

Microsoft.AI-102.v2026-05-27.q184

Exam Code:	AI-102
Exam Name:	Designing and Implementing a Microsoft Azure AI Solution
Certification Provider:	Microsoft
Free Question Number:	184
Version:	v2026-05-27
# of views:	159
# of Questions views:	1840
https://www.freepdfdumps.com/Microsoft.AI-102.v2026-05-27.q184.html	

NEW QUESTION: 1

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

Winch attributes should you enable for the field?

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

Answer: (SHOW ANSWER)

Searchable # Field content is tokenized and added to the full-text index (supports free-text queries).

Retrievable # Field values can be included in search results.

Sortable # Used for ordering results, not required here.

Facetable # Used for aggregation/filtering by categories, not required here.

Filterable # Used for structured filters (e.g., WHERE conditions), not required here.

Reference: Azure Cognitive Search field attributes

The Answer: A

NEW QUESTION: 2

You have a video recording of a company meeting that includes presentations by multiple presenters.

You plan to use Azure AI Video Indexer to process the video and generate a separate text file for each presenter ' s presentation.

You need to identify each presenter in the video and attribute each text file to a presenter.

Which insight should you extract?

- A. speaker indexing
- B. keyframe detection
- C. topic indexing
- D. face detection

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 3

You have an Azure Cognitive Search instance that indexes purchase orders by using Form Recognizer You need to analyze the extracted information by using Microsoft Power BI. The solution must minimize development effort.

What should you add to the indexer?

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

Answer: A ([LEAVE A REPLY](#))

- * To connect extracted structured data directly to Power BI, you must project the enriched data into Azure Table Storage using a table projection.
- * This allows Power BI to query and visualize data with minimal coding.
- * Object projection # JSON objects in Blob storage, useful for apps, not BI.
- * File projection # Saves images or files, not structured data.
- * Projection group # A grouping concept but not directly queried by BI.

The answer: A

Reference: Knowledge store projections in Azure Cognitive Search

NEW QUESTION: 4

You are planning the product creation project.

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

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

NOTE: Each correct selection is worth one point.

Answer Area

api.cognitive.microsofttranslator.com

api-nam.cognitive.microsofttranslator.com

westus.tts.speech.microsoft.com

wwics.cognitiveservices.azure.com/translator

/detect

/languages

/text-to-speech

/translate

?api-version=3.0&to=es&to=pt

Answer:

Answer Area

Microsoft

api.cognitive.microsofttranslator.com

api-nam.cognitive.microsofttranslator.com

westus.tts.speech.microsoft.com

wwics.cognitiveservices.azure.com/translator

/detect

/languages

/text-to-speech

/translate

?api-version=3.0&to=es&to=pt

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>

NEW QUESTION: 5

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:

Actions	Answer Area
Add prebuilt domain models as required.	Enable active learning.
Validate the utterances logged for review and modify the model.	Validate the utterances logged for review and modify the model.
Migrate authoring to an Azure resource authoring key.	
Enable active learning.	Train and republish the Language Understanding model.
Enable log collection by using Log Analytics.	
Train and republish the Language Understanding model.	

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>

NEW QUESTION: 6

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

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

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

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

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

< product > ?

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

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

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

Does this meet the goal?

A. Yes

B. No

Answer: (SHOW ANSWER)

Creating an entity for price is a feature of Conversational Language Understanding (CLU/LUIS), not Question Answering. Custom question answering does not use entities for extraction or intent resolution; it matches user utterances to Q/A pairs using alternates, synonyms, and ranking. Therefore, defining an "entity for price" would not make the chatbot correctly answer "How much does < product > cost?". The correct fix is to author additional alternate phrasings (as in Q266).

References

Custom question answering capabilities (no entity modeling).

<https://learn.microsoft.com/azure/ai-services/language-service/question-answering/overview> CLU entities vs. QnA authoring-separate features and workflows.

<https://learn.microsoft.com/azure/ai-services/language-service/conversational-language-understanding/overview>

NEW QUESTION: 7

You are building a chatbot.

You need to use the Content Moderator API to identify aggressive and sexually explicit language.

Which three settings should you configure? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

Image

Text

Content Moderator - Moderate

Text - Screen

The operation detects profanity in more than 100 languages and match against custom and shared blacklists.

Host

Name

uksouth.api.cognitive.microsof

Query parameters

autocorrect

Value

Remove parameter

pii

Value

Remove parameter

listid

Value

Remove parameter

classify

false

Remove parameter

language

Value

Remove parameter

Add parameter

Headers

Content-Type

text/plain

Remove header

Ocp-Apim-Subscription-Key

Value

Add header



Microsoft

Answer:

Microsoft

Cognitive Services

Image

Text

Content Moderator - Moderate

Text - Screen

The operation detects profanity in more than 100 languages and match against custom and shared blacklists.

Host

Name

uksouth.api.cognitive.microsr

Query parameters

autocorrect

Value

Remove parameter

PII

Value

Remove parameter

listid

Value

Remove parameter

classify

false

Remove parameter

language

Value

Remove parameter

Add parameter

Headers

Content-Type

text/plain

Remove header

Ocp-Apim-Subscription-Key

Value

Add header

Microsoft

Explanation:

Answer Area



Content Moderator - Moderate

Text - Screen

The operation detects profanity in more than 100 languages and match against custom and shared blacklists.

Host

Name	[resource name].cognitiveser
Resource Name	

Query parameters

autocorrect	Value	✕ Remove parameter
Pii	Value	✕ Remove parameter
listid	Value	✕ Remove parameter
classify	Value	✕ Remove parameter
language	Value	✕ Remove parameter
+ Add parameter		

To detect aggressive language (e.g., profanity/abuse) and sexually explicit content in text using the Content Moderator Text - Screen API:

classify = true

Enables machine-assisted text classification that returns category scores (e.g., sexually explicit/suggestive categories). This is required to identify sexual content.

language = < code > (for example, eng)

Specifies the language so profanity/abusive terms are detected correctly. The profanity detection in Text- Screen is language dependent.

autocorrect = true

Helps normalize intentionally misspelled or obfuscated words so that profanity or aggressive terms are still detected (e.g., "f@#\$" # "fuck").

Other parameters in the blade are not necessary for this goal:

Pii targets personal data extraction, not aggression/sexuality.

listId is for matching against custom blocklists.

Headers like Ocp-Apim-Subscription-Key and Content-Type are required to call the API, but they are not the content-analysis settings that control detection behavior asked for in this question.

Microsoft Azure AI References

Azure AI Content Moderator - Screen text API parameters (autocorrect, language, classify) and category outputs.
Azure AI Content Moderator - Text moderation concepts: profanity detection and classification categories.

NEW QUESTION: 8

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: ([SHOW ANSWER](#))

This question is about selecting valid training files for a custom Form Recognizer model in Azure AI Document Intelligence (formerly Form Recognizer).

Supported File Formats for Training

According to Microsoft documentation, Form Recognizer supports the following for training and analysis:

PDF (# 500 pages, # 50 MB for training)

JPEG/JPG (# 50 MB)

PNG (# 50 MB)

TIFF (# 50 MB)

Unsupported:

MP4 # Video not supported.

GIF # Not supported.

Maximum File Size Rules (for training input files):

Images (JPG, PNG, TIFF): # 50 MB

PDFs: # 50 MB

Note: 500 pages limit applies to PDFs.

File-by-File Analysis

File1 - PDF - 20 MB

Type supported.

Size within limit (< 50 MB).

Valid #

File2 - MP4 - 100 MB

MP4 is not supported.

Invalid #

File3 - JPG - 20 MB

Type supported.

Size within limit (< 50 MB).

Valid #

File4 - PDF - 100 MB

Type supported.

Size exceeds the 50 MB maximum limit.

Invalid #

File5 - GIF - 1 MB

GIF format not supported.

Invalid #

File6 - JPG - 40 MB

Type supported.

Size within limit (< 50 MB).

Valid #

The Answer:

A). File1, C. File3, F. File6

Microsoft References

Azure AI Document Intelligence (Form Recognizer) - Supported file formats Train a custom model with Azure Form Recognizer

NEW QUESTION: 9

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.

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

Answer:

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

Explanation:

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

Code Recap

```
static void GetKeyPhrases ( TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.ExtractKeyPhrases(text);
    Console.WriteLine( " Key phrases: " );
    foreach ( string keyphrase in response.Value)
    {
        Console.WriteLine( $ " \t {keyphrase} " );
    }
}
```

Then called with:

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

Statement Analysis

"The call will output key phrases from the input string to the console." Yes.

The method calls ExtractKeyPhrases , iterates over response.Value , and writes each key phrase to the console.

"The output will contain the following words: the, cat, sat, on, and mat." No.

The Text Analytics key phrase extraction returns meaningful phrases, not stop words or every word.
Likely result: " cat " , " mat " , maybe " sat " depending on language model. But not " the " , " on " .
"The output will contain the confidence level for key phrases."
No.
The ExtractKeyPhrases method in Text Analytics API returns only key phrases as strings, not confidence scores.

Final Answer
The call will output key phrases from the input string to the console # Yes The output will contain the following words: the, cat, sat, on, and mat # No The output will contain the confidence level for key phrases # No
Microsoft References Azure Text Analytics - Key Phrase Extraction TextAnalyticsClient.ExtractKeyPhrases

NEW QUESTION: 10

You are building an app that will perform speech translation by using the Azure ai Language service.
You need to ensure that the language input to the app is supported.
How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE; Each correct selection is worth one point.

Answer Area

Microsoft

```
...
public static void CheckLanguage(TextTranslationClient client, string language)
{
    Response<GetLanguagesResult> languagesResponse = await
        client. GetLanguagesAsync (scope:"translation").ConfigureAwait(false);
    GetLanguagesResult languages = languagesResponse.Value;
    if (languages.Translation. ContainsKey (language))
    {
        return true;
    }
    else
    {
        Console.WriteLine($"Sorry, {language} is not a supported language.");
        return false;
    }
}
```

GetLanguagesAsync
GetLanguages
GetLanguagesAsync
TranslationRecognizer

ContainsKey
Contains
ContainsKey
GetValueOrDefault
Where

Answer:

```

...
public static void CheckLanguage(TextTranslationClient client, string language)
{
    Response<GetLanguagesResult> languagesResponse = await
        client. GetLanguagesAsync (scope:"translation").ConfigureAwait(false);
    GetLanguagesResult languages = languagesResponse.Value;
    if (languages.Translation. ContainsKey (language))
    {
        return true;
    }
    else
    {
        Console.WriteLine($"Sorry, {language} is not a supported language.");
        return false;
    }
}

```

Explanation:

Answer Area

```

...
public static void CheckLanguage(TextTranslationClient client, string language)
{
    Response<GetLanguagesResult> languagesResponse = await
        client. GetLanguagesAsync (scope:"translation").ConfigureAwait(false);
    GetLanguagesResult languages = languagesResponse.Value;
    if (languages.Translation. ContainsKey (language))
    {
        return true;
    }
    else
    {
        Console.WriteLine($"Sorry, {language} is not a supported language.");
        return false;
    }
}

```



To validate that an input language is supported for translation, call the Translator SDK's GetLanguagesAsync with scope: " translation " to retrieve the supported languages. The call returns a GetLanguagesResult, whose Translation property is a dictionary keyed by language codes (e.g., " en " , " it "). You can then check membership with ContainsKey(language) . If the key exists, the language is supported; otherwise, it isn't.

This matches the .NET Translator client design:

TextTranslationClient.GetLanguagesAsync(...) fetches supported languages.

GetLanguagesResult.Translation is a dictionary of translation-capable languages, so ContainsKey is the correct check.

Microsoft References

TextTranslationClient.GetLanguagesAsync (overview and signature).

Languages REST reference - use scope=translation to get languages supported for translation. Microsoft Learn GetLanguagesResult class - Translation dictionary property.

NEW QUESTION: 11

Select the answer that correctly completes the sentence.

Answer Area

A data analyst

A data engineer

A data scientist

A database administrator

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

Answer:

Answer Area

A data analyst

A data engineer

A data scientist

A database administrator

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

Explanation:

A data analyst

The statement is:

"_____ is responsible for creating visuals and charts that help a company make informed decisions." Let's analyze the options:

A data analyst

Main responsibility: interpreting data, creating visuals, dashboards, and reports using tools like Power BI, Excel, or Tableau.

Helps stakeholders make business decisions from data.

Correct answer.

A data engineer

Focuses on designing, building, and maintaining data pipelines, ETL processes, and storage systems.

Not primarily responsible for creating visuals.

A data scientist

Specializes in advanced analytics, machine learning, and statistical modeling.

They may present insights but their focus is on predictions and algorithms, not routine visuals.

A database administrator (DBA)

Responsible for managing databases, ensuring security, availability, backups, and performance.

Not responsible for creating visuals or dashboards.

The Answer: A data analyst
Microsoft Reference
Roles in data analytics
Microsoft Learn: Data Analyst role description - responsible for creating reports and visualizations to support decision-making.

NEW QUESTION: 12

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

"faceListId"

"LargeFaceListId"

"matchFace"

"matchPerson"

Answer Area

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

Answer:

Values

"faceListId"

"LargeFaceListId"

"matchFace"

"matchPerson"

Answer Area

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

Explanation:

Values

"faceListId"

"LargeFaceListId"

"matchFace"

"matchPerson"

Answer Area

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

You are calling the Face API to find similar faces from an existing list called employeefaces.

Key facts from the scenario:

The list contains 60,000 images.

In Face API:

faceListId # limited to 1,000 faces.

largeFaceListId # designed for very large lists, up to 1,000,000 faces.

Since employeefaces has 60,000 images, you must use LargeFaceListId, not faceListId.

Parameters in the Request Body

" faceId " : " 18c51a87-3a69-47a8-aedc-a54745f708a1 "

This is the face to search for.

" largeFaceListId " : " employeefaces "

Must use LargeFaceListId because the list has 60,000 images.

" maxNumOfCandidatesReturned " : 1

Limits the number of matches returned.

" mode " : " matchFace "

Two options exist:

" matchFace " # returns the most similar individual faces.

" matchPerson " # groups faces by person.

Since the requirement is find similar faces, use " matchFace " .

Completed Answer

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

Correct Values to Drag and Drop

LargeFaceListId # " employeefaces "

matchFace # " mode "

Final Answer:

First blank: LargeFaceListId

Second blank: matchFace

Microsoft References

Face API - Find Similar

Face API - Large Face List operations

NEW QUESTION: 13

You have a web app that uses Azure AI search.


When reviewing activity, you see greater than expected search query volumes. You suspect that the query key is compromised.

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

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



Answer:



Explanation:



Comprehensive Detailed Explanation

You are dealing with Azure AI Search (formerly Azure Cognitive Search). The issue is that the query key is compromised, and you need to ensure minimal downtime while keeping users with read-only access. Azure AI Search uses two types of keys: Admin keys (full control - manage indexes, data sources, etc.) and Query keys (read-only access for client applications). If a query key is compromised, the remediation must be safe and fast, without breaking the application unnecessarily.

Step 1 - Add a new query key

Instead of immediately deleting the compromised key, first generate a new query key in the Azure portal or via API. This ensures you have a replacement ready before cutting off the old key.

Step 2 - Change the app to use the new key

Update your web app configuration (for example, connection strings or environment variables) so that it authenticates against Azure AI Search using the newly created query key. This step ensures a smooth transition with minimal downtime.

Step 3 - Delete the compromised key

Finally, once your app is verified to work with the new query key, remove the old compromised key to prevent any further unauthorized access.

Correct Order:

Add a new query key.

Change the app to use the new key.

Delete the compromised key.

Microsoft References

Manage admin and query keys in Azure AI Search

Best practices for key management

NEW QUESTION: 14

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 ([LEAVE A REPLY](#))

To index millions of PDFs/images quickly:

Partition the blob data into virtual folders (prefixes) and create multiple indexers , one per folder, so they can run in parallel .

Increase search units (SUs) to raise indexing throughput.

Schedule the indexers with the same cadence so they run concurrently.

Creating multiple services (A, C) is unnecessary and costly. Scheduling sequential execution (B) defeats the goal of "as quickly as possible." References Azure Cognitive Search indexer performance guidance: parallelizing blob indexing using multiple indexers over virtual folders/prefixes and scaling SUs for higher throughput.

NEW QUESTION: 15

You use the Custom Vision service to build a classifier.

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

Which two metrics are available for review? Each correct answer presents a complete solution. (Choose two.) NOTE: Each correct selection is worth one point.

A. recall

B. F-score

C. weighted accuracy

D. precision

E. area under the curve (AUC)

Answer: A,D ([LEAVE A REPLY](#))

The question is about evaluating a Custom Vision classifier after training.

When you train an image classifier in Azure Custom Vision, the service automatically calculates performance metrics to help you evaluate the quality of the model.

Metrics Provided by Custom Vision:

Precision:

The percentage of correct positive predictions out of all positive predictions made.

Formula: Precision = $TP / (TP + FP)$

Helps determine how reliable positive predictions are.

Recall:

The percentage of actual positives that were correctly predicted.

Formula: Recall = $TP / (TP + FN)$

Helps determine how many of the true positives were captured.

These two metrics are explicitly shown in the Custom Vision portal and via API.

Option Analysis

A). Recall

Yes, Custom Vision reports recall.

Correct

B). F-score

Not directly reported in the Custom Vision portal.

Although it can be derived from precision and recall, it is not provided as a direct metric.

Incorrect

C). Weighted accuracy

Not reported by Custom Vision.

Incorrect

D). Precision

Yes, Custom Vision reports precision.

Correct

E). Area under the curve (AUC)

Not reported by Custom Vision.

More common in ROC curve analysis, not part of Custom Vision output.

Incorrect

The Answer: A. recall, D. precision

Microsoft References

Evaluate the prediction performance of your classifier in Custom Vision Custom Vision training and evaluation

NEW QUESTION: 16

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

A. a five-minute wav or mp3 file of the voce talent describing the kiosk system

B. a five-minute .flac audio file and the associated transcript as a w file

C. a .wav or mp3 file of the voice talent consenting to the creation of a synthetic version of their voice

D. a .zip file that contains 10-second .wav files and the associated transcripts as .txt files

Answer: (SHOW ANSWER)

When creating a voice talent profile for Custom Neural Voice in Azure Cognitive Services Speech, you must upload a consent recording .

The consent recording is a .wav or .mp3 file where the voice talent explicitly states that they consent to the creation of a synthetic version of their voice.

Audio samples and transcripts (training data) are uploaded later for training the model.

The profile itself requires the consent recording to comply with Microsoft's responsible AI principles.

Reference: Custom Neural Voice requirements
The Answer: C

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (416 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 17

You have an Azure subscription that contains an Azure AI Document Intelligence resource named DM.
You create a PDF document named Test.pdf that contains tabular data.
You need to analyze Test.pdf by using DM.
How should you complete the command? To answer, select the appropriate options in the answer area NOTE: Each correct selection is worth one point.

Answer Area

```
curl -v -i POST "{endpoint}/formrecognizer/documentModels/
```

prebuilt-layout
prebuilt-contract
prebuilt-document
prebuilt-layout
prebuilt-read

```
:analyze?api-version=2023-07-31" -H "Content-Type: applica
```

-H Ocp-Apim-Subscription-Key

Key1
Ocp-Apim-Subscription-Key
Secret
Subscription-Key

```
: {yourkey}" --data-ascii '{"urlSource': 'test.pdf'}"
```

Answer:

Answer Area

```
curl -v -i POST "{endpoint}/formrecognizer/documentModels/
```

prebuilt-layout
prebuilt-contract
prebuilt-document
prebuilt-layout
prebuilt-read

```
:analyze?api-version=2023-07-31" -H "Content-Type: applica
```

-H Ocp-Apim-Subscription-Key

Key1
Ocp-Apim-Subscription-Key
Secret
Subscription-Key

```
: {yourkey}" --data-ascii '{"urlSource': 'test.pdf'}"
```

Explanation:



Comprehensive Detailed Explanation

You're analyzing a PDF (Test.pdf) that "contains tabular data." In Azure AI Document Intelligence, the Layout prebuilt model is the right choice when your primary goal is to extract document structure such as tables, lines, words, selection marks, and paragraphs. Microsoft's documentation clearly states that the Layout model "extracts text, tables, selection marks, and document structure," which is exactly what you need for tabular content. Microsoft Learn While the General document (prebuilt-document) model can also return tables, it is intended for broader extraction (including key-value pairs, entities, etc.). For a document whose key requirement is table extraction, Microsoft guidance recommends starting with Layout for structure-first scenarios.

For authentication in the REST call, the correct header to pass your resource key is Ocp-Apim-Subscription-Key . This is explicitly shown in the official REST API reference for the Analyze Document operation under Security (Type: apiKey, In: header). Microsoft Learn Putting it together, the command shape is:

```
curl -v -i POST "{endpoint}/formrecognizer/documentModels/prebuilt-layout:analyze?api-version=2023-07-31" \
```

```
-H "Content-Type: application/json" \
```

```
-H "Ocp-Apim-Subscription-Key: {yourkey}" \
```

```
--data-ascii "{ \"urlSource\" : \"test.pdf\" }"
```

This uses a single analyze call to the prebuilt-layout model to extract the tables and authenticates with the Ocp-Apim-Subscription-Key header-matching the two selections in the answer area. The model matrix and overview pages also confirm Layout supports tables. Microsoft Learn References (Microsoft Documentation) Document layout analysis (Layout model) - capabilities include extracting tables, text, selection marks, and structure. Microsoft Learn Choose the best Document Intelligence model - guidance on when to use Layout (structure-first; tables).

Microsoft Learn

General document (prebuilt-document) model - broader extraction including key-value pairs and tables.

Microsoft Learn

Model overview / matrix - shows Layout supports Tables. Microsoft Learn Analyze Document (REST API) - Ocp-Apim-Subscription-Key header for authentication. Microsoft Learn

NEW QUESTION: 18

You are building a social media messaging app.

You need to identify in real time the language used in messages.

Which SDK package should you install?

A. Azure.AI.Translation.Text

B. Microsoft.CognitiveServices.Speech

C. Azure.AI.Translation.Document

D. Azure.AI.Translation.Speech

Answer: A (LEAVE A REPLY)

For a social media messaging app that must identify the language of text messages in real time, the most direct SDK is the Azure AI Translator - Text Translation SDK. Its client libraries (e.g., Azure.AI.Translation.Text in .NET) include a language detection operation that returns the detected language code and a confidence score, which is exactly what you need for real-time text scenarios.

Other options are not a fit:

Microsoft.CognitiveServices.Speech is the Speech SDK for speech-to-text, TTS, and speech translation (audio), not text-only language detection.

Azure.AI.Translation.Document targets asynchronous document translation workflows, not short text message detection. Microsoft Learn Azure.AI.Translation.Speech isn't the package used; speech translation is covered by the Speech SDK package above. Microsoft Learn Microsoft References Text Translation SDK overview (supports language detection): Azure AI Translator. Microsoft Learn DetectedLanguage in

Azure.AI.Translation.Text (returned by detect operation). Microsoft Learn Quickstart: Translator text client libraries. Microsoft Learn Install/use the Speech SDK (contrast with text scenario). Microsoft Learn Document Translation SDK overview (not applicable to short text detection). Microsoft Learn

NEW QUESTION: 19

You plan to deploy an Azure OpenAI resource by using an Azure Resource Manager (ARM) template.

You need to ensure that the resource can respond to 600 requests per minute.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity": 600,
    "count": 1,
    "maxValue": 100,
    "size": 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
      ...
    }
  }
}
```

Answer:

Answer Area

 Microsoft

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity": 600,
    "count": 1,
    "maxValue": 100,
    "size": 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
      ...
    }
  }
}
```

freeedtdumps.com

Explanation:

When deploying an Azure OpenAI model using the Microsoft.CognitiveServices/accounts/deployments resource in an ARM template, the throughput for a Standard deployment is set on the SKU object.

Specifically:

Use the capacity field of the SKU to specify the per-deployment rate.

Provide the integer value that represents the requests per minute (RPM) you want the deployment to handle.

Therefore, to meet the requirement of 600 requests per minute, complete the template as:

```
"sku": {
  "name": "Standard",
  "capacity": 600
}
```

This uses a single deployment SKU with a defined capacity. The ARM schema for accounts/deployments explicitly includes sku.capacity for resources that support scale (Azure OpenAI deployments do), and Azure OpenAI quota/rate-limit guidance confirms that rate limits are assigned per deployment.

Microsoft Azure AI Solution References

ARM template-Microsoft.CognitiveServices/accounts/deployments (2023-05-01) : property reference showing sku.capacity. Microsoft Learn Resource type reference-latest accounts/deployments : confirms SKU usage on deployments. Microsoft Learn Azure OpenAI quota and rate limits : explains assigning deployment-level limits (RPM/TPM). Microsoft Learn

NEW QUESTION: 20

You have an Azure subscription that contains an Azure AI Language resource named Resource1.
You run the following cURL command, and then play the Output.mp3 file.

```
curl --location --request POST "https://eastus.tts.speech.microsoft.com/cognitiveservices/v1" ^
--header "Ocp-Apim-Subscription-Key: 3795c4011f714f5aa66469e573109e4f" ^
--header "Content-Type: application/ssml+xml" ^
--header "X-Microsoft-OutputFormat: audio-16khz-128kbitrate-mono-mp3" ^
--header "User-Agent: curl" ^
--data-raw "<?xml version='1.0' xml:lang='en-US'>
  <voice xml:gender='Female' name='en-US-JennyNeural'>
    Welcome to the Azure Text-to-Speech demonstration.
  </voice>
  <voice xml:gender='Male' name='en-GB-RyanNeural'>
    This service allows you to convert text into natural-sounding speech.
  </voice>
  <voice xml:gender='Male' name='en-US-ChristopherNeural'>
    <mstts:express-as style='advertisement_upbeat' styledegree='2'>
      It's easy to integrate, customizable, and supports multiple languages and voices.
    </mstts:express-as>
  </voice>
</?xml>" --output output.mp3
```

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

Statements	Yes	No
You hear three sentences in different voices.	☐	☐
You hear three sentences in different accents.	☐	☐
You hear three sentences expressed in a neutral tone.	☐	☐

Answer:

Statements	Yes	No
You hear three sentences in different voices.	☒	☐
You hear three sentences in different accents.	☒	☐
You hear three sentences expressed in a neutral tone.	☐	☒

Explanation:

Answer Area

Statements	Yes	No
You hear three sentences in different voices.	☒	☐
You hear three sentences in different accents.	☒	☐
You hear three sentences expressed in a neutral tone.	☐	☒

The SSML payload contains three < voice > elements with different voice names:

name= " en-US-JennyNeural " (US English) - neutral delivery.

name= " en-GB-RyanNeural " (UK English) - neutral delivery.

name= " en-US-ChristopherNeural " (US English) - wrapped in < mstts:express-as style= " advertisement_upbeat " styledegree= " 2 " > , which applies an upbeat, non-neutral style.

