

Microsoft

AI-102 Exam

Designing and Implementing a Microsoft Azure AI Solution

Product Questions: 114

Version: 8.0

Topic 1, Wide World Importers

Case study

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

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

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

To start the case study

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

Overview

Existing Environment

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

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

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

Applications

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

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

Azure Resources

You have the following resources:

An Azure Active Directory (Azure AD) tenant

The tenant supports internal authentication.

All employees belong to a group named AllUsers.

Senior managers belong to a group named LeadershipTeam.

An Azure Functions resource

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

An Azure Cosmos DB account

The account uses the Core (SQL) API.

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

An Azure Storage account

The account contains blob containers for assets related to products.

The assets include images, videos, and PDFs.

An Azure Cognitive Services resource named wwics

A Video Indexer resource named wwivi

Requirements

Business Goals

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

Planned Changes

Wide World Importers plans to start the following projects:

A product creation project: Help employees create accessible and multilingual product entries, while expediting product entry creation.

A smart e-commerce project: Implement an Azure Cognitive Search solution to display products for customers to browse.

A shopping on-the-go project: Build a chatbot that can be integrated into smart speakers to support customers.

Business Requirements

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

Provide a multilingual customer experience that supports English, Spanish, and Portuguese.
Whenever possible, scale based on transaction volumes to ensure consistent performance.
Minimize costs.

Governance and Security Requirements

Wide World Importers identifies the following governance and security requirements:

Data storage and processing must occur in datacenters located in the United States.
Azure Cognitive Services must be inaccessible directly from the internet.

Accessibility Requirements

Wide World Importers identifies the following accessibility requirements:

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

Product Creation Requirements

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

Minimize how long it takes for employees to create products and add assets.
Remove the need for manual translations.

Smart E-Commerce Requirements

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

Ensure that the Cognitive Search solution meets a Service Level Agreement (SLA) of 99.9% availability for searches and index writes.
Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products.
Support autocompletion and autosuggestion based on all product name variants.
Store all raw insight data that was generated, so the data can be processed later.
Update the stock level field in the product index immediately upon changes.
Update the product index hourly.

Shopping On-the-Go Requirements

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

Answer common questions.

Support interactions in English, Spanish, and Portuguese.

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

Provide all employees with the ability to edit Q&As. Only senior managers must be able to publish updates.

Support purchases by providing information about relevant products to customers. Product displays must include images and warnings when stock levels are low or out of stock.

Product JSON Sample

You have the following JSON sample for a product.

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

Question: 1

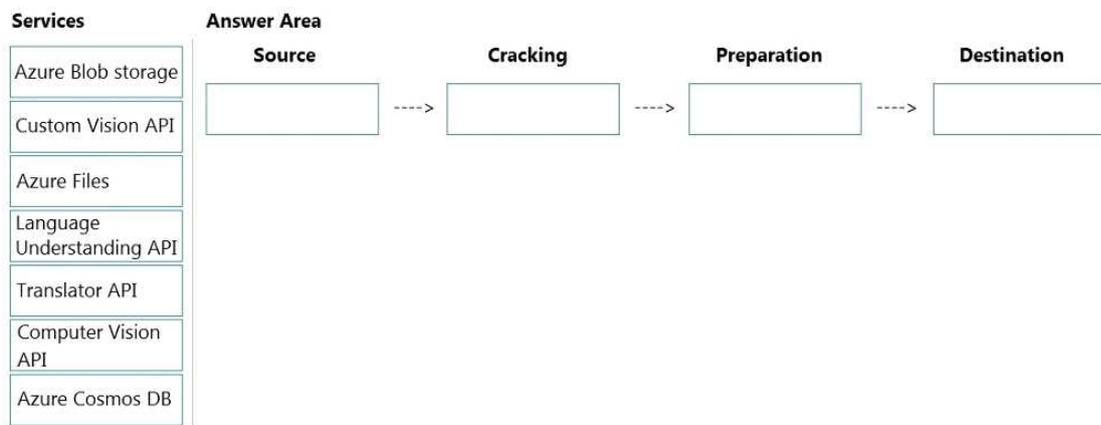
DRAG DROP

You are developing the smart e-commerce project.

