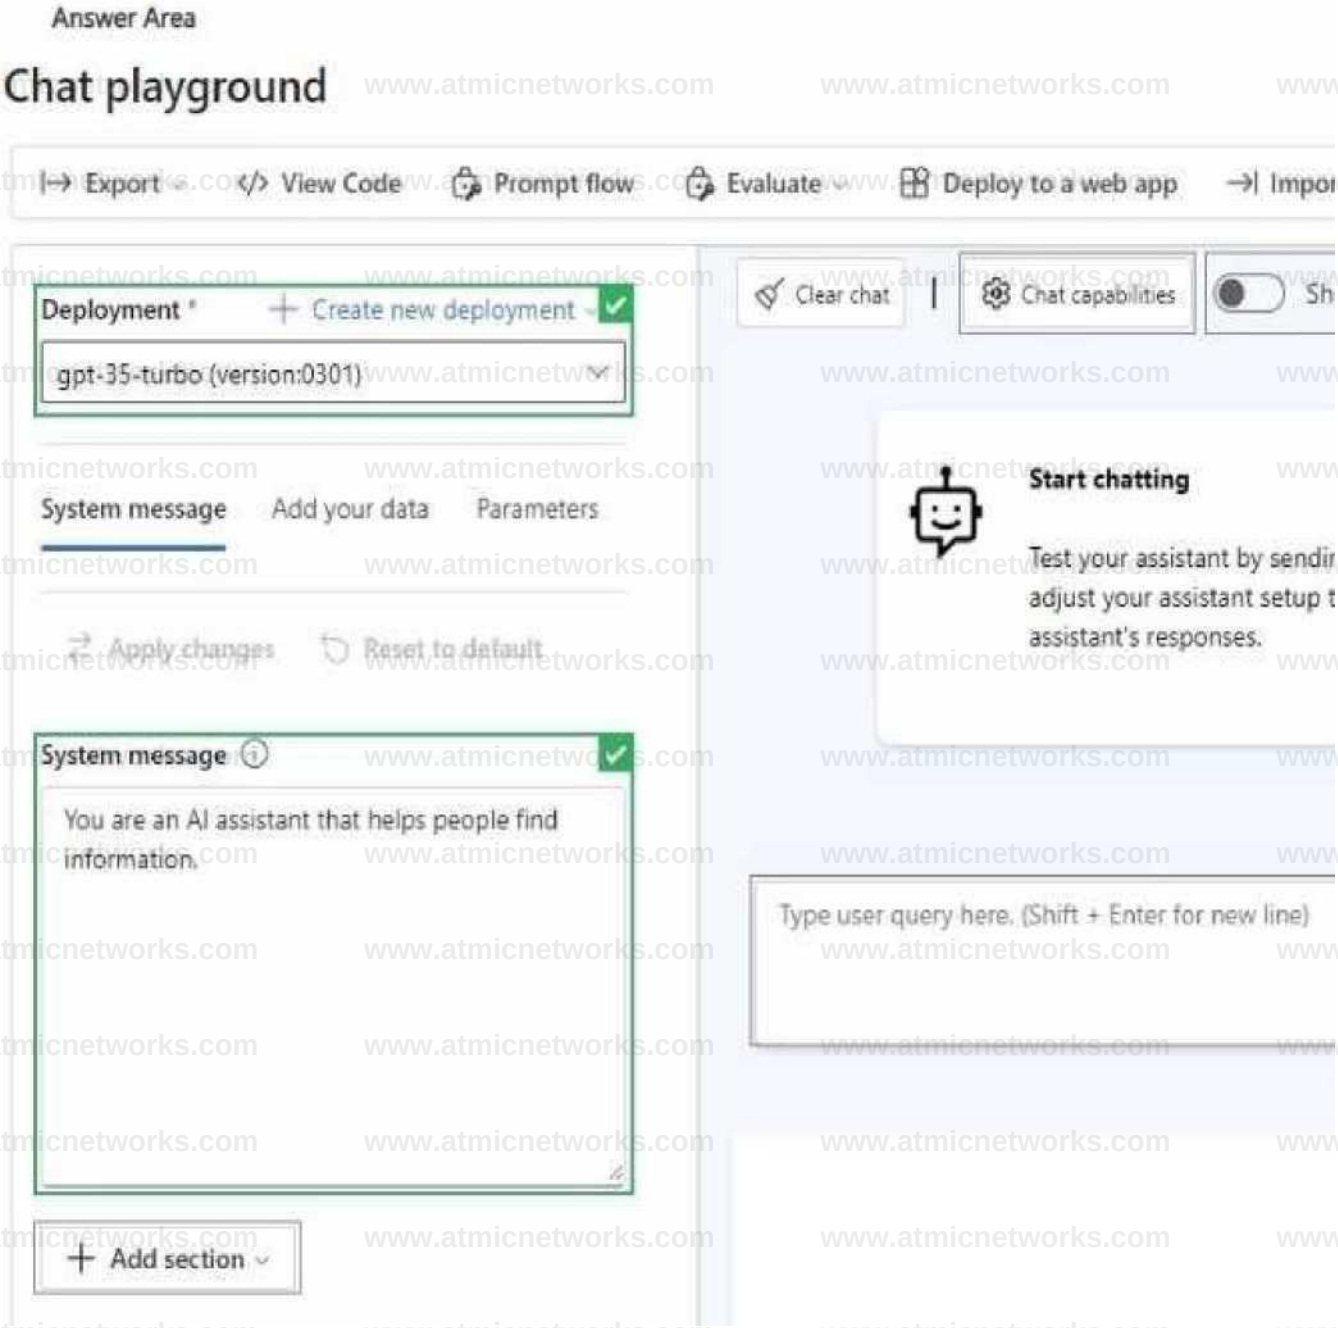

Question: 320

HOTSPOT

You have an Azure subscription that contains an Azure AI Video Indexer account.

You need to add a custom brand and logo to the indexer and configure an exclusion for the custom brand. How should you complete the REST API call? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Answer:

Explanation:

Answer Area

Question: 321

HOTSPOT

You are developing an application that will use the Azure Al Vision client library. The application has the

following code.

```
def analyze_ircage(local_image):
    with open(local image, "rb") as image^stream:
    image_analy5is = client.analy:e_image_in_5tream(
      irnage=itnage_stream,
      visual_features[
        Vis ualFeatureTypes.tags,
        VisualFeatureTypes.description
      ]
    )
    for caption in imageanalysis.description.captions:
      print(f"\n{captlon.text) with confidence (caption.confidence}")
    for tag in image_analysis,tags:
      print(f"\n{tag.name} with confidence {tag.confidence}")
```

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The code will perform face recognition.	O	Q
The code will list tags and their associated confidence.	£	0
The code will read an image file from thelocal file system,	r	O

Answer:

Explanation:

Answer Area

Statements

The code will perform face recognition.

The code will list tags and their associated confidence

The code will read an image file from the local file system.

Question: 322

You have an Azure subscription that contains an Azure Al Content Safety resource named CS1.

You create a test image that contains a circle.

You submit the test image to CS1 by using the curl command and the following command-line parameters.

```
data raw [
*iiMEe": {
   "content"; "(base 64 strlng>“


    "violence”
b
"outputType”: "E ightSeveritytevels*
```

What should you expect as the output?

A. 0
B. 0.0
C. 7
D. 100

Answer: A,C

Explanation:

Question: 323

You have an Azure subscription.

You are building a social media app that will enable users to share images.

You need to ensure that inappropriate content uploaded by the users is blocked. The solution must minimize development effort

What are two tools that you can use? Each correct answer presents a complete solution

NOTE: Each correct selection is worth one point.