Therefore:

Different voices # Yes (Jenny, Ryan, Christopher are three distinct voices).

Different accents # No (voices 1 and 3 are both en-US; only voice 2 is en-GB).

Neutral tone # No (the third sentence uses an expressive style instead of neutral).

References (Microsoft Azure Speech / Text-to-Speech)

Use SSML to control Speech Synthesis (voices, < voice > , < speak > , and expressive styles).

Neural voice styles and roles (e.g., mstts:express-as with styles such as advertisement_upbeat).

Text-to-Speech REST API v1 (synthesis endpoint and headers like Ocp-Apim-Subscription-Key , X- Microsoft-OutputFormat).

NEW QUESTION: 21

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

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

Which index attributes should you configure for the category field?

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

Answer: (SHOW ANSWER)

For the category field in Azure Cognitive Search:

- * searchable # allows users to search by words in the category field.
- * facetable # enables drill-down filtering (facets).
- * retrievable # ensures the field appears in search results.
- * filterable/sortable/key are not required here based on the scenario.

The answer: A

Reference: Index attributes in Azure Cognitive Search

NEW QUESTION: 22

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

findsimilars

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:

Values

detect

findsimilars

group

identify

matchFace

matchPerson

verify

Answer Area

POST {Endpoint}/face/v1.0/

findsimilars

Request Body

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

matchPerson

"
}

Explanation:

Box 1: findsimilars

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

Box 2: matchPerson

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

Reference:

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

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

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: A (LEAVE A REPLY)

Scenario Recap:

Application: Flower species classification using Custom Vision .

New images are received.

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

Analysis:

The Smart Labeler tool in Custom Vision can help automatically generate suggested labels for newly uploaded images.

This is a valid way to extend a classifier with new images.

It doesn't complete the entire pipeline (you'd still retrain afterward), but using Smart Labeler does meet the goal of adding new species efficiently.

NEW QUESTION: 24

You have an Azure subscription that contains an Azure OpenA1 resource named All.

You plan to develop a console app that will answer user questions.

You need to call All and output the results to the console.

How should you complete the code? To answer, select the appropriate options in the answer area NOTE: Each correct selection is worth one point.

Microsoft Azure

```
OpenAIClient client =  
    new OpenAIClient(new Uri(endpoint), new AzureKeyCredential(key));  
Response<Completions> response =  
    client.  
        (deploymentName, "What is Microsoft Azure?")
```

Get Completions
Get Completions
Get Embeddings
Get Image Generations

Console.WriteLine
Microsoft

(response.Value.Choices[0].Text);
(response.Value.Choices[0].Text);
(response.Value.Id);
(response.Value.PromptFilterResults);

Answer:



Explanation:



You are using the Azure OpenAI .NET SDK to send a prompt ("What is Microsoft Azure?") to a text model deployment. For text generation, you must call the Completions endpoint; hence `client.GetCompletions(...)` (or the async variant) is correct.

The `Response < Completions >` object returns one or more choices; the generated text is contained in `response.Value.Choices[0].Text`, which you can write to the console.

`GetEmbeddings` would return vector embeddings, not natural-language text, and `GetImageGenerations` is for DALL E image generation, so neither meets the requirement to "answer user questions" and print text.

Microsoft References

Azure OpenAI Service (.NET) - Completions quickstart and sample showing `GetCompletions` and reading `Choices[0].Text`.

Azure OpenAI Service - Embeddings overview (returns vectors, not text).

Azure OpenAI Service - Image generation overview (DALL E, not for text answers).

NEW QUESTION: 25

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: (SHOW ANSWER)

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.



Comprehensive Detailed Explanation

You trained an object detection model obj1 inside project proj1 on the acvdev Custom Vision resource and need to move it to the acvprod resource.

Moving a model between Custom Vision resources is done by exporting the project (not an iteration) from the source resource and importing it into the target resource.

Step 1: GetProjects (acvdev)

Retrieve the list of projects on the development resource to obtain the project ID of proj1 . The export operation requires the project ID.

Step 2: ExportProject (acvdev)

Call the Export Project API on the development resource to generate the portable project package (includes tags, images' metadata, iterations, settings). This is specifically intended for moving/cloning a project across resources or regions.

Step 3: ImportProject (acvprod)

Use the Import Project API on the production resource to create a new project from the exported package.

This recreates the project (and associated iterations/metadata) under acvprod .

Notes:

ExportIteration is for exporting a trained iteration to formats for edge/container or mobile use, not for transferring to another Custom Vision resource.

UpdateProject is for changing metadata (name, description, domain) of an existing project and is not part of the move workflow.

References (Microsoft Azure AI Solution)

Custom Vision REST API - Projects: Export/Import

<https://learn.microsoft.com/azure/ai-services/custom-vision-service/export-import-project> Custom Vision REST API - Projects API (Get projects)

<https://learn.microsoft.com/rest/api/customvision/training/projects>

Custom Vision - Export model vs. Export project guidance

<https://learn.microsoft.com/azure/ai-services/custom-vision-service/export-your-model> (describes iteration export purpose versus project portability)

NEW QUESTION: 27

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

Microsoft

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:

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

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>

NEW QUESTION: 28

You have the following files:

* File1.pdf

* File2.jpg

* File3.docx

* File4.webp

* File5.png

Which files can you analyze by using Azure AI Content Understanding?

A. File1.pdf, File2.jpg, and File3.docx only

B. File1.pdf and File3.docx only

- C. File1.pdf, File2.jpg, File3.docx, File4.webp, and File5.png
 - D. File1.pdf, File2.jpg, File3.docx and File5.png only
 - E. File1.pdf, File2.jpg, and FileS.png only
- Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 29

You have an app that uses an Azure AI Foundry Content Safety blocklist.
You need to remove an entry from the blocklist. The solution must minimize the impact on existing entries on the list.
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

Microsoft

freepdfdumps.com

```
client = BlocklistClient(endpoint, AzureKeyCredential(key))
blocklist_name = "TestBlocklist"
block_item_id = 1
client.  

    add_or_update_blocklist_items  

    create_or_update_text_blocklist  

    delete_text_blocklist  

    remove_blocklist_items  

    blocklist_name=blocklist_name, options= (blocklist_item_ids=[block_item_id])  

    )  

print(f"\nRemoved blockItem: {add_result.blockItem}")  

...

```

Answer:

ANSWER AREA

```
client = BlocklistClient(endpoint, AzureKeyCredential(key))

blocklist_name = "TestBlocklist"

block_item_id = 1

client.
  add_or_update_blocklist_items
  create_or_update_text_blocklist
  delete_text_blocklist
  remove_blocklist_items

blocklist_name=blocklist_name, options=
  AddOrUpdateTextBlocklistItemsOptions
  BlocklistClient
  ContentSafetyClient
  RemoveTextBlocklistItemsOptions
) (blocklist_item_ids=[block_item_id])

print(f"\nRemoved blockItem: {add_result.blocklist_items[0].blocklist_item_id}")

...
```

Explanation:

```
Answer Area

client = BlocklistClient(endpoint, AzureKeyCredential(key))

blocklist_name = "TestBlocklist"

block_item_id = 1

client. remove_blocklist_items (
  blocklist_name=blocklist_name, options= RemoveTextBlocklistItemsOptions (blocklist_item_ids=[block_item_id])
)

print(f"\nRemoved blockItem: {add_result.blocklist_items[0].blocklist_item_id}")

...
```

NEW QUESTION: 30

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

Actions

Pull an image from Docker Hub.

Provision an on-premises Kubernetes cluster that has internet connectivity.

Provision an Azure Kubernetes Service (AKS) resource.

Run the container and specify an App ID and Client Secret.

Provision an on-premises Kubernetes cluster that is isolated from the internet.

Pull an image from the Microsoft Container Registry (MCR).

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Answer Area

Answer:

Actions

Pull an image from Docker Hub.

Provision an on-premises Kubernetes cluster that has internet connectivity.

Provision an Azure Kubernetes Service (AKS) resource.

Run the container and specify an App ID and Client Secret.

Provision an on-premises Kubernetes cluster that is isolated from the internet.

Pull an image from the Microsoft Container Registry (MCR).

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Answer Area

Provision an on-premises Kubernetes cluster that is isolated from the internet.

Pull an image from the Microsoft Container Registry (MCR).

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Explanation:

Actions

Pull an image from Docker Hub.

Provision an on-premises Kubernetes cluster that has internet connectivity.

Provision an Azure Kubernetes Service (AKS) resource.

Run the container and specify an App ID and Client Secret.

Answer Area

1 Provision an on-premises Kubernetes cluster that is isolated from the internet.

2 Pull an image from the Microsoft Container Registry (MCR).

3 Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Comprehensive Detailed Explanation

You need to analyze confidential documents with the Language service while ensuring those documents never leave your premises. The supported pattern for Azure AI Services (Cognitive Services) is to run the Language service container on-premises and send the app's requests to that local container endpoint. The container performs inference locally; only metering/billing calls are sent to Azure.

To achieve this:

Provision an on-premises Kubernetes cluster that has internet connectivity.

Azure AI service containers running on-prem require outbound access to Azure for metering and billing unless you're using the special disconnected model (not implied here). Therefore, the cluster must have internet connectivity so the container can reach the Azure billing endpoint while keeping document content local.

Pull an image from the Microsoft Container Registry (MCR).

Official Azure AI service containers (including Language) are published in MCR, not Docker Hub. You must pull the Language (Text Analytics/Language) container image from MCR.

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

When you start the container, you configure it with the API key and billing (endpoint) URL of your Azure Cognitive Services (multi-service or Language) resource. This enables the container to authenticate and report usage to Azure, while your data stays on-prem.

Why not the other options?

AKS would place workloads in Azure, which isn't required when the requirement is to keep documents on- prem.

On-prem cluster isolated from the internet would block the container's required calls to Azure's billing endpoint.

Pull from Docker Hub is incorrect because Azure AI service containers are published to MCR.

App ID and Client Secret are not the standard parameters for Azure AI service containers; they require API key + billing endpoint.

Microsoft Azure AI Solution References

Use Azure AI services in containers (billing/endpoint and API key requirements, data remains local):

Microsoft Learn - Install and run containers for Azure AI services and Use containers with Azure AI services (metering and billing) .

<https://learn.microsoft.com/azure/ai-services/containers/>

Language service container images on Microsoft Container Registry (MCR):

Microsoft Learn - Deploy language containers (Text Analytics/Language) .

<https://learn.microsoft.com/azure/ai-services/language-service/how-to/use-containers> Networking requirements (internet connectivity for metering):

Microsoft Learn - Configure containers and network access for Azure AI services containers .

<https://learn.microsoft.com/azure/ai-services/containers/container-requirements-and-limits> These documents collectively show that: (a) containers are pulled from MCR, (b) they are configured with API key + endpoint for billing, and (c) internet egress is required for metering while payload data remains on- premises.

NEW QUESTION: 31

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: ([SHOW ANSWER](#))

Scenario Recap:

* Application: Flower classification.

* Solution: Create a new model, and then upload the new images and labels.

Analysis:

* Creating a new model from scratch is not the correct approach for adding species to an existing classifier.

* The goal is to extend the existing model, not reset and rebuild a new one.

* This would waste training data and is not aligned with Custom Vision usage.

The answer: B. No

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (**416** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 32

You are building a phone call handling solution that will use the Azure AI Speech service and a custom neural voice.

You need to create a custom speech model.

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

Actions	Answer Area
☐ Upload a consent statement for the voice talent as a WAV file.	
☐ Upload speech samples as WMA files.	
☐ Create a custom voice project.	
☐ Upload a consent statement for the voice talent as a signed PDF file.	
☐ Analyze the quality of the audio data and resolve identified issues.	
☐ Upload speech samples as MP3 files.	
☐ Train the model by using a neural training method.	

Answer:

ACTIONS

Upload a consent statement for the voice talent as a WAV file.

Upload speech samples as WMA files.

Create a custom voice project.

Upload a consent statement for the voice talent as a signed PDF file.

Analyze the quality of the audio data and resolve identified issues.

Upload speech samples as MP3 files.

Train the model by using a neural training method.

ANSWER AREA

Create a custom voice project.

Upload a consent statement for the voice talent as a signed PDF file.

Analyze the quality of the audio data and resolve identified issues.

Upload speech samples as MP3 files.

Train the model by using a neural training method.

Explanation:

Actions

Upload a consent statement for the voice talent as a WAV file.

Upload speech samples as WMA files.

Answer Area

1 Create a custom voice project.

2 Upload a consent statement for the voice talent as a signed PDF file.

3 Analyze the quality of the audio data and resolve identified issues.

4 Upload speech samples as MP3 files.

5 Train the model by using a neural training method.

Comprehensive Detailed Explanation

A custom neural voice (professional voice) project in Azure AI Speech requires that you follow Responsible AI and consent steps before any model training. The correct workflow in Speech Studio / Azure AI Foundry is:

Create a custom voice project - This is the container that will hold the voice talent consent, your training datasets, models, and deployments. Each project is scoped to a locale and other attributes.

Collect and upload consent - Microsoft requires explicit authorization from the voice talent. In practice, you add the voice talent consent to the project. The platform requires the verbal consent recording (WAV); many guidance and exam items also include the signed consent document (PDF) as part of consent capture to meet Responsible AI expectations. The consent step is performed before uploading training data.

Upload the voice talent's verbal consent as a WAV file - The consent audio verifies that the person who granted consent is the same as the speaker in your training data. The documentation shows consent is a recorded statement uploaded to the project.

Microsoft Learn Analyze the quality of the audio data and fix issues - After data upload, Speech Studio runs validation and provides a per-utterance analysis (format checks, sampling rate, script match, pronunciation/noise scoring).

You must resolve issues and run Analyze data before the dataset can be used for training.

Microsoft Learn Train the model using a neural training method - Start model training from the project and choose Neural as the training type for custom neural voice.

Microsoft Learn Why not "upload speech samples as WMA/MP3 files"?

For the standard, recommended Individual utterances + matching transcript data type used to fine-tune a professional neural voice, the audio files must be WAV/RIFF PCM; the docs explicitly reference ".wav" throughout data checks and corrections. WMA isn't supported, and MP3 isn't accepted for this core data type.

(Long audio / audio-only workflows can accept MP3 for preprocessing, but the canonical training dataset for custom neural voice is WAV.) Microsoft Azure AI References Create a project for professional voice (Speech Studio/Azure AI Foundry): Project contains consent, datasets, models, endpoints. Microsoft Learn Add voice talent consent to the professional voice project: Upload the recorded consent statement for the voice talent. Microsoft Learn Professional voice fine-tuning data - required formats: Individual utterances + transcript require WAV/RIFF PCM; guidance on sample rate and transcript format. Microsoft Learn Add a professional voice training dataset & resolve data issues: Data validation, pronunciation/noise checks, and the Analyze data step prior to training. Microsoft Learn Train your professional voice model: Start training and select Neural method. Microsoft Learn

NEW QUESTION: 33

You are building an app that will process scanned expense claims and extract and label the following data:

- . Merchant information
- . Time of transaction
- . Date of transaction
- . Taxes paid
- . Total cost

You need to recommend an Azure AI Document Intelligence model for the app. The solution must minimize development effort.

What should you use?

- A.** a custom neural model
- B.** the prebuilt receipt model
- C.** a custom template model
- D.** the prebuilt Read model

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 34

You have an Azure subscription that contains an Azure Cognitive Service for Language resource. You need to identify the URL of the REST interface for the Language service. Which blade should you use in the Azure portal?

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

Answer: B ([LEAVE A REPLY](#))

In the Azure portal, the Keys and Endpoint blade for a Cognitive Service resource provides:

API keys (for authentication)

Endpoint URL (the REST interface base URL)

Identity blade relates to managed identities.

Properties blade shows metadata (resource ID, location, subscription ID).

Networking blade controls network access and firewall rules.

The Answer: B. Keys and Endpoint

Reference: Manage keys and endpoints for Azure Cognitive Services

NEW QUESTION: 35

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

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

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

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

NOTE: Each correct selection is worth one point.

Answer Area

Generate alt text:

https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image

https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate

https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image

https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description

https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description

Detect inappropriate content:

https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description

https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate

https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image

https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description

https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1

Answer:

Answer Area

Generate alt text:

https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image

https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate

https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image

https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description

https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description

Detect inappropriate content:

https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description

https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate

https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image

https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description

https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1

Explanation:

Generate alt text

To automatically suggest descriptive captions for images (usable as alt text), use the Computer Vision Describe Image API. It returns natural-language captions summarizing the image content. Endpoint: /vision/v3.2/describe . Setting maxCandidates=1 gives a single best caption.

Detect inappropriate content

To block adult/racy/gory images with minimal effort, use Computer Vision Analyze Image with the Adult visual feature. This returns flags (isAdult , isRacy , isGory) and confidence scores. The provided option includes Adult,Description which still returns the needed adult/racy/gore signals, satisfying the requirement.

These choices minimize development effort by using built-in Computer Vision capabilities without training a custom model or wiring up multiple services.

References

Microsoft Learn - Describe images (Computer Vision): "The Describe Image API returns a list of natural language sentences that describe the image."

<https://learn.microsoft.com/azure/ai-services/computer-vision/overview-image-analysis#describe-images> Microsoft Learn - Analyze images (Adult/Racy/Gore detection): "Use the Adult feature to detect adult, racy, and gory content."

<https://learn.microsoft.com/azure/ai-services/computer-vision/overview-image-analysis#adult-racy-and-gory-content-detection>

NEW QUESTION: 36

You have an Azure subscription that contains an Azure OpenAI resource. You configure a model that has the following settings:

- * Temperature: 1
- * Top probabilities: 0.5
- * Max response tokens: 100

You ask the model a question and receive the following response.

```
{
  "choices": [
    {
      "finish_reason": "stop",
      "index": 0,
      "message": {
        "content": "The founders of Microsoft are Bill Gates and Paul Allen. They co-founded the company in 1975.",
        "role": "assistant"
      }
    }
  ],
  "created": 1679014554,
  "id": "chatcmpl-6usfn2yyjkbmESe364jaQR6bDSc01",
  "model": "gpt-3.5-turbo-0301",
  "object": "chat.completion",
  "usage": {
    "completion_tokens": 86,
    "prompt_tokens": 37,
    "total_tokens": 123
  }
}
```

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

NOTE: Each correct selection is worth point.

Answer Area

Statements	Yes	No
The subscription will be charged 86 tokens for the execution of the session.	☐	☐
The text completion was truncated because the Max response tokens value was exceeded.	☐	☐
The prompt_tokens value will be included in the calculation of the Max response tokens value.	☐	☐

Answer:

Answer Area

Statements	Yes	No
The subscription will be charged 86 tokens for the execution of the session.	☐	☒
The text completion was truncated because the Max response tokens value was exceeded.	☐	☐
The prompt_tokens value will be included in the calculation of the Max response tokens value.	☐	☐

Explanation:

The subscription will be charged 86 tokens for the execution of the session. No The text completion was truncated because the Max response tokens value was exceeded. No The prompt_tokens value will be included in the calculation of the Max response tokens value. No

From the response payload:

```
usage.prompt_tokens = 37
usage.completion_tokens = 86
usage.total_tokens = 123
finish_reason = " stop "
```

Billing is for total tokens (input + output), not just completion tokens.

Azure OpenAI bills on both prompt (input) and completion (output) tokens. Here, the total is 123 tokens, so it is not just 86. Hence No.

Not truncated by max response tokens.

If a response hits the maximum allowed output tokens, finish_reason would be " length " . Here it is " stop " and the model produced only 86 completion tokens while Max response tokens was set to 100, so the limit was not reached. Hence No.

max response tokens applies only to the generated output, not the prompt.

max_tokens (Max response tokens) limits the number of tokens the model may generate in the completion; it does not include prompt_tokens . Hence No.

Microsoft References

Azure OpenAI Service - Chat Completions API (parameters, max_tokens , response usage , finish_reason).

<https://learn.microsoft.com/azure/ai-services/openai/reference#chat-completions> Azure OpenAI pricing and token usage model (charges for input and output tokens).

<https://learn.microsoft.com/azure/ai-services/openai/overview#pricing>

Token counting and limits (context vs. max output tokens).

<https://learn.microsoft.com/azure/ai-services/openai/concepts/tokenizers#token-limits>

NEW QUESTION: 37

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 (LEAVE A REPLY)

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>

NEW QUESTION: 38

You are building a transcription service for technical podcasts.

Testing reveals that the service fails to transcribe technical terms accurately.

You need to improve the accuracy of the service.

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

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

Actions

Create a Speaker Recognition model.

Create a Conversational Language Understanding model.

Create a Custom Speech project.

Create a speech-to-text model.

Upload training datasets.

Train the model.

Deploy the model.

Answer Area

Answer:

Actions

Create a Speaker Recognition model.

Create a Conversational Language Understanding model.

Create a Custom Speech project.

Create a speech-to-text model.

Upload training datasets.

Train the model.

Deploy the model.

Answer Area

Create a Custom Speech project.

Create a speech-to-text model.

Upload training datasets.

Train the model.

Deploy the model.

Explanation:

Actions

Create a Speaker Recognition model.

Create a Conversational Language Understanding model.

>

<

Answer Area

1

Create a Custom Speech project.

2

Create a speech-to-text model.

3

Upload training datasets.

4

Train the model.

5

Deploy the model.

↑

↓

<https://learn.microsoft.com/en-us/azure/cognitive-services/speech-service/custom-speech-overview#how-does-it-work> With Custom Speech, you can upload your own data, test and train a custom model, compare accuracy between models, and deploy a model to a custom endpoint.

- Create a project and choose a model. Use a Speech resource that you create in the Azure portal. If you will train a custom model with audio data, choose a Speech resource region with dedicated hardware for training audio data.
- Upload test data. Upload test data to evaluate the speech to text offering for your applications, tools, and products.
- Train a model. Provide written transcripts and related text, along with the corresponding audio data. Testing a model before and after training is optional but recommended.
- Deploy a model. Once you're satisfied with the test results, deploy the model to a custom endpoint. With the exception of batch transcription, you must deploy a custom endpoint to use a Custom Speech model.

NEW QUESTION: 39

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:

You have a Computer Vision resource named contoso1 in West US region.

You must use this resource to resize an image using smart cropping.

Let's analyze the key points:

Endpoint

Each Azure AI service resource has a regional endpoint.

Since the resource name is contoso1 and is deployed in West US, the correct endpoint will be:

https://contoso1.cognitiveservices.azure.com

The older projectoxford.ai and generic westus.api.cognitive.microsoft.com are legacy endpoints and are not recommended.

Therefore, the correct choice is https://contoso1.cognitiveservices.azure.com .

API Method

The task is to generate a different size of a product photo with smart cropping.

This is done using the Generate Thumbnail API in Computer Vision.

The path for this is:

/vision/v3.1/generateThumbnail

Other options like detect (for object detection) or areaOfInterest (for finding regions of interest) do not generate thumbnails.

The Answer:

Endpoint: https://contoso1.cognitiveservices.azure.com

API: /vision/v3.1/generateThumbnail

Microsoft References

Computer Vision API reference - Generate Thumbnail

Computer Vision REST API documentation

Azure AI services endpoints

NEW QUESTION: 40

You have a SQL query that combines customer data and order data. The query includes calculated columns.

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

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

Answer: B ([LEAVE A REPLY](#))

- * A view is a virtual table based on the result set of a query.
- * Views allow you to encapsulate complex queries (including joins and calculated columns) and make them reusable by other users.
- * Index improves performance but does not store or rerun queries.
- * Scalar function returns a single value, not a tabular result.
- * Table stores raw data, not queries.

Reference: Create views in SQL Server

NEW QUESTION: 41

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:

Azure Databricks is an Apache Spark-based analytics platform. # Yes

Azure Analysis Services is used for transactional workloads. # No

Azure Data Factory orchestrates data integration workflows. # Yes

Let's break down each statement:

Azure Databricks is an Apache Spark-based analytics platform.

True.

Azure Databricks is a fast, easy, and collaborative Apache Spark-based big data analytics platform .

It is designed for data engineering, machine learning, and AI workloads.

The Answer: Yes

Azure Analysis Services is used for transactional workloads.

False.

Azure Analysis Services is an analytical service for semantic data modeling, OLAP (online analytical processing), and BI (business intelligence).

Transactional workloads belong to OLTP (online transaction processing) databases such as Azure SQL Database.

The Answer: No

Azure Data Factory orchestrates data integration workflows.

True.

Azure Data Factory (ADF) is a cloud-based ETL and data integration service that orchestrates and automates data movement and transformation across on-premises and cloud systems.

NEW QUESTION: 42

You have an Azure subscription that contains a Microsoft Foundry resource.

You need to build an app that will suggest product names from a given product description.

Which Foundry model should you use?

A. DALL-E

B. Whisper

C. embeddings

D. GPT-4

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 43

You are building an app that uses a Language Understanding model to analyze text files. You need to ensure that the app can detect the following entities:

- * Temperatures
- * Currency values
- * Email addresses
- * Telephone numbers

The solution must minimize development effort.

Which model capability should you use?

- A. list entities
- B. learned entities
- C. utterances
- D. regular expression components
- E. pre-built entity components

Answer: E ([LEAVE A REPLY](#))

To detect common data types such as temperatures, currency values, email addresses, and telephone numbers with minimal development effort, you should use pre-built entity components in Conversational Language Understanding (CLU). Prebuilt components are ready-made recognizers that automatically extract well- known patterns from utterances-no training, labeling, lists, or regex authoring required.

Specifically, CLU offers the following relevant prebuilt components:

Quantity.Temperature - extracts temperatures (e.g., "34 degrees celsius").

Quantity.Currency - extracts monetary amounts and currency units (e.g., "\$20", "£15.50").

Email - extracts email addresses.

Phone Number - extracts phone numbers.

These components are designed to be dropped into your entities and immediately provide predictions, which minimizes development effort compared to:

List entities (you must curate values/synonyms yourself),

Learned entities (require labeling utterances and training),

Regular expression components (you must craft and maintain regex patterns), or Utterances (training data, not a capability for detection by itself).

Microsoft References

Supported prebuilt entity components in CLU (shows Quantity.Temperature, Quantity.Currency, Email, Phone Number, etc.). Microsoft Learn Entity components in CLU (explains prebuilt, learned, list, and regex components and why prebuilt components are automatically detected). Microsoft Learn

NEW QUESTION: 44

You have an Azure subscription that contains an Azure AI Content Safety resource named Resource1.

You create the following cURL command.

```
curl -X POST "https://resources.cognitiveservices.azure.com/contentSafety/text:detectProtectedMaterial?api-version=2024-09-01" \
-H "Content-Type: application/json" \
-H "Ocp-Apim-Subscription-Key: <your_subscription_key>" \
-d '{ "text": "<your_content>" }'
```

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

NOTE: Each correct selection is worth point.

