



Microsoft

Exam Questions AI-901

Microsoft Azure AI Fundamentals (Updated Version)

NEW QUESTION 1

You need to convert written customer notifications into natural-sounding spoken audio that can be played over a phone system. Which Azure Speech in Foundry Tools capability should you use?

- A. speaker recognition
B. speech synthesis
C. speech recognition
D. speech translation

Answer: B

Explanation:

The requirement is to convert written customer notifications into natural-sounding spoken audio. This is speech synthesis, also known as text to speech. Microsoft's Azure Speech documentation describes text to speech as a capability that converts text into natural-sounding synthesized speech. Therefore, for playing written notifications over a phone system, the correct Azure Speech capability is speech synthesis.

Why the other options are incorrect:

- * A. speaker recognitionidentifies or verifies speakers by voice.
- * C. speech recognitionconverts spoken audio into text.
- * D. speech translationtranslates spoken audio between languages.

NEW QUESTION 2

You are developing an application that extracts fields from PDFs by using Azure Content Understanding in Foundry Tools. You need to use the Python SDK to submit a PDF for analysis and retrieve the extraction results. What should you do?

- A. Call `begin_analyze()`, and then call `poller.result()` to retrieve the results.
- B. Submit the PDF to an analyzer and read the results from the request headers.
- C. Use optical character recognition (OCR) to extract text from the PDF and map the fields in the code.
- D. Call `analyze()` to return the extracted fields synchronously in the same request.

Answer: A

NEW QUESTION 3

HOTSPOT

You are developing an application that converts text into spoken audio and saves the synthesized audio to a file by using Azure Speech in Foundry Tools. How should you complete the Python code? To answer, select the appropriate option in the answer area.

NOTE:Each correct selection is worth one point.

```
import azure.cognitiveservices.speech as speechsdk
...
speech_config = speechsdk.SpeechConfig(subscription=
audio_config = speechsdk.audio.
synthesizer = speechsdk.SpeechSynthesizer(
    speech_config=speech_config,
    audio_config=audio_config
)
...
```

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

```
AudioOutputConfig(filename="output.wav")
```

The question specifically states the application must save the synthesized audio to a file. In the Azure Speech SDK for Python, `speechsdk.audio.AudioOutputConfig(filename="output.wav")` directs the synthesizer to write the generated speech output directly to a WAV file on disk — which is exactly the requirement.

Why the other options are wrong:

AudioOutputConfig(stream)— This routes audio output to an in-memory audio stream object, not a file. It is used when you want to process or play the audio programmatically without saving it to disk.

`AudioStreamFormat(wave_stream_format=AudioStreamWaveFormat.PCM)`— This class defines the format of an audio stream (e.g., PCM encoding, sample rate). It is used when configuring custom audio streams, not when specifying a file output destination. It is not a valid argument for `speechsdk.audio.` in this context.

The correct and complete line is:

```
audio_config = speechsdk.audio.AudioOutputConfig(filename="output.wav")
```

NEW QUESTION 4

NEW QUESTION 4

You are using the Azure Speech SDK to develop a Python application that supports real-time spoken conversations. Which Azure speech class should you use to configure the connection to the Azure Speech service?

- A. AudioConfig
- B. SpeechSynthesizer
- C. AuditOutputConfig

Answer: C

NEW QUESTION 5

What are two purposes of instructions when prompting a generative AI model? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

- A. defines constraints on the model's responses
- B. defines the agent's role and behavior
- C. defines the Azure region where inference occurs
- D. selects which model to use
- E. defines the tokens per minute (TPM) allocation for the model

Answer: AB

NEW QUESTION 6

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

Answer Area

Statements	Yes	No
System prompts can be used to authorize users.	☐	☐
A system prompt is used to reduce tokens per minute (TPM).	☐	☐
A system prompt guides the behavior of a generative AI model.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

A system prompt, also called a system message, is used to guide how a generative AI model behaves, including its tone, role, response style, and output format. Microsoft states that system messages help steer an Azure OpenAI chat model toward the desired behavior, tone, and output format.
Statement 1: System prompts can be used to authorize users. = No
Authorization is handled by application identity, access control, authentication, and security configuration. A system prompt can guide model behavior, but it is not an authorization mechanism.
Statement 2: A system prompt is used to reduce tokens per minute (TPM). = No
TPM is a quota or throughput setting for model deployments. A system prompt does not reduce or configure TPM.
Statement 3: A system prompt guides the behavior of a generative AI model. = Yes
This is correct. Microsoft documentation says system messages help steer the model's behavior, tone, and output format.