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

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

NOTE: Each correct selection is worth one point.



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**Answer Area**

Index the video by using the Video Indexer API.

Upload the video to blob storage.

Analyze the video by using the Computer Vision API.

Extract the transcript from Microsoft Stream.

Send the transcript to the Language Understanding API as an utterance.

Extract the transcript from the Video Indexer API.

Translate the transcript by using the Translator API.

Upload the video to file storage.

Answer:

Explanation:

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

Upload the video to blob storage.

Index the video by using the Video Indexer API.

Extract the transcript from the Video Indexer API.

Translate the transcript by using the Translator API.

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

		<code>?api-version=3.0&to=es&to=pt</code>
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 autocomplete 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 autocomplete 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

AllUsers:

	▼
Cognitive Service User	
Contributor	
Owner	
QnA Maker Editor	
QnA Maker Read	

LeadershipTeam:

	▼
Cognitive Service User	
Contributor	
Owner	
QnA Maker Editor	
QnA Maker Read	

Answer:

Explanation:

Answer Area

AllUsers:

▼
Cognitive Service User
Contributor
Owner
QnA Maker Editor
QnA Maker Read

LeadershipTeam:

▼
Cognitive Service User
Contributor
Owner
QnA Maker Editor
QnA Maker Read

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: 6

HOTSPOT

You need to develop code to upload images for the product creation project. The solution must meet the accessibility requirements.

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

NOTE: Each correct selection is worth one point.

```

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

Dictionary
stream
string

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

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

Answer:

Explanation:

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

Reference:

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

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": "${
      if(language == 'en', 'en', name)
      name
      name.en
      name[language]
    },
  {
    "type": "TextBlock",
    "$when": "${stockLevel != 'OK'}"
    "$when": "${stockLevel == 'OK'}"
    "$when": "${stockLevel.OK}"
    color : Attention
  },
  {
    "type": "Image",
    "url": "${image.uri}",
    "size": "Medium",
    "altText": "${
      image.altText.en
      image.altText.language
      image.altText["language"]
      image.altText[language]
    }"
  }
]
}

```

Explanation:

Answer:

```

version": "1.3",
"body": [
  {
    "type": "TextBlock",
    "size": "Medium",
    "weight": "Bolder",
    "text": "${
      if(language == 'en', 'en', name)
      name
      name.en
      name[language]
    }"
  },
  {
    "type": "TextBlock",
    "${when": "${stockLevel != 'OK'}"
    "${when": "${stockLevel == 'OK'}"
    "${when": "${stockLevel.OK}"
    color : Attention
  },
  {
    "type": "Image",
    "url": "${image.uri}",
    "size": "Medium",
    "altText": "${
      image.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 roles shown in the following table.

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

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-[Level]-[Role]
- [Level]-[Role]

Intellectual Property

Contoso has the intellectual property shown in the following table.

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

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

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 Form Recognizer API for analysis, and the API returns the following JSON.

```
"documentResults":[
  {
    "docType":"prebuilt:receipt",
    "pageRange":[
      1,
      1
    ],
    "fields":{
      "ReceiptType":{
        "type":"string",
        "valueString":"Itemized",
        "confidence":0.672
      },
      "MerchantName":{
        "type":"string",
        "valueString":"Tailwind",
        "text":"Tailwind",
        "boundingBox":[],
        "page":1,
        "confidence":0.913,
        "elements":[
          "#/readResults/0/lines/0/words/0"
        ]
      }
    }
  },
  ...
]
```

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

- A. `documentResults.docType == "prebuilt:receipt"`
- B. `documentResults.fields.".confidence < 0.7`
- C. `documentResults.fields.ReceiptType.confidence > 0.7`
- D. `documentResults.fields.MerchantName.confidence < 0.7`

Answer: C

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

```

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

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: [Error! Hyperlink reference not valid.](#)

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

Management-Accountants	<div style="border: 1px solid black; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid black; margin-bottom: 5px;">▼</div> Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor </div>
Consultant-Accountants	<div style="border: 1px solid black; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid black; margin-bottom: 5px;">▼</div> Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor </div>
Agent-CustomerServices	<div style="border: 1px solid black; padding: 5px;"> <div style="text-align: right; border-bottom: 1px solid black; margin-bottom: 5px;">▼</div> Owner Contributor Cognitive Services User Cognitive Services QnA Maker Read Cognitive Services QnA Maker Editor </div>

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. 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.