A. Microsoft Defender for Cloud Apps
B. Azure Al Custom Vision
C. Azure Al Vision
D. Azure Al Content Safety
E. Azure Al Document Intelligence

Answer: B,D

Explanation:

Question: 324

You have an Azure subscription.

You need to build an app that will compare documents for semantic similarity. The solution must meet the following requirements:

- Return numeric vectors that represent the tokens of each document.
- Minimize development effort.

Which Azure OpenAI model should you use?

A. GPT-3.5
B. embeddings
C. DALL-E
D. GPT-4

Answer: B

Explanation:

Question: 325

You have an Azure subscription that contains an Azure OpenAI resource.

You deploy the GPT-4 model to the resource.

You need to ensure that you can upload files that will be used as grounding data for the model. Which two types of resources should you create? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Azure AI Bot Service
B. Azure SQL
C. Azure AI Document Intelligence
D. Azure Blob Storage
E. Azure AI Search

Answer: D, E

Explanation:

Question: 326

HOTSPOT

You have 100,000 images.

You need to build an app that will perform the following actions:

- Identify road signs in the images and extract the text on the signs.
- Analyze the text to identify well-known locations.

The solution must minimize development effort.

What should you use for each action? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Extract the text: Azure AI Vision / Azure AI Document Intelligence / Azure AI Language / Azure AI Search / Azure AI Vision

Identify well-known locations: Azure AI Search / Azure AI Document Intelligence / Azure AI Language / Azure AI Search / Azure AI Vision

Answer:

Explanation:

Answer Area

Extract the text: Azure AI Vision

Identify well-known locations: Azure AI Search

Question: 327

HOTSPOT

You have an Azure subscription that contains an Azure AI Content Safety resource named Resource1.

You create the following cURL command.

```
curl -X POST "https://resource!.cognitiveservices.azure.com/contentsafety/text:detectProtectedMaterial?api- version=2024-09-01" \
-H "Content-Type: application/json" \
-H "Ocp-Apim-Subscription-Key: <your_subscription_key>" \
-d '{ "text": "<your_content>" }'
```

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

Answer Area

Statements	Yes	No
The text value must use JSON formatted text	Q	o
The command will analyze inputted text and identify whether the text contains published song lyrics.	Q	
The Ocp-Apim-Subscription-Key value must contain the ID of the Azure subscription that hosts Resource!.		

Answer:

Explanation:

Answer Area Statements	Yes	No
The text value must use JSON formatted text	®	O
The command will analyze inputted text and identify whether the text contains published song lyrics.	0	•
The Ocp-Apim-Subscription-Key value must contain the ID of the Azure subscription that hosts Resource!.	0	8

Question: 328

HOTSPOT

You have an Azure subscription that contains an Azure AI Document Intelligence resource named DM.

You create a PDF document named Test.pdf that contains tabular data.

You need to analyze Test.pdf by using DM.

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

NOTE: Each correct selection is worth one point.

Answer Area

curl -v -i POST "{endpointJ/formrecognizer/documentModels/ | prebuilt-ia'/out I analyze?api-version=2023-07-31“ -H “Content-Type: applies

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

-H Ocp-Apim-Subscription-Key ’ : {yourkey}" --data-ascii "{'urlSource': 'test.pdf'}" iKeyl I

Ocp-ftpim-Subscriptign-Key
I Secret
Subscription-Key

Answer:

Explanation:

Answer Area ^

curl -v -i POST "{endpointj/formrecognizer/documentlfodels/ prebuilt-layout * :analyze?api-version=2023-07-31“ -H “Content-Type: applies

-H Ocp-Apim-Subscnption-Key ^ : {yourkey}" --data-ascii "{'urlSource': 'test.pdf'}"

Question: 329

You are building a social media messaging app.
You need to identify in real time the language used in messages.
Which SDK package should you install?

A. Azure.AI.Translation.Text
B. Microsoft.CognitiveServices.Speech
C. Azure.AI.Translation.Document
D. Azure.AI.Translation.Speech

Answer: A

Explanation:

Question: 330

You have a training dataset that contains 10,000 PDF documents. The documents contain scanned books, comics, and magazines.
You are building a solution that will use Azure AI and a custom model.
You need to train the model by using Language Studio. The solution must meet the following requirements:
• Tag each item as a book, comic, or magazine.

Answer: B

Explanation:

Question: 331

DRAG DROP
You are building a phone call handling solution that will use the Azure AI Speech service and a custom neural voice.
You need to create a custom speech model.
Which five actions should you perform in sequence from Speech Studio? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

* Minimize development effort.
What should you use?

A. a custom extraction model
B. a multi label classification project
C. a custom named entity recognition (NER) project
D. a multi label image classification model

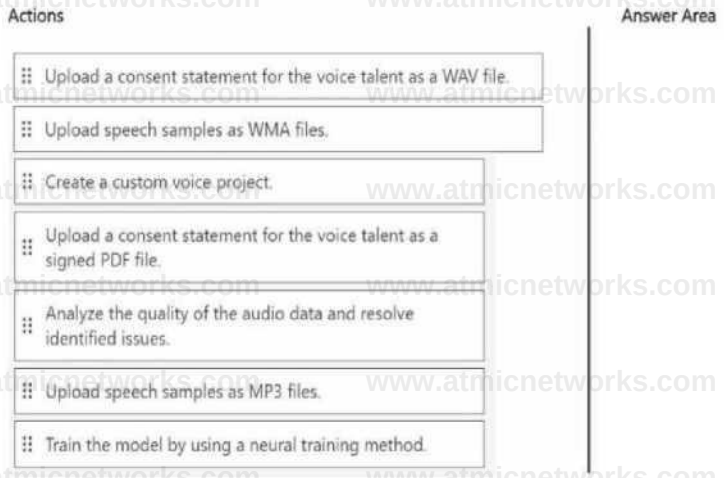

Answer

Explanation:

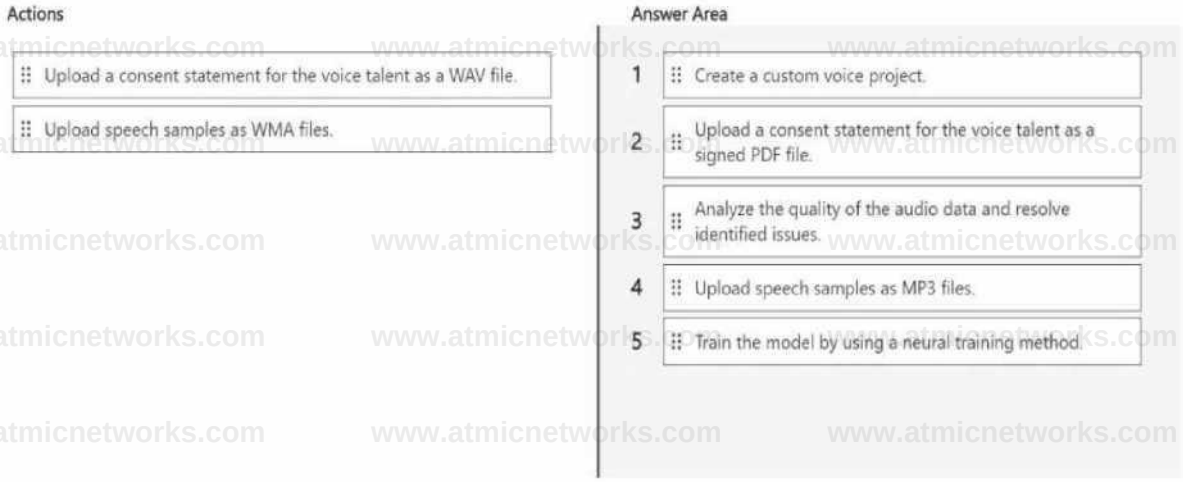

Question: 332

You are building an app that will analyze documents by using the Azure AI Language service.