Answer Area

Statements	Yes	No
The <code>text</code> value must use JSON-formatted text.	☐	☐
The command will analyze inputted text and identify whether the text contains published song lyrics.	☐	☐
The <code>Ocp-Apim-Subscription-Key</code> value must contain the ID of the Azure subscription that hosts Resource1.	☐	☐

Answer:

Answer Area

Statements	Yes	No
The <code>text</code> value must use JSON-formatted text.	☐	☒
The command will analyze inputted text and identify whether the text contains published song lyrics.	☒	☐
The <code>Ocp-Apim-Subscription-Key</code> value must contain the ID of the Azure subscription that hosts Resource1.	☐	☒

Explanation:

The text value must use JSON-formatted text. No

The command will analyze inputted text and identify whether the text contains published song lyrics. Yes

The Ocp-Apim-Subscription-Key value must contain the ID of the Azure subscription that hosts Resource1. No

Comprehensive Detailed Explanation

The cURL calls the Protected Material detection endpoint:

```
POST {endpoint}/contentsafety/text:detectProtectedMaterial?api-version=2024-09-01.
```

This API expects a JSON body with a text property whose value is a plain string (the content to check). The body is JSON, but the text value itself is not required to be "JSON-formatted text"; it's just ordinary text inside a JSON payload. Hence, statement 1 is No .

Microsoft Learn The Protected Material (text) feature "flags any known text content ... such as song lyrics , articles, recipes, and selected web content." Therefore, this command will analyze the input and can identify whether it contains published song lyrics. Statement 2 is Yes .

Ocp-Apim-Subscription-Key is the API (resource) key used to authenticate requests to Azure AI services or APIs fronted by API Management. It is not the Azure subscription ID . Authentication guidance explicitly states the header carries a resource key or token; API Management documentation identifies Ocp-Apim- Subscription-Key as the subscription key header. Therefore, statement 3 is No .

Microsoft Azure AI References

Detect Text Protected Material - REST API (endpoint, request/response schema): Microsoft Learn Protected material detection (concepts; includes song lyrics examples): Microsoft Learn What's new in Azure AI Content Safety (mentions protected material detection for lyrics): Microsoft Learn Azure AI services authentication (resource key in headers): Microsoft Learn API Management subscriptions (meaning of Ocp-Apim-Subscription-Key): Microsoft Learn

NEW QUESTION: 45

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.	☐	☐
The code will list tags and their associated confidence.	☐	☐
The code will read a file from the local file system.	☐	☐

Answer:

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

Explanation:

The code will perform face recognition # No

The code will list tags and their associated confidence # Yes

The code will read a file from the local file system # Yes

The given method uses the Computer Vision client library:

It specifies the features:

```
List < VisualFeatureTypes > features = new List < VisualFeatureTypes > ()  
{  
    VisualFeatureTypes.Description,  
    VisualFeatureTypes.Tags,  
};
```

That means it will return descriptions (captions) and tags from the image.

The image is opened from the local file system:

```
using (Stream imageStream = File.OpenRead(localImage))
```

The analysis call:

```
ImageAnalysis results = await client.AnalyzeImageInStreamAsync(imageStream, features);
```

 It prints captions and tags with their confidence values.

Statement Analysis

"The code will perform face recognition."

No.

The requested features are Description and Tags , not Faces .

Face recognition would require VisualFeatureTypes.Faces .

"The code will list tags and their associated confidence."

Yes.

The foreach (var tag in results.Tags) loop outputs tag name and confidence.

"The code will read a file from the local file system."

Yes.

It uses File.OpenRead(localImage) which reads from disk.

Final Answer

The code will perform face recognition # No

The code will list tags and their associated confidence # Yes

The code will read a file from the local file system # Yes

Microsoft References

Analyze an image with the Computer Vision API

VisualFeatureTypes enum reference

NEW QUESTION: 46

You are building a chatbot.

You need to configure the bot to guide users through a product setup process.

Which type of dialog should you use?

- A.** action
- B.** component
- C.** waterfall
- D.** adaptive

Answer: C ([LEAVE A REPLY](#))

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (416 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 47

You are developing a text processing solution.

You have the function shown below.

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

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

For the second argument, you call the function and specify the following string.

Our tour of Paris included a visit to the Eiffel Tower.

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

NOTE: Each correct selection is worth one point.

Answer:

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

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

Explanation:

The output will include the following words: our and included. No

The output will include the following words: Paris, Eiffel, and Tower. Yes The function will output all the key phrases from the input string to the console. No The output will include the following words: our and included.

No The output will include the following words: Paris, Eiffel, and Tower. Yes The function will output all the key phrases from the input string to the console. No The function uses RecognizeEntities to identify entities in the text, which typically includes named entities like locations (e.g., Paris, Eiffel Tower). It then prints the text of each recognized entity. The input string " Our tour of Paris included a visit to the Eiffel Tower " contains the entities " Paris " and " Eiffel Tower, " which will be broken down into individual words like " Paris, " " Eiffel, " and " Tower. " However, common words like " our " and " included " are not typically recognized as entities, and the function does not output all key phrases, only the recognized entities.

NEW QUESTION: 48

You are building an app that will use the Azure Video Indexer service.

You plan to train a language model to recognize industry-specific terms.

You need to upload a file that contains the industry-specific terms.

Which file format should you use?

A. PDF

B. XML

C. TXT

D. XLS

Answer: C ([LEAVE A REPLY](#))

* To extend Video Indexer with industry-specific terms, you upload a plain text (TXT) file containing the vocabulary.

* PDF, XML, XLS are not supported for custom vocabulary upload.

The answer: C

Reference: Customize Video Indexer with custom vocabulary

NEW QUESTION: 49

You have an Azure subscription that contains an Azure OpenAI resource named OpenAI1 and a user named User1.

You need to ensure that User1 can upload datasets to OpenAI1 and finetune the existing models. The solution must follow the principle of least privilege.

Which role should you assign to User1?

A. Cognitive Services Contributor

B. Contributor

C. Cognitive Services OpenAI User

D. Cognitive Services OpenAI Contributor

Answer: ([SHOW ANSWER](#))

A user User1 must be able to upload datasets and fine-tune existing models in Azure OpenAI.

The role must follow principle of least privilege .

Role analysis:

A). Cognitive Services Contributor : Too broad; allows management of all Cognitive Services resources, not least privilege.

B). Contributor : Overly permissive; allows full management rights on the entire resource, not just OpenAI.

C). Cognitive Services OpenAI User : Grants permissions only to use models (query endpoints). Does not allow fine-tuning or dataset upload.

D). Cognitive Services OpenAI Contributor : Correct role. This provides the ability to manage OpenAI resources, including uploading training data and fine-tuning models.

The Answer: D. Cognitive Services OpenAI Contributor

Microsoft Reference:

Azure RBAC roles for Azure OpenAI

Built-in roles: Cognitive Services OpenAI Contributor

NEW QUESTION: 50

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: (SHOW ANSWER)

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

QnA Maker is a cloud-based Natural Language Processing (NLP) service that allows you to create a natural conversational layer over your data. It is used to find the most appropriate answer for any input from your custom knowledge base (KB) of information.

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

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

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

NEW QUESTION: 51

You have an Azure subscription that contains an Azure AI Video Indexer account.

You need to add a custom brand and logo to the indexer and configure an exclusion for the custom brand.

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

NOTE: Each correct selection is worth one point.

```
{
  "referenceUrl": "https://www.contoso.com/Contoso",
  "id": 97974,
  "name": "Contoso",
  "accountId": "ContosoAccountId",
  "lastModifierUserName": "SampleUserName",
  "created": "2023-04-25T14:59:52.7433333",
  "lastModified": "2023-04-25T14:59:52.7433333",
```

"enabled":	false
"enabled":	["Excluded"]
"state":	["Included"]
"tags":	false
"useBuiltin":	true

Answer:

Answer Area

```
{
  "referenceUrl": "https://www.contoso.com/Contoso",
  "id": 97974,
  "name": "Contoso",
  "accountId": "ContosoAccountId",
  "lastModifierUserName": "SampleUserName",
  "created": "2023-04-25T14:59:52.7433333",
  "lastModified": "2023-04-25T14:59:52.7433333",
```

"enabled":	false
"enabled":	["Excluded"]
"state":	["Included"]
"tags":	false
"useBuiltin":	true

Explanation:

Answer Area

```
{
  "referenceUrl": "https://www.contoso.com/Contoso",
  "id": 97974,
  "name": "Contoso",
  "accountId": "ContosoAccountId",
  "lastModifierUserName": "SampleUserName",
  "created": "2023-04-25T14:59:52.7433333",
  "lastModified": "2023-04-25T14:59:52.7433333",
  "enabled": false
}
```



Comprehensive Detailed Explanation

The scenario:

You are working with Azure AI Video Indexer.

You need to add a custom brand/logo and also configure an exclusion so that this brand is not indexed.

In Video Indexer's REST API for brands, the configuration JSON includes:

"enabled" : true/false # Determines whether the brand should be indexed.

Setting "enabled" : false means the brand is excluded.

The other properties in the dropdown (state , tags , useBuiltin) are not correct in this context for configuring exclusions.

Therefore, to exclude a custom brand:

"enabled" : false

The Answer:

Attribute: "enabled"

Value: false

Microsoft References

Video Indexer - Brands customization

Azure Video Indexer API - Brands

NEW QUESTION: 52

You are building an app that will share user images.

You need to configure the app to perform the following actions when a user uploads an image:

* Categorize the image as either a photograph or a drawing.

* Generate a caption for the image.

The solution must minimize development effort.

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

NOTE: Each correct selection is worth one point.

A. object detection in Computer Vision

- B. content tags in Computer Vision
- C. image descriptions in Computer Vision
- D. image type detection in Computer Vision
- E. image classification in Custom Vision

Answer: C,D (LEAVE A REPLY)

According to the Microsoft documentation, Computer Vision is a cloud-based service that provides developers with access to advanced algorithms for processing images and returning information. By uploading an image or specifying an image URL, Computer Vision algorithms can analyze visual content in different ways based on inputs and user choices.

According to the Microsoft documentation, image type detection is one of the features of Computer Vision that can categorize an image as either a photograph or a drawing. You can use the image type detection feature by calling the Analyze Image API with the visualFeatures parameter set to ImageType. The API will return a JSON response with an imageType field that indicates whether the image is a photo or a clipart.

According to the Microsoft documentation, image descriptions is another feature of Computer Vision that can generate a caption for an image. You can use the image descriptions feature by calling the Analyze Image API with the visualFeatures parameter set to Description. The API will return a JSON response with a description field that contains a list of captions for the image, each with a confidence score.

Therefore, by using these two features of Computer Vision, you can achieve your app requirements with minimal development effort. You don't need to use any other services, such as object detection, content tags, or Custom Vision, which are designed for different purposes.

NEW QUESTION: 53

You have a flow that you plan to deploy to Azure Machine Learning managed online endpoints for real-time inference.

You need to modify the YAML configuration file for the deployment to ensure that trace data and system metrics are collected and sent to an Application Insights resource named my-app-insights.

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

NOTE: Each correct selection is worth one point.

Answer Area

Add a property named:

app_insights

app_insights_enabled

application_insights

azureml

environment_variables

Set the property value to:

true

my-app-insights

The name of the deployment

The instrumentation key of my-app-insights

The fully qualified resource ID of my-app-insights

Answer:

Answer Area



Microsoft

Add a property named:

app_insights

app_insights_enabled

application_insights

azureml

environment_variables

Set the property value to:

true

my-app-insights

The name of the deployment

The instrumentation key of my-app-insights

The fully qualified resource ID of my-app-insights

Explanation:

Answer Area



Microsoft

Add a property named:

app_insights_enabled

Set the property value to:

true

NEW QUESTION: 54

You have a question answering project in Azure Cognitive Service for Language.
You need to move the project to a Language service instance in a different Azure region.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

From the new Language service instance, import the project file.

From the new Language service instance, enable custom text classification.


From the new Language service instance, train and publish the project.

From the original Language service instance, export the existing project.

From the new Language service instance, regenerate the keys.

From the original Language service instance, train and publish the model.

Answer Area



Microsoft


⏪

⏩

⏪

⏩

Answer:



Comprehensive Detailed Explanation

To move a Question Answering project to a different Azure region, you migrate the project data and then re- deploy it in the target region:

Export the existing project (source region).

In Language Studio (Question Answering), you first export the project. This produces a package (JSON/ZIP) containing your knowledge bases, sources, synonyms, and settings. Exporting is the supported way to move a project between resources/regions.

Import the project into the new Language service instance (target region).

In the target region's Language service instance, use Language Studio to import the previously exported file.

This recreates the project (KBs and configuration) under the new resource.

Train and publish from the new instance.

After import, you must train (build) and then publish the project so that an endpoint is created under the new regional resource. Clients can then point to the new endpoint and key.

Actions like "enable custom text classification" or "regenerate the keys" are not required for moving a Question Answering project, and training/publishing from the original instance won't deploy to the target region.

Microsoft Azure AI Solution References

Export and import Question Answering projects in Language Studio (move between resources/regions):

Microsoft Learn - Export and import projects (Question Answering) .

<https://learn.microsoft.com/azure/ai-services/language-service/question-answering/how-to/export-import- project> Train and publish a Question Answering project:

Microsoft Learn - Create, train, and publish projects (Question Answering) .

<https://learn.microsoft.com/azure/ai-services/language-service/question-answering/how-to/create-test-deploy> Language service overview and regional deployment concepts:

Microsoft Learn - What is Azure AI Language?

<https://learn.microsoft.com/azure/ai-services/language-service/overview>

NEW QUESTION: 55

You have a computer that contains the files shown in the following table.

Name	Format	Length (mins)	Size (MB)
File1	MP4	34	1,500
File2	AVI	500	1,700
File3	MP3	300	980
File4	MP4	350	2,800

Which files can you upload and analyze by using Azure AI Video Indexer?

- A. File1 only
- B. File3only
- C. File1 andFile3 only
- D. File1, File2, and File3 only
- E. File1, File2, File3, and File4

Answer: C (LEAVE A REPLY)

Azure AI Video Indexer (VI) imposes limits on both file duration and file size:

Duration: Up to 6 hours for all presets; Basic Audio preset supports up to 12 hours.

Size: When uploading from your device, the limit is 2 GB (uploading by URL allows up to 30 GB).

Formats: MP4, AVI (video) and MP3 (audio codec/file) are supported.

Evaluating each file:

File1 - MP4, 34 mins, 1,500 MB: Within 6-hour duration and under 2 GB # Supported.

File2 - AVI, 500 mins (#8h20m), 1,700 MB: Exceeds the 6-hour limit # Not supported.

File3 - MP3, 300 mins (5h), 980 MB: Under 6 hours (and also under the 12-hour Basic Audio limit) and under 2 GB # Supported.

File4 - MP4, 350 mins (#5h50m), 2,800 MB: Duration is under 6 hours but size exceeds 2 GB for device uploads # Not supported (unless uploaded by URL, which the question doesn't state).

Therefore, the files you can upload and analyze (under standard device upload constraints) are File1 and File3 only.

Microsoft References

Support matrix & limits - file duration (6h for all presets; 12h for Basic Audio), device upload size limit 2 GB; supported formats (MP4, AVI, WAV, etc.) and audio codecs including MP3. Microsoft Learn Release notes - increase of duration limit from 4h to 6h (and 12h for Basic Audio). Microsoft Learn Upload guidance - device upload 2 GB vs 30 GB via URL.

NEW QUESTION: 56

You are building a model to detect objects in images.

The performance of the model based on training data is shown in the following exhibit.



Answer:



Explanation:

The percentage of false positives is [0].

The value for the number of true positives divided by the total number of true positives and false negatives is [100].

The screenshot shows the Performance tab of a Custom Vision object detection model after training.

Two key performance metrics are in focus:

False Positives (FP):

False positives occur when the model incorrectly predicts an object that isn't present.

In the provided exhibit, the model shows 0% false positives.

This means it did not incorrectly classify any non-object as an object.

Recall (True Positives ÷ (True Positives + False Negatives)):

This measures how many of the actual objects were correctly detected.

Formula: Recall = TP / (TP + FN).

In the performance report, recall is 100%, meaning the model correctly detected all the actual objects (no false negatives).

Evaluating the answer choices:

False positives: 0%

Recall (TP ÷ (TP + FN)): 100%

Correct Answers:

False positives # 0

Recall # 100

Microsoft References

Custom Vision performance metrics

Precision, recall, and accuracy in machine learning

NEW QUESTION: 57

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 = "FF956C68B83B21838691ABD200A4C686"
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
...

```



```
... , str(body), headers)
```

Answer:
Answer Area

```
...
api_key = "FF956C68B83B21838691ABD200A4C686"
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
...

```



```
... , str(body), headers)
```

Explanation:

API Endpoint Selection

The constraint states that the content of the text must remain within the Americas Azure geography .

" api.cognitive.microsofttranslator.com " : This is the global endpoint , which doesn ' t guarantee data residency in a specific geography.

" api-apc.cognitive.microsofttranslator.com " : This endpoint refers to the Asia Pacific geography .

" api-nam.cognitive.microsofttranslator.com " : This endpoint refers to the North and South America (NAM) geography .

Therefore, to meet the data residency requirement for the Americas, you must select " api-nam.cognitive.

microsofttranslator.com " .

2. Path/Query String Selection

The requirement is to translate the text to a single language . The input text can be in " one of many languages, " meaning the source language isn ' t fixed and should be detected automatically. The language to translate to is implied to be English (en) based on the options available and common translation scenarios.

" /translate?from=en " : This is incorrect because it specifies the source language as English, which contradicts the " text can be in one of many languages " requirement (implying automatic source language detection).

" /translate?suggestedFrom=en " : This is not the standard way to initiate a translation request and seems to suggest a source language, which is unnecessary if auto-detection is desired/possible.

" /translate?to=en " : This is the correct and standard way to request a translation where the source language is automatically detected , and the target language is English (en) . This fulfills the translation requirement.

" /detect?to=en " / " /detect?from=en " : These endpoints are used for language detection and not for the translation itself. The prompt explicitly asks to develop code to translate the text.

NEW QUESTION: 58

You have 100,000 documents.

You are building an app that will identify city names in each document by using Azure AI Language.

You need to test the detection client.

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

NOTE: Each correct selection is worth one point.

```

static void EntityRecognitionExample(
{
    var response = client.RecognizeEntities("I had a great experience visiting Contoso in Redmond.");
    Console.WriteLine("Named Entities:");
    foreach (var entity in response.Value)
    {
        if (entity.
        Console.WriteLine($"\\tText: {entity.Text},\\tResult: {entity.ConfidenceScore}");
    }
}
}

```

TextAnalyticsClient client)

DocumentAnalysisClient client)

FormRecognizerClient client)

QuestionAnsweringClient client)

TextAnalyticsClient client)

Category

Category

KeyValuePairs

SubCategory

TextAppearance



Answer:



Microsoft

```
static void EntityRecognitionExample(  
    TextAnalyticsClient client)  
{  
    var response = client.RecognizeEntities("I had a great experience visiting Contoso in Redmond."  
    Console.WriteLine("Named Entities:");  
    foreach (var entity in response.Value)  
    {  
        if (entity.  
            Console.WriteLine($"{entity.Text}, {entity.ConfidenceScore}");  
    }  
}
```

TextAnalyticsClient client)
DocumentAnalysisClient client)
FormRecognizerClient client)
QuestionAnsweringClient client)
TextAnalyticsClient client)

Category
Category
KeyValuePair
SubCategory
TextAppearance

Explanation:



To identify city names with Azure AI Language, you use the Text Analytics client and its Named Entity Recognition (NER) capability. In the .NET SDK, create a `TextAnalyticsClient` and call `RecognizeEntities(...)`.

This returns a `CategorizedEntityCollection`, where each `CategorizedEntity` has properties such as `Text`, `Category`, `SubCategory`, and `ConfidenceScore`.

City mentions are recognized under the broader `Location` category, and the more specific type is surfaced in `SubCategory` (commonly "City"). Therefore, iterating the returned entities and filtering with `entity.SubCategory.Contains("City")` correctly locates city names.

Using `DocumentAnalysisClient` or `FormRecognizerClient` would be incorrect because those are for document/form extraction scenarios, not general-language NER. `QuestionAnsweringClient` is for Q&A over knowledge sources. Hence the correct client is `TextAnalyticsClient`, and the correct discriminator is `SubCategory`.

Microsoft Azure AI References (titles only)

- * Azure AI Language - Text Analytics client library: Recognize Entities
- * Azure AI Language - Named Entity Recognition concepts and categories
- * Azure AI Language - CategorizedEntity model (Category, SubCategory, ConfidenceScore)

NEW QUESTION: 59

You are building an app that will use Azure AI to monitor workspaces for safety regulation compliance.

You need to recommend a service that meets the following requirements:

Generates alerts when employees enters high-risk areas

Monitors video feeds in real time

Minimizes development effort

What should you recommend?

- A. Object detection in Azure Custom Vision
- B. Azure Vision in Foundry Tools Spatial Analysis
- C. Azure AI Video Indexer
- D. Azure Vision in Foundry Tools Image Analysis

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 60

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

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

Answer: A ([LEAVE A REPLY](#))

* In Power BI, reports are interactive, visual, and support KPIs, drill-through, and drill-down functionality.

* Power Apps is for app development, not analytics.

* A view is a database object, not a visualization solution.

* Dataflow is for ETL/ELT (data preparation), not reporting.

The answer: A

Reference: Power BI Reports

NEW QUESTION: 61

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 ([LEAVE A REPLY](#))

Scenario Recap:

* Goal: Ensure app1 (on VM1 in vnet1) can connect to service1 (Azure Cognitive Search) without routing over the public internet.

* Proposed Solution: Deploy service1 with a public endpoint, and configure an IP firewall rule.

Analysis:

* IP firewall rules restrict which public IPs can access a service.

* However, traffic still routes over the public internet. The firewall only filters access, it does not provide private connectivity.

* To avoid internet traffic, you must use Private Link (private endpoint).

The answer: B. No

Reference: Secure access to Cognitive Search with Private Link

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (**416** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 62

You have a local folder that contains the files shown in the following table.

Name	Format	Length (mins)	Size (MB)
File1	WMV	34	400
File2	AVI	90	1,200
File3	MOV	300	980
File4	MP4	80	1,800

You need to analyze the files by using Azure Ai Video Indexer. Which files can you upload to the Video Indexer website?

- A. File1,FileZ and File4 only
- B. File1, and File2 only
- C. File1, File2, and File3 only
- D. File1, File2. File3 and Fi1e4
- E. File1, and File3 only

Answer: E (LEAVE A REPLY)

You have files with the following formats:

File1: WMV (34 mins, 400 MB)

File2: AVI (90 mins, 1,200 MB)

File3: MOV (300 mins, 980 MB)

File4: MP4 (80 mins, 1,800 MB)

You must determine which files can be uploaded to Azure AI Video Indexer.

Key Rules (Video Indexer upload limitations):

Supported formats: MP4, MOV, WMV, AVI, M2TS.

Maximum file length: 4 hours (240 minutes).

Maximum file size: 2 GB (2048 MB) for website uploads.

Check each file:

File1 (WMV, 34 min, 400 MB): Supported format, within time and size # # Allowed.

File2 (AVI, 90 min, 1,200 MB): Supported format, within time and size # # Allowed.

File3 (MOV, 300 min, 980 MB): Supported format, but exceeds 240 min limit # # Not Allowed.

File4 (MP4, 80 min, 1,800 MB): Supported format, within size (under 2 GB), within time # # Allowed.

The Answer:

A). File1, File2, and File4 only

Microsoft Reference:

Video Indexer upload limitations

NEW QUESTION: 63

You plan to build an agent that will combine and process multiple files uploaded by users.

You are evaluating whether to use the Azure AI Agent Service to develop the agent.

What is the maximum size of all the files that can be uploaded to the service?

- A.** 1 GB
- B.** 10 GB
- C.** 100 GB
- D.** 1 TB

Answer: C ([LEAVE A REPLY](#))

For Azure AI Agent Service, the documented quota for the maximum total size of all files uploaded for agents has been listed as 100 GB. Individual files are limited to 512 MB each. These limits are relevant when you plan an agent that must combine and process multiple user-uploaded files.

(Note: Quotas can evolve; check the current "Quotas and limits" page for your region.) References (Microsoft Docs / Q & A):

Microsoft Q & A (Agent Service limits: "Max size for all uploaded files for agents 100 GB"; per-file 512 MB). Microsoft Learn File Search tool how-to (per-file 512 MB guidance). Microsoft Learn

NEW QUESTION: 64

Which scenario is an example of a streaming workload?

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

Answer: (SHOW ANSWER)

* A streaming workload means data is generated continuously and ingested in near real time.

* Telemetry from edge/IoT devices is a classic example of streaming workloads because the devices constantly push data (sensor readings, logs, metrics).

* A. sending transactions daily # batch workload (not streaming).

* B. sending metadata every 30 minutes # periodic batch, not continuous streaming.

* C. sending old transactions to archive # cold data archival, not streaming.

The answer: D

Reference: Azure Stream Analytics overview

NEW QUESTION: 65

You have an Azure AI agent solution.

You plan to create an agent-based app named App1 that will analyze and summarize data for users and generate data-driven recommendations. App1 will be used by non-technical business users and must adapt to new and unforeseen business challenges and improve its performance over time.

You need to identify which type of agent to use in App1. The solution must meet the following requirements:

* Adapt and improve the agents ' performance over time based on user feedback.

* Provide tailored recommendations to help users make informed decisions.

* Provide the best possible performance of the app.

Which agent type should you identify?

- A.** prompt-and-response
- B.** autonomous
- C.** task automation
- D.** cognitive

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 66

You develop a Conversational Language Understanding model by using Language Studio. During testing, users receive incorrect responses to requests that do NOT relate to the capabilities of the model.

You need to ensure that the model identifies spurious requests.

What should you do?

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 67

You have an Azure OpenAI model.

You have 500 prompt-completion pairs that will be used as training data to fine-tune the model.

You need to prepare the training data.

Which format should you use for the training data file?

- A.** XML
- B.** JSONL
- C.** CSV
- D.** TSV

Answer: B ([LEAVE A REPLY](#))

For Azure OpenAI fine-tuning, training data must be provided as a JSON Lines (.jsonl) file. Each line represents one example. Because your data is described as prompt-completion pairs, the required shape is typically: { " prompt " : " < your prompt > " , " completion " : " < ideal completion > " } (Chat model fine-tuning instead uses a messages array per line, but the file format is still JSONL.) JSONL is the only accepted format for upload to the fine-tuning job.

Microsoft References

Azure OpenAI Service - Fine-tuning: data preparation and formats (JSONL; prompt-completion and chat messages).

Azure OpenAI Service - How to fine-tune models (training files must be .jsonl).

NEW QUESTION: 68

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 ([LEAVE A REPLY](#))

Your app uses the primary admin key for index management. If it's suspected compromised, the least- downtime rotation is:

Regenerate the secondary admin key to ensure it's known only to you.

Switch the app to use the secondary key (no downtime-primary still works until you cut it off).

Regenerate the primary key to invalidate the compromised credential.

Query keys are read-only and cannot manage indexes, so switching to a query key would not meet the requirement. Regenerating the primary first would break the app until you update it.

References (Microsoft Azure Cognitive Search)

Admin (primary/secondary) vs. query keys and key rotation; key regeneration invalidates only the selected key.

Guidance to use the alternate admin key while rotating the other to avoid downtime.

NEW QUESTION: 69

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 (LEAVE A REPLY)

Obtain values to connect your bot to the knowledge base

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

Reference:

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

NEW QUESTION: 70

Select the answer that correctly completes the sentence.

Answer Area

Microsoft

A data analyst

A data engineer

A data scientist

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

Answer:

Answer Area

Microsoft

A data analyst

A data engineer

A data scientist

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

Explanation:

Data Analyst.

When dealing with data roles in an organization, responsibilities are divided among data analysts, data engineers, and data scientists. Let's review them:

Data Analyst

Focuses on understanding business requirements and ensuring that the right rules and logic are applied to data so it becomes meaningful for decision-making. They work closely with stakeholders to define business rules such as data validation rules, aggregation logic, and reporting metrics. This makes them responsible for identifying which business rules must be applied to data.

Data Engineer

Responsible for designing, building, and maintaining the infrastructure and pipelines that move and process data. They ensure data availability, reliability, and scalability, but do not usually define business rules.

Data Scientist

Focuses on advanced analytics, predictive modeling, and machine learning. They use data prepared by engineers and analysts but are not typically the ones who define the business rules that govern how raw data should be interpreted. Therefore, the correct answer is Data Analyst.

The Answer: A data analyst

Microsoft References

Data roles and responsibilities in Azure

What is a Data Analyst?

Responsibilities of Data Engineers vs Data Scientists vs Data Analysts

NEW QUESTION: 71

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

Values

addTargetLanguage

speechSynthesisLanguage

speechRecognitionLanguage

voiceName

Answer Area