<p>Values</p> <ul style="list-style-type: none"> AddPhraseListAsync Phraselist PhraselistCreateObject Phrases SavePhraselistAsync UploadPhraseListAsync 	<p>Answer Area</p> <pre>var phraselistId = await client.Features. <input type="text"/> (appId, versionId, new <input type="text"/> { EnabledForAllModels = false, IsExchangeable = true, Name = "PL1", Phrases = "item1,item2,item3,item4,item5" });</pre>
--	---

Answer:

Explanation:

Values

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

Answer Area

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

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
V1.1	2020-10-01	None
V1.0	2020-09-01	2020-09-15

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

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

Actions

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

Export the model by using the Export as JSON option.

Select v1.1 of app1.

Run a container and mount the model file.

Select v1.0 of app1.

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

Select v1.2 of app1.

Answer Area

Answer:

Explanation:

Actions

- Run a container that has `version` set as an environment variable.
- Export the model by using the Export as JSON option.
- Select v1.1 of app1.
- Run a container and mount the model file.
- Select v1.0 of app1.
- Export the model by using the Export for containers (GZIP) option.
- Select v1.2 of app1.

Answer Area

- Export the model by using the Export for containers (GZIP) option.
- Select v1.1 of app1.
- Run a container and mount the model file.

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 data. It is used to find the most appropriate answer for any input from your custom knowledge base (KB) of information.

Text Analytics: Mine insights in unstructured text using natural language processing (NLP)—no machine learning expertise required. Gain a deeper understanding of customer opinions with sentiment analysis. The Language Detection feature of the Azure Text Analytics REST API evaluates text input

Incorrect Answers:

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

Reference:

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

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

Question: 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

```

https://management.azure.com/subscriptions/xxxxxxx-xxxx-
PATCH
POST
PUT
xxxx-xxxx-
xxxxxxxxxxxx/resourceGroups/RG1/providers/Microsoft.CognitiveServices/
accounts/CS1?api-version=2017-04-18
{
  "location": "West US",
  "kind": "
    CognitiveServices
    ComputerVision
    TextAnalytics
  ",
  "sku": {
    "name": "S0"
  },
  "properties": {},
  "identity": {
    "type": "SystemAssigned"
  }
}

```

Answer:

Explanation:

Answer Area

PATCH
POST
PUT

xxxx-xxxx-

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

{
 "location": "West US",
 "kind": " ",
 "sku": {
 "name": "S0"
 },
 "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.com
mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment

```
Eula=accept \
```

Billing=
▼
\

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

```
ApiKey=xxxxxxxxxxxxxxxxxxxx
```

Answer:

Explanation:

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

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. Error! Hyperlink reference not valid.
```

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

{ENDPOINT_URI} is the endpoint for accessing the Text Analytics API: [Error! Hyperlink reference not valid.](#) 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 create_resource(CognitiveServicesManagementClient client, string
resource_name, string kind, string account_tier, string location)
{
    CognitiveServicesAccount parameters =
        new CognitiveServicesAccount(null, null, kind, location, resource_name,
new CognitiveServicesAccountProperties(), new Sku(account_tier));
    var result = client.Accounts.Create(resource_group_name, account_tier,
parameters);
}
```

You need to call the method to create a free Azure resource in the West US Azure region. The

resource will be used to generate captions of images automatically.

Which code should you use?

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

Answer: A

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: D

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

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

NOTE: Each correct selection is worth one point.

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

Answer: A, D, E

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.

Values

-
-
-
-

Answer Area

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

Answer:

Explanation:

Values

-
-
-
-

Answer Area

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

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.