You need to identify industry-specific technical terms in the documents. The solution must minimize development effort.

What should you use?

A. key phrase extraction
B. custom named entity recognition (NER)
C. conversational language under-standing (CLU)
D. language detection

Answer: B

Explanation:

Question: 333

You have an Azure subscription that contains an Azure AI Language custom question answering project named QA1.

You need to import question and answer pairs to QA1.

Which two file formats can you use? Each correct answer presents a complete solution

NOTE; Each correct selection is worth one point.

A. Excel
B. TSV
C. JSON
D. LU
E. CSV

Answer: B, E

Explanation:

Question: 334

DRAG DROP

You Have a chatbot that uses the Azure AI Language custom question answering service. The model used by the service was trained by using an internal support FAQ document.

You discover that the chatbot fails to provide correct answers to common questions.

You need to increase the accuracy of the responses provided by the chatbot. The solution must minimize development effort.

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

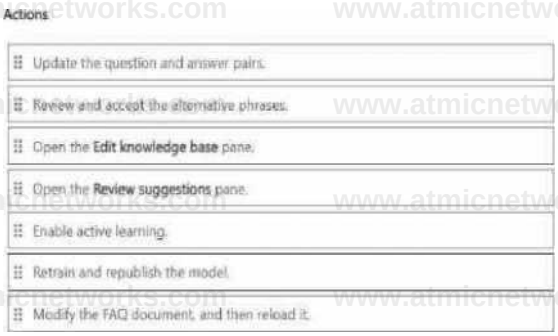

Answer:

Explanation:

Question: 335

You have an Azure subscription that contains an Azure Al Document Intelligence resource named Aldoc1.
You have an app named App1 that uses Aldoc1. App1 analyzes business cards by calling business card model v2.1.
You need to update App1 to ensure that the app can interpret QR codes. The solution must minimize administrative effort.
What should you do first?

A. Deploy a custom model.
B. Implement the read model.
C. Upgrade the business card model to v3.0
D. Implement the contract model

Answer: C

Explanation:

Question: 336

DRAG DROP
You have an Azure subscription that contains an Azure Al Search resource named AS1.
You implement a custom skill in AS1 that performs language and sentiment analysis of documents.
You ate evaluating the use of AS1 as part of an enrichment pipeline.
In which order will AS1 index the documents? To answer, move all indexing stages from the list of stages to the answer area and arrange them in the correct order.

Answer:

Explanation:

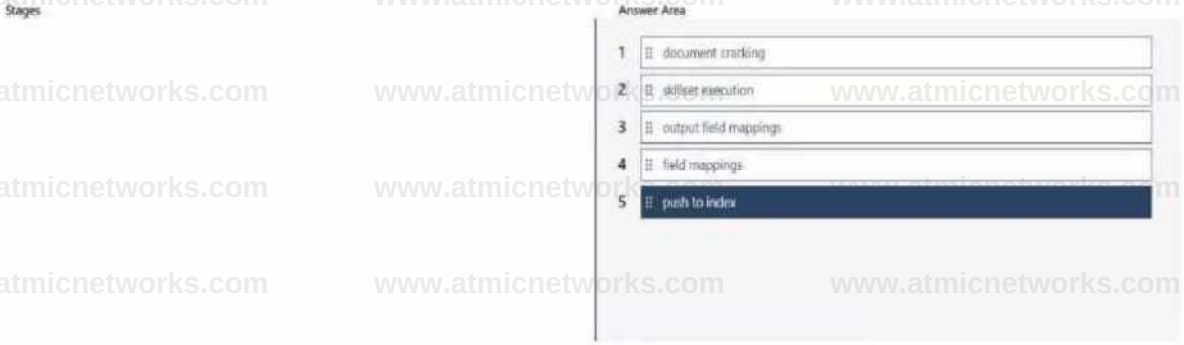

Question: 337

You have a product support manual.

You need to build a product support chatbot based on the manual. The solution must minimize development effort and costs.

What should you use?

A. Azure Al Phi-3-medium with fine-tuning
B. Azure A1 Language Custom question answering
C. Azure OpenA1 GPI-4 with grounding data that uses Azure Al Search
D. Azure Al Document Intelligence

Answer: B

Explanation:

Question: 338

You plan to build an agent that will combine and process multiple files uploaded by users.

You are evaluating whether to use the Azure Al Agent Service to develop the agent. What is the maximum size of all the files that can be uploaded to the service?

A. 1 GB

B. 10 GB
C. 100 GB
D. 1 TB

Answer: C

Explanation:

Question: 339

You are building an app by using the Semantic Kernel.

You need to include complex objects in the prompt templates of the app. The solution must support objects that contain sub-properties.

Which two prompt templates can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Liquid
B. JSONL
C. Handlebars
D. YAML
E. Semantic Kernel

Answer: A, E

Explanation:

Question: 340

HOTSPOT

You are building an agent that will retrieve the current time at a given location by using a custom API.

You need to test the functionality of the custom API.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
curl https://cuj.servlces.ai.azure.coio/ assistants          ᵀ ?api-version«2e24 u-01 prev
                                          assistants
                                          completions embeddings threads
-       H "Authorization: Bearer JAZUR£_M_TOKEN" \
-       H “Content Type: application/json" \
-       a '(
  "instructions": "You are a weather bot. Use the provided functions to answer
questions.”,
  “model": "gpt-Ao-mini",
   contented *
   content=[(
   functions-H
   tool_resources=[( tools=[(
   "type": "function",
   "function": {
     "name”: “get current time”,
     "description": "Get the current time in location",
     “parameters": {
       "type": "object”,
       "properties": (
         "location": ("type": "string", "description": "The city name, for example Seattle")
       }.
       "required": ("location”]
```

Answer:

Explanation:

Answer Area

```
curl https://cui.servlces.ai.a2ure.coa1/ assistants                    ^ ?api version-2024 12 01 preview \
-H "Authorization: Bearer JAZUHE.AI.TOKEH" \
■H "Content-Type: application/json“ \
                                                                       '(
  "instructions": "You are a weather bot. Use the provided functions to answer questions.",
  "model": ”gpt-4o-mini",
    content={{
    “type": “function",
    "function": {
      "name": "get_current_tlaie".
    “description": "Get the current time in location",
    "parameters": {
      "type": "object",
      "properties”: (
        "location": ("type": "string", "description": "the city name, for example Seattle")