```
var translationConfig = SpeechTranslationConfig.FromSubscription(SPEECH__SUBSCRIPTION__KEY, SPEECH__SERVICE__REGION);  
translationConfig. = "en-US";  
translationConfig. ("de");
```

Answer:

Values

addTargetLanguage

speechSynthesisLanguage

speechRecognitionLanguage

voiceName

Answer Area

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

Explanation:

Values

addTargetLanguage

speechSynthesisLanguage

speechRecognitionLanguage

voiceName

Answer Area

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



You are configuring speech-to-speech translation in an app. The workflow is:

Recognize spoken input language (in this case, English).

Translate the recognized text into the target language (in this case, German).

(Optional) Synthesize speech in the target language.

Step 1 - SpeechRecognitionLanguage

This property defines the input language for speech recognition.

Since the input is English, it should be set to " en-US " .

Step 2 - AddTargetLanguage

This method specifies the output translation language(s).

Since the requirement is to translate into German, use " de " .

Other options (incorrect here):

speechSynthesisLanguage # used if you want to output translated text as synthesized speech in the target language voice. Not required in this snippet.

voiceName # specifies the actual voice to use for speech synthesis, not translation.

Correct Completed Code Snippet:

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

Microsoft References Speech Translation with Azure AI Speech SDK SpeechTranslationConfig Class Reference

NEW QUESTION: 72

You have an Azure subscription that contains an Azure AI Content Understanding resource named cu1.

You need to create a custom analyzer that will analyze documents.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
curl -i -X POST "https://cu1.services.ai.azure.com/contentunderstanding/analyzers/myInvoice?api-version=2024-12-01-preview"
```



Answer:



Explanation:



Explanation:

curl -i -X: The POST method is appropriate because you are creating a new custom analyzer, which involves sending a request to create a resource. POST is used for creating new resources in REST APIs.

https:// < cu1 > .services.ai.azure.com/: The base URL should be updated to include the contentunderstanding endpoint, as this is the specific service within the Azure AI Content Understanding resource (cu1) that handles analyzer-related operations. Replace < cu1 > with the actual resource name or endpoint provided by your Azure subscription.

/analyzers/myInvoice?api-version=2024-12-01-preview: This part of the URL is correct as it specifies the endpoint for creating a custom analyzer (/analyzers) with the name myInvoice and the API version 2024-12-01-preview. No changes are needed here, assuming the analyzer name and version are intended as shown.

NEW QUESTION: 73

You are building an agent by using custom question answering in the Azure Language in Foundry Tools service. You upload a product catalog and price list and train the model.

The agent responds correctly when a user asks the following question: " What is the price of Product1? " The agent responds incorrectly when a user asks the following question: " Which colors of Product1 are available? "

You need to improve the accuracy of the agent. What should you do?

A. Modify the original documents and retrain the model.

B. Add alternative phrasing to the question and response pair.

C. Modify the system prompt.

D. Add a new question and response pair.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 74

You are examining the Language service output of an application.

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

The response contains the data shown in the following table.

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

Which Language service API is used to analyze the Text?

- A.** Entity Linking
- B.** Named Entity Recognition

C. Key Phrase Extraction

D. Sentiment Analysis

Answer: B ([LEAVE A REPLY](#))

The given text:

" Our tour guide took us up the Space Needle during our trip to Seattle last week. " The Language service output shows:

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

Why the correct answer is Named Entity Recognition (NER):

NER identifies named entities (people, places, dates, events, organizations, etc.) in text and classifies them into categories with confidence scores.

In the table, we see categories such as PersonType, Location, Event, DateTime # these are classic outputs of NER.

Why not the others:

A). Entity Linking # Identifies entities and links them to a knowledge base (like Wikipedia). Output would include links/IDs rather than just categories.

C). Key Phrase Extraction # Returns important phrases ("tour guide", "Space Needle", "trip", "Seattle", "last week") but does not assign categories like PersonType, Location, DateTime.

D). Sentiment Analysis # Returns sentiment scores (positive/neutral/negative), not entity categories.

The Answer:

B). Named Entity Recognition

Microsoft References

Named Entity Recognition in Azure AI Language

Entity Linking vs Named Entity Recognition


NEW QUESTION: 75

You have an app that uses the AI Language custom question answering service.


You need to ad alternatives for the word testing by using the Authoring API.

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

NOTE: Each correct selection is worth one point.



Answer:



Explanation:



In Custom Question Answering (part of Azure AI Language), if you want to define alternative words or spellings for a specific term (e.g., testing , trials , evaluate), you must use the alterations property. Each alteration object contains a list of alternative values for the same meaning.

Example:

```
{
  "alterations" : [
    {
      "alterations" : [
        "testing",
        "trials",
        "evaluate"
      ]
    }
  ]
}
```

phrases is used for pattern matching in question-answer pairs.

synonyms is used for synonyms within LUIS (Language Understanding), not for QnA alterations.

value is used inside synonyms definitions in some contexts but not here.

Correct JSON Payload Completion:

The JSON should be:

```
{
  "alterations" : [
    {
      "alterations" : [
        "testing " ,
        " trials " ,
        " evaluate "
      ]
    }
  ]
}
```

Verified Answer:

Left box: " alterations "

Right box: " alterations "

Microsoft References:

Manage alterations in Custom Question Answering (Azure AI Language)

Authoring API reference for alterations

NEW QUESTION: 76

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 ([LEAVE A REPLY](#))

Scenario Recap:

Application: Flower classification.

Solution: Add new images and labels to the existing model. Retrain and publish the model.

Analysis:

This is the best practice workflow:

Add new labeled training images.

Retrain the model with updated dataset.

Publish the retrained model.

This clearly meets the goal.

The Answer: A. Yes

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (**416** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 77

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

You need to configure the chatbot to present a list of available options. The solution must ensure that an image is provided for each option.

Which two features should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. an Azure function
- B. an adaptive card
- C. an entity
- D. a dialog
- E. an utterance

Answer: (SHOW ANSWER)

- * To present a list of options with images, you use Adaptive Cards, which allow for rich content such as images, buttons, and formatted text in the conversation.
- * These cards are typically surfaced within the context of a dialog, which controls the flow of the conversation.
- * Azure functions are unrelated here.
- * Entities and utterances belong to natural language understanding, not UI presentation.

The answer: B and D

Reference: Adaptive Cards in Bot Framework Composer

NEW QUESTION: 78

You are building an app that will automatically translate speech from English to French, German, and Spanish by using Azure AI service.

You need to define the output languages and configure the Azure AI Speech service.

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

NOTE: Each correct selection is worth one point.

Answer Area



Microsoft

```
static async Task TranslateSpeechAsync()
{
    var config = SpeechTranslationConfig.FromSubscription(SPEECH_SUBSCRIPTION_KEY, SPEECH_SERVICE_REGION);
    var languages = new List<string>
    {
        ["fr","de","es"],
        ["en-GB"],
        ["en","fr","de","es"],
        ["fr","de","es"],
        ["French","German","Spanish"]
    };

    languages.ForEach(config.AddTargetLanguage);

    using var recognizer = new
    {
        TranslationRecognizer,
        IntentRecognizer,
        SpeakerRecognizer,
        SpeechSynthesizer,
        TranslationRecognizer
    };
    ...
}
```

Answer:

Answer Area



Microsoft

```
static async Task TranslateSpeechAsync()
{
    var config = SpeechTranslationConfig.FromSubscription(SPEECH_SUBSCRIPTION_KEY, SPEECH_SERVICE_REGION);
    var languages = new List<string>
    {
        ["fr","de","es"],
        ["en-GB"],
        ["en","fr","de","es"],
        ["fr","de","es"],
        ["French","German","Spanish"]
    };

    languages.ForEach(config.AddTargetLanguage);

    using var recognizer = new
    {
        TranslationRecognizer,
        IntentRecognizer,
        SpeakerRecognizer,
        SpeechSynthesizer,
        TranslationRecognizer
    };
    ...
}
```

Explanation:

You are asked to build an app that translates speech from English into French, German, and Spanish using Azure AI Speech Translation.

Step 1 - Define target languages

The SpeechTranslationConfig object requires language codes (not full names like " French " , " German " , " Spanish ").

Valid language codes:

French # " fr "

German # " de "

Spanish # " es "

Therefore, the correct syntax for the output languages is:

```
var languages = new List < string > { " fr " , " de " , " es " };
```

Step 2 - Choose the recognizer type

IntentRecognizer : Used for intent recognition with LUIS, not translation.

SpeakerRecognizer : Used for speaker identification/verification, not translation.

SpeechSynthesizer : Used for text-to-speech, not translation.

TranslationRecognizer : Used specifically for speech-to-speech or speech-to-text translation.

Therefore, the correct recognizer is TranslationRecognizer.

The Answer:

Output Languages: { " fr " , " de " , " es " }

Recognizer: TranslationRecognizer

Microsoft References

Speech Translation with Azure AI Speech

TranslationRecognizer Class (C# API)

Supported language codes

NEW QUESTION: 79

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 ([**LEAVE A REPLY**](#)**)**

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>

NEW QUESTION: 80

You have an Azure subscription.

You plan to build an app that will use the Azure AI DALL-E model.

You need to deploy the model.

What should you use?

A. Azure OpenAI Studio and Azure Command-Line Interface (CU)

B. the Azure portal and Microsoft Graph API

C. the Azure SDK for JavaScript and Azure Machine Learning Studio

D. the Azure SDK for Python and PowerShell cmdlets

Answer: A ([**LEAVE A REPLY**](#)**)**

NEW QUESTION: 81

You have an Azure DevOps pipeline named Pipeline1 that is used to deploy an app. Pipeline1 includes a step that will create an Azure AI services account.

You need to add a step to Pipeline1 that will identify the created Azure AI services account. The solution must minimize development effort.

Which Azure Command-Line interface (CLI) command should you run?

- A.** Az resource link
- B.** Az account list
- C.** Az cognitivesservices account network-rule
- D.** As cognitiveservices account show

Answer: D ([LEAVE A REPLY](#))

- * The requirement is to identify the created Azure AI services account within a DevOps pipeline.
- * The az cognitiveservices account show command returns the details (name, id, location, etc.) of an existing Azure AI service account.
- * az resource link creates or manages links between resources, not what we need.
- * az account list lists Azure subscriptions, not resources.
- * az cognitiveservices account network-rule manages network access rules but does not return account identity.

Microsoft References:

- * az cognitiveservices account show

NEW QUESTION: 82

You are designing a content management system.

You need to ensure that the reading experience is optimized for users who have reduced comprehensive and learning differences, such as dyslexia.

Which Azure service should you include in the solution?

- A.** Azure AI Translator
- B.** Azure AI Document Intelligence
- C.** Azure AI Immersive Reader
- D.** Azure AI Language

Answer: ([SHOW ANSWER](#)**)**

- * The requirement is to optimize the reading experience for users with dyslexia and learning differences.
- * Azure AI Immersive Reader is specifically designed to make text more accessible, improving comprehension by enabling features like text-to-speech, line focus, and translation.
- * Translator (A) is for translating text, not accessibility.
- * Document Intelligence (B) is for extracting information from documents.
- * Azure AI Language (D) is for NLP tasks, not improving reading accessibility.

Microsoft References:

- * Azure AI Immersive Reader overview

NEW QUESTION: 83

You have an Azure subscription that contains an Azure AI Content Safety resource named resource1.

You need to add a custom category to resource1.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
curl -X PUT "<endpoint>/contentsafety/text/categories/learning-advice?api-version=2024-02-15-preview" \
-H "Ocp-Apim-Subscription-Key: <api_key>" \
-H "Content-Type: application/json" \
-d "{
  \"categoryName\": \"learning-advice\",
  \"prompt\": \"text prompts about learning advice while preparing for an exam.\",
  \"definition\": \"\",
  \"sampleBlobUrl\": \"https://<your-azure-storage-url>/container/learning-advice.jsonl\"
}
```

\"prompt\":

\"definition\":

\"prompt\":

\"text\":

sampleBlobUrl

blocklistName

outputType

sampleBlobUrl

Answer:

Answer Area

```
curl -X PUT "<endpoint>/contentsafety/text/categories/learning-advice?api-version=2024-02-15-preview" \
-H "Ocp-Apim-Subscription-Key: <api_key>" \
-H "Content-Type: application/json" \
-d "{
  \"categoryName\": \"learning-advice\",
  \"prompt\": \"text prompts about learning advice while preparing for an exam.\",
  \"definition\": \"\",
  \"sampleBlobUrl\": \"https://<your-azure-storage-url>/container/learning-advice.jsonl\"
}
```

\"prompt\":

\"definition\":

\"prompt\":

\"text\":

sampleBlobUrl

blocklistName

outputType

sampleBlobUrl

Explanation:

Use " definition " for the first field and " sampleBlobUrl " for the second field When creating a custom text category in Azure AI Content Safety, you call the PUT /contentsafety/text /categories/{categoryName} endpoint and provide a JSON body that includes:

categoryName : the identifier of your category.

definition : a natural-language description of what the category covers (e.g., "text prompts about learning advice while preparing for an exam").

sampleBlobUrl : a pointer to labeled sample data (typically a JSONL file in Azure Storage) used to help the service understand the category.

Keys like prompt or text are used when analyzing content (for example, text:detect), not when defining a category. Likewise, blocklistName applies to blocklist operations, and outputType is not part of the category-creation schema. Therefore, the correct payload fields to complete the cURL are definition and sampleBlobUrl .

Microsoft Azure AI References (titles only)

Azure AI Content Safety - Create or update custom text categories (REST) Azure AI Content Safety - Custom categories concepts and training data (JSONL in blob storage) Azure AI Content Safety - Text analysis vs. category management endpoints

NEW QUESTION: 84

Select the answer that correctly completes the sentence.

Answer Area

The

ALTER
JOIN
SET
WHERE

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



Answer:

Answer Area

The

ALTER
JOIN
SET
WHERE

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



Explanation:

Answer Area

The

WHERE

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



In SQL, DML (Data Manipulation Language) statements include SELECT , INSERT , UPDATE , and DELETE . Often, you need to filter which rows are affected or returned. That filtering is done with the WHERE clause.

Let's review the options:

ALTER

Used in Data Definition Language (DDL), not DML.

Alters the structure of a database object (e.g., add/remove columns, modify constraints).

Not used to filter rows.

ON

Used in JOIN operations to specify the relationship between tables.

Does not specify row-level filtering criteria in general DML.

SET

Used in the UPDATE statement to assign new values to columns.

Does not specify which rows to update.

WHERE

Used in DML statements to specify criteria for rows that must match.

Examples:

SELECT * FROM Customers WHERE Country = ' USA ' ;

DELETE FROM Orders WHERE OrderDate < ' 2020-01-01 ' ;

UPDATE Products SET Price = Price * 1.1 WHERE CategoryID = 5;

Thus, the correct clause is WHERE.

The Answer: WHERE

Microsoft References

SQL WHERE clause

SQL UPDATE with WHERE
SQL DELETE with WHERE

NEW QUESTION: 85

You have a Language Understanding solution that runs in a Docker container.
You download the Language Understanding container image from the Microsoft Container Registry (MCR).
You need to deploy the container image to a host computer.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

From the Language Understanding portal, retrain the model.

From the host computer, run the container and specify the input directory.

From the Language Understanding portal, export the solution as a package file.

From the host computer, move the package file to the Docker input directory.

From the host computer, build the container and specify the output directory.

Answer Area

Answer:

Actions

From the Language Understanding portal, retrain the model.

From the host computer, run the container and specify the input directory.

From the Language Understanding portal, export the solution as a package file.

From the host computer, move the package file to the Docker input directory.

From the host computer, build the container and specify the output directory.

Answer Area

Explanation:

Actions

From the Language Understanding portal, retrain the model.

From the host computer, run the container and specify the input directory.

Answer Area

1

From the Language Understanding portal, export the solution as a package file.

2

From the host computer, move the package file to the Docker input directory.

3

From the host computer, build the container and specify the output directory.

You are deploying a Language Understanding (LUIS) container. When running LUIS in Docker, the model must first be exported from the Language Understanding portal, and then provided to the container at runtime.

Step-by-step reasoning:

From the Language Understanding portal, export the solution as a package file.

The trained LUIS model must be exported from the portal into a .json package file.

This is required because the container cannot access the hosted service directly.

From the host computer, move the package file to the Docker input directory.

Containers expect the model to be available locally.

The package file is placed into the input directory that the container maps for models.

From the host computer, run the container and specify the input directory.

When starting the container, you specify --volume < local_input > :/input so the container has access to the package file.

This makes the exported model available inside the container for processing.

Why not the other options?

Retain the model in the portal is not sufficient; the container cannot pull directly from the cloud.

Build the container and specify the output directory is not required; the container image is already available from MCR and is not custom-built for this step.

Correct Answer Order:

Export the solution as a package file.

Move the package file to the Docker input directory.

Run the container and specify the input directory.

Microsoft References

Run LUIS containers

Use containers with Azure AI services

NEW QUESTION: 86

You have a mobile app that manages printed forms.

You need the app to send images of the forms directly to Forms Recognizer to extract relevant information.

For compliance reasons, the image files must not be stored in the cloud.

In which format should you send the images to the Form Recognizer API endpoint?

A. raw image binary

B. form URL encoded

C. JSON

Answer: A ([LEAVE A REPLY](#))

The requirement is to process images of forms with Azure Form Recognizer (now part of Azure AI Document Intelligence) while ensuring that the image files are not stored in the cloud.

When sending images directly to the Form Recognizer REST API endpoint, you have two options:

Send the file as binary data in the request body (for local file upload).

Send a URL reference (if the file is stored in Azure Blob Storage or another public location).

Since the requirement explicitly states that the image files must not be stored in the cloud, you cannot send a URL. Instead, you must send the raw binary content of the file directly in the request body.

Option A - raw image binary

Correct. This allows you to directly send the local image file to the API without uploading/storing it in the cloud.

Option B - form URL encoded

Incorrect. Form Recognizer does not accept URL-encoded form data.

Option C - JSON

Incorrect. JSON is only used when providing the API with the URL of a cloud-stored document, not when sending the actual file.

The Answer: A. raw image binary

Microsoft References

Form Recognizer REST API - Analyze a document

Quickstart: Extract layout and text - Document Intelligence

" If your document is available locally, pass the file as binary data in the request body. "

NEW QUESTION: 87

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);
        }
    }
});
```

File
Stream
Uri
Url

AddFaceFromStreamAsync
AddFaceFromUrlAsync
CreateAsync
GetAsync



Answer:

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 (
            File
            Stream
            Uri
            Url
            t = File.OpenRead(imagePath))
        {
            await faceClient.PersonGroupPerson.
                AddFaceFromStreamAsync
                AddFaceFromUrlAsync
                CreateAsync
                GetAsync
                (personGroupId, personId, t);
        }
    }
});
```



Explanation:

Comprehensive Detailed Explanation

You're adding multiple images for each person into a PersonGroup using the Face API SDK. File.OpenRead(imagePath) returns a Stream, so the using declaration must declare Stream t = File.OpenRead (imagePath) . To upload an image from a stream to a person in a person group, call faceClient.PersonGroupPerson.AddFaceFromStreamAsync(personGroupId, personId, t); Putting it together: 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 (Stream t = File.OpenRead(imagePath))
        {
            await faceClient.PersonGroupPerson.AddFaceFromStreamAsync(
                personGroupId, personId, t);
        }
    }
});
```

This loops the images for each person, opens each as a stream, and adds the face to the specified person in the person group.

References

Microsoft Docs - Face API: Add a face to a person (AddFaceFromStreamAsync , AddFaceFromUrlAsync) in PersonGroupPerson operations.

https://learn.microsoft.com/azure/ai-services/face/ (see SDK/PersonGroupPerson methods) NET File.OpenRead returns Stream .

https://learn.microsoft.com/dotnet/api/system.io.file.openread

NEW QUESTION: 88

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 AI Search REST API. What should you do?

- A. Ingest the data in Logs into Azure Data Explorer.
- B. Migrate the data in HR to Azure Blob storage.
- C. Ingest the data in Logs into Azure Sentinel1.
- D. Export the data in Finance to Azure Data Lake Storage.

Answer: (SHOW ANSWER)

NEW QUESTION: 89

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

Upload the images by category.

Get suggested tags.

Upload all the images.

Group the images locally into category folders.

Review the suggestions and confirm the tags.

Tag the images manually.

Answer Area

←


→

↑

↓

Microsoft

Answer:



Explanation:

Comprehensive Detailed Explanation

You have 1,000 untagged images and want to retrain the Custom Vision model as quickly as possible. The key is to avoid manually tagging every image, since that is time-consuming.

Step 1: Upload all the images

First, bulk upload all the untagged images into the Custom Vision project.

At this stage, they will have no labels.

Step 2: Get suggested tags

The Custom Vision portal offers an auto-suggest feature that uses the current trained model to predict tags for new, untagged images.

This dramatically reduces manual tagging effort.

Step 3: Review the suggestions and confirm the tags

Auto-suggestions are not always 100% accurate, so you must confirm or correct them.

Once confirmed, the dataset is labeled properly and ready for retraining.

Why not the other options?

Upload the images by category / group into category folders: This is only helpful if the images already have known categories, but here they arrive untagged.

Tag the images manually: This would take much longer and contradicts the requirement to minimize time.

Correct Sequence

Upload all the images.

Get suggested tags.

Review the suggestions and confirm the tags.

Microsoft References

Custom Vision - Tag images

Custom Vision - Suggested tags with predictions

NEW QUESTION: 90

You have an Azure Cognitive Search resource named Search 1 that is used by multiple apps You need to secure Search 1. The solution must meet the following requirements:

* Prevent access to Search1 from the internet.

* Limit the access of each app to specific queries
What should you do? To answer, select the appropriate options in the answer area NOTE Each correct answer is worth one point.

Answer Area

To prevent access from the internet:

- Create a private endpoint.
- Configure an IP firewall.
- Create a private endpoint.
- Use Azure roles.

To limit access to queries:

- Use Azure roles.
- Create a private endpoint.
- Use Azure roles.
- Use key authentication.

Answer:

Answer Area

To prevent access from the internet:

- Create a private endpoint.
- Configure an IP firewall.
- Create a private endpoint.
- Use Azure roles.

To limit access to queries:

- Use Azure roles.
- Create a private endpoint.
- Use Azure roles.
- Use key authentication.

Explanation:

Answer Area

To prevent access from the internet: Create a private endpoint.

To limit access to queries: Use Azure roles.

You are securing an Azure Cognitive Search resource with two specific requirements:

Prevent access to Search1 from the internet

The correct way is to use a private endpoint.

A private endpoint connects the Cognitive Search resource directly into your virtual network (VNet), blocking public internet access.

Configuring an IP firewall can restrict access to specific IP ranges but does not fully prevent internet access.

Azure roles and authentication do not address network-level restrictions.

Limit the access of each app to specific queries

This is done via Azure role-based access control (RBAC).

You can assign different Azure roles or API keys with query rights to control what each app can do.
Key authentication alone only provides coarse-grained access (admin/query keys) but does not allow fine- grained per-app restrictions.
RBAC ensures principle of least privilege is enforced.

- The Answer:
- Prevent access from internet # Create a private endpoint
 - Limit access to queries # Use Azure roles
- Microsoft References
- Secure access to Azure Cognitive Search using private endpoints
 - Role-based access control in Azure Cognitive Search

NEW QUESTION: 91

You are building a language learning solution.

You need to recommend which Azure services can be used to perform the following tasks:

- * Analyze lesson plans submitted by teachers and extract key fields, such as lesson times and required texts.
- * Analyze learning content and provide students with pictures that represent commonly used words or phrases in the text

The solution must minimize development effort.

Which Azure service should you recommend for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Analyze lesson plans: Azure AI Document Intelligence

Analyze learning content: Immersive Reader

Answer:

Answer Area

Analyze lesson plans: Azure AI Document Intelligence

Analyze learning content: Immersive Reader

Explanation:

Answer Area

Analyze lesson plans: Azure AI Document Intelligence

Analyze learning content: Immersive Reader

Analyze lesson plans and extract key fields

Use Azure AI Document Intelligence (Form Recognizer) to automatically extract structured data (dates/times, required texts, names, etc.) from uploaded lesson-plan documents (PDFs/images). It provides prebuilt and custom extraction with minimal code via Document Intelligence Studio.

Analyze learning content and show pictures for common words/phrases

Use Immersive Reader, which includes features like Picture Dictionary and text comprehension supports directly in apps with minimal development. It can highlight words and show illustrative images to aid vocabulary learning-exactly the requirement.