Values

- detect
- findsimilar
- group
- identify
- matchFace
- matchPerson
- verify

Answer Area

```
POST {Endpoint}/face/v1.0/ 
Request Body
{
  "faceId": "c5c24a82-6845-4031-9d5d-978df9175426",
  "largeFaceListId": "sample_list",
  "largeFaceListId": "sample_list",
  "maxNumOfCandidatesReturned": 10,
  "mode": "  "
}
```

Answer:

Explanation:

Box 1: findsimilar

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

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

NOTE: Each correct selection is worth one point.

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

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))
            {
                
                await faceClient.PersonGroupPerson.
                    
            }
        (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 TaskAnalyzeImage(ComputerVisionClient client, string localImage)
{
    List<VisualFeatureTypes> features = new List<VisualFeatureTypes>()
    {
        VisualFeatureTypes.Description,
        VisualFeatureTypes.Tags,
    };
    using (Stream imageStream = File.OpenRead(localImage))
    {
        try
        {
            ImageAnalysis results = await client.AnalyzeImageInStreamAsync(imageStream, features);

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

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

Answer:

Explanation:

Box 1: No

Box 2: Yes

Box 3: Yes

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

    ReadOperationResult results;

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

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

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

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

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

NOTE: Each correct selection is worth one point.

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

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: CDF

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**Answer Area**

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:

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

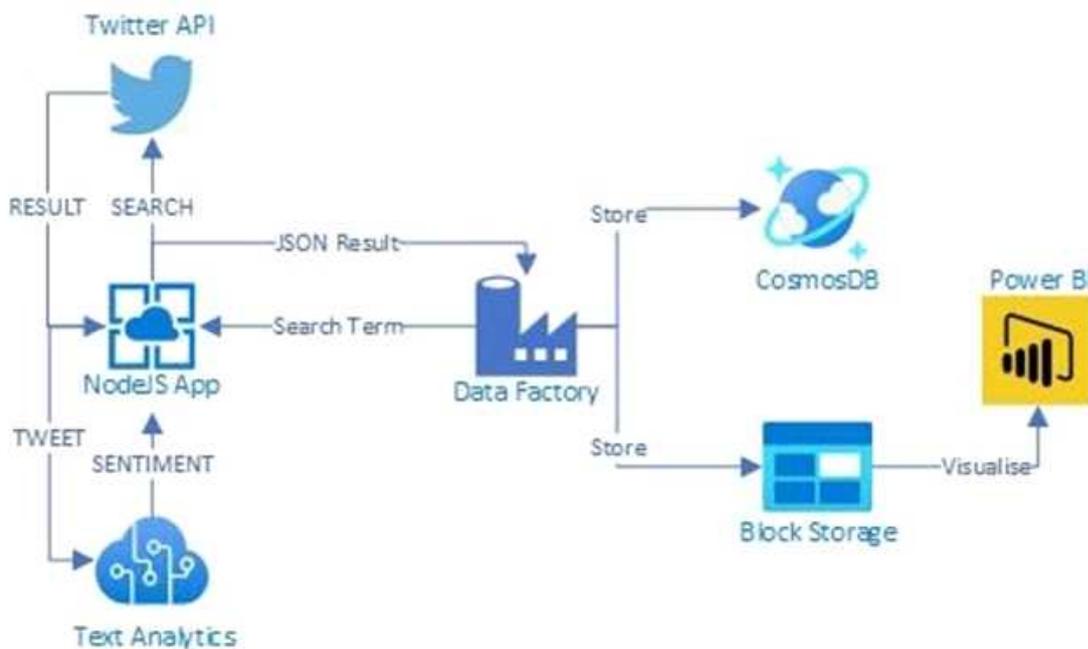
D. Form Recognizer

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: B

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

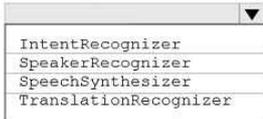
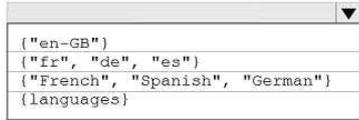
```
static async Task TranslateSpeechAsync()
{
    var config = SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");

    var lang = new List<string>
    {
        "en-GB"
    };

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

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

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



Answer:

Explanation:

Answer Area

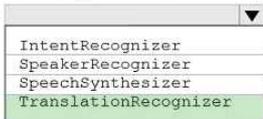
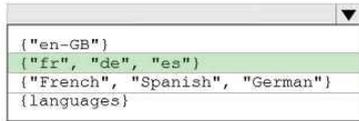
```
static async Task TranslateSpeechAsync()
{
    var config = SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");

    var lang = new List<string>
    {
        "en-GB"
    };