    "required": ["location"]
```

Question: 341

DRAG DROP

You have an Azure subscription that contains an Azure Al Search instance named AlSearch1.

AlSearch1 contains an index that includes a vector.

You need to perform the following actions:

- Deploy a new agent by using the Azure Al Agent Service.
- Connect the AlSearch1 index to the new agent.
- Validate the integration of the index and the agent.

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.

Actions

|Create an agent and enable the code interpreter tool

:: Download the files generated by the code interpreter too[1]

.. Create an Azure AI Client and retrieve the connection ID of the " Azure AI Search resource.

:: Create an agent and enable the Azure AI Search tool

:: Configure the Azure AI Search tool

- : Ask the agent questions about the data in the AISearch! index

Answer Area

Answer:

Explanation:

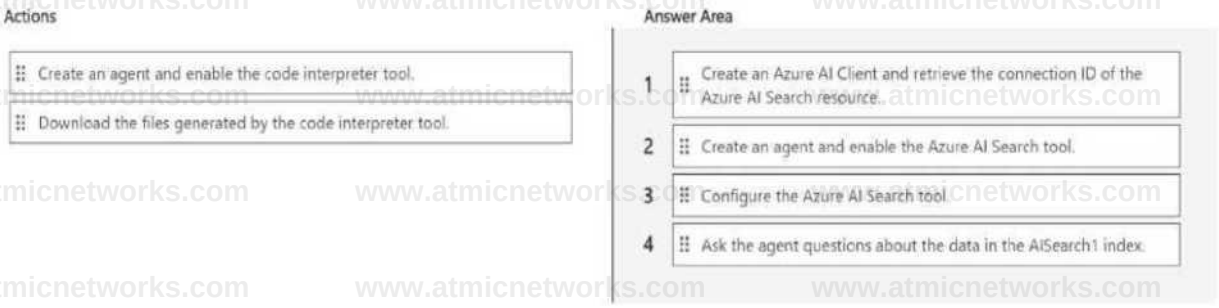

Question: 342

You are building an app that uses a Language Understanding model to analyze text files. You need to ensure that the app can detect the following entities:

- Temperatures
- Currency values
- Email addresses
- Telephone numbers

The solution must minimize development effort.

Which model capability should you use?

A. list entities
B. learned entities
C. utterances
D. regular expression components
E. pre-built entity components

Answer: E

Explanation:

Question: 343

HOTSPOT

You have an Azure subscription that contains an Azure AI Content Safety resource named Resource1 and a storage account named storage1.

You create a blob container named container1 and upload a sample set of image files to container1.

You need to validate whether Resource1 can identify images that contain potential violence.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
curl -location -request POST
'https://resourcel.cognitiveservices.a2ure.com Hcgntentsafety 3 /inugeanalyze ᵛ 'api-version^zejd-oo-er \
J /contentmoderator |.; /detect
/contentsafety _____ | /imageanalyze
|/Vision ]/Processimage
header 'ocp-Apim-subscriptlon Key: 33esdfi6c33f43as942926d«a*ac32b6' \
--header 'Content-Type: application/ison' \
```

data raw '(

"image": (

"blobUrl": “https://storagel.blob.core.Kindows.net/rontainerl

r

Answer:

Explanation:

Answer Area

curl -location -request POST

'https://resourcel.cognitiveservices.a2ure.co1> /contentsafety ' /imaqe:analyze ˅ ?apl-version«2e24-09-ei' \

header 'ocp-Apim-subscrlptlon Key: 33e5dflbc33f43aS942926dsa8ac3i , . /imaqeanalyze * ?api-version-2P24-09 ei' \

header Content Type: appJicatlon/json \

--data-raw {

- sHcation/ison' \

■image": (

"blobUrl": "https://stoi

1

' igel.blob.core.windows,net/container1

Question: 344

You are building a social media messaging app.
You need to identify in real time the language used in messages.
Which service should you use?

A. Azure Al Speech
B. Azure Al Content Safety
C. Azure Al Translator
D. Azure Al Language

Answer: D

Explanation:

Question: 345

HOTSPOT

You are building an app that will analyze text by using the Azure AI Language service.
You need to configure the app to mask the telephone number and email details in a given document. How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

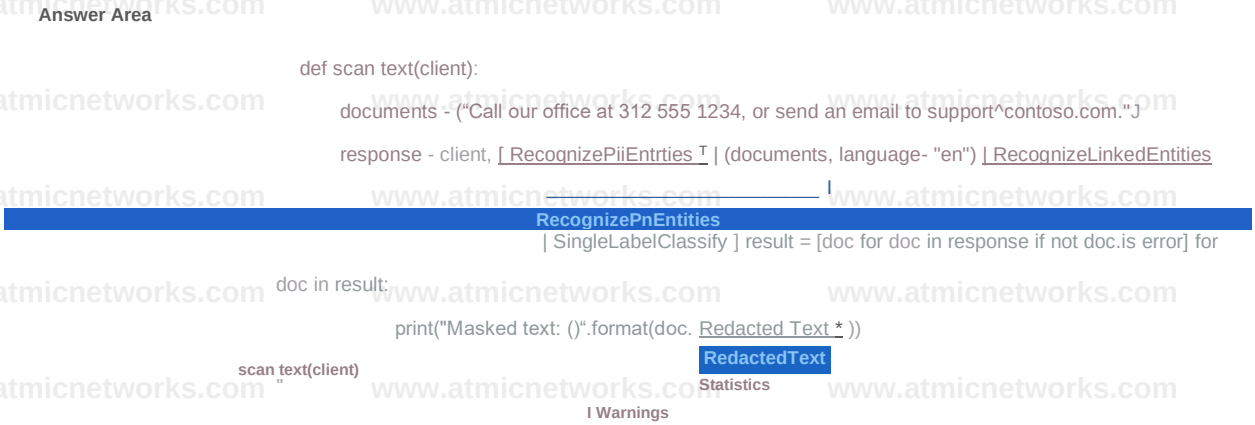

Answer:

Explanation:

Question: 346

HOTSPOT

You are building an app that will translate speech by using the Azure AI Language service.
You need configure the app to translate the speech from English to Italian.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

speech translation config - speechsdk.translation.SpeechTranslationConfig(subscriptian^os.environ.get('SPEECH KEY')

region-os.environ.get('SPEECH REGION'))

speech translation ronfig. . speechrecognition lanquage addJargetJanguage J -"enUS"

region

speechjecognitionjanguage

voice name

speechtranslationconfig. addUrgetlanquaoe

addtargetjanguage

set speech synthesis _ output Jo rmat speech recognitionJanguage voice name

audlo_conflg - speechsdk.audio.AudioConfig(use_default_«lcrophone'True)

recognizer † sperchsdk.trans 1 at ion.TranslationHecognizer(translation_config=speech trans 1 ationcnnfig, audioronfig^audio config)

Answer:

Explanation:

Answer Area

speech translation config - speechsdk.translation.SpeechtranslationConfig(subscriptionM>s.environ.get('SPEECH KEY'),

region-os.environ.get('SPEECH.REGION'))

speech translation config. speech recoqnition language ^ -"en-US"

speech translation config. add target language . ^ ("it")

audlo_conflg - speechsdk.audio.Audioconflg(use_default_«lcrophone'True)

recognizer * speechsdk.translation.lranslationRecogni-er(translation_config-speech translation config, audioconfig-audio config)

Question: 347

HOTSPOT

You have 100,000 images.

You need to build an app that will perform the following actions:

† Analyze the descriptions to generate a report about the different types of road signs and how often each type occurred.

The solution must minimize costs.

What should you use for each action? To answer, select the appropriate options in the answer area. NOTE; Each correct selection is worth one point.

Answer Area

Azure Ai Vision

Azure AI Document Intelligence Azure AI Phi-3-mini

Azure AI Vision

| Azure OpenAI GPT-4-TurbO'

Azure OpenAI GPT-4Turbo

Azure AI Document Intelligence

Azure AI Language

Azure AI Phi-3-mini

Azute OpenAI GPT-4-Turbo

HOTSPOT

You have a custom analyzer named analyzer1 that performs the following functions

• The transcription of video content

• Identify road signs in the images and generate a short description of each road sign.

Answer:

Explanation:

Answer Area

Identify the road signs and generate a short description of each road sign Azure AI Vision

Analyze the descriptions to geneiate a teport about the different types of Azuie OpenAI GPT 4 Turbo road signs and how often each type occurred

Identify the road signs and generate a short description of each road sign

Analyze the descriptions to genet ate a report about the different types of road signs and how often each type occurred:

Question: 348

You have a library that contains 1,000 video files.