Other options are not a fit:

Azure AI Custom Vision is for training image classifiers/detectors, not text comprehension.

Azure AI Search indexes and searches content; it doesn't provide picture dictionary/reading experience.

Microsoft Azure AI References

Document Intelligence (Form Recognizer) overview & extraction capabilities: [learn.microsoft.com/azure/ai- services/document-intelligence/overview](https://learn.microsoft.com/azure/ai-services/document-intelligence/overview) Immersive Reader - Picture Dictionary & reading features:

[learn.microsoft.com/azure/ai-services/immersive- reader/overview](https://learn.microsoft.com/azure/ai-services/immersive-reader/overview)

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (416 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 92

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: [
```

Attachment
ConversationUpdate
Typing

```
][Delay=3000]
```

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

```
[AttachmentLayout=
```

adaptivecard
carousel
thumbnail

```
[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]
```

```
user: I like the first one.
bot: Sure, pulling up more information.
bot: [Attachment=cards\watchProfileCard.json
```

```
user: That's nice! Thank you.
```

```
bot: Sure, you are most welcome!
```

adaptivecard
carousel
list

Answer:

user=User1

bot=watchbot

user: I want a new watch.

bot: [][Delay=3000]

Attachment
ConversationUpdate
Typing

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

bot: Here's what I found.

bot:

[AttachmentLayout= ]



adaptivecard
carousel
thumbnail

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

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

user: I like the first one.

bot: Sure, pulling up more information.

bot: [Attachment=cards\watchProfileCard.json

user: That's nice! Thank you.

bot: Sure, you are most welcome!

]

adaptivecard
carousel
list

Explanation:

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

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

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

Comprehensive Detailed Explanation along with All References available from Microsoft Azure AI Solution at the end of the explanation Typing: In a .chat transcript, you can simulate the bot's typing indicator with the Typing activity, optionally followed by a [Delay=...] . This matches the intent to pause before sending the next message.

AttachmentLayout=carousel: When sending multiple attachments (the two watch images), set the layout to carousel so the Emulator displays the attachments horizontally as a card carousel.

Adaptive Card attachment: The JSON file cards\watchProfileCard.json represents an Adaptive Card; specify adaptivecard so the Emulator renders the card.

Resulting .chat snippet (essentials):

```
user =User1
bot=watchbot
user : I want a new watch.
bot: [Typing][Delay= 3000 ]
bot: I can help you with that! Let me see what I can find.
bot: Her e ' s what I found.
bot:
[AttachmentLayout=carousel]
[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 adaptivecard]
user: That ' s nice! Thank you.
bot: Sure, you are most welcome!
```

References

Bot Framework Emulator transcript (.chat) activities and typing indicator: Bot Framework Activity schema; Emulator supports typing and delay in transcripts.

<https://learn.microsoft.com/azure/bot-service/bot-service-debug-emulator>

[https://learn.microsoft.com/azure/bot-service/rest-api/bot-framework-rest-connector-api-reference#activity-](https://learn.microsoft.com/azure/bot-service/rest-api/bot-framework-rest-connector-api-reference#activity-object-attachment-layout) object Attachment layout (carousel / list) for multiple attachments:

<https://learn.microsoft.com/azure/bot-service/nodejs/bot-builder-nodejs-send-rich-cards#carousel-layout> Adaptive Cards in Bot Framework:

<https://learn.microsoft.com/adaptive-cards/get-started/bots>

<https://learn.microsoft.com/azure/bot-service/bot-builder-howto-add-media-attachments#adaptive-cards>

NEW QUESTION: 93

You have an Azure subscription.

You need to build an app that will compare documents for semantic similarity. The solution must meet the following requirements:

- * Return numeric vectors that represent the tokens of each document.

- * Minimize development effort.

Which Azure OpenAI model should you use?

A. GPT-3.5

B. embeddings

C. DALL-E

D. GPT-4

Answer: B ([LEAVE A REPLY](#))

Analysis:

GPT-3.5 (A): A generative model; not specialized for embeddings # #

Embeddings (B): Purpose-built models (like text-embedding-ada-002) return numeric vector representations of text, ideal for similarity comparisons # # DALL-E (C): Used for image generation, not semantic similarity # #

GPT-4 (D): Powerful LLM, but not efficient for embeddings tasks # # The Answer:

B). embeddings

Reference:

Azure OpenAI embeddings

NEW QUESTION: 94

You have an Azure subscription that contains an Azure OpenAI resource named All and a user named User1 You need to ensure that User1 can perform the following actions in Azure OpenAI Studio;

- * Identify resource endpoints.

- * View models that are available for deployment.

- * Generate text and images by using the deployed models

The solution must follow the principle of least privilege Which role should you assign to User1?

A. Cognitive Services OpenAI User

B. Cognitive Services Contributor

C. Contributor

D. Cognitive Services OpenAI Contributor

Answer: A ([LEAVE A REPLY](#))

The Cognitive Services OpenAI User role grants least-privilege read/use permissions: view resource endpoints, list available models, use existing deployments in Azure OpenAI Studio (e.g., generate text and images).

It cannot create or manage deployments (that would require Cognitive Services OpenAI Contributor), satisfying least-privilege.

References

Azure OpenAI built-in roles: OpenAI User vs OpenAI Contributor .

<https://learn.microsoft.com/azure/ai-services/openai/how-to/role-based-access-control> Azure RBAC for Azure AI services.

<https://learn.microsoft.com/azure/ai-services/role-based-access-control?tabs=azure-portal>

NEW QUESTION: 95

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

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

Answer: A ([LEAVE A REPLY](#))

Comprehensive Detailed Explanation

When you develop a custom question answering project (formerly QnA Maker, now part of Azure Cognitive Service for Language - Question Answering), you can configure it for multi-turn conversations.

Follow-up prompts

Used to create multi-turn conversations.

They allow the system to present additional questions to the user after an initial question, guiding the conversation flow.

Example:

User: " Tell me about Azure Storage. "

Bot: " Which type of storage: Blob, File, Queue, or Table? "

These follow-up options are configured using follow-up prompts.

Correct choice for this question.

Active learning

Helps improve the knowledge base by suggesting alternative questions from user queries.

Not related to multi-turn conversation flow.

Alternate questions

Used to add different phrasings for the same question.

Example: "What is Azure Blob Storage?" and "Explain Blob Storage in Azure." This improves recognition but does not enable multi-turn conversations.

Chit-chat

Adds small talk and casual conversation handling (e.g., greetings, jokes).

Not related to multi-turn Q & A about a knowledge base.

The Answer: A. Add follow-up prompts

Microsoft References

Multi-turn conversations in Question Answering

Add follow-up prompts in Question Answering

NEW QUESTION: 96

You are building a social media messaging app.

You need to identify in real time the language used in messages.

- A.** Azure AI Foundry Content Safety
- B.** Azure AI Speech
- C.** Azure AI Language
- D.** Azure AI Translator

Answer: D (LEAVE A REPLY)

For a social media messaging app that needs to determine the language of short text in real time, the most appropriate service is Azure AI Translator. Translator provides a low-latency language detection operation for text, returning the detected language code and a confidence score, and it's optimized for interactive scenarios like chat and messaging. You can use detection on its own or feed the result directly into translation if needed.

Why not the others:

Azure AI Speech is for audio (speech-to-text, text-to-speech, speech translation), not text-only language detection.

Azure AI Language has language detection within its Text Analytics features, but Translator's detect endpoint is the recommended, streamlined choice for real-time translation pipelines and chat apps.

Azure AI Foundry Content Safety is for moderation (toxicity, harassment, hate), not language identification.

Microsoft References

Azure AI Translator - Text Translation overview (includes language detection).

<https://learn.microsoft.com/azure/ai-services/translator/>

Text Translation REST API v3.0 - Detect operation.

<https://learn.microsoft.com/azure/ai-services/translator/reference/v3-0-detect> Quickstart: Use the Translator text client libraries (shows detect + translate flow).

<https://learn.microsoft.com/azure/ai-services/translator/quickstart-text-sdk>

NEW QUESTION: 97

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 (LEAVE A REPLY)

The HTTP request provided is:

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 " } Key Observations:

The request is made against the Azure Cognitive Services Management API (Microsoft.CognitiveServices /accounts/.../regenerateKey).

The regenerateKey operation is specifically designed to regenerate one of the two subscription keys used to authenticate Cognitive Services API calls.

Cognitive Services resources always have two keys: Key1 (primary) and Key2 (secondary).

This design allows key rotation without downtime: you can regenerate one key while using the other in production.

Body Payload:

```
{ " keyName " : " Key2 " }
```

This explicitly tells Azure to regenerate Key2 (the secondary subscription key).

After the call, the secondary subscription key value changes, while Key1 (primary) remains unaffected.

Option Analysis

A). A key for Azure Cognitive Services was generated in Azure Key Vault Incorrect. The operation is not integrated with Key Vault; it only regenerates Cognitive Services subscription keys.

B). A new query key was generated

Incorrect. Query keys are related to Azure Cognitive Search (not general Cognitive Services).

The request clearly targets CognitiveServices/accounts.

- C). The primary subscription key and the secondary subscription key were rotated Incorrect. The request regenerates only the specified key, not both.
- D). The secondary subscription key was reset

Correct. The payload specifies " Key2 " , so only the secondary subscription key is regenerated.

The Answer: D. The secondary subscription key was reset

Microsoft References

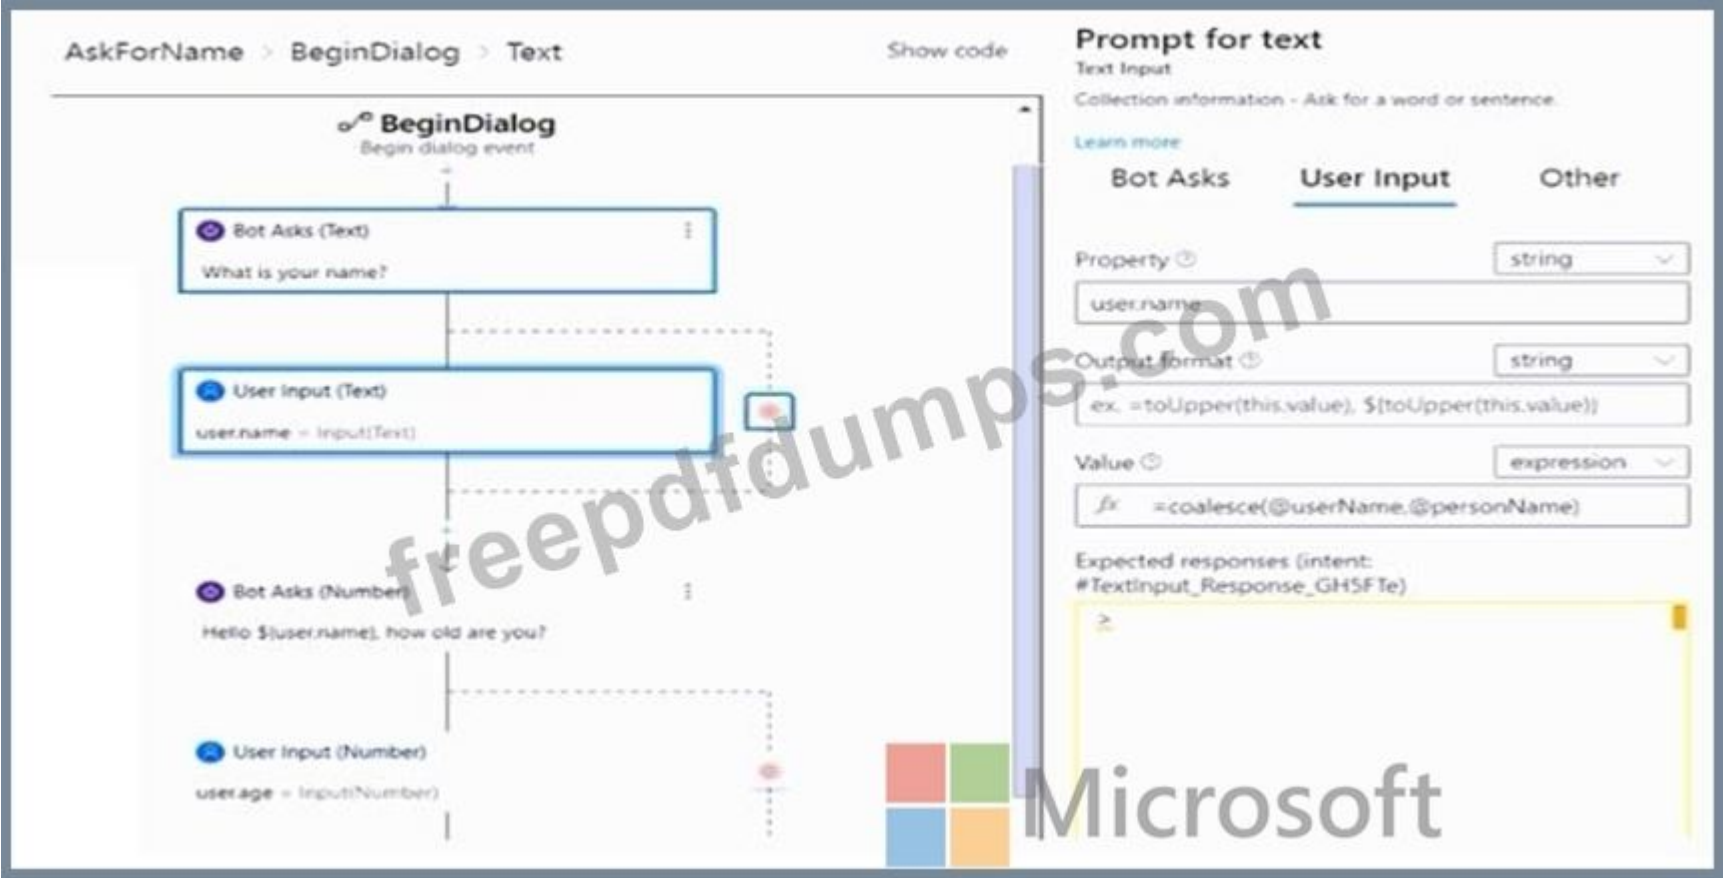
Azure REST API - Cognitive Services regenerate key

Authenticate requests to Azure AI services with keys and endpoint

Manage Cognitive Services keys (key1/key2)

NEW QUESTION: 98

You are building a chatbot by using the Microsoft Bot Framework Composer. You have the dialog design shown in the following exhibit.



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

NOTE: Each correct selection Is worth one point.

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

Answer:

Answer Area

Statements

Microsoft

YesNo

user.name is an entity.

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

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

Explanation:

Answer Area

Statements

Microsoft

YesNo

user.name is an entity.

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

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

Comprehensive Detailed Explanation

The screenshot shows a Bot Framework Composer dialog that collects user input:

user.name is an entity.

False.

user.name is a property in bot state (a variable), not an entity.

Entities are part of LUIS/NLU models (e.g., PersonName, DateTime).

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

True.

The dialog explicitly has:

A Bot Asks (Text) step # assigns user's response to user.name .

A Bot Asks (Number) step # assigns response to user.age .

So the values are stored in those properties.

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

True.

On the right-hand side, the Value expression is:

=coalesce(@userName, @personName)

coalesce returns the first non-null value.

This means if userName entity is available, it is used; otherwise, personName is used.

That result is assigned to user.name .

Correct Answers:

No

Yes

Yes

Microsoft References

Bot Framework Composer - Manage variables and properties

Bot Framework Composer - Recognizers and entities

Coalesce function in Composer expressions

NEW QUESTION: 99

You have a custom agent named Agent1.

You need to control access to and monitor activity for Agent1 by using Microsoft Foundry.

What should you do first?

- A.** Add Agent1 to a Microsoft Foundry project.
- B.** Create a Microsoft Foundry project.
- C.** Provision an Application Insights resource.
- D.** Add Agent1 to the Microsoft Foundry playground.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 100

You are building a social media messaging app.

You need to identify in real time the language used in messages.

Which service should you use?

- A.** Azure AI Speech
- B.** Azure AI Content Safety
- C.** Azure AI Translator
- D.** Azure AI Language

Answer: ([SHOW ANSWER](#)**)**

For a social media messaging app, you need real-time language identification on short text messages. The Azure AI Translator service provides a dedicated Detect API that returns the language and a confidence score for each input text, designed for low-latency scenarios such as chat and in-app communication. This aligns directly with "identify in real time the language used in messages." Microsoft Azure +2 Why the other options are not the best fit:

Azure AI Speech (A) targets spoken audio (speech-to-text, translation, diarization). It's not intended for detecting the language of text messages. Azure AI Azure AI Content Safety (B) focuses on safety/moderation (harmful or unsafe content), not language identification.

Azure AI Language (D) also offers a Language Detection feature, but it's typically used within broader text analytics workflows. For real-time messaging and potential downstream translation, Translator's Detect endpoint is the targeted and streamlined choice.

References (Microsoft Learn)

Translator - Detect method (returns detected language + confidence). Microsoft Learn Translator - Product overview and real-time translation scenarios. Microsoft Azure Translator - Language support (Auto Language Detection). Microsoft Learn Azure AI Language - Language Detection (feature overview/how-to) for context.

NEW QUESTION: 101

You are building a chatbot.

You need to use the Content Moderator service to identify messages that contain sexually explicit language.

Which section in the response from the service will contain the category score, and which category will be assigned to the message? To answer, select the appropriate options in the answer area, NOTE: Each correct selection is worth one point.

Answer Area



Section:

Category:

Answer:

Answer Area



Section:

Category:

Explanation:

Answer Area



Section:

Category:

Comprehensive Detailed Explanation

The scenario:

You are using Azure Content Moderator to filter chatbot messages.

You must detect sexually explicit language.

The question asks: which section of the response provides the category score, and which category number corresponds to sexual content.

Step 1 - Understand Content Moderator response structure

Content Moderator returns several sections:

Classification # Provides scores for categories of potentially unsafe or inappropriate content.

PII # Identifies personally identifiable information (phone numbers, email, addresses, etc.).

Terms # Identifies terms from a custom or built-in blocklist.

Since the requirement is sexually explicit language detection, the correct section is Classification.

Step 2 - Identify the correct Category

Content Moderator's Classification section returns three default categories:

Category 1: Potentially sexually explicit or adult content.

Category 2: Potentially sexually suggestive content.

Category 3: Offensive or derogatory language.

Since the question specifies sexually explicit language, the correct category is 1.

Correct Answer Mapping:

Section: Classification

Category: 1

Microsoft References

Content Moderator - Text moderation (classification categories explained) Content Moderator overview and response schema

NEW QUESTION: 102

Select the answer that correctly completes the sentence.

Answer Area

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

distributes processing across compute nodes.

distributes processing across control nodes.

redirects client connections across compute nodes.

redirects client connections across control nodes.

Microsoft

Answer:

Answer Area

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

distributes processing across compute nodes.

distributes processing across control nodes.

redirects client connections across compute nodes.

redirects client connections across control nodes.

Microsoft

Explanation:

Answer Area

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

distributes processing across compute nodes.

Microsoft

Azure Synapse Analytics (formerly Azure SQL Data Warehouse) uses a massively parallel processing (MPP) architecture for big data and analytics workloads.

Control Node:

Acts as the brain of the system.

Receives queries, optimizes them, and then coordinates execution.

Does not perform large-scale data crunching itself.

Compute Nodes:

Perform the actual parallel execution of queries.

Data is distributed across these nodes.

The MPP engine breaks down queries into smaller tasks and distributes them to compute nodes for fast, parallel execution.

Review of Answer Options:

Distributes processing across compute nodes # Correct. This is the essence of MPP.

Distributes processing across control nodes # Incorrect. The control node does not process queries in parallel; it coordinates them.

Redirects client connections across compute nodes # Incorrect. Client connections go to the control node, not directly to compute nodes.

Redirects client connections across control nodes # Incorrect. There is typically one control node per Synapse instance.

The Answer: Distributes processing across compute nodes

Microsoft References

Azure Synapse Analytics architecture

Distributed query processing in Azure Synapse

NEW QUESTION: 103

You Build a bot in JavaScript.

From the Azure Command-Line interface (CLI), you run the following command.

az bot prepare-deploy

You need to deploy the bot to Azure.

Which three Azure CLI commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the client order.

Microsoft

Commands

az ad app credential

az ad app create

az deployment group create

az webapp deployment source config-zip

az ad app update

>

<

Answer Area

<

>

Answer:

Microsoft

Commands

az ad app credential

az ad app create

az deployment group create

az webapp deployment source config-zip

az ad app update

>

<

Answer Area

az deployment group create

az webapp deployment source config-zip

az ad app update

<

>

Explanation:

To deploy the bot to Azure, you should run the following three Azure CLI commands in sequence:

az deployment group create This command will create the Azure resources for your bot using an ARM template and a parameters file. You need to specify the resource group name, the template file path, and the parameters file path. For example:

az deployment group create --resource-group myResourceGroup --template-file "deploymentTemplates\template-with-preexisting-rg.json" --parameters "deploymentTemplates\parameters- for-template-BotApp-with-rg.json"

This command will also output the app ID and password of your bot, which you will need for the next command1.

az webapp deployment source config-zip This command will deploy your bot code to the app service that you created in the previous step. You need to specify the resource group name, the app service name, and the zip file path of your bot code. For example:

```
az webapp deployment source config-zip --resource-group myResourceGroup --name myBotAppService --src  
"code.zip"
```

This command will also output the URL of your bot endpoint, which you will need for the next command2.

az ad app update This command will update your bot registration with the endpoint URL of your bot. You need to specify the app ID of your bot and the endpoint URL. For example:

```
az ad app update --id myBotAppId --set replyUrls="https://myBotAppService.azurewebsites.net/api  
/messages"
```

This command will complete the deployment process and make your bot ready to be tested3.

NEW QUESTION: 104

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

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

Answer: B (LEAVE A REPLY)

To automatically delete blobs in Azure Blob Storage, you can configure a lifecycle management policy .

This policy allows you to:

Transition blobs between hot, cool, and archive tiers.

Delete blobs automatically after a specified number of days.

Other options:

Change feed # Tracks all changes to blobs, but does not delete them.

Soft delete # Protects blobs from accidental deletion by allowing recovery, not automated deletion.

Archive storage # Storage tier for rarely accessed data, not a deletion mechanism.

The Answer: a lifecycle management policy

Reference: Azure Blob Storage lifecycle management

NEW QUESTION: 105

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 (LEAVE A REPLY)

The question is about creating an Azure Cognitive Services resource programmatically with C#.

Key requirements from the question:

Free resource # This requires the F0 (Free) pricing tier.

West US region # Location should be " westus " .

Generate captions of images automatically # This requires the Computer Vision service.

Option Analysis

A). create_resource(client, " res1 " , " ComputerVision " , " F0 " , " westus ") Uses " ComputerVision " kind # Correct, because Computer Vision provides the image captioning feature.

Tier is " F0 " # Correct, free tier.

Location " westus " # Correct.

This is the correct answer.

B). create_resource(client, " res1 " , " CustomVision.Prediction " , " F0 " , " westus ") Custom Vision is for building and deploying custom image classifiers.

The Prediction resource is specifically for running inference, not general captioning.

Wrong service for the requirement.

C). create_resource(client, " res1 " , " ComputerVision " , " S0 " , " westus ")

" ComputerVision " is correct.

But " S0 " is a paid tier, not the required free tier.

Does not meet the requirement.

D). create_resource(client, " res1 " , " CustomVision.Prediction " , " S0 " , " westus ") Wrong service (captioning is not handled by Custom Vision).

Wrong tier (paid instead of free).

Incorrect.

The Answer: A. create_resource(client, " res1 " , " ComputerVision " , " F0 " , " westus ") Microsoft References Azure AI Computer Vision documentation Create a Cognitive Services resource in Azure Computer Vision capabilities - Image analysis and captioning

NEW QUESTION: 106

You create five bots by using Microsoft Bot Framework Composer.

You need to make a single bot available to users that combines the bots. The solution must support dynamic routing to the bots based on user input.

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

NOTE: Each correct selection is worth one point.

A. Create an Orchestrator model.

B. Change the Recognizer/Dispatch type.

C. Create a composer extension.

D. Enable WebSockets.

E. Create a custom recognizer JSON file.

F. Install the Orchestrator package.

Answer: (SHOW ANSWER)

You created five bots with Microsoft Bot Framework Composer, but want to make them appear as one unified bot for users, with dynamic routing depending on user input.

This is exactly the Bot Framework Orchestrator scenario. Orchestrator is the recognizer that allows you to dispatch incoming utterances to the correct skill (bot) or dialog.

Step 1 - Install Orchestrator package

To use Orchestrator inside Bot Framework Composer, you first install the Orchestrator package.

This provides the functionality for language understanding and routing.

Correct answer: F.

Step 2 - Change Recognizer/Dispatch type

By default, a bot project might use LUIS recognizer or Regex recognizer.

To use Orchestrator, you need to change the recognizer type to Orchestrator in Composer.

Correct answer: B.

Step 3 - Create an Orchestrator model

Orchestrator requires a trained model that maps utterances to intents or routes them to the correct skill/bot.

You must create an Orchestrator model for dispatch.

Correct answer: A.

Other Options (Incorrect):

C). Create a Composer extension # Not required for orchestrating multiple bots.

D). Enable WebSockets # Useful for streaming scenarios, not for bot orchestration.

E). Create a custom recognizer JSON file # Applies to building custom recognizers, not needed with Orchestrator.

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (**416** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 107

You are building an agent that will retrieve the current time at a given location by using a custom API.

You need to test the functionality of the custom API.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
curl https://cu1.services.ai.azure.com/assistant/completions?api-version=2024-12-01-prev

-H "Authorization: Bearer $AZURE_AI_TOKEN" \
-H "Content-Type: application/json" \
-d '{
  "instructions": "You are a weather bot. Use the provided functions to answer questions.",
  "model": "gpt-4o-mini",
  "content": [{
    "type": "function",
    "function": {
      "name": "get_current_time",
      "description": "Get the current time in location",
      "parameters": {
        "type": "object",
        "properties": {
          "location": {"type": "string", "description": "The city name, for example Seattle"}
        },
        "required": ["location"]
      }
    }
  },
  ...
}
```

assistants

assistants

completions

embeddings

threads

content=[{

content=[{

functions=[{

tool_resources=[{

tools=[{



Answer:

Answer Area

```
curl https://cu1.services.ai.azure.com/ assistants ?api-version=2024-12-01-prev  
-H "Authorization: Bearer $AZURE_AI_TOKEN" \  
-H "Content-Type: application/json" \  
-d '{  
  "instructions": "You are a weather bot. Use the provided functions to answer questions.",  
  "model": "gpt-4o-mini",  
  "content": {  
    "content": {  
      "functions": {  
        "tool_resources": {  
          "tools": {  
            "type": "function",  
            "function": {  
              "name": "get_current_time",  
              "description": "Get the current time in location",  
              "parameters": {  
                "type": "object",  
                "properties": {  
                  "location": {"type": "string", "description": "The city name, for example Seattle"}  
                },  
                "required": ["location"]  
              }  
            }  
          }  
        }  
      }  
    }  
  }  
}
```



Explanation:

Answer Area

```
curl https://cu1.services.ai.azure.com/ assistants ?api-version=2024-12-01-preview \
-H "Authorization: Bearer $AZURE_AI_TOKEN" \
-H "Content-Type: application/json" \
-d '{
  "instructions": "You are a weather bot. Use the provided functions to answer questions.",
  "model": "gpt-4o-mini",
  "content": [[
    {
      "type": "function",
      "function": {
        "name": "get_current_time",
        "description": "Get the current time in location",
        "parameters": {
          "type": "object",
          "properties": {
            "location": {"type": "string", "description": "The city name, for example Seattle"}
          },
          "required": ["location"]
        }
      }
    }
  ]
}
```

NEW QUESTION: 108

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:

Actions

Change the model domain.