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

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

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



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/stt-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($"{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.	<input type="radio"/>	<input type="radio"/>
The output will contain the following words: the, cat, sat, on, and mat.	<input type="radio"/>	<input type="radio"/>
The output will contain the confidence level for key phrases.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

Answer Area

Statements	Yes	No
The call will output key phrases from the input string to the console.	<input checked="" type="radio"/>	<input type="radio"/>
The output will contain the following words: the, cat, sat, on, and mat.	<input type="radio"/>	<input checked="" type="radio"/>
The output will contain the confidence level for key phrases.	<input type="radio"/>	<input checked="" type="radio"/>

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:

	▼
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/*"
    }
  ],
  "outputs": [
    {
      "name": "companyDescription",
    }
  ]
}
```

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
CompanyDescription is available for indexing.	<input type="radio"/>	<input type="radio"/>
The definition calls a web API as part of the enrichment process.	<input type="radio"/>	<input type="radio"/>
The enrichment step is called only for the first organization under "/document/organizations".	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

Answer Area

Statements	Yes	No
CompanyDescription is available for indexing.	<input checked="" type="radio"/>	<input type="radio"/>
The definition calls a web API as part of the enrichment process.	<input checked="" type="radio"/>	<input type="radio"/>
The enrichment step is called only for the first organization under "/document/organizations".	<input type="radio"/>	<input checked="" type="radio"/>

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:

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

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

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

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

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**

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

Answer Area

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:

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.

The screenshot shows a dialog flow for 'AskForName > BeginDialog > Text'. The flow starts with a 'BeginDialog' event, followed by a 'Bot Asks (Text)' step with the prompt 'What is your name?'. This leads to a 'User Input (Text)' step where 'user.name = Input(Text)'. The flow then continues to a 'Bot Asks (Number)' step with the prompt 'Hello \$(user.name), how old are you?'. This leads to a 'User Input (Number)' step where 'user.age = Input(Number)'. The right-hand pane is titled 'Prompt for text' and shows configuration for 'Text input'. It includes a 'Property' dropdown set to 'string' with the value 'user.name', an 'Output Format' dropdown set to 'string' with the example 'ex. =toUpper(this.value), \$(toUpper(this.value))', and a 'Value' dropdown set to 'expression' with the value 'fx = coalesce(@user.Name, @personName)'. The 'Expected responses' section shows the intent '#TextInput_Response_GH5FTe'.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

Answer:

Explanation:

Answer Area

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

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/coalescefuction>

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, ITurnContext<IConversationUpdateActivity> turnContext,
Cancellation token 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.	<input type="radio"/>	<input type="radio"/>
When a new user joins the conversation, the existing users in the conversation will see the chatbot greeting.	<input type="radio"/>	<input type="radio"/>
OnMembersAddedAsync will be initialized when a user sends a message.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

Answer Area

Statements	Yes	No
OnMembersAddedAsync will be triggered when a user joins the conversation.	<input checked="" type="radio"/>	<input type="radio"/>
When a new user joins the conversation, the existing users in the conversation will see the chatbot greeting.	<input checked="" type="radio"/>	<input type="radio"/>
OnMembersAddedAsync will be initialized when a user sends a message.	<input type="radio"/>	<input checked="" type="radio"/>

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 <code>UserProfile</code> object in the underlying storage layer.	<input type="radio"/>	<input type="radio"/>
The code will create and maintain the <code>ConversationData</code> object in the underlying storage layer.	<input type="radio"/>	<input type="radio"/>
The <code>UserProfile</code> and <code>ConversationData</code> objects will persist when the Bot Framework runtime terminates.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

Statements	Yes	No
The code will create and maintain the <code>UserProfile</code> object in the underlying storage layer.	<input checked="" type="radio"/>	<input type="radio"/>
The code will create and maintain the <code>ConversationData</code> object in the underlying storage layer.	<input checked="" type="radio"/>	<input type="radio"/>
The <code>UserProfile</code> and <code>ConversationData</code> objects will persist when the Bot Framework runtime terminates.	<input type="radio"/>	<input checked="" type="radio"/>

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
Amsterdam



San Francisco
Amsterdam

SFO
AMS

SFO
AMS

Non-Stop

Fri, Jun 2, 2017 11:55 PM

San Francisco
Amsterdam



San Francisco
Amsterdam

SFO
AMS

SFO
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
<code>\${user.name}</code> retrieves the user name by using a prompt.	<input type="radio"/>	<input type="radio"/>
<code>Greet ()</code> is the name of the language generation template.	<input type="radio"/>	<input type="radio"/>
<code>\${Greeting () }</code> is a reference to a template in the language generation file.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

Statements	Yes	No
<code>\${user.name}</code> retrieves the user name by using a prompt.	<input type="radio"/>	<input checked="" type="radio"/>
<code>Greet ()</code> is the name of the language generation template.	<input type="radio"/>	<input checked="" type="radio"/>
<code>\${Greeting () }</code> is a reference to a template in the language generation file.	<input checked="" type="radio"/>	<input type="radio"/>

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: A

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.

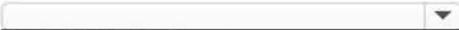
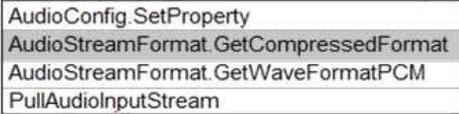
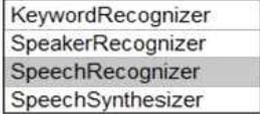
NOTE: Each correct selection is worth one point.