You need to perform sentiment analysis on the videos by using an Azure AI Content Understanding project. The solution must minimize development effort.

Which type of template should you use for the project?

A. Video shot analysis
B. Media asset management
C. Advertising

Answer: B

Explanation:

Question: 349

Answer Area

Statements	Yes	No
analyzer 1 is in a ready state.	0	0
analyzer 1 finished analyzing the file.	Q	0
The key frames of Videol .mp4 can be retrieved by using the data in the command and the output		○

Answer:

Explanation:

Answer Area

Statements	Ves	No
analyzer 1 is in a ready state.	0	«
analyzer 1 finished analyzing the file.	○	®
The key frames of Videol mp4 can be retrieved by using the data in the command and the output	W	0

Question: 350

You have a computer that contains the files shown in the following table.

Name	Format	Length (mins)	Size (MB)
Ftlel	MP4	34	1.500
Ftle2	AVI	ri in	1,700
File3	MP3	300	980
File4	MP4	350	2,800

Which files can you upload and analyze by using Azure Al Video Indexer?

A. File1 only
B. File3only
C. File1 andFile3 only
D. File1, File2, and File3 only
E. File1, File2, File3, and File4

Answer: D

Explanation:

Question: 351

HOTSPOT

You have an Azure Al Speech service resource named Resource1.

You call Resource1 by running the following Python code.

```
def synthesize_speech(input):
    speechconfig • speechsdk.speechconf ig( subscript ion-os. environ, get ('VourSpeechKey'),
region=os.environ.get('VourspeechHegion'))
    audio_config = speechsdk. audio. AudioOutputconf ig(filenai»e='path/to/file. wav")
    speechsynthesizer = speechsdk.SpeechSynthesizer{speech config=speech config,audio config=audie config)
    speechsynthesizer.speak_text_async(input)
```

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

Answer Area

Statements	
The function will fail tf there is an existing file named File wav	
The function will sample File wav to use as a synthesized voice.	
The function will generate an audio file based on the input text	

Answer

Explanation:

Answer Area

Statements	Yes	No
The function will fail if there is an existing file named File wav.	0	
The function will sample File wav to use as a synthesized voice.		
The function will generate an audio file based on the input text		

Question: 352

HOTSPOT

You have a chatbot that uses the Azure Al Language custom question answering service.

You need to test the chatbot. The solution must ensure that the chatbot responds only when an answer has a confidence score of at least 95 percent.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
cur] -x POST -M "Ocp-Apim-Subscription-Key: $LANGUA6E_KEV" -H "Content-Type: application/json" -d '(
"question": "How much energy does ay phone have Left?",
confidenceScoreThreshold ᵛl : "0.95", confidencescore
confidenceSaxeThfeshold
. context
| rankerType
'SLANGUAGE EUJPOIKT.apt.cognitive.Microsoft.COM/language/:query-knowledgebases?   ProitCtName — i -ChatBot-prcjcctkopi-version-2021-10-01
                                                                                  deploymentName
                                                                                  kbld
                                                                                  knowledgebases
                                                                                  project Name
```

Answer:

Explanation:

Answer Area

```
curl -x POST -H "Ocp-Apim-Subscription-Key: $LAN6UAGE_KEV" -H "Content-Type: application/jsoft" -d '{
"question": "How much energy does my phone have left?", confidenceScoreThreshold ᵛ : "e.«",
)- 'SLANGUAGE.ENOPOXNT. epi. cognitive. Microsoft. co«/language/:queryknouledgebases?   prqjectName   ' -ChatBot-projectSopi-versior-2021-ie-GI
```

Question: 353

HOTSPOT

You are building an app that will perform translations by using the Azure Al Translator service.

You need to ensure that the app will translate user-inputted text.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
while (true)
Console.WriteLine("Enter text to translate or 'exit's");
string input = Console.Readline();
if (input.ToLower() -• "exit") break;
Re5ponse<IReadOnlyList<TranslatedTevtltern*> translationResponse «
await client.TranslateAsync (TranslateAsync(tarqetLanquaq€. input) ^"   • ConfLgureAwalt(false);
                            TranslalefUrgetLarguagej^^ ]
                            IranslaleAsyMIaigen^igi^^
                            TranslationRecognizerO|
IRea50nlyList<TranslatefITextlteo» translations • rranslat* ^Response .Value;
TranslatedTextltern translation • translations[8];
string source Language • translation?.Detectedanguage?.language;
Console.*riteLine(S"';input)' translated from (sourceLanguage; to (translation?.Translations^].To.- as | (translation*Translators^!)]7 Text) 7 translations Value'.' ■
                                                                translationsfO] text 7 (TransiaiionsIW/T
                                                                (lmnsLition?Translations?(0|?Teit|'.'):
```

Answer:

Explanation:

Answer Area

```
while (true)
{
Console.WriteLine("Enter text to translate or 'exit':");
```

```
string input « Console.Readiine();
if (input.ToLower() -- "exit") break;
Respon5e<IRead0nlyList<TranslatedTextItem>> trar,slationRespanse -
    await client.TransiateAsync TranslateAsyncdargetLanguage. input) ^ .conflgureAwalt(false);
IReadOnlyList<TranslatedTextItem> translations • translationResponse.Value;
TranslatedTextItew translation = translations^];
string sourceLanguage ■ translation?.DetectedLanguage*.Language;
Console.writeLine($' {Input}' translated from {sourceLanguage} to {translation³.Translations^®].To} as * (ttanslabor^TfansiationsTIOf Texl}.")
```

Question: 354

HOTSPOT

You are building an app that will perform speech translation by using the Azure ai Language service.

You need to ensure that the language input to the app is supported.

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

NOTE; Each correct selection is worth one point.

Answer Area

```
...
public static void CheckLanguage(TextTranslationClient client, string language)
{
  Response<GetLanguagesResult> languagesResponse = await
    client.[GetLanguagesAsync ▼] (scope:"translation").ConfigureAwait(false);
  GetLanguagesResult languages = languagesResponse.Value;
  if (languages.Translation.[ContainsKey ▼] (language))
  {
    return true;
  }
  else
  {
    Console.WriteLine($"Sorry, {language} is not a supported language.");
    return false;
  }
```

GetLanguages
GetLanguagesAsync
TranslationRecognizer

Contains
ContainsKey
GetValueOrDefaul
Where

Answer:

Explanation:

Answer Area

```
public static void QiecKLanguBge(TextTran51#tionClient client, string language¹
```

```
ltS5ons65etLJ«gusse$R«ult> languagesR«ponse = await
  client. GeUanquaqesAsync * (scope:"transiation" j.ConfigureAwalti-false);
GetiangjagesPesult languages = languagesSesponse.Value;
if (languages.Translation. ContamsKey              ^ (languageJ
  return true;
else
  Console.WrltetineCS^orry, ^language? is not a supported language.");
  return false;
```

Question: 355

You have an Azure subscription that contains an Azure OpenAI resource.

You plan to build an agent by using the Azure Ai Agent Service. The agent will perform the following actions:

- Interpret written and spoken questions from users.
- Generate answers to the questions.
- Output the answers as speech.

You need to create the project for the agent.

What should you use?

A. Language Studio
B. Azure Al Foundry
C. Speech studio
D. the Azure portal

Answer: C

Explanation:

Question: 356

HOTSPOT

You have an Azure subscription that contains an Azure Al Language resource named Resource1.

