PASS QUESTION

PRACTICE QUESTIONS

EXAM PREPARATION GUIDE

Learn Smarter

Practice Effectively

Achieve Your Goals

https://www.passquestion.com

One Year of Content Updates Included

Exam : **AI-901**

Title : Microsoft Azure AI Fundamentals (Updated Version)

Version : DEMO

1.HOTSPOT

Select the answer that correctly completes the sentence.

When content is submitted to Azure Content Understanding in Foundry Tools,

the analysis is [ ].

- synchronous
- asynchronous
- returned only as unstructured plain text
- limited to optical character recognition (OCR)-only processing

Answer:

When content is submitted to Azure Content Understanding in Foundry Tools,

the analysis is [ ].

- synchronous
- asynchronous
- returned only as unstructured plain text
- limited to optical character recognition (OCR)-only processing

Explanation:

When content is submitted to Azure Content Understanding in Foundry Tools, the analysis is asynchronous. This means the service does not return results immediately within the same HTTP request. Instead, it uses the standard Azure long-running operation (LRO) pattern — you call begin_analyze() to submit the content, which immediately returns a poller object, and then call poller.result() to wait for processing to complete and retrieve the structured extraction results.

Why the other options are wrong:

Synchronous is incorrect — the analysis pipeline involves multiple AI steps (OCR, speech transcription, schema mapping) that take time; a blocking synchronous call is not supported.

Returned only as unstructured plain text is incorrect — Azure Content Understanding returns richly structured JSON output with named fields mapped to your defined schema, not plain unstructured text.

Limited to OCR-only processing is incorrect — Content Understanding goes far beyond OCR; it supports document, audio, image, and video analyzers, and performs semantic field extraction using AI, not just character recognition.

This asynchronous design is consistent across all Azure AI services that perform complex, multi-step content processing.

2.HOTSPOT

Select the answer that correctly complete the sentence.

Answer Area

To develop an application that analyzes text by using Azure Language in Foundry Tools, use the [ ] package.

- Azure AI Projects client library
- Azure OpenAI client library
- Azure Speech SDK
- Azure Text Analytics client library

Answer:

Answer Area

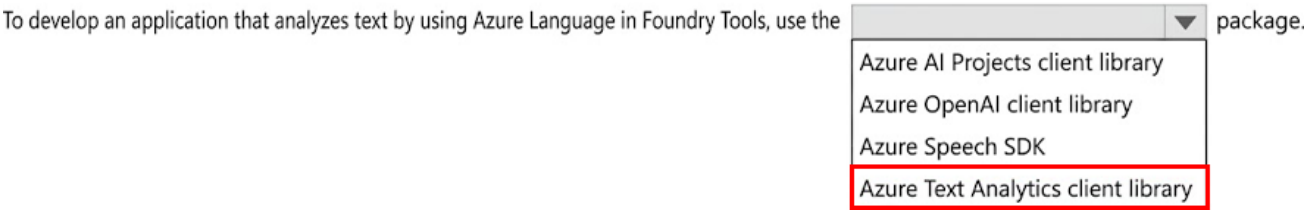

Explanation:

Azure Text Analytics client library

Comprehensive and Detailed Explanation with all Azure AI documents: =

The completed sentence is:

To develop an application that analyzes text by using Azure Language in Foundry Tools, use the Azure Text Analytics client library package.

Azure AI Language text analysis features, such as sentiment analysis, named entity recognition, key phrase extraction, language detection, and text summarization, are accessed in code by using the Azure Text Analytics client library.

The other options are incorrect:

Azure AI Projects client library is used for Foundry project and agent-related workflows, not specifically Azure Language text analysis.

Azure OpenAI client library is used to call OpenAI model deployments, not Azure Language text analytics APIs.

Azure Speech SDK is used for speech-to-text and text-to-speech workloads, not text analysis.

3.HOTSPOT

Select the answer that correctly completes the sentence.