```

var audioFormat =  (AudioStreamContainerFormat.MP3);
var speechConfig = SpeechConfig.FromSubscription("18c51a87-3a69-47a8-aedc-a54745f708a1", "westus");
var audioConfig = AudioConfig.FromStreamInput(pushStream, audioFormat);
using (var recognizer = new  (speechConfig, audioConfig))
{
    var result = await recognizer.RecognizeOnceAsync();
    var text = result.Text;
}
    
```

Answer:

Explanation:

```
var audioFormat =  (AudioStreamContainerFormat.MP3);  
  
var speechConfig = SpeechConfig.FromSubscription("18c51a87-3a69-47a8-aedc-a54745f708a1", "westus");  
var audioConfig = AudioConfig.FromStreamInput(pushStream, audioFormat);  
using (var recognizer = new  (speechConfig, audioConfig))  
  
{  
    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.

```
curl -H "Ocp-Apim-Subscription-Key: xxx" /
-o "sample.png" -H "Content-Type: application/json" /
/vision/v3.1/
?width=100&height=100&smartCropping=true" /
-d "{\"url\": \"https://upload.litwareinc.org/litware/bicycle.jpg\"}"
```

https://api.projectoxford.ai

https://contoso1.cognitiveservices.azure.com

https://westus.api.cognitive.microsoft.com

areaOfInterest

detect

generateThumbnail

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.

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

Answer:

Explanation:

Cognitive Insights Widget

https://www.videoindexer.ai/embed/insights/<accountId>/<videoId>/?widgets= controls=

Player Widget

https://www.videoindexer.ai/embed/player/<accountId>/<videoId>/? showcaptions= captions=

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

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

Answer Area

⬅
➡

⬆
⬇

Answer:

Explanation:

In user want to change to deploy offline 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("Ocp-Apim-Subscription-Key", subscription_key);
    var endpointURI = $"https://westus.api.cognitive.microsoft.com/face/v1.0/persongroups/{personGroupId}/persons/{personId}/persistedFaces";
    HttpResponseMessage response;
    var body = "{ \"url\": \"" + imageURI + "\"}";
    var content = new StringContent(body, Encoding.UTF8, "application/json");
    var response = await client.PutAsync(endpointURI, content);
}
```

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

NOTE: Each correct selection is worth one point.

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

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.

Actions	Answer Area
Use the <code>ExportProject</code> endpoint on acvdev.	
Use the <code>GetProjects</code> endpoint on acvdev.	
Use the <code>ImportProject</code> endpoint on acvprod.	⬅
Use the <code>ExportIteration</code> endpoint on acvdev.	➡
Use the <code>GetIterations</code> endpoint on acvdev.	
Use the <code>UpdateProject</code> endpoint on acvprod.	⬆
	⬇

Answer:

Explanation:

Use the `GetProjects` endpoint on `acvdev`.

Use the `ExportProject` endpoint on `acvdev`.

Use the `ImportProject` endpoint on `acvprod`.

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.

Actions	Answer Area
<input type="text" value="Upload the images by category."/>	
<input type="text" value="Get suggested tags."/>	
<input type="text" value="Upload all the images."/>	⬅️
<input type="text" value="Group the images locally into category folders."/>	➡️
<input type="text" value="Review the suggestions and confirm the tags."/>	⬆️
<input type="text" value="Tag the images manually."/>	⬆️

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: A

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 WebJob that will create knowledge bases from an array of URLs.

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

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

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

NOTE: Each correct selection is worth one point.

- A. Create a list of FileDTO objects that represents data from the WebJob.

- B. Call the client. 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.

```

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


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

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

Answer:

Explanation:

```

. . .
var endpoint =  ;

var apiKey = "FF956C68883821B38691ABD200A4C606";
var text = getTextToBeTranslated();
var body = '[{"Text":"' + text + '"}]';
var client = new HttpClient();
client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", apiKey);

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

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

```

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: B

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. lceyPhrases
- 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.

Actions		Answer Area
Add a new query key.		
Regenerate the secondary admin key.		
Change the app to use the secondary admin key.	⊞	⊞
Change the app to use the new key.	⊞	⊞
Regenerate the primary admin key.		
Delete the compromised key.		

Answer:

Explanation:

Add a new query key.

Change the app to use the new key.

Delete the compromised key.

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 {@DirectionalReference=bottom left}
- choose {@DirectionalReference=top right}
- i like {@DirectionalReference=left} one

## SelectNone
- none

@m1 DirectionalReference
@prebuilt personName
```

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

NOTE: Each correct selection is worth one point.

SelectItem is [answer choice].

a domain
an entity
an intent
an utterance

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

a domain
an entity
an intent
an utterance

Answer:

Explanation:

SelectItem is [answer choice].

a domain
an entity
an intent
an utterance

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

a domain
an entity
an intent
an utterance

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.

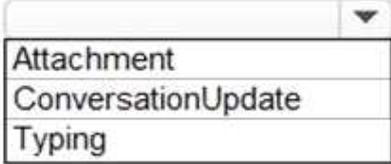
How should you complete the .chat file? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

user=User1

bot=watchbot

user: I want a new watch.

bot: [][Delay=3000]

bot: I can help you with that! Let me see what I can find.

bot: Here's what I found.

bot:

[AttachmentLayout= ]

[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]

[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]

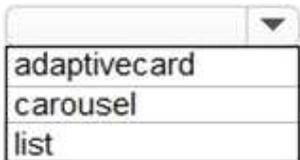
user: I like the first one.

bot: Sure, pulling up more information.

bot: [Attachment=cards\watchProfileCard.json

user: That's nice! Thank you.

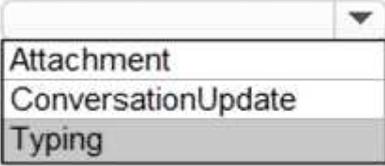
bot: Sure, you are most welcome!

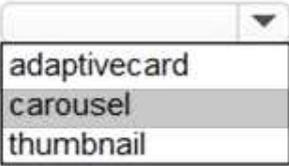
]

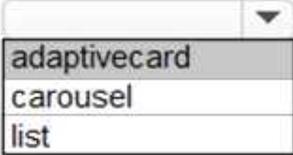
Answer:

Explanation:

```

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

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

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

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. tenanId
- C. appld
- D. objecId
- 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 bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

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

Answer:

Explanation:

Detect the incoming language:	<input type="text" value="Text Analytics"/>
Respond in the callers' own language:	<input type="text" value="Translator"/>

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

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(receiptUri)).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

Answer Area

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



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.