You run the following cURL command, and then play the Output.mp3 file.

```
curl --location --request POST "https://cdstus.tts.speech.micro50ft.com/cognitlve5ervices/vl" --
header "Ocp-Apim-Subscription-Key: 3795c4eilf714f5aa6W9e573ie9e4f' '
--beader "Content-Type: applicatlon/ssml+xiil" *
--header "X-Mlcrosoft-OutputFonnat: audio-16khi-128kbitrate-mono-mp3" ^
--header "User-Agent: curl" *
--data-raw *<speak version^'l.e' w»l:lang='en-US'>
  cvolce xml:gender='Female' name='en-U5-3ennyMeural> welcome to the Axure Text-to-Speech
```

```
demonstration. </volce>
cvoice xml:gender='Hale' nane='en-G8-aydnMeural'>
    This service allows you to convert text into natural-sounding speech.
</voices cvolce xml:gender='Male' name='en-US-Chrlstopherfieural> cmittsiexpress-as
style^advertlsement-Upbeat" styledegree="2">
    It's easy to integrate, customizable, and supports multiple languages and voices.
    C/mstts:express-as >

    </votce>

  </speak>" --output output.mp3
```

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

Answer Area

Statements	Yes	No
You hear three sentences in different voices.	Q	O
You hear three sentences in different accents.	Q	0
You hear three sentences expressed in a neutral tone.		G

Answer:

Explanation:

Answer Area

Statements	Yes	No
You hear three sentences	in different voices.t	O
You hear three sentences	in different accents. I	
You hear three sentences	expressed in a neutral tone.	

Question: 357

You have a blog that allows users to append feedback comments. Some of the feedback comments contain harmful content that includes discriminatory language.

You need to create a prototype of a solution that will detect the harmful content. The solution must minimize development effort.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Sign in to Content Safety Studio and select Moderate text content.
B. From the Azure portal, create an Azure Al Content Safety resource.
C. From the Azure portal, create an Azure OpenAI resource.
D. Sign in to Azure Al Foundry and select Safety + security.
E. Sign in to Content Safety Studio and select Protected material detection for text.

Answer: A, B

Explanation:

Question: 358

You are designing a solution that will answer questions about human resources (HR) policies stored in the PDF format.

You need to ensure that the identical answer to a specific question is returned every time. The solution must minimize development effort

Which service should you include in the solution?

A. Azure Al Language
B. Azure Machine Learning
C. Azure OpenAI
D. Azure Al Document Intelligence

Answer: A

Explanation:

Question: 359

You have the following files:

- File1.pdf
- File2.jpg
- File3.docx
- File4.webp
- FileS.png

Which files can you analyze by using Azure Al Content Understanding?

A. File1 .pdf and File3.docx only
B. File1.pdf, File2jpg, and File5.png only
C. File1.pdf, File2.jpg. and File3.docx only
D. File1.pdf, File2.jpg. File3.docx, and FileS.png only
E. File1.pdf. File2.jpg, File3.docx. File4.webp, and File5.png

Answer: D

Explanation:

Question: 360

You are building an app that will use Azure AI Language to extract meaning from text messages.
You need to provide additional context by adding Reference to supporting articles in Wikipedia. What should you use?

A. entity linking
B. custom entity extraction recognition (NER)
C. Azure AI Content Safety
D. key phrase extraction

Answer: A

Explanation:

Question: 361

HOTSPOT

You are building an app that will provide users with definitions of common AJ terms.
You create the following C# code.

```
OpenAXClient client =
  new OpenAXClient(new Uri(endpoint), new AiureKeyCredentiel(key));
ChatCompletionsOptions options ■ new ChatCo«pletionsOptions()

    Messages •

        new Charlessage(Cn3tAole.System, "You are a helpful assistant."), new
        ChatMessage(Chat«ole.User, “What is an LLM?")
};

ChatCowpietions response • client.GetChatCompletions(
    deployuentNMej options);
Chafes sage completion * response.Cholces[0].Message;
Console.WriteLine($"Chatbot: {completion.Content}*);
```

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

Answer Area

Statements	Yes	No
The response will contain an explanation of large language models (LLMs) that has a high degree of certainty.		
Changing "What is an LUI?" to "What Is an Lin in the context of AI models?' will produce the intended response.		
Changing "You are a helpful assistant." to "You must answer only within the context of AI language models." will oive a hioher likelihood of oroducina the intended		

Answer:

Explanation:

Answer Area

Statements	Yes	No
The response will contain an explanation of large language models (LLMs) that has a high degree of certainty.	O	•
Changing 'what is an tut?" to "what is an LUI in the context of AI models?' Will produce the intended response.	9	
Changing "You are a helpful assistant." to "You must answer only within the context of AI language models.' will oive a hioher likelihood of oroducina the intended	®	

Question: 362

You have an Azure OpenAI model.

You have 500 prompt-completion pairs that will be used as training data to fine-tune the model.

You need to prepare the training data.

Which format should you use for the training data file?

A. XML
B. JSONL
C. CSV
D. TSV

Answer: B

Explanation:

Question: 363

You have an Azure subscription.
You need to deploy an Azure AI Search resource that will recognize geographic locations.
Which built-in skill should you include in the skillset for the resource?

A. AzureOpenAIEmbeddingSkill
B. Document Extract ionSkill
C. EntityLinkingSkill
D. EntityRecognitionSkill

Answer: D

Explanation:

Question: 364

HOTSPOT

You have an Azure subscription that contains an Azure AI Content Understanding resource named cu1.
You need to create a custom analyzer that will analyze documents.

How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer

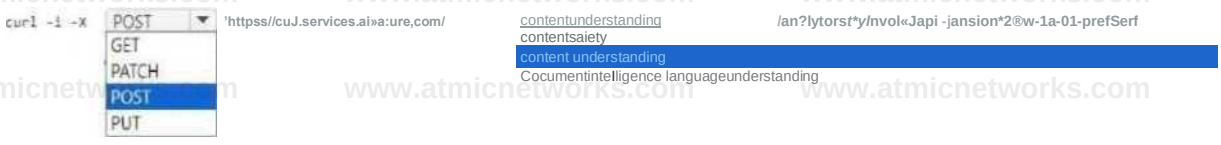

Answer:

Explanation:

Answer Area

curl -1 -X POST pᵛ "https://cul.services.ai.a2ure.cOM/ COntentunderstandmg ᵛ /analyxers/#ylnvoice?api-version-2024-12-®l-preview"

-H "Ocp-Apia-Subscription-Key: {key}' -H "Content-Type: application/json" -d freqbody.json

Question: 365

You have an Azure Al Search resource named Search1.

You have an app named App1 that uses Search1 to index content.

You need to add a custom skill to App1 to ensure that the app can recognize and retrieve properties from invoices by using Search1.

What should you include in the solution?