Retrain the model.

Test the model.

Export the model.

Answer Area

Change the model domain.

Retrain the model.

Export the model.

⬅️

➡️

⬆️

⬆️

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>

NEW QUESTION: 109

You are designing an Azure AI solution to identify defective products on a production line.

You have a real-time video feed and an image library of sample products that are approved or rejected manually.

You need to recommend a service that meets the following requirements:

Monitors the video feed and identifies the defective products.

Can train a new model by using the image library.

Minimizes development effort.

What should you recommend?

- A. Azure AI Custom Vision
- B. Azure Machine Learning
- C. Azure AI Video Indexer
- D. Azure Vision in Foundry Tools

Answer: A (LEAVE A REPLY)

NEW QUESTION: 110

You have a factory that produces cardboard packaging for food products. The factory has intermittent internet connectivity.

The packages are required to include four samples of each product.

You need to build a Custom Vision model that will identify defects in packaging and provide the location of the defects to an operator. The model must ensure that each package contains the four products.

Which project type and domain should you use? To answer, drag the appropriate options to the correct targets.

Each option may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content NOTE: Each correct selection is worth one point.

Options

Food

General

General (compact)

Image classification

Logo

Object detection

Answer Area

Project type:

Domain:

Answer:

Options

Food

General

General (compact)

Image classification

Logo

Object detection

Answer Area

Project type: Object detection

Domain: General

Explanation:

Project type: Object detection

Domain: General (compact)

The scenario requires building a Custom Vision model for a factory with intermittent internet connectivity.

The solution must:

Detect defects in packaging

Locate the defects (provide bounding box information)

Ensure each package contains four products

Support offline use due to intermittent connectivity

Step 1 - Determine Project Type

Image Classification: Assigns a label to an image, but cannot locate where objects/defects are in the image.

Object Detection: Identifies what objects are present and their location (via bounding boxes).

Since the requirement is to identify location of defects and count products, the correct choice is Object detection.

Step 2 - Determine Domain

General domain: Standard general-purpose domain, optimized for cloud inference.

General (compact) domain: A lightweight version designed for exporting models to run offline on edge devices (iOS, Android, ONNX, etc.). Perfect for scenarios with intermittent internet connectivity.

Food domain: Specialized for food image classification, not defect detection in packaging.

Logo domain: Specialized for logo recognition, irrelevant here.

Since the requirement includes offline capability, the correct domain is General (compact).

The Answer:

Project type: Object detection

Domain: General (compact)

Microsoft References

Custom Vision project types: Classification vs Object detection

Domains in Custom Vision (General, General compact, Food, Logo, etc.)

Exporting compact models for offline use

NEW QUESTION: 111

You are designing a solution that will answer questions about human resources (HR) policies stored in the PDF format.

You need to ensure that the identical answer to a specific question is returned every time. The solution must minimize development effort Which service should you include in the solution?

A. Azure AI Language

B. Azure Machine Learning

C. Azure OpenAI

D. Azure AI Document Intelligence

Answer: A ([LEAVE A REPLY](#))

For consistent, deterministic answers about policy PDFs with minimal effort, use Azure AI Language - Question Answering (Custom Question Answering). You can ingest PDFs into a knowledge base and define /curate the exact answers. The service then returns the same answer for the same question every time (subject to your KB content and configuration), avoiding the variability of generative models and custom ML.

Microsoft Azure+2

Microsoft References:

* What is Custom Question Answering? (Language service). Microsoft Learn

* Azure AI Language - Question Answering product page. Microsoft Azure

* Azure AI Language overview (includes Question Answering feature)

NEW QUESTION: 112

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

You need to configure the bot to support voice interactions. The solution must support multiple client apps.

Which type of channel should you use?

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

Answer: (SHOW ANSWER)

The requirement is for supporting voice interactions across multiple client apps.

Cortana is deprecated and limited to Cortana channels only.

Microsoft Teams is app-specific and not suited for multiple client apps.

Direct Line Speech is designed to connect bots with Speech SDK-enabled clients, enabling voice-first bots that can work across multiple applications and devices.

Therefore, Direct Line Speech is the correct channel.

NEW QUESTION: 113

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

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

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

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

Actions

Register the bot with the Azure Bot Service.

Open the Bot Framework Composer.

Build and run the bot.

Open the Bot Framework Emulator.

Connect to the bot endpoint.

Answer Area

Answer:

Actions

Register the bot with the Azure Bot Service.

Open the Bot Framework Composer.

Build and run the bot.

Open the Bot Framework Emulator.

Connect to the bot endpoint.

Answer Area

Build and run the bot.

Open the Bot Framework Emulator.

Connect to the bot endpoint.

Explanation:

Actions

Register the bot with the Azure Bot Service.

Open the Bot Framework Composer.

Answer Area

1 Build and run the bot.

2 Open the Bot Framework Emulator.

3 Connect to the bot endpoint.

Comprehensive Detailed Explanation

The question asks about testing a bot interactively on a local machine when using the Microsoft Bot Framework SDK.

Steps to test locally:

Build and run the bot

Compile the bot application code and run it locally (e.g., using Visual Studio or Node.js).

This exposes the bot's endpoint on localhost (commonly `http://localhost:3978/api/messages`).

Open the Bot Framework Emulator

The Bot Framework Emulator is a desktop app for testing and debugging bots built with the Bot Framework SDK.

It allows you to send messages to the bot and see the responses.

Connect to the bot endpoint

Use the Emulator to connect to the locally running bot by entering its endpoint URL.

Once connected, you can interact with the bot as if you were a real user.

Actions Not Required in This Case

Register the bot with the Azure Bot Service # Needed for cloud deployment, not local testing.

Open the Bot Framework Composer # Only needed if designing the bot with Composer, but this scenario uses the SDK directly.

Correct Sequences (multiple possible valid orders):

Build and run the bot # Open the Bot Framework Emulator # Connect to the bot endpoint Microsoft References Test a bot locally with Bot Framework Emulator Bot Framework SDK overview

Final Answer Sequence:

Build and run the bot

Open the Bot Framework Emulator

Connect to the bot endpoint

NEW QUESTION: 114

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
Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than infrastructure as a service (IaaS) database offerings.	☐	☐
Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.	☐	☐
All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.	☐	☐

Answer:

Answer Area



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

Explanation:

Answer Area

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



Platform as a service (PaaS) database offerings in Azure require less setup and configuration effort than Infrastructure as a Service (IaaS) database offerings.

Yes

PaaS abstracts away the OS, patching, backups, and much of the infrastructure management. IaaS requires users to configure and manage VMs and databases manually.

Platform as a service (PaaS) database offerings in Azure provide end users with the ability to control and update the operating system version.

No

In PaaS, the OS and platform management are handled by Azure. Users cannot directly control or update the OS version. This is a responsibility difference between IaaS and PaaS.

All relational and non-relational platform as a service (PaaS) database offerings in Azure can be paused to reduce costs.

No

Not all PaaS database services support pausing. For example, Azure SQL Database serverless can auto-pause, but Azure Cosmos DB and provisioned SQL databases cannot. Hence this statement is incorrect.

Correct Responses:

Yes

No

No

Microsoft References

Azure SQL Database - PaaS overview

What is PaaS?

Serverless compute tier in Azure SQL Database (auto-pause)

NEW QUESTION: 115

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



Microsoft

\$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:

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	

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>

You are developing an app that will use the Decision and Language APIs.

You need to provision resources for the app. The solution must ensure that each service is accessed by using a single endpoint and credential.

Which type of resource should you create?

- A.** Speech
- B.** Content Moderator
- C.** Language
- D.** Azure Cognitive Services

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 117

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

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

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

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

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

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

Does this meet the goal?

- A.** Yes
- B.** No

Answer: B ([LEAVE A REPLY](#))

Scenario Recap:

* Proposed Solution: Deploy service1 with a public endpoint, and configure a Network Security Group (NSG) for vnet1.

Analysis:

- * NSGs only control inbound/outbound traffic at the VNet/subnet or NIC level.
- * Configuring NSGs does not create private connectivity to Cognitive Search.
- * Traffic still flows over the public endpoint.

The answer: B. No

NSGs filter traffic, but they don't prevent routing via the public internet. Only Private Link does.

NEW QUESTION: 118

You are building an agent by using the Azure AI Foundry Agent Service.

You need to ensure that the agent can access publicly accessible data that was released during the past 90 days.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
...
conn_id = "search_connection_id"
search = AzureAISearchTool (connection_id=conn_id)
with project_client:
    agent = project_client.agents.create_agent(
        model="gpt-4o",
        name="my-assistant",
        instructions="You are a helpful assistant",
        search.definitions,
        metadata=
    )
    print(tool_resources=
          tools=
```

Answer:

```
...
conn_id = "search_connection_id"
search = BingGroundingTool (connection_id=conn_id)
with project_client:
    agent = project_client.agents.create_agent(
        model="gpt-4o",
        name="my-assistant",
        instructions="You are a helpful assistant",
        search.definitions,
        metadata=
    )
    print(tool_resources=
          tools=
```

Explanation:

```
...
conn_id = "search_connection_id"
search = BingGroundingTool (connection_id=conn_id)
with project_client:
    agent = project_client.agents.create_agent(
        model="gpt-4o",
        name="my-assistant",
```

```

agent = aiohttp_client_agents.create_agent(
    model="gpt-4o",
    name="my-assistant",
    instructions="You are a helpful assistant",
    tools=[search_definitions],
)
print(f"Created agent, ID: {agent.id}")

```

To let an agent use public web data with recency controls (e.g., limit to content from the last 90 days), you add the Bing grounding tool. In the Azure AI Foundry Agent Service SDK, you create a connection (e.g., `search_connection_id`) and then construct a `BingGroundingTool` with that connection. The tool exposes a list of definitions (tool operations) that the agent can call at runtime. When creating the agent, you pass tool definitions via the `tools=` argument (a list of `ToolDefinition`). The `tool_resources=` argument is used when the tool requires additional resource descriptors (e.g., Azure AI Search index configuration). For Bing grounding, providing the connection and passing the definitions as `tools=` is sufficient. This enables the agent to call Bing with date filters to retrieve only recent, publicly accessible content.

Microsoft Azure AI References (titles only)

Azure AI Agent Service - Bing web grounding tool

Azure AI Agent Service - Create an agent (agent creation parameters including tools and tool_resources) Azure AI Agent Service - Tool connections and definitions

NEW QUESTION: 119

You have a file share that contains 5,000 images of scanned invoices.

You need to analyze the images. The solution must extract the following data:

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

What should you use?

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

Answer: D (LEAVE A REPLY)

- * The requirement is to extract structured information (invoice items, sales amounts, customer details) from scanned invoice images.
- * Form Recognizer (now part of Azure AI Document Intelligence) is specifically designed for extracting structured data from documents such as invoices, receipts, and business forms.
- * Custom Vision # used for image classification or object detection, not for text/data extraction.
- * Computer Vision # can extract text with OCR, but it doesn't provide structured field-level extraction for invoices.
- * Immersive Reader # for accessibility (reading content aloud, translating text), not document data extraction.

The answer: D

Reference: Azure AI Document Intelligence (Form Recognizer)

NEW QUESTION: 120

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

```
def analyze_image(local_image):  
    with open(local_image, "rb") as image_stream:  
        image_analysis = client.analyze_image_in_stream(  
            image=image_stream,  
            visual_features=[  
                VisualFeatureTypes.tags,  
                VisualFeatureTypes.description  
            ]  
        )  
    for caption in image_analysis.description.captions:  
        print(f"\n{caption.text} with confidence {caption.confidence}")  
    for tag in image_analysis.tags:  
        print(f"\n{tag.name} with confidence {tag.confidence}")
```

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

Answer Area

Statements	Yes	No
The code will perform face recognition.	☐	☐
The code will list tags and their associated confidence.	☐	☐
The code will read an image file from the local file system.	☐	☐

Answer:

Answer Area

Statements	Yes	No
The code will perform face recognition.	☐	☒
The code will list tags and their associated confidence.	☒	☐
The code will read an image file from the local file system.	☒	☐

Explanation:

Answer Area

Statements	Yes	No
The code will perform face recognition.	☐	☒
The code will list tags and their associated confidence.	☒	☐
The code will read an image file from the local file system.	☒	☐

Comprehensive Detailed Explanation

The given code is using the Azure AI Vision client library with the method:

```
image_analysis = client.analyze_image_in_stream(  
image=image_stream,  
visual_features=[  
VisualFeatureTypes.tags,  
VisualFeatureTypes.description  
]  
)
```

Breakdown of the functionality:

VisualFeatureTypes.tags # Extracts tags (labels) describing the image (e.g., "dog", "tree", "outdoor"), including confidence scores.

VisualFeatureTypes.description # Generates captions (natural-language sentences) that describe the image.

The code explicitly iterates through tags and captions , printing both with their confidence values.

The image is read from the local file system using:

with open (local_image, " rb ") as image_stream:

Face recognition is not included because the code does not request VisualFeatureTypes.faces .

Correct Answer Matrix:

Face recognition # No

Tags with confidence # Yes

Read image from local file system # Yes

Microsoft References

Azure AI Vision features (tags, description, faces, etc.)

analyze_image_in_stream method

NEW QUESTION: 121

You have an Azure Cognitive Services model named Model that identifies the intent of text input.

You develop an app in C# named App1.

You need to configure App1 to use Model1.

Which package should you add to App1?

- A.** Xamarin.Cognitive.Speech
- B.** Azure.AI.Language.Conversations
- C.** Universal.Microsoft.CognitiveServices.Speech
- D.** SpeechServicesToolkit

Answer: ([SHOW ANSWER](#))

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (**416** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 122

You have an Azure subscription.
You are building a chatbot that will use an Azure OpenAI model.
You need to deploy the model.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

⋮ Apply for access to Azure OpenAI.

⋮ Deploy the embeddings model.

⋮ Provision Azure API Management.

⋮ Provision an Azure OpenAI resource.

⋮ Deploy the GPT model.

⋮ Deploy the DALL-E model.

Answer Area



Answer:

ACTIONS

⋮ Apply for access to Azure OpenAI.

⋮ Deploy the embeddings model.

⋮ Provision Azure API Management.

⋮ Provision an Azure OpenAI resource.

⋮ Deploy the GPT model.

⋮ Deploy the DALL-E model.

ANSWER AREA

⋮ Provision an Azure OpenAI resource.

⋮ Deploy the GPT model.

⋮ Deploy the DALL-E model.

Explanation:

Actions

⋮ Apply for access to Azure OpenAI.

⋮ Deploy the embeddings model.

⋮ Provision Azure API Management.

Answer Area

1 ⋮ Provision an Azure OpenAI resource.

2 ⋮ Deploy the GPT model.

3 ⋮ Deploy the DALL-E model.

You are building a chatbot that will use an Azure OpenAI model.
You need to deploy the model in your Azure subscription.
Key Steps to Deploy an Azure OpenAI Model for a Chatbot:

Apply for access to Azure OpenAI

Azure OpenAI is a gated service and requires approval before you can provision resources.

Provision an Azure OpenAI resource

After approval, you must provision an Azure OpenAI resource in your subscription.

Deploy the GPT model

Since the chatbot requires natural language understanding and conversation, you deploy a GPT model (e.g., GPT-4, GPT-3.5).

Actions That Are Not Required for this Question:

Provision Azure API Management # Not needed for basic deployment (used only for advanced API management scenarios).

Deploy the embeddings model # Used for search/RAG, not necessary for chatbot baseline.

Deploy the DALL-E model # Used for image generation, not chatbot text responses.

Correct Order:

Apply for access to Azure OpenAI.

Provision an Azure OpenAI resource.

Deploy the GPT model.

Verified Answer (Sequence):

Apply for access to Azure OpenAI

Provision an Azure OpenAI resource

Deploy the GPT model

Microsoft References:

Quickstart: Get started with Azure OpenAI

Deploying models in Azure OpenAI

NEW QUESTION: 123

You plan to build a chatbot to support task tracking.

You create a Language Understanding service named lu1.

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

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

Actions

Train the application.

Publish the application.

Add a new application.

Add example utterances.

Add the prebuilt domain ToDo.

Answer Area


freepdfdumps.com

Answer:

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

Explanation:

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

NEW QUESTION: 124

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: ([SHOW ANSWER](#))

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>

NEW QUESTION: 125

You are building an app that will provide users with definitions of common AI terms.
You create the following C# code.

```
...
OpenAIClient client =
    new OpenAIClient(new Uri(endpoint), new AzureKeyCredential(key));
ChatCompletionsOptions options = new ChatCompletionsOptions()
{
    Messages =
    {
        new ChatMessage(ChatRole.System, "You are a helpful assistant."),
        new ChatMessage(ChatRole.User, "What is an LLM?")
    }
};

ChatCompletions response = client.GetChatCompletions(
    deploymentName, options);
ChatMessage completion = response.Choices[0].Message;
Console.WriteLine($"Chatbot: {completion.Content}");
...
```

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

Statements	Yes	No
The response will contain an explanation of large language models (LLMs) that has a high degree of certainty.	☐	☐
Changing "What is an LLM?" to "What is an LLM in the context of AI models?" will produce the intended response.	☐	☐
Changing "You are a helpful assistant." to "You must answer only within the context of AI language models." will give a higher likelihood of producing the intended	☐	☐

Answer:

Answer Area

Statements	Yes	No
The response will contain an explanation of large language models (LLMs) that has a high degree of certainty.	☐	☒
Changing "What is an LLM?" to "What is an LLM in the context of AI models?" will produce the intended response.	☒	☐
Changing "You are a helpful assistant." to "You must answer only within the context of AI language models." will give a higher likelihood of producing the intended	☒	☐

Explanation:

Answer Area

Statements	Yes	No
The response will contain an explanation of large language models (LLMs) that has a high degree of certainty.	☐	☒
Changing "What is an LLM?" to "What is an LLM in the context of AI models?" will produce the intended response.	☒	☐
Changing "You are a helpful assistant." to "You must answer only within the context of AI language models." will give a higher likelihood of producing the intended	☒	☐

High degree of certainty - No.

The code uses default generation parameters for chat completions (e.g., default temperature # 1.0), which means responses are probabilistic and not guaranteed to be highly certain or consistent. Azure OpenAI guidance explains that parameters like temperature and top_p control randomness; to increase determinism /certainty you would lower temperature and/or adjust top_p, which this snippet does not do.

Clarifying the user question - Yes.

Prompt engineering guidance emphasizes that clear, specific instructions lead to better outputs. Rephrasing the user message from "What is an LLM?" to "What is an LLM in the context of AI models?" adds context and intent, improving the chance of getting the exact explanatory answer you want.

Constrain with a stronger system message - Yes.

In chat completions, the system message sets behavior and has the highest priority. Tightening it to "You must answer only within the context of AI language models." narrows scope and reduces topic drift, increasing the likelihood that the model's response stays on-task (definitions of AI terms, here specifically LLMs).

Microsoft References (Azure AI / Azure OpenAI Service)

Azure OpenAI Service - Chat Completions: roles (system, user), message ordering, and usage.

Azure OpenAI Service - Prompt engineering: write clear, specific prompts; add constraints to guide responses.

Azure OpenAI Service - Sampling parameters: temperature and top_p control variability/creativity vs. determinism.

Azure OpenAI Service - Best practices for grounding, controllability, and reducing model drift with stronger system instructions.

NEW QUESTION: 126

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 (LEAVE A REPLY)

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

NEW QUESTION: 127

You are building an app that will use Azure AI Language to extract meaning from text messages. You need to provide additional context by adding references to supporting articles in Wikipedia. What should you use?

- A.** entity linking
- B.** custom entity extraction recognition (NER)
- C.** Azure AI Content Safety
- D.** key phrase extraction

Answer: ([SHOW ANSWER](#))

Entity linking in Azure AI Language identifies entities in text and links them to entries in a knowledge base such as Wikipedia, providing precisely the additional context and references you need. Custom NER extracts entities but doesn't link to Wikipedia; Content Safety moderates content; Key phrase extraction surfaces topics without external links.

Microsoft References:

- * Entity linking overview (links entities to Wikipedia). Microsoft Learn
- * How to call the Entity Linking API (returns entities with Wikipedia links)

NEW QUESTION: 128

You have an Azure IoT hub that receives series data from machinery. You need to build an app that will perform the following actions:

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

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

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

Answer: A ([LEAVE A REPLY](#))

- * Azure Metrics Advisor builds on Anomaly Detector but adds root cause analysis, correlation across multiple signals, incident detection, and alerting out-of-the-box.
- * D. Anomaly Detector # Works for single-signal anomaly detection but lacks correlation and root cause analysis.
- * C. Azure Machine Learning # Would work but requires custom model development (more dev effort).
- * B. Form Recognizer # Used for document processing, not sensor telemetry.

The answer: A

Reference: Azure Metrics Advisor overview

NEW QUESTION: 129

You have a factory that produces food products.

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

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

Which service should you use?

- A.** Face

B. Computer Vision

C. Azure Video Analyzer for Media (formerly Video indexer)

Answer: ([SHOW ANSWER](#))

You need to detect PPE compliance (masks or safety glasses) on a periodic schedule (every 15 minutes) while minimizing development effort and cost.

Why Computer Vision (Azure AI Vision)

Azure AI Vision (formerly "Computer Vision") lets you analyze still images with prebuilt Image Analysis and-when needed-train a lightweight custom object detection model directly in Vision Studio without building ML pipelines. You can simply grab a frame (snapshot) from your camera feed every 15 minutes and call the Vision API-this keeps costs low because you pay per image analyzed rather than continuously processing video.

Microsoft's current Vision stack supports image analysis and custom object detection via Vision Studio and v4.0 APIs, designed for scenarios like PPE detection. Microsoft Azure +1 Why not Face The Face service is intended for face detection/recognition workflows. While Face can return some accessory attributes (e.g., glasses/mask confidence), Microsoft has scaled back several face attributes under Responsible AI guidance, and

Face is not positioned as a general PPE-compliance detector across a scene (multiple people, full body PPE). It's therefore a poorer fit and can increase false decisions for PPE checks beyond the face region. Azure

Docs +3 Why not Azure Video Analyzer for Media (Video Indexer) Video Indexer is optimized for extracting insights from full videos (speech, people, labels, brands). Running it just to check PPE every 15 minutes is

overkill and more expensive than analyzing a single still frame with Vision. In addition, Microsoft has been adjusting video services post-Azure Media Services retirement; using snapshot + Vision avoids those

complexities and keeps costs minimal.

How this meets the requirements

Identify masks/safety glasses: Use Image Analysis or a quick custom object detector in Vision Studio trained on a small set of PPE images. Microsoft Learn Every 15 minutes: Capture a single frame from the camera

and call the Vision endpoint-simple scheduling with a timer-triggered Function or Logic App, minimal code. (General Azure pattern; Vision API is stateless per image.) Microsoft Azure Minimize development effort: Vision

Studio provides no-code/low-code training and testing for custom detectors if the prebuilt model isn't sufficient. Microsoft Learn Minimize costs: Pay-per-call image analysis is cheaper than indexing entire video

segments; no continuous streaming required. Microsoft Azure The Answer: B. Computer Vision Microsoft References Azure AI Vision (product overview): image and video analysis with OCR, object detection, Image

Analysis.

Microsoft Azure

What's new in Azure AI Vision: custom image classification/object detection in Vision Studio and v4.0 APIs.

Microsoft Learn

Face detection & attributes (accessories including glasses and mask)-not the right tool for scene-level PPE compliance. Microsoft Learn Face Detect REST API (attributes list, including glasses/mask). Microsoft Learn

Responsible AI updates indicating changes to facial analysis attributes; facial detection/accessories remain but are limited in scope. Azure Docs Azure AI Video Indexer release notes/updates (video-centric use cases,

ongoing adjustments post-AMS retirement).

NEW QUESTION: 130

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 ([LEAVE A REPLY](#))

The scenario involves a Language Understanding (LUIS) or Conversational Language Understanding (CLU) model that has:

An intent called FindContact (used for searching a contact list).

Training phrases like:

"Find contacts in London"

"Who do I know in Seattle?"

"Search for contacts in Ukraine"

Key Requirement

The model must recognize locations (London, Seattle, Ukraine) inside the FindContact intent.

Why the Solution is Incorrect

The proposed solution is to create a new intent for location.

This is not correct because:

Intents represent what the user wants to do (the action, e.g., FindContact).

Entities represent details within the user's utterance (the data, e.g., location).

Here, the intent is already defined: FindContact .

Location should be modeled as an entity (Location entity), not as a separate intent.

Creating a new intent for each location (London, Seattle, Ukraine) would fragment the model and make it harder to generalize. Instead, the model should capture location as a parameter of the single intent FindContact .

Correct Approach

Keep the existing FindContact intent.

Add an entity for location (prebuilt GeographyV2 in LUIS or "Location" entity in CLU).

Train the model so that London , Seattle , Ukraine are recognized as entity values inside the FindContact intent.

The Answer: B. No

Microsoft References

Intents and entities in Language Understanding

Prebuilt entity: GeographyV2

NEW QUESTION: 131

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

A. Revoke

B. UPDATE

C. DROP

D. CREATE

Answer: B ([LEAVE A REPLY](#))

In SQL, commands are categorized into different language groups depending on their purpose:

Data Manipulation Language (DML):

Deals with manipulating the data stored in existing database objects.

Includes:

SELECT # query data

INSERT # add new records

UPDATE # modify existing records

DELETE # remove records

Therefore, UPDATE is a DML statement.

Data Definition Language (DDL):

Used to define and manage database schema objects (tables, indexes, schemas).

Includes: CREATE , ALTER , DROP , TRUNCATE .

Data Control Language (DCL):

Used to control access and permissions.

Includes: GRANT , REVOKE .

Now let's map the options:

A). REVOKE # DCL (not DML).

B). UPDATE # DML # (correct).

C). DROP # DDL (removes database objects).

D). CREATE # DDL (creates database objects).

Microsoft References

SQL Language Elements - DDL, DML, DCL

UPDATE (Transact-SQL)

NEW QUESTION: 132

You have a Microsoft OneDrive folder that contains a 20-GB video file named FileVavi. You need to index File1.avi by using the Azure Video Indexer website. What should you do?

A. Upload File1.avi to the www.youtube.com webpage. and then copy the URL of the video to the Azure AI Video Indexer website

B. From OneDrive. create a download link, and then copy the link to the Azure AI Video Indexer website.

C. From OneDrive, create a sharing link for File1.avi and then copy the link to the Azure AI Video Indexer website.

D. Download File1 avi to a local computer, and then upload the file to the Azure AI Video Indexer website.

Answer: (SHOW ANSWER)

Azure Video Indexer can ingest a video by URL as long as the URL is publicly accessible and directly downloadable (no authentication or HTML landing page). Creating a download link in OneDrive produces a direct-file URL that Video Indexer can fetch and index without you re-uploading the 20-GB file through your browser. A generic OneDrive "sharing link" often lands on a web page (not a direct file) and may require authentication, which Video Indexer cannot crawl. Uploading locally is possible but unnecessary and time- consuming for a 20-GB video.

Microsoft References

Upload/index videos in Video Indexer (supports uploading from a public URL).

<https://learn.microsoft.com/azure/azure-video-indexer/upload-index-videos> Video Indexer quotas and limits (file size limits; uploading by URL).

NEW QUESTION: 133

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: B,C (LEAVE A REPLY)

Comprehensive Detailed Explanation

The question focuses on monitoring a public-facing sales system that processes video and text, ensuring that the system provides equitable results regardless of user location or background.

Two Microsoft Responsible AI principles apply directly here:

Fairness

AI systems should treat all people fairly and avoid bias.

Monitoring ensures that outcomes do not disadvantage people based on location, background, or demographics.

This directly addresses the requirement for equitable results.

Inclusiveness

AI systems should empower everyone and engage people of all backgrounds.

Monitoring for inclusiveness ensures that the system works well across different regions, languages, and cultures, making the system accessible and beneficial for all users.

Why not the other options?

A). Transparency

Important for explaining how AI systems work, but does not directly address equitable results across demographics.

D). Reliability and safety

Focuses on consistent and safe system operation, not equity across users.

E). Privacy and security

Critical for data protection, but not related to equitable treatment of users.

The Answer:

B). fairness

C). inclusiveness

Microsoft References

Microsoft Responsible AI Principles

Fairness in AI systems

Inclusiveness in AI systems

NEW QUESTION: 134

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

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

```

"documentResults":[
  {
    "docType":"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: (SHOW ANSWER)

The requirements state:

All AI solution responses must have a confidence score # 70%.