Project Types: ▼

Classification
Object Detection

Classification Types: ▼

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

Domains: ▼

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

Answer:

Explanation:

Project Types: ▼

Classification
Object Detection

Classification Types: ▼

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

Domains: ▼

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

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.	<input type="radio"/>	<input type="radio"/>
The container logging provider will write log data.	<input type="radio"/>	<input type="radio"/>
Going to http://localhost:5000/swagger will provide the details to access the documentation for the available endpoints.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

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.	<input checked="" type="radio"/>	<input type="radio"/>
The container logging provider will write log data.	<input checked="" type="radio"/>	<input type="radio"/>
Going to http://localhost:5000/swagger will provide the details to access the documentation for the available endpoints.	<input checked="" type="radio"/>	<input type="radio"/>

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.

Entity Types**Answer Area**

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

Answer:

Explanation:

Paris:	GeographyV2
email@domain.com:	Email
2 audit business:	Machine learned

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 screenshot displays the Bot Framework Designer interface. On the left, a dialog flow is shown for 'GetUserDetails > BeginDialog > Text'. The flow starts with a 'BeginDialog' event, followed by a 'Bot Asks (Text)' activity with the prompt 'What is your name?'. This leads to a 'User input (Text)' activity with the code '(SCOPE).name = Input(Text)'. On the right, the configuration pane for the 'Text Input' control is visible. It has three tabs: 'Bot response', 'User input', and 'Other'. The 'User input' tab is selected. Under this tab, there are three properties: 'Property' (set to '(SCOPE).name'), 'Output format' (set to 'string'), and 'Value' (set to 'string'). Below these is the 'Expected responses' section, which shows the intent '#TextInput_Response_FuyF4'.

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

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

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

Which scope should you assign to name?

- A. dialog
- B. user
- C. turn
- D. conversation

Answer: A

Explanation:

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

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.

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:

For the knowledge base, select Show active learning suggestions.

Approve and reject suggestions.

Save and train the knowledge base.

Publish the knowledge base.

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	PersonType	0.45
Space Needle	Location	0.38
Trip	Event	0.78
Seattle	Location	0.78
Last week	DateTime	0.80

Which Text Analytics API is used to analyze the text?

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

Answer: C

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.

Answer Area

```

...
api_key = "FF956C68B83B21B38691ABD200A4C606"
text = get_text_to_be_translated()
headers = {
    'Content-Type': 'application/json',
    'Ocp-Apim-Subscription-Key': api_key
}
body = {
    'Text': text
}
conn = httplib.HTTPSConnection
conn.request("POST",
response = conn.getresponse()
    conn.request("POST",
response = conn.getresponse()
response_data = response.r
...

```

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": "unrelateddoclayout",
          "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]

- not be saved.
- be saved to a blob container.
- be saved to file storage.
- be saved to an Azure Data lake.

Answer:

Explanation:

Answer Area

There will be [answer choice].

Images will [answer choice]

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.

```

static async Task DetectFaces(string imagePath)
{
    HttpClient client = new HttpClient();
    DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", subscriptionKey);
    string requestParameter = "detectionModel=detection_01&returnFaceId=true&returnFaceLandmarks=false";
    string uri = endpoint + "/face/v1.0/detect?" + requestParameters;
    HttpResponseMessage response;
    byte[] byteData = GetImagesAsByteArray(imageFilePath);
    using (ByteArrayContent content = new ByteArrayContent(byteData))
    {
        Headers.ContentType = new MediaTypeHeaderValue("application/octet-stream");
        response = await PostAsync(uri, content);
        string contentString = await Content.ReadAsStringAsync();
        ProcessDetection(contentString);
    }
}

```

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.

```

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

```

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

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

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

Answer: A

Explanation:

Thank You for Purchasing AI-102 PDF