A. Azure OpenAI
B. Azure Al Immersive Reader
C. Azure Al Document Intelligence
D. Azure Custom Vision

Answer: C

Explanation:

Question: 366

HOTSPOT

You need to build an app that will use the Azure Al Speech service to translate audio files.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
static void Output$peechRecognitionResult(TransiationRecognitionResult translatlonRecognitionResult)

switch (translatlonRecognitionResult. j reason
```

reason / text / translations

```
case ResultReason.
```

TranslatedSpeech / NoMatch / RecognizingSpeech / TranslatedSpeech

```
console.Writeline(*"RECOGNIZED: Te«f{translationRecognitionResult.Text}');

foreach (var element in translationRecognitiorResult.Translations)

    Console.Writeline($"TRANSIATEO Into '{element.Key}': {element.Value)");

  break;
case ResultReason.NoMatch:
  Console.WrlteLine($"5peech could not be recognised.");
  break;
```

Answer:

Explanation:

Answer Area

```
static void OutputSpeechRecognitionResult:TranslationRecognitionResult translatior.RecognitionResult { switch
(translationRecognitionResult. reason                    ˅

  { case ResultReason. TranslatedSpeech ^ Console.WriteiIne($'RECOGNIZED:
    Text»{translationRecognitionResult.Text}*); foreach (var element in translationRecognitionResult.Translations)

                Console.Uriteiine S'TRANSLATED into '{elenent.key}': {element.Value}*); break; case

            ResultReason.NoMatch: Console.Writeline($*Speech could not be recognized.*); break;
```

Question: 367

HOTSPOT

You have 100,000 documents.

You are building an app that will identify city names in each document by using Azure Al Language.

You need to test the detection client.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
static void EntityRecognitionExample(      TextAnalyticsClient client)
                                           DocumentAnalysisClient client)
                                           FormRecognizerClient client)
                                           QuestionAnsweringClient client)
                                           TextAnalyticsClient client)

  ar response • client.RecogniteEntitles('I rad a great experience isiting Cor to so in Redmond.');
Cons ole. Write Line ("Named Entities:");

foreach (var entity in response.Value)
{
  if (entity, | Category ᵀ I .ContainsC'City")) {
                Category
                KeyValuePairs
                SubCategory
                TextAppearance

    console.WriteLine(S'\tText: (entity.Text},XtResult: (entity.ConfidenceScore}");
```

Answer:

Explanation:

Answer Area

```
static void EntityRecognitionEx*»ple( TextAnalyticsClient Client)

{
  ar response • client.RecogniieEntitiesfl had a great experience .isiting Contoso in Redmond.')
 Console.UriteLine(*Ma*ed Entities: ");
 foreach (var entity in response.value) {
  if (entity. Category                ^ .Contains('City')) {
      Console.WriteLine(J"\tText: (entity.Text},\tResult: (entity.ConfIdenceScore}");

 }
```

Question: 368

HOTSPOT

You have an Azure subscription that contains an Azure AI Content Safety resource named resource1.

You need to add a custom category to resource1.

How should you complete the cURL statement? To answer, select the appropriate options in the

answer area.

NOTE: Each correct selection is worth one point.

Answer Area

curl -X PVT ■<endpoint>/contentsafety/text/categories/learning-advice?api-version-2024-02-15-preview" X

-H "Ocp-Apie-Subscription-Key: <api_key>“ \

-H "Content-Type: application/json" \

-d "{

\"categoryM«»e\“: \"learnIng- adviceV,

Vprompty: \"text prompts about learning advice while preparing for an axM.V>

VdefinrtionV:

VpromptV:

VtextV:

SampleBIpblUrl ˅ \"https ://<your-azure- storage url>; container /learning-advice. jsonl\" **blocklistName**

outputType _____

sampteBloblUri

Answer:

Explanation:

Answer Area

curl -X PUT ’<endpoint>/content$afety/text/categorif$/learning-advice?api-version«2024-02-15-preview" \

-H "Ocp*Apfa-Subscription*Key: <api_key>" \

•H "Content-Type application/json’ \

VcategoryNMeV: ^'learning - advIceV,

VpromptV: ^ Vtext prompts about learning advice while preparing for an exaa.V,

sampleBloblUrl ˅ Vhttps://<yQur-»ture-5torege-url>/container/le»rning-»dvice. jsonlV

Question: 369

HOTSPOT

You are building an agent by using the Azure AI Agent Service.

You need to ensure that the agent can access publicly accessible data that was released during the past 90 days.

How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

```
var connectlonld ■ “blngConnectionld";
AgentsClient agentclient * projectclient.GetAgentsClient();
Var connectionList * new ToolConnectionList

     ConnectionList * ( new Toolconnection(connectionjd) }
h
```

```
var grounding • new BingGrOundingToolDefiniTIon ’ (connectionlist);
                    AzureAISearchResource
                    BingGioundingToolDefinition
                    | ToofResources _____ |
Response<Agent> agentResponse = await agentclient.CreateAgentAsync(
    model: "gpt-4o“,
    name: "my-assistant*,
    instructions: "You are a helpful assistant.*,
    ToolReSQUrceS ᵛ I ne- List<ToolDefinition> ( grounding )); metadata;
(
    ToolResources:
    | tools:                |
    Agent agent = agentResponse.Value;
```

Answer:

Explanation:

Answer Area

```
var connectionid « "blngConnectIonId";
AgentsClient agentclient * projectclient.GetAgent»Client();
Var connectIonList • new ToolConnectionList {
    ConnectionList • { new ToolConnection(connectionId) }
};
var grounding • new BingGrourwiingToolDefotion ᵛ (connectionList);
Response*Agent> agentSesponse * await agentclient.CreateAgentAsync(
    ■odel: “gpt-4o",
    name: “ry-assistant",
    instructions: "You are a helpful assistant.%
        TOOlResOUrces.' ^ new List<ToolDefinition> ( grounding });
Agent agent « agentResponse.Value;
```

Question: 370

HOTSPOT

You are building an agent by using the Semantic Kernel. The agent will use a custom plugin. You need to ensure that the agent meets the following requirements:

- The agent must use function calling.
- All functions that match the instructions must be triggered.
- All required parameters in the function must be requested by the agent if the user fails to provide them.

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

a. NOTE: Each correct selection is worth one point.

Answer Area

```
dgentKernel. CreatePluginFromFunctions ᵀ I ("Prompt instructions")}
```

CreateFunctionFromMethod

CreateFunctionFrom Prompt
CreatePluqinFromFunctions

```
new ChatCompletlonAgent;
    Name • “<agent names',
    Instructions • "ragent instructions>",
    Kernel • agentKernel,
    Arguments • new KernelArguments(
        new OpenAIPromptExecutionSettingsi ; {
FunctionchoiceBehavior« FunctionChoiceBetiaviour.Required
```

FunctionChoiceBehaviour^uto
FunctionChoiceBehaviour.None
FuncttonChoiceBehaviour.Requited

Answer:

Explanation:

Answer Area

```
agentKernel. CreatePluginFromFunctlOns T ("Prompt instructions');
new ChatCcwpletionAjent!){
    Wane - "cagent names'.
    Instructions - "cagent instructions)".
    Kernel - agentKernel,
    Arguments ■ new KernelArguments(
        new OpenAIPromptExecutionSettingsf) {
```

Funttionchoicesehavior « FunctionChoiceBehaviour,Required ˅

Question: 371

You are building a social media messaging app.
You need to identify in real time the language used in messages.

A. Azure AI Foundry Content Safety
B. Azure AI Speech
C. Azure AI Language
D. Azure AI Translator

Answer: D

Explanation:

Question: 372

HOTSPOT

You are building an agent by using the Azure Al Foundry Agent Service.
You have the following code.

```
file = project_client.agents.upload_file_andjoll(
     fili_p»ths"./d«i/filel.:jp",
     purpose=FilePurpose.AGENTS

vector store = project client.agents.create vector store and DOIK
```

Answer Area

Statements	Yes	No
The agent will reason over the uploaded file	O	
The code will create a run and check the output	J	
The code will create an agent and enable a file search	O	

Answer:

Explanation:

Answer Are*

Statements	S^s	No
The agent will reason over the uploaded file	0	
The rode will create a run and check the output		
The code Will create an agent and enable a file search		

Question: 373

HOTSPOT

You are building an agent by using the Azure Al Foundry Agent Service.

You need to ensure that the agent can access publicly accessible data that was released during the past 90 days.