If the response confidence score is < 70%, the response must be improved by human input.

Members of the Consultant-Bookkeeper group must be able to process financial documents, which includes performing manual review when the AI confidence is below threshold.

Looking at the provided JSON:

```

" fields " : {
  " ReceiptType " : {
    " type " : " string " ,
    " valueString " : " Itemized " ,
    " confidence " : 0.672
  },
  " MerchantName " : {
    " type " : " string " ,
    " valueString " : " Tailwind " ,

```

" confidence " : 0.913

}

}

ReceiptType.confidence = 0.672 # this is below 0.7.

MerchantName.confidence = 0.913 # this is above 0.7.

Therefore, the correct condition for triggering manual review is:

documentResults.fields. < field > .confidence < 0.7

Now, evaluating the options:

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

Always true for receipts, does not check confidence. Not correct.

B). documentResults.fields.*.confidence < 0.7

This is the correct general expression: trigger manual review whenever any field confidence is below 0.7.

C). documentResults.fields.ReceiptType.confidence > 0.7

This would bypass manual review when ReceiptType has high confidence. The requirement is to trigger review when confidence < 0.7, so this is the opposite.

D). documentResults.fields.MerchantName.confidence < 0.7

This only checks one field (MerchantName). In the JSON, MerchantName has confidence 0.913 (> 0.7), so this condition would not trigger, but ReceiptType clearly needs review. Too narrow.

The Answer: B. documentResults.fields.*.confidence < 0.7

Microsoft References

Azure AI Document Intelligence - Confidence scores

Human-in-the-loop for Document Intelligence

NEW QUESTION: 135

You are developing an application that will recognize faults in components produced on a factory production line. The components are specific to your business.

You need to use the Custom Vision API to help detect common faults.

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


Actions

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

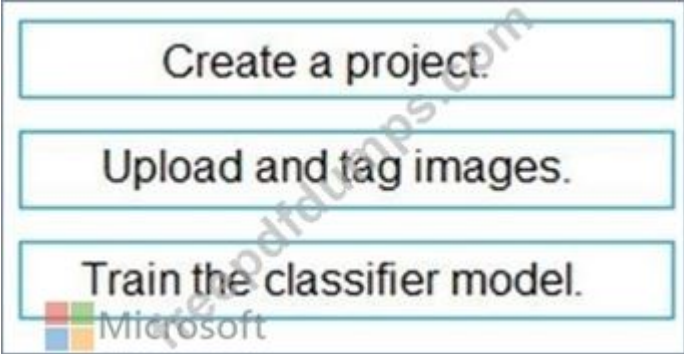
Answer:

Answer Area





Explanation:



Comprehensive Detailed Explanation

You are tasked with building a solution using Custom Vision to detect faults in factory components. Custom Vision follows a defined workflow:

Create a project

Before anything else, you must create a Custom Vision project in the Azure AI portal.

This project specifies the type of model you are building: classification (classifying entire images) or object detection (locating items within images).

Since the question is about detecting faults in specific components (a business-specific scenario), it aligns with classification unless explicitly stated otherwise.

Upload and tag images

You must provide training data.

Upload a set of images showing both faulty and non-faulty components.

Tag each image correctly so that the model can learn from labeled examples.

Train the classifier model

Once images are tagged, you run training on the dataset to generate the classifier model.

The model can then be published and used for inference via API calls.

Why not the other options?

Initialize the training dataset: This is not a direct step in Custom Vision; instead, uploading and tagging images inherently prepares the training dataset.

Train the object detection model: Object detection is used if you need bounding boxes around items within an image. Since the question is about recognizing faults in a component (faulty vs non-faulty), a classifier model is appropriate.

Correct Sequence:

Create a project

Upload and tag images

Train the classifier model

Microsoft References

Quickstart: Build an image classification project with the Custom Vision portal Custom Vision overview Train and test a Custom Vision model

NEW QUESTION: 136

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]
    }"
  }
]
```

Answer:

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:



Comprehensive Detailed Explanation

The requirement is to build an Adaptive Card for the chatbot that supports multilingual product display and meets accessibility requirements.

Product Name Display

The product JSON stores localized names like " name " : { " en " : " ... " , " es " : " ... " , " pt " : " ... " } .

To dynamically render based on the current language, you must use name[language].

This ensures the product name shown in the card adapts automatically to the user's preferred language.

Stock Level Warning

Business rules require warnings when stock is low or out of stock.

This means you display the stock warning only when stockLevel is not ' OK ' .

Therefore, the correct condition is:

" \$when " : " \${stockLevel != ' OK ' } "

Image Alt Text for Accessibility

Accessibility requirements specify that all images must have alt text in English, Spanish, and Portuguese.

In the product JSON, " alttext " is also localized:

" altText " : { " en " : " Bicycle " , " es " : " Bicicleta " , " pt " : " Bicicleta " } To render the correct localized alt text dynamically, use image.altText[language].

Correct Selections:

First blank: name[language]

Second blank: " \$when " : " \${stockLevel != ' OK ' } "

Third blank: image.altText[language]

Microsoft References

Adaptive Cards Templating

Adaptive Card \$when property

Multilingual data in Azure Cognitive Search and Adaptive Cards

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (416 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 137

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: A,B,F (LEAVE A REPLY)

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>

NEW QUESTION: 138

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: B,D ([LEAVE A REPLY](#))

Comprehensive Detailed Explanation

The scenario:

You are building an Azure WebJob to create knowledge bases from an array of URLs.

You already have a QnAMakerClient object instantiated with the relevant API keys.

You need to know what steps are required to create a knowledge base programmatically.

Step 1 - Create a CreateKbDTO object

CreateKbDTO is the request payload for creating a knowledge base.

It contains the name of the knowledge base and its sources, which may include:

A list of URLs.

A list of QnA pairs (QnADTO).

A list of files (FileDTO).

In this case, since the WebJob provides URLs, they will be passed inside the CreateKbDTO.Sources property.

Step 2 - Call client.Knowledgebase.CreateAsync

Once the CreateKbDTO object is defined, you use the QnAMakerClient to send the request:

```
var response = await client.Knowledgebase.CreateAsync(createKbDTO);
```

This triggers the asynchronous creation of a new knowledge base in the QnA Maker service.

Why not A or C?

A). Create a list of FileDTO objects: Not required here, because the question specifies URLs, not files.

C). Create a list of QnADTO objects: Also not needed, since you're not directly supplying QnA pairs; the knowledge base will be built from URLs.

Correct Answer

B). Call the client.Knowledgebase.CreateAsync method.

D). Create a CreateKbDTO object.

Microsoft References

QnA Maker - Create Knowledge Base API

QnAMakerClient Class - Azure SDK for .NET

CreateKbDTO Class

NEW QUESTION: 139

You are developing a text processing solution.

You have the following function.

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

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

You call the function and use the following string as the second argument Our tour of London included a visit to Buckingham Palace What will be the output of the function?

- A. Tour and visit only
- B. Our tour of London included a visit to Buckingham Palace
- C. London and Tour only
- D. London and Buckingham Palace only

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 140

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

What should you include in the solution?

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

Answer: C ([LEAVE A REPLY](#))

The system must analyze multiple sensor data points (rotation speed, angle, temperature, pressure) and generate alerts on anomalies.

Application Insights # used for monitoring applications, not IoT sensor data.

Metric alerts in Azure Monitor # works with single metrics but does not handle complex correlations.

Univariate Anomaly Detection # only works on a single time series (e.g., just temperature).

Multivariate Anomaly Detection # specifically built to detect anomalies across multiple correlated variables (sensor data).

Thus, the correct choice is Multivariate Anomaly Detection.

NEW QUESTION: 141

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:

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

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


Answer:

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

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

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

Explanation:



Scenario Recap

You are building a remote training monitoring solution.

Requirement: Use video and audio feeds to detect if a learner is present, paying attention, and talking.

Services available: Face, Speech, Text Analytics.

Analysis

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

The Face API can detect and identify faces in a video feed.

It can tell if a person is present and recognized, fulfilling the requirement.

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

Again, the Face API provides facial expression and emotion recognition (happiness, anger, neutral, etc.).

This can be mapped to "paying attention vs. distracted."

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

The Speech service detects spoken input and can determine if speech is present.

Text Analytics works on text (not raw audio) and is therefore not appropriate here.

Final Answer (Answer Area Selections)

From a learner's video feed, verify whether the learner is present: Face From a learner's facial expression in the video feed, verify whether the learner is paying attention: Face From a learner's audio feed, detect whether the learner is talking: Speech Microsoft References Face API - Face detection & identification Face API - Emotion recognition Azure Speech service

NEW QUESTION: 142

You have an Azure subscription that contains an Azure AI Content Safety resource named CS1.

You need to call CS1 to identify whether a user request contains hateful language.

How should you complete the command? To answer, select the appropriate options in the answer area NOTE: Each correct selection is worth one point

Answer Area

```
curl --location --request POST 'https://cs1.cognitiveservices.azure.com/
--header 'Ocp-Apim-Subscription-Key: <your_subscription_key>' \
--header 'Content-Type: application/json' \
--data-raw '{
  "text": "What is the weather forecast for Seattle",
  "categories": ["Hate"]
  "blocklistNames": [
    "string"
  ],
}
```

language/ completions/ contentsafety/ healthinsights/ language/ embeddings completions embeddings text:analyze text/blocklists ?api-version=2023-10-01



Answer:

Answer Area

```
curl --location --request POST 'https://cs1.cognitiveservices.azure.com/
--header 'Ocp-Apim-Subscription-Key: <your_subscription_key>' \
--header 'Content-Type: application/json' \
--data-raw '{
  "text": "What is the weather forecast for Seattle",
  "categories": ["Hate"]
  "blocklistNames": [
    "string"
  ],
}
```

language/ completions/ contentsafety/ healthinsights/ language/ embeddings completions embeddings text:analyze text/blocklists ?api-version=2023-10-01



Explanation:

Answer Area

```
curl --location --request POST 'https://cs1.cognitiveservices.azure.com/
--header 'Ocp-Apim-Subscription-Key: <your_subscription_key>' \
--header 'Content-Type: application/json' \
--data-raw '{
  "text": "What is the weather forecast for Seattle",
  "categories": ["Hate"]
  "blocklistNames": [
    "string"
  ],
  "haltOnBlocklistHit": true,
  "outputType": "FourSeverityLevels"
}'
```

language/ embeddings ?api-version=2023-10-01



NEW QUESTION: 143

You have an Azure Subscription that contains an Azure OpenAI resource named AI1 and a user named User1.

You need to ensure that User1 can add custom data sources to AI1. The solution must follow the principle of least privilege.

Which role should you assign to User1?

- A. Search Service Contributor
- B. Cognitive Services OpenAI Contributor
- C. Cognitive Services Contributor
- D. Search index Data Contributor

Answer: B (LEAVE A REPLY)

- * The requirement is to allow User1 to add custom data sources to an Azure OpenAI resource (AI1).
- * Following the principle of least privilege means granting only the specific rights required.
- * The Cognitive Services OpenAI Contributor role allows management of Azure OpenAI resources, including attaching custom data sources for RAG (retrieval-augmented generation).
- * The Cognitive Services Contributor role is broader and grants access to all cognitive services, which violates least privilege.
- * Search Service Contributor and Search Index Data Contributor are roles for Azure AI Search, not for managing data sources on Azure OpenAI.

Microsoft References:

- * Azure built-in roles: Cognitive Services OpenAI Contributor

NEW QUESTION: 144

You plan to use a Conversational Language Understanding application named app1 that is deployed to a container. App1 was developed by using a Conversational 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.

Actions

⋮ Export the model by using the Export as JSON option.

⋮ Select v1.2 of app1.

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

⋮ Select v1.0 of app1.

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

⋮ Select v1.1 of app1.

⋮ Run a container and mount the model file.

Answer Area



Microsoft

Answer:

⋮

Export the model by using the Export as JSON option.

⋮

Select v1.2 of app1.

⋮

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

⋮

Select v1.0 of app1.

⋮

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

⋮

Select v1.1 of app1.

⋮

Run a container and mount the model file.

⋮

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

⋮

Select v1.1 of app1.

⋮

Run a container and mount the model file.

Explanation:
Actions

⋮

Export the model by using the Export as JSON option.

⋮

Select v1.2 of app1.

⋮

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

⋮

Select v1.0 of app1.

1

⋮

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

2

⋮

Select v1.1 of app1.

3

⋮

Run a container and mount the model file.

NEW QUESTION: 145

You are building an Azure web app named App1 that will translate text from English to Spanish. You need to use the Text Translation REST API to perform the translation. The solution must ensure that you have data sovereignty in the United States. How should you complete the URI? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

https://

api-nam.cognitive.microsofttranslator.com

api.cognitive.microsofttranslator.com

api-nam.cognitive.microsofttranslator.com

api-nam.cognitiveservices.azure.com

eastus.api.cognitive.microsoft.com

translate

detect


languages

text-to-speech

translate

?api-version=3.0&to=e

Answer:



Explanation:

Answer Area

https:// / ?api-version=3.0&to=es

1. api-nam.cognitive.microsofttranslator.com

2. translate

<https://learn.microsoft.com/en-us/azure/cognitive-services/Translator/reference/v3-0-reference#base-urls> Requests to Translator are, in most cases, handled by the datacenter that is closest to where the request originated. If there's a datacenter failure when using the global endpoint, the request may be routed outside of the geography.

To force the request to be handled within a specific geography, use the desired geographical endpoint. All requests are processed among the datacenters within the geography.

- United States

api-nam.cognitive.microsofttranslator.com

<https://learn.microsoft.com/en-us/azure/cognitive-services/translator/reference/rest-api-guide>

- translate

Translate specified source language text into the target language text.

NEW QUESTION: 146

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

You have the dialog design shown in the following exhibit.

AskForName > BeginDialog > Text [Show code](#)

BeginDialog
Begin dialog event

Bot Asks (Text)
What is your name?

User Input (Text)
user.name = Input(Text)

Bot Asks (Number)
Hello \$(user.name), how old are you?

User Input (Number)
user.age = Input(Number)

Prompt for text
Text input
Collection information - Ask for a word or sentence.
[Learn more](#)

Bot Asks **User Input** **Other**

Property ⓘ string
user.name

Output Format ⓘ string
ex. \$toUpper(this.value), \$(toUpper(this.value))

Value ⓘ expression
fx =coalesce(@user.Name.@personName)

Expected responses (intent:
#TextInput_Response_GH5FTe)

>

Microsoft

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

Answer:

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

Explanation:

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

Dialog Recap

The Bot Framework Composer dialog shown:

Bot asks: "What is your name?"

User input stored in user.name .

Property is user.name with a value expression: =coalesce(@user.Name, @personName) Bot asks: "Hello \${user.name}, how old are you?" User input stored in user.age .

Statement Analysis

user.name is an entity.

No.

user.name here is not an entity. It is a property in the bot's memory (user state).

Entities come from LUIS/CLU recognition, while this is a variable property.

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

Yes.

The dialog explicitly prompts for a text input (name) and a number input (age) and assigns them to user.name and user.age .

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

Yes.

The expression =coalesce(@user.Name, @personName) ensures that if @user.Name exists, it's used; otherwise, it tries @personName . This is exactly first-non-null entity resolution.

Final Answer

user.name is an entity. # No

The dialog asks for a user name and a user age and assigns appropriate values to the user.name and user.age properties. # Yes The chatbot attempts to take the first non-null entity value for userName or personName and assigns the value to user.name. # Yes Microsoft References Bot Framework Composer - Memory and Properties Expressions and coalesce function Reference:

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

NEW QUESTION: 147

You are developing a custom analyzer by using Azure Content Understanding in Foundry Tools.

After creating the schema, you label the data as shown in the exhibit. (Click the Exhibit tab.) 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 schema defines the type of the Benefits field as [answer choice].

fixed table

string

table

The schema contains [answer choice] fields.

2

5

10

Microsoft

Answer:

Answer Area

Microsoft

The schema defines the type of the Benefits field as [answer choice].

fixed table

string

table

The schema contains [answer choice] fields.

2

5

10

Explanation:

Answer Area

The schema defines the type of the Benefits field as [answer choice]. table

The schema contains [answer choice] fields. 5


NEW QUESTION: 148

You have a chatbot.


You need to test the bot by using the Bot Framework Emulator. The solution must ensure that you are prompted for credentials when you sign in to the bot.

Which three settings should you configure? To answer, select the appropriate settings in the answer area.


NOTE Each correct selection is worth one point.



Answer:



Explanation:



When testing a bot with the Bot Framework Emulator, especially if the bot uses authentication (OAuth), you need the Emulator to mimic real-world authentication flows. To do that:

Use version 1.0 authentication tokens

Ensures that authentication tokens are generated in the same format as Bot Framework v1.0, which is compatible with most OAuth flows.

Use a sign-in verification code for OAuthCard

Prompts the user with a verification code during sign-in flows, ensuring the emulator behaves like a real client app (e.g., mobile or Teams).

This is required to test sign-in prompts and enforce credential entry.

Use your own user ID to communicate with the bot

Ensures that the Emulator identifies you uniquely and triggers authentication properly.

Without this, the bot may treat sessions generically, and authentication prompts might not behave as expected.

Why not the others?

Application Updates (automatically download/install) # Related to emulator updates, not authentication.

Data Collection # Only for telemetry; does not affect authentication.

Ngrok settings # Used for tunneling remote bot connections; unrelated to credential prompts.

Correct Answer Selections:

Use version 1.0 authentication tokens

Use a sign-in verification code for OAuthCard

Use your own user ID to communicate with the bot

Microsoft References

Bot Framework Emulator - Authentication

Testing bots with Bot Framework Emulator

NEW QUESTION: 149

You are developing an application that will detect faulty components produced on a factory production line.

The components are specific to your business.

You need to use the Azure 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

Upload and tag images.

Train the object detection model.

Initialize the training dataset.

Train the classifier model.

Create a project.

Answer Area

Answer:

Actions

Upload and tag images.

Train the object detection model.

Initialize the training dataset.

Train the classifier model.

Create a project.

Answer Area

Create a project.

Initialize the training dataset.

Upload and tag images.

Explanation:

Answer Area

1 Create a project.

2 Initialize the training dataset.

3 Upload and tag images.

NEW QUESTION: 150

You have an app named App1 that uses a custom Azure AI Document Intelligence model to recognize contract documents. You need to ensure that the model supports an additional contract format. The solution must minimize development effort. What should you do?

- A. Lower the confidence score threshold of App1.
- B. Lower the accuracy threshold of App1.
- C. Add the additional contract format to the existing training set. Retrain the model.
- D. Create a new training set and add the additional contract format to the new training set.
- E. Create and train a new custom model.

Answer: (SHOW ANSWER)

Analysis of Options:

A). Lower the confidence score threshold of App1

Would just accept more low-confidence predictions, but won't add support for a new format. # B). Lower the accuracy threshold of App1 Similar reasoning: affects prediction acceptance, not model capabilities. # C). Add the additional contract format to the existing training set and retrain the model Correct approach. You extend the training data to include the new format, retrain, and reuse the same model. # D). Create a new training set and add the additional contract format This means starting from scratch, which requires more effort. Not minimal. # E). Create and train a new custom model Same as D, higher effort than necessary. # The Answer: C). Add the additional contract format to the existing training set. Retrain the model.

Microsoft Reference:

Azure AI Document Intelligence - Custom models

NEW QUESTION: 151

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

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

Which Azure Cognitive Services service should you use?

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

Answer: B (LEAVE A REPLY)

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>

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (**416** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 152

You are developing an app that will use the Speech and Language APIs.

You need to provision resources for the app. The solution must ensure that each service is accessed by using a single endpoint and credential.

Which type of resource should you create?

- A.** Azure AI Language
- B.** Azure AI Foundry service
- C.** Azure AI Speech
- D.** Azure AI Foundry Content Safety

Answer: ([SHOW ANSWER](#))

The scenario requires:

- * An app that uses both Speech and Language APIs.
- * A single endpoint and single credential for accessing both services.

Let's analyze the options:

- * Option A - Azure AI Language
- * Provides only language-related capabilities (text analytics, QnA, summarization, etc.).
- * Does not include Speech APIs.
- * So, this would not meet the requirement.
- * Option B - Azure AI Foundry service #
- * Previously known as Azure Cognitive Services (multi-service account).
- * Allows you to provision one unified resource that gives access to multiple Azure AI services (Vision, Speech, Language, Content Safety, etc.).
- * Provides a single endpoint and single API key/credential across these services.
- * Perfectly fits the requirement of the question.
- * Option C - Azure AI Speech
- * Provides only speech-related capabilities (speech-to-text, text-to-speech, speech translation, speaker recognition).
- * Does not include Language APIs.
- * Therefore, not suitable.
- * Option D - Azure AI Foundry Content Safety
- * Focused on safety and responsible AI (moderation of text and images).
- * Does not include general speech or language APIs.
- * Not applicable here.

Microsoft References

- * Azure AI Foundry documentation (formerly Cognitive Services)
- * Azure AI services overview
- * Provisioning Azure AI services as a single multi-service resource

NEW QUESTION: 153

You are designing two Azure AI agents named Agent 1 and Agent2 to assist different departments at your company. Each agent will have access to sensitive corporate data. The design must meet the following requirements:

- * Agent1 must process specific user queries and provide relevant answers in a conversational format
- * Agent2 must perform the following actions:
 - o Adapt actions based on patterns in user behavior,
 - o Understand and learn from historical data.
 - o Provide proactive recommendations.

Which type of agent should you recommend for each agent? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area



Microsoft

Agent1:

Autonomous

Cognitive

Prompt-and-response

Agent2:

Autonomous

Cognitive

Prompt-and-response

Answer:

Answer Area



Microsoft

Agent1:

Autonomous

Cognitive

Prompt-and-response

Agent2:

Autonomous

Cognitive

Prompt-and-response

Explanation:

Answer Area



Microsoft

Agent1:

Cognitive

Agent2:

Autonomous

NEW QUESTION: 154

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: ([SHOW ANSWER](#))

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>

NEW QUESTION: 155

You have an Azure Cognitive Search solution and a collection of handwritten letters stored as JPEG files.

You plan to index the collection. The solution must ensure that queries can be performed on the contents of the letters.

You need to create an indexer that has a skillset.

Which skill should you include?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 156

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 ([LEAVE A REPLY](#))

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>

NEW QUESTION: 157

You have a flow that you plan to deploy to Azure Machine Learning managed online endpoints for real-time inference.
You need to modify the YAML configuration file for the deployment to ensure that trace data and system metrics are collected and sent to an Application Insights resource named my-app-insights.
How should you modify the file? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Identify the road signs and generate a short description of each road sign:

Azure Document Intelligence in Foundry Tools
Azure AI Phi-3-mini
Azure Vision in Foundry Tools
Azure OpenAI GPT-4-Turbo

Analyze the descriptions to generate a report about the different types of road signs and how often each type occurred:

Azure Document intelligence in Foundry Tools
Azure Language in Foundry Tools
Azure AI Phi-3-mini
Azure OpenAI GPT-4-Turbo

Microsoft

Answer:

Answer Area

Identify the road signs and generate a short description of each road sign:

Analyze the descriptions to generate a report about the different types of road signs and how often each type occurred:

Microsoft

Azure Document Intelligence in Foundry Tools
 Azure AI Phi-3-mini
 Azure Vision in Foundry Tools
 Azure OpenAI GPT-4-Turbo

Azure Document intelligence in Foundry Tools
 Azure Language in Foundry Tools
 Azure AI Phi-3-mini
 Azure OpenAI GPT-4-Turbo

Explanation:

Answer Area

Identify the road signs and generate a short description of each road sign:

Analyze the descriptions to generate a report about the different types of road signs and how often each type occurred:

Microsoft

Azure Vision in Foundry Tools

Azure Document Intelligence in Foundry Tools

NEW QUESTION: 158

You have an Azure subscription.

You plan to build a solution That will analyze scanned documents and export relevant fields to a database.

You need to recommend which Azure AI service to deploy for the following types of documents:

- * Internal expenditure request authorization forms
- * Supplier invoices

The solution must minimize development effort.

What should you recommend for each document type? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point



Answer:



Explanation:



Azure AI Document Intelligence offers pre-built models for common document types like invoices, which extract key fields (e.g., vendor, amount, date) with minimal setup. For custom forms like internal authorization requests, the pre-built general document model handles layout and text extraction, but to minimize development effort, use the pre-built invoice model where applicable and a custom model only if pre-built options fall short. However, the standard recommendation prioritizes pre-built for both to avoid training.

NEW QUESTION: 159

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

You plan to deploy the model as part of a mobile app for Android phones.
You need to prepare the model for deployment.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Retrain the model.

Test the model.

Change the model domain.

Export the model.

Answer Area

Answer:

Actions

Retrain the model.

Test the model.

Change the model domain.

Export the model.

Answer Area

Change the model domain.

Retrain the model.

Test the model.

Explanation:

Actions

Export the model.

Answer Area

1 Change the model domain.

2 Retrain the model.

3 Test the model.

For deployment to a mobile app (Android), Custom Vision requires a compact domain so the model can be exported and run on-device (e.g., as TensorFlow Lite or ONNX). The model was trained using the Retail domain, which is not an exportable compact domain for object detection. Therefore:
Change the model domain to an exportable Object Detection (compact) domain (e.g., General [compact]).
Retrain the model so weights are learned under the new domain-domain changes require retraining.
Export the model for Android (typically TensorFlow Lite or ONNX) and integrate it into the app.
This sequence prepares the model for on-device mobile deployment with minimal extra development.
References (Microsoft Azure AI Solution)
Custom Vision: Only compact domains support model export for edge/mobile.
Changing a domain requires retraining the iteration.
Exporting object detection models supports TensorFlow/TensorFlow Lite/ONNX formats for Android integration.

You have an Azure subscription that contains a Content Safety in Foundry Control Plane resource.

You are building a social media messaging app.

You need to build a solution that will analyze messages and flag messages that contain explicit content.

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

NOTE: Each correct selection is worth one point.

Answer Area