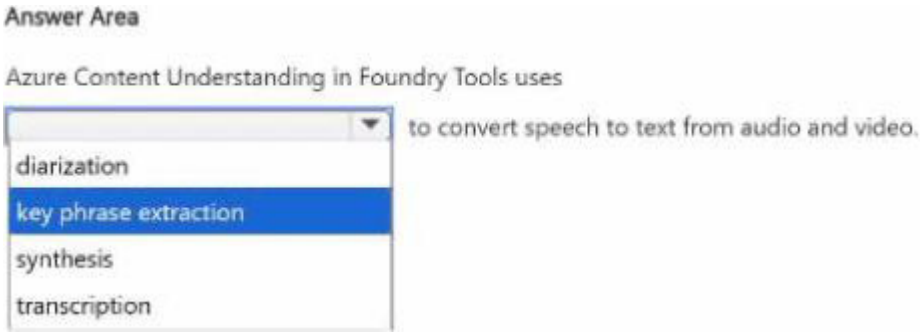

Answer:

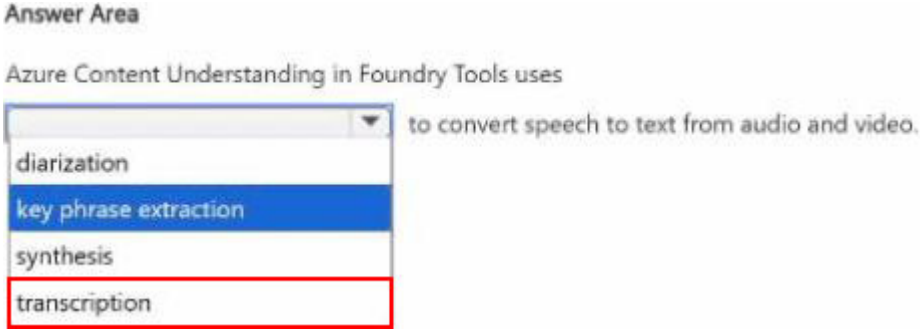

Explanation:

Azure Content Understanding in Foundry Tools uses transcription to convert speech to text from audio

and video.
Transcription is the process of converting spoken audio into written text. In Azure Content Understanding, audio and video analyzers can use transcription to produce text transcripts from spoken content.
The other options are incorrect:
diarization identifies or separates speakers, such as “Speaker 1” and “Speaker 2.” key phrase extraction identifies important phrases from text after text exists.
synthesis converts text into spoken audio, which is the opposite direction of speech-to-text.

4.You have a Microsoft Foundry project that contains a vision-enabled model deployment.
You need to develop an application that sends a message containing text and an image URL. The solution must ensure the quickest response time.
Which message structure should you include in the request?
A. a user message that includes only a text item and sends an image item in a separate request
B. a system message that includes both a text item and an image item in the content array
C. a system message that includes only a text item and sends an image item in a separate request
D. a user message that includes both a text item and an image item in the content array
Answer: D
Explanation:
For a vision-enabled model, Microsoft documentation shows that multimodal prompts can include both text and image content in the same user message content array. Azure OpenAI vision models accept multimodal image-and-text input and return natural language responses.
The correct request structure is a user message with a content array containing both items, for example:
{
"role": "user", "content": [
{"type": "input_text", "text": "What is in this image?"}, {"type": "input_image", "image_url": image_url}
]
}
This avoids sending separate requests and is therefore the best choice for quickest response time.
A and C are incorrect because they send the text and image separately.
B is incorrect because the image/question should be sent as a user message, while system messages are used for instructions and behavior.

5.HOTSPOT
Select the answer that correctly completes the sentence.

Answer Area

When developing an AI system for self-driving cars, the Microsoft [▼] principle for responsible AI should be applied to ensure consistent operation of the system during its lifespan.

- Fairness
- reliability and safety
- accountability
- inclusiveness

Answer:

Answer Area

When developing an AI system for self-driving cars, the Microsoft [ ] principle for responsible AI should be applied to ensure consistent operation of the system during its lifespan.

- Fairness
- reliability and safety
- accountability
- inclusiveness

Explanation:

The completed sentence is:

When developing an AI system for self-driving cars, the Microsoft reliability and safety principle for responsible AI should be applied to ensure consistent operation of the system during use.

Self-driving cars are safety-critical AI systems. The key requirement in the question is consistent operation of the system during use, which maps directly to the Microsoft responsible AI principle of reliability and safety.

The other options are incorrect:

Fairness focuses on reducing bias and treating people equitably.

Accountability focuses on ensuring people remain responsible for AI systems and their outcomes.

Inclusiveness focuses on designing AI systems that work for and empower diverse users, including people with different needs or abilities.

Therefore, the correct answer is reliability and safety.