How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

nMt^oy-iisiitaflt".

agent » project-Client.agents.create_ag«nt!

eodel="gpt-4o".

na«'e«"iey-assi stent".

instructions "You me a helpful assistant".

J search.definitions.

metadata-

prin' tool resources- gent.id}')

tools=

Answer:

Explanation:

Answer Area

```
                                wnnjd = *search_connection_id"

                                search * BmgGroundingTool ^ (connection_id*conn_id) with

                                proj»ct_client:
                          agent ■ project_client.agents.create_agent(

                                        i»d»li"gpt-4o“,

                                         name="»y-assistant",

agent - praject_client.agents.create_agent(

    model="gpt-4o",

    name:"my-assistant",

    instructions*"You are a helpful assistant",

     tools-              ᵛ search.definitions,


print(f"Created agent, ID; (agent.id)”)
```

Question: 374

DRAG DROP

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

You plan to deploy the model as part of a mobile 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.

Actions

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

Answer Area

Answer:

Explanation:

Actions

:: Export the model.

Answer Area

1 :: Change the model domain.
2 :: Retrain the model.
3 :: Test the model.

Question: 375

You have a 20-GB video file named File1.avi that is stored on a local drive.

You need to index File1.avi by using the Azure AI Video Indexer website. What should you do first?

A. Upload File1.avi to the Azure AI Video Indexer website.
B. Upload File1.avi to the www.youtube.com webpage.
C. Upload File1.avi to Microsoft OneDrive.
D. Upload File1.avi to an Azure Storage queue.

Answer: A

Explanation:

Question: 376

You are processing text by using the Azure AI Language service.

You need to identify music band names in the text. The solution must minimize development effort. What should you use?

A. Key phrase extraction
B. Conversational Language Understanding (CLU)
C. Entity linking
D. Custom named entity recognition (NER)

Answer: C

Explanation:

Question: 377

You are developing an app that will use the Speech and Language APIs.

You need to provision resources for the app. The solution must ensure that each service is accessed by using a single endpoint and credential.

Which type of resource should you create?

A. Azure AI Language
B. Azure AI Foundry service

C. Azure AI Speech
D. Azure AI Foundry Content Safety

Answer: B

Explanation:

Question: 378

HOTSPOT

You have an Azure subscription that contains an Azure AI Foundry Content Safety resource named resource1.

You are building an app that will analyze text by using resource1.

You need to identify text that contains hateful content.

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

NOTE: Each correct selection is worth one point.

Answer Area

Answer:

Explanation:

Answer Area

clint • COTtartl*f«tyCll>M(»M>alrrt. An^nCre^muH**?)) request * AnalyzaTaxtOptions(text?'sou aniauls ere aoro equal than others') •

Analyse text try:

response • client. analyie_text(request)

except HTtpAe^ponseFrror at e:

print('Analyte text •ailed.') raise

hate result • aaxt(itaa tor Ites in reaponse. (fIteQprm^rwbyi'i * if item,category •• Ceit Category. HATE) W MU^rtasUt

print(e'Mate seventy; (hatejesutLseventy]') *

Question: 379

HOTSPOT

You have an Azure OpenAl resource named Al1 that hosts three deployments of the GPT 3.5 model.

Each deployment is optimized for a unique workload.

You plan to deploy three apps. Each app will access Al1 by using the REST API and will use the deployment that was optimized for the app's intended workload.

You need to provide each app with access to Al1 and the appropriate deployment. The solution must ensure that only the apps can access Al1.

What should you use to provide access to Al1, and what should each app use to connect to its appropriate deployment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Provide access to All by using

An API key

Connett to the deployment by using

A bearer token

A shared access signature (SAS) token

Connect to the deployment by using:

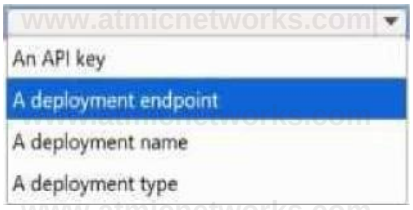

Answer:

Explanation:

Answer Area

Provide access to All by using A bearer token

Connect to the deployment by using: A deployment endpoint

Provide access to Al1 by using: A bearer token

Connect to the deployment by using: A deployment name

To ensure only your apps can access the Azure OpenAI resource, authenticate with Microsoft Entra ID (Azure AD) and obtain a Bearer token (ideally via managed identity). This avoids sharing long-lived API keys and lets you scope access by app identity and Azure RBAC. Microsoft's guidance recommends token-based auth for Azure OpenAI and shows how to acquire a bearer token/managed identity token for calls

When calling the REST API, Azure OpenAI requires you to specify the deployment name to select which model deployment to use. The REST route includes {deployment-id} (deployment name) in the path, e.g., .../openai/deployments/{deployment-id}/chat/completions. This is explicitly documented as a key difference from OpenAI's public API, which takes only a model.

Therefore:

Access method: Bearer token (via Entra ID/managed identity)

Deployment selector: Deployment name in the REST path

Microsoft Reference

Authenticate to Azure OpenAI with Microsoft Entra ID / Managed identity (bearer token): guidance and how-to

Azure OpenAI request header recommends token-based authentication (Bearer).

Azure OpenAI requires deployment name in API calls; REST path uses

Question: 380

HOTSPOT

You have an Azure subscription that contains an Azure Al Language service resource named Resourcel. You query Resourcel by running a cURL command and receive the following response.

```
"text"! "Mateo 6o«2", "category": "Person", "offset": 6, “length": 11,
"confidencescore": 1

"warnings": []

Jr
"errors": [], "modelVersion": "2023-09’01"
}
```

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

Answer Area

Statements	Yes	No

Resourcel was queried by using Personally Identifiable Information (PH) detection,	Q	
The response indicates that Resource! has low confidence in the accuracy of the identified entity.	Q	0
The request URL includes the following string https://resourcel.cognitiveservices.azure.coa/contentsafety	O	0

Answer:

Explanation:

Answer Area

Statements	Yes	No
Resource! was queried by using Personally Identifiable Information (PII)detection.	•	O
The response indicates that Resource! has low confidence in the accuracy of theidentified entity.	0	®
The request URL includes the following string https://resourcel.cognitiveservices.azure.com/contentsafety.	O	®

Question: 381

You are developing a text processing solution.

You have the following function.

```
static void GetKeyword5(TextAnalyticsClient textAnalyticsClient, string text)
    var response • textAnalyticsClient.SecognizeEntities (text);
    Console.WriteLine("Key words:");

    foreach (CategorizedEntity entity in response.Value)
        Consol#.wrltetine(S"\t(entity.Text}');
    }
}
```

You call the function and use the following string as the second argument Our tour of London included a visit to Buckingham Palace What will be the output of the function?

A. Our tour of London included a visit to Buckingham Palace

B. London and Tour only

C. Tour and visit only

D. London and Buckingham Palace only

Answer: D

Explanation:

• The extraction of key frames from videos

You run the following command.

```
curl -1 x POST "https://contoso-at.openai.azure.co»/contentunderstanding/analy2ers/analyzerl?api-version.2024-12 01-preview* \
  ■H "Ocp Apl» Subscription Key: 285d057el5e6419eaaSdl7Sb9291eccl" \
  -H 'Content-Type: application/json* \
  -d '(\"url\":\"https://www.contoso.con/videos/videol.»p4\")'
```

You receive the following output

```
202 Accepted
Operation location: https://contoso
ai. openai. azure, con/contentunderstanding/analyzers/analyzerl/resuits/zcf fcb70f71e.ilJeb2193f6653d8028d?api
version-2024-12-01-preview
```

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