```
public static Analyze(dynamic client, string text)
{
    var request = new
    {
        AnalyzeText(text);
        AnalyzeTextOptions(text);
        AnalyzeImage(text);
        client.AnalyzeText(text);
    };

    return
    {
        AnalyzeImage(request);
        client.AnalyzeImage(request);
        client.AnalyzeText(request);
        request.AnalyzeImage(client);
    };
}
```

Answer:

Answer Area



```
public static Analyze(dynamic client, string text)
{
    var request = new
    {
        AnalyzeText(text);
        AnalyzeTextOptions(text);
        AnalyzeImage(text);
        client.AnalyzeText(text);
    };

    return
    {
        AnalyzeImage(request);
        client.AnalyzeImage(request);
        client.AnalyzeText(request);
        request.AnalyzeImage(client);
    };
}
```

Explanation:

Answer Area



```
public static Analyze(dynamic client, string text)
{
    var request = new
    {
        AnalyzeTextOptions(text);
        client.AnalyzeImage(request);
    };
}
```

NEW QUESTION: 161

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: ([SHOW ANSWER](#))

Scenario Recap:

IoT data from 100 machines × 50 sensors × 1-minute intervals for 1 year = 5,000 time series datasets .

Requirement: Detect unusual values (outliers) to predict failures.

Correct Service:

Azure Anomaly Detector (A) is designed for time-series anomaly detection.

It learns temporal patterns and detects anomalies in sensor/IoT data streams.

Other options:

Cognitive Search = full-text search

Form Recognizer = extract text/fields from documents

Custom Vision = image classification

The Answer: A. Anomaly Detector

Reference: Anomaly Detector overview

NEW QUESTION: 162

You have an Azure subscription that contains an Azure AI Content Safety resource named CS1. You plan to build an app that will analyze user-generated documents and identify obscure offensive terms. You need to create a dictionary that will contain the offensive terms. The solution must minimize development effort.

What should you use?

- A.** a text classifier
- B.** text moderation
- C.** language detection
- D.** a blacklist

Answer: D ([LEAVE A REPLY](#))

Azure AI Content Safety supports custom blocklists -dictionaries of terms/phrases you define to detect obscure or domain-specific offensive terms in user content. Creating a blocklist is the minimal-effort, built-in way to meet this requirement; you don't need to build a classifier.

References

Create and use blocklists in Azure AI Content Safety.

<https://learn.microsoft.com/azure/ai-services/content-safety/concepts/blocklists> Text moderation overview (using built-in classifiers with optional blocklists).

<https://learn.microsoft.com/azure/ai-services/content-safety/concepts/text-moderation>

NEW QUESTION: 163

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:

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)

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>

NEW QUESTION: 164

You have an app that analyzes images by using the Computer Vision API.

You need to configure the app to provide an output for users who are vision impaired. The solution must provide the output in complete sentences.

Which API call should you perform?

- A. readInputStreamAsync
- B. describeImageInStreamAsync
- C. tagImageInStreamAsync
- D. analyzeImageByDomainInStreamAsync

Answer: B (LEAVE A REPLY)

* To generate complete sentences describing an image (for accessibility for visually impaired users), you must call the Describe Image API.

* describeImageInStreamAsync analyzes the content of the image and generates human-readable descriptions in natural language sentences.

* readInputStreamAsync is for OCR (extracting text).

* tagImageInStreamAsync is for generating tags (keywords), not full sentences.

* analyzeImageByDomainInStreamAsync is for domain-specific models (e.g., celebrities, landmarks).

The answer: B

Reference: Computer Vision Describe Images

NEW QUESTION: 165

You have an Azure subscription that contains an Azure OpenAI resource named AM.
You build a chatbot that uses All to provide generative answers to specific questions.
You need to ensure that questions intended to circumvent built-in safety features are blocked.
Which Azure AI Content Safety feature should you implement?

- A.** Protected material text detection
- B.** Jailbreak risk detection
- C.** Monitor online activity
- D.** Moderate text content

Answer: B ([LEAVE A REPLY](#))

- * The requirement: Ensure the chatbot blocks questions that attempt to circumvent built-in safety features.
- * This is known as a jailbreak attempt (e.g., prompt injection, role-play tricking).
- * Azure AI Content Safety recently added Jailbreak risk detection for precisely this purpose.
- * Other options:
- * Protected material detection # for copyright/IP detection.
- * Monitor online activity # not related to OpenAI safety.
- * Moderate text content # general harmful content moderation (violence, hate, sexual, self-harm), but not jailbreaks.

Microsoft Reference:

Azure AI Content Safety jailbreak risk detection

NEW QUESTION: 166

You have an Azure subscription and 10,000 ASCII files.
You need to identify files that contain specific phrases. The solution must use cosine similarity.
Which Azure OpenAI model should you use?

- A.** text-embedding-ada-002
- B.** GPT-4
- C.** GPT-35 Turbo
- D.** GPT-4-32k

Answer: A ([LEAVE A REPLY](#))

- You need to identify files containing specific phrases using cosine similarity.
Cosine similarity is used in vector embeddings, not in generative LLM completion.
- * text-embedding-ada-002 # Correct model for embeddings and cosine similarity search.
 - * GPT-4, GPT-3.5 Turbo, GPT-4-32k # Generative models, not designed for embedding similarity.

The answer: A. text-embedding-ada-002

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (416 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 167

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 (LEAVE A REPLY)

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>

NEW QUESTION: 168

You need to build an app that will use the Azure AI Speech service to translate audio files.

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

NOTE: Each correct selection is worth one point.

Answer Area



Microsoft

```
static void OutputSpeechRecognitionResult(TranslationRecognitionResult translationRecognitionResult)
{
    switch (translationRecognitionResult.Reason)
    {
        case ResultReason.TranslatedSpeech:
            Console.WriteLine($"RECOGNIZED: Text={translationRecognitionResult.Text}");
            foreach (var element in translationRecognitionResult.Translations)
            {
                Console.WriteLine($"TRANSLATED into '{element.Key}': {element.Value}");
            }
            break;
        case ResultReason.NoMatch:
            Console.WriteLine($"Speech could not be recognized.");
            break;
    }
}
```


reason
reason
text
translations

TranslatedSpeech
NoMatch
RecognizingSpeech
TranslatedSpeech

Answer:

Answer Area

```
static void OutputSpeechRecognitionResult(TranslationRecognitionResult translationRecognitionResult)
{
    switch (translationRecognitionResult.reason)
    {
        case ResultReason.TranslatedSpeech:
            Console.WriteLine($"RECOGNIZED: Text={translationRecognitionResult.Text}");
            foreach (var element in translationRecognitionResult.Translations)
            {
                Console.WriteLine($"TRANSLATED into '{element.Key}': {element.Value}");
            }
            break;
        case ResultReason.NoMatch:
            Console.WriteLine($"Speech could not be recognized.");
            break;
    }
}
```



Explanation:

QUESTION 169

```
static void OutputSpeechRecognitionResult(TranslationRecognitionResult translationRecognitionResult)
{
    switch (translationRecognitionResult.Reason)
    {
        case ResultReason.TranslatedSpeech:
            Console.WriteLine($"RECOGNIZED: Text={translationRecognitionResult.Text}");
            foreach (var element in translationRecognitionResult.Translations)
            {
                Console.WriteLine($"TRANSLATED into '{element.Key}': {element.Value}");
            }
            break;
        case ResultReason.NoMatch:
            Console.WriteLine($"Speech could not be recognized.");
            break;
    }
}
```

When using the Speech SDK's speech translation, a TranslationRecognitionResult includes a Reason property (type ResultReason) indicating the outcome. For successful translation results, the reason is ResultReason.TranslatedSpeech. In that case, you can read the recognized source text from translationRecognitionResult.Text and iterate the translated outputs from translationRecognitionResult.Translations (a dictionary keyed by target language code).

RecognizingSpeech corresponds to interim hypotheses raised on the Recognizing event and is not the final result reason used here; NoMatch is the branch already handled for unrecognized audio.

Microsoft Azure AI References

- * Azure AI Speech SDK - Speech translation overview and TranslationRecognizer usage (C#): result handling with ResultReason.TranslatedSpeech, reading Text and Translations.
- * Speech SDK API reference - TranslationRecognitionResult (Reason, Text, Translations) and ResultReason enum (values include TranslatedSpeech, NoMatch, Canceled).

NEW QUESTION: 169

Match the types of workloads to the appropriate scenarios.

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

NOTE: Each correct match is worth one point.

Workload Types	Answer Area
Batch	Workload type Data for a product catalog will be loaded every 12 hours to a data warehouse.
Streaming	Workload type Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time.
	Workload type Updates to inventory data will be loaded to a data warehouse every 1 million transactions.

Answer:

Workload Types

Batch

Streaming

Answer Area

Batch

Streaming

Batch

Data for a product catalog will be loaded every 12 hours to a data warehouse.

Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time.

Updates to inventory data will be loaded to a data warehouse every 1 million transactions.

Explanation:

Answer Area

Batch

Streaming

Batch

Data for a product catalog will be loaded every 12 hours to a data warehouse.

Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time.

Updates to inventory data will be loaded to a data warehouse every 1 million transactions.

Comprehensive Detailed Explanation

There are two main types of data workloads: Batch and Streaming.

Batch Workload

Processes data in large groups (batches) at scheduled intervals.

Common when real-time is not required.

Examples: daily sales reports, hourly data refreshes, periodic ETL (Extract, Transform, Load).

Streaming Workload

Processes data in real time or near real time as it arrives.

Used for scenarios where immediate insights or actions are required.

Examples: fraud detection, IoT sensor monitoring, live purchase tracking.

Applying to the scenarios:

"Data for a product catalog will be loaded every 12 hours to a data warehouse." This is scheduled, not continuous. # Batch

"Thousands of data sets per second for online purchases will be loaded into a data warehouse in real time." This requires real-time processing. # Streaming

"Updates to inventory data will be loaded to a data warehouse every 1 million transactions." Data is grouped and loaded after reaching a threshold (1 million). # Batch Correct Matches:

Product catalog (12 hours) # Batch

Online purchases (real time) # Streaming

Inventory data (1 million transactions) # Batch

Microsoft References

Batch processing vs. Stream processing in Azure

Azure Stream Analytics overview

Data processing approaches: Batch vs Real-time

NEW QUESTION: 170

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

NOTE: Each correct selection is worth one point.

Answer Area

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

Answer:

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

Explanation:

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

Normalization involves eliminating relationships between database tables.

No

Normalization does not eliminate relationships. In fact, it often creates relationships between tables by splitting data into smaller, related tables.

Normalizing a database reduces data redundancy.

Yes

The primary goal of normalization is to minimize duplicate data and ensure that each fact is stored only once.

Normalization improves data integrity.

Yes

By reducing redundancy and ensuring data is stored in a structured manner, normalization enforces consistency and integrity across the database.

Correct Responses:

No

Yes

Yes

Microsoft References

Database normalization basics

Database design and normalization

NEW QUESTION: 171

You are building content for a video training solution.

You need to create narration to accompany the video content. The solution must use Custom Neural Voice.

What should you use to create a custom neural voice, and which service should you use to generate the narration? To answer, select the appropriate options in the answer area.

NOTE: Each correct answer is worth one point.

Answer Area

Custom neural voice:

Narration:

Narration:

Answer:

Answer Area

Custom neural voice:

Narration:

Narration:

Explanation:

Answer Area

Custom neural voice:

Narration:

To build narration for video training content using Custom Neural Voice, we must know where each step is done:

Creating a Custom Neural Voice

A Custom Neural Voice is a text-to-speech (TTS) feature within Azure AI Speech.

It allows you to create a unique, natural-sounding synthetic voice.

The correct tool to build and configure a custom voice is the Speech Studio portal.

Not the Language Understanding portal (LUIS), Azure portal directly, or Bot Framework Composer.

Generating narration

Narration requires converting written text into speech.

This is done via the Text-to-Speech service in Azure AI Speech.

Other options:

Language Understanding # For intent/utterance processing, not narration.

Speaker Recognition # For identifying/authenticating speakers, not narration.

Speech-to-text # Converts spoken words to text (the opposite direction).

Thus:

Custom Neural Voice # Speech Studio portal

Narration # Text-to-speech

Correct Answers:

Custom neural voice: The Speech Studio portal

Narration: Text-to-speech

Microsoft References

Custom Neural Voice in Speech Studio

Text-to-Speech overview

NEW QUESTION: 172

You have a Language service resource that performs the following:

* Sentiment analysis

* Named Entity Recognition (NER)

* Personally Identifiable Information (PII) identification

You need to prevent the resource from persisting input data once the data is analyzed. Which query parameter in the Language service API should you configure?

A. loggingOptOut

B. piiCategories

C. showStats

D. Model-version

Answer: (SHOW ANSWER)

By default, Azure AI Language may log request and response data for monitoring.

Setting the loggingOptOut parameter to true ensures that input data is not persisted after analysis.

piiCategories specifies which PII types to detect.

showStats enables request/response statistics.

model-version specifies which model version to use.

The Answer: A

Reference: Azure AI Language API parameters

NEW QUESTION: 173

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:

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"  
  }  
}
```

Explanation:

▼

https://management.azure.com/subscriptions/xxxxxxxx-xxxx

PATCH

POST

PUT

xxxx-xxxx-
xxxxxxxxxxxx/resourceGroups/RG1/providers/Microsoft.CognitiveServices
accounts/CS1?api-version=2017-04-18
{
 "location": "West US",
 "kind": "

▼

CognitiveServices

ComputerVision

TextAnalytics

",
 "sku": {
 "name": "S0"
 },
 "properties": {},
 "identity": {
 "type": "SystemAssigned"
 }
}

To create a single multi-service Cognitive Services resource (now Azure AI Services) that provides one key /endpoint for multiple services , consolidates billing, and allows adding services like Text Analytics (sentiment analysis) and Computer Vision (OCR) from the same resource, you must: Use the Azure Resource Manager (ARM) REST API with PUT to create or update the account:
PUT /subscriptions/{subld}/resourceGroups/{rg}/providers/Microsoft.CognitiveServices/accounts/{name}?api-version=2017-04-18
Set " kind " : " CognitiveServices " to create the multi-service account (not a single-service kind like ComputerVision or TextAnalytics).
This satisfies the requirements: single key/endpoint, consolidated billing, and future Computer Vision use.

Microsoft Azure AI Solution References

ARM REST for Cognitive Services accounts (create/update via PUT , multi-service " kind " : " CognitiveServices ").
<https://learn.microsoft.com/azure/ai-services/management/resource-manager-create-account> Azure AI Services multi-service resource (single endpoint/key across services).
<https://learn.microsoft.com/azure/ai-services/multi-service-resource>

NEW QUESTION: 174

You have an Azure subscription that contains an AI enrichment pipeline in Azure Cognitive Search and an Azure Storage account that has 10 G8 of scanned documents and images. You need to index the documents and images in the storage account. The solution must minimize how long it takes to build the index. What should you do?

A. From the Azure portal, configure parallel indexing.
B. Create a text-based indexer by using the REST API.
C. From the Azure portal, configure scheduled indexing.

D. Configure field mappings by using the REST API.

Answer: (SHOW ANSWER)


- * You need to minimize indexing time for 10 GB of scanned docs and images.
- * Azure Cognitive Search supports parallel indexing to increase throughput by partitioning work.
- * A text-based indexer or scheduled indexing will not reduce build time.
- * Field mappings help transform fields but don't affect indexing speed.
- * Therefore, the correct way to speed up index building is parallel indexing.

The answer: A

NEW QUESTION: 175

Select the answer that correctly completes the sentence.

Answer Area



Answer:

Answer Area



Explanation:

account

The statement is:

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

account # Correct. When you create a Cosmos DB account, you must choose the API type (e.g., SQL API, MongoDB API, Cassandra API, Table API, Gremlin API). This decision defines the data model and query language. # container # A container is created inside a database (tables, collections, graphs) but does not determine the API type.

database # Exists inside an account and inherits the API type from the account.

item # Represents individual records/documents; not where the API is chosen.

Thus, the API choice is made at the account level.

The Answer: account

Microsoft References

Azure Cosmos DB account overview

Choose the right API for Cosmos DB

NEW QUESTION: 176

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

- A. durability
- B. atomicity
- C. consistency

D. isolation

Answer: ([SHOW ANSWER](#))

The D in ACID stands for Durability .

Durability guarantees that once a transaction is committed, the changes are permanent, even in the case of power loss, crashes, or OS restarts.

Atomicity = all-or-nothing.

Consistency = valid state transitions.

Isolation = transactions run independently without interference.

The Answer: A. durability

NEW QUESTION: 177

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: ([SHOW ANSWER](#))

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>

NEW QUESTION: 178

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: B (LEAVE A REPLY)

<https://learn.microsoft.com/en-us/azure/cognitive-services/language-service/named-entity-recognition>

/overview Named Entity Recognition (NER) is one of the features offered by Azure Cognitive Service for Language, a collection of machine learning and AI algorithms in the cloud for developing intelligent applications that involve written language. The NER feature can identify and categorize entities in unstructured text. For example: people, places, organizations, and quantities.

NEW QUESTION: 179

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

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

You need to run model1 on Host1.

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

Actions

Configure disk logging.

Export model1 to Host1.

Request approval to run the container.

Retrain the model.

Run the container.

Answer Area

Answer:

Configure disk logging.

Export model1 to Host1.

Request approval to run the container.

Retrain the model.

Run the container.

Export model1 to Host1.

Run the container.

Configure disk logging.

Explanation:

According to the course AI-102T00: Designing and Implementing a Microsoft Azure AI Solution1, the correct sequence of actions to run model1 on Host1 is:

Export model1 to Host1

Run the container

Configure disk logging

The explanation and references are as follows:

Export model1 to Host1: This step is required to deploy the custom speech-to-text model to the Docker host. You can use the Azure portal or the Azure CLI to export the model as a container image2.

Run the container: This step is required to start the container and run the model on the Docker host. You can use the Docker CLI to run the container image3.

Configure disk logging: This step is optional but recommended to monitor the performance and health of the container. You can use the Docker CLI to configure disk logging for the container4.

NEW QUESTION: 180

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: A,B,C ([LEAVE A REPLY](#))

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>

NEW QUESTION: 181

You are developing an app that will use the Speech and language APIs.

You need to provision resources for the app. The solution must ensure that each service is accessed by using a single endpoint and credential Which type of resource should you create?

A. Azure AI Content Safety

B. Azure AI service

C. Azure AI Speech

D. Azure AI Language

Answer: B ([LEAVE A REPLY](#))

To access multiple Azure AI capabilities (e.g., Speech and Language) using a single endpoint and a single set of credentials, you should provision an Azure AI multi-service resource (called Azure AI service in the portal; formerly "Cognitive Services" multi-service account). This resource provides one endpoint and one key (or Microsoft Entra ID) that authorizes calls across supported services, including Speech and Language. Creating separate "Azure AI Speech" or "Azure AI Language" resources would give you distinct endpoints/keys for each, which does not meet the requirement. Azure AI Content Safety is unrelated.

Microsoft References

Azure AI services multi-service resource (single endpoint/credential across services).

<https://learn.microsoft.com/azure/ai-services/multi-service-resource>

What are Azure AI services (overview of consolidated access model).

<https://learn.microsoft.com/azure/ai-services/what-are-ai-services>

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (416 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 182

You have an Azure subscription that contains an Azure AI Foundry Content Safety resource named resource1.

You are building an app that will analyze text by using resource1.

You need to identify text that contains hateful content.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
'''
client = ContentSafetyClient(endpoint, AzureKeyCredential(key))
request = AnalyzeTextOptions(text="Some animals are more equal than others")
# Analyze text
try:
    response = client.analyze_text(request)
except HttpResponseError as e:
    print("Analyze text failed.")
    raise
hate_result = next(item for item in response.items() if item.category == TextCategory.HATE)
if hate_result:
    print(f"Hate severity: ")
'''
except HttpResponseError as e:
    print("Analyze text failed: {hate_result.blocklist_item_id}")
    raise
hate_result = next(item for {hate_result.category}) if item.category == TextCategory.HATE)
if hate_result:
    print(f"Hate severity: {hate_result.severity}")
```

blocklist_match

categories_analysis

content

{hate_result.blocklist_item_text}

{hate_result.severity}

Answer:

Answer Area

```
...
client = ContentSafetyClient(endpoint, AzureKeyCredential(key))

request = AnalyzeTextOptions(text="Some animals are more equal than others")

# Analyze text

try:

    response = client.analyze_text(request)
except HttpResponseError as e:
    print("Analyze text failed.")
    raise

hate_result = next(item for item in response.categories_analysis if item.category == TextCategory.HATE)

if hate_result:
    print(f"Hate severity: {hate_result.severity}")
...

except HttpResponseError as e:
    print("Analyze text failed")
    raise

hate_result = next(item for item in response.categories_analysis if item.category == TextCategory.HATE)

if hate_result:
    print(f"Hate severity: {hate_result.severity}")
```

Explanation:

```
Answer Area

...
client = ContentSafetyClient(endpoint, AzureKeyCredential(key))
request = AnalyzeTextOptions(text="Some animals are more equal than others")
# Analyze text
try:
    response = client.analyze_text(request)
except HttpResponseError as e:
    print("Analyze text failed.")
    raise
hate_result = next(item for item in response.categories_analysis if item.category == TextCategory.HATE)
if hate_result:
    print(f"Hate severity: {hate_result.severity}")
```

To detect hateful content with Azure AI Foundry Content Safety - Text , you should read the category analysis results returned by the `analyze_text` call. The response exposes a collection named `categories_analysis`, where each item contains the analyzed category (e.g., `TextCategory.HATE`) and its severity level.

Hence:

Iterate over `response.categories_analysis`, filter the item where `item.category == TextCategory.HATE`.

Print the severity using `hate_result.severity`.

This uses category analysis (not blocklist matches), which is the correct mechanism for detecting hate content with a severity score.

Microsoft References

Azure AI Content Safety - Analyze text (Python quickstart/samples): iterate response.categories_analysis, check TextCategory.HATE, read item.severity.
Azure AI Content Safety - Text categories and severity levels: explains categories (Hate, Violence, Sexual, Self-harm) and severity output.

NEW QUESTION: 183

You have an Azure subscription that contains an Azure AI Search resource named AS1.
You implement a custom skill in AS1 that performs language and sentiment analysis of documents.
You are evaluating the use of AS1 as part of an enrichment pipeline.
In which order will AS1 index the documents? To answer, move all indexing stages from the list of stages to the answer area and arrange them in the correct order.

Stages

document cracking

skillset execution

push to index

output field mappings

field mappings

Answer Area

Answer:

Stages

document cracking

skillset execution

push to index

output field mappings

field mappings

Answer Area

document cracking

skillset execution

output field mappings

field mappings

push to index

Explanation:

Stages

Answer Area

1 document cracking

2 skillset execution

3 output field mappings

4 field mappings

5 push to index

Comprehensive Detailed Explanation
When an indexer runs in Azure AI Search, it processes each document through a fixed series of stages:

Document cracking

The indexer opens the item (file/row) and extracts text, images, and metadata, creating the root of the enriched document (/document). This is always the first stage.

Field mappings

Next, source fields are mapped directly to target index fields before any transformations . Field mappings occur after document cracking but before skillset execution . Microsoft Learn Skillset execution Built-in or custom skills (for example, language and sentiment analysis in your scenario) run and write outputs into the in-memory enriched document. Microsoft Learn Output field mappings Because skill outputs live only in memory, you must define output field mappings to route selected nodes of the enriched document into index fields. This step comes after skillset execution.

Push to index

Finally, the indexer hands off raw and enriched content to the search index (the physical data structures) so it becomes searchable-i.e., the content is "pushed to the index." Microsoft Learn This order- document cracking # field mappings # skillset execution # output field mappings # push to index -ensures raw content is extracted, direct source-to-index mappings are applied, AI enrichment is performed, enriched outputs are projected to fields, and the completed documents are written into the index.

Microsoft Azure AI References

Indexers in Azure AI Search - Stages of indexing (Document cracking; Field mappings; Skillset execution; Output field mappings, with field mappings occurring before transformations). Microsoft Learn Skillset concepts (Output of skillset is routed to the index via output field mappings after skillset execution).

Microsoft Learn

AI enrichment concepts (Field mappings move raw source content to the index; output field mappings move enriched content; indexing ingests raw and enriched content into the physical index). Microsoft Learn

Annotations & enriched document tree (Enriched document is created during document cracking at

/document). Microsoft Learn

NEW QUESTION: 184

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: C,D,F (LEAVE A REPLY)

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>

Valid AI-102 Dumps shared by Actual4test.com for Helping Passing AI-102 Exam! Actual4test.com now offer the **newest AI-102 exam dumps**, the Actual4test.com AI-102 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AI-102 dumps with Test Engine here: https://www.actual4test.com/AI-102_examcollection.html (**416** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)