NEW QUESTION 7

HOTSPOT
You are developing a voice application that listens for spoken commands and converts them into text by using Azure Speech in Foundry Tools. How should you complete the Python code? To answer, select the appropriate option in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
import azure.cognitiveservices.speech as speechsdk
...
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
recognizer = speechsdk.SpeechRecognizer(speech_config=speech_config)
recognizer.recognize_once()
recognizer.speak_text_async("Ready")
recognizer.start_continuous_recognition()
recognizer.start_keyword_recognition()
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The correct Python method is:

`recognizer.recognize_once()`

Completed code:

```
import azure.cognitiveservices.speech as speechsdk
speech_config = speechsdk.SpeechConfig(subscription=key, region=region)
recognizer = speechsdk.SpeechRecognizer(speech_config=speech_config)
recognizer.recognize_once()
```

Microsoft's Azure Speech documentation for Foundry Tools explains that Speech to text is used for real-time speech recognition and converting spoken audio into text. The Python Speech SDK uses a `SpeechRecognizer` object for recognition from audio input, such as a microphone.

Why the other options are incorrect:

`recognizer.speak_text_async("Ready")` is incorrect because speaking text is text-to-speech, not speech-to-text recognition.

`recognizer.start_continuous_recognition()` can be used for continuous recognition, but the code shown is asking for the basic method to recognize spoken input and convert it to text from the `SpeechRecognizer`.

`recognizer.start_keyword_recognition()` is used for keyword/wake-word recognition, not general speech-to-text transcription of spoken commands.

Therefore, the correct answer is:

`recognizer.recognize_once()`

NEW QUESTION 8

You need to build an AI solution that generates marketing email drafts based on a short description of a product and its target audience.

Which AI workload should you use?

- A. generative AI
- B. computer vision
- C. text classification
- D. speech recognition

Answer: A

NEW QUESTION 9

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

After deploying a vision-enabled GPT model in Microsoft Foundry,

you can configure an application to send requests to the

endpoint of the model
 evaluation pipeline of model
 Foundry playground
 training dataset of the model


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

After deploying a vision-enabled GPT model in Microsoft Foundry,
you can configure an application to send requests to the



NEW QUESTION 10

You are developing an application that analyzes voicemail recordings by using Azure Content Understanding in Foundry Tools. You need to extract a transcript and structured information from the recordings. Which type of analyzer should you use?

- A. document analyzer
- B. video analyzer
- C. audio analyzer
- D. image analyzer

Answer: C

NEW QUESTION 10

You are developing a web app that processes invoices to calculate expenses. You need to extract structured fields, including nested values, from the invoices by using a defined schema. What should you use?

- A. a transcription workflow in Azure Speech in Foundry Tools
- B. an optical character recognition (OCR)-only document processing pipeline
- C. an analyzer in Azure Content Understanding in Foundry Tools
- D. an Azure AI Search service

Answer: C

NEW QUESTION 14

DRAG DROP

You have a Microsoft Foundry project that contains a model deployment. You are developing an application that sends an image and a user question to the model. You need to send both text and image content in the same request so the model can return an answer. How should you complete the Python code? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

Values

input_image

input_text

input_url

output_image

output_text

Answer Area

```
input=[
    {
        "role": "user",
        "content": [
            {"type": "image_url", "image_url": input_image},
            {"type": "text", "text": "What is in this image? Provide 3 bullet points."},
        ],
    },
]
response = client.chat.completions.create(model=model_id, messages=input)
print(response.output_text)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Values

input_image

input_text

input_url

output_image

output_text

Answer Area

```
input=[
  {
    "role": "user",
    "content": [
      {
        "type": "text",
        "text": "What is in this image? Provide 3 bullet points."
      },
      {
        "type": "image_url",
        "image_url": image_url
      }
    ]
  }
],
print(response.output_text)
```

NEW QUESTION 17

HOTSPOT

Select the answer that correctly complete the sentence.

Answer Area

The

Model catalog

Monitor page

Service endpoints page

Solution templates page

 is used for comparing and deploying a wide range of models for generative AI development in Microsoft Foundry.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

The

Model catalog

Monitor page

Service endpoints page

Solution templates page

 is used for comparing and deploying a wide range of models for generative AI development in Microsoft Foundry.

NEW QUESTION 20

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": "text",
      "text": "What is in this image?"
    },
    {
      "type": "image_url",
      "image_url": "https://example.com/image.jpg"
    }
  ]
}
```

```
{ "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.

NEW QUESTION 23

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

When using the OpenAI Responses API and a vision-enabled model, you can include an image in a request by providing the image as

a base64-encoded image data
a CSV file attachment
an MP4 video stream
a shared access signature (SAS) token

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

When using the OpenAI Responses API and a vision-enabled model, you can include an image in a request by providing the image as

a base64-encoded image data
a CSV file attachment
an MP4 video stream
a shared access signature (SAS) token

NEW QUESTION 27

You are developing an AI-powered customer support application.
Which task is an example of the Microsoft responsible AI principle of inclusiveness?

- A. Provide explanations about how predictions are generated.
- B. Design the interface to support multiple languages and screen readers.
- C. Evaluate model outputs across demographic groups to reduce bias.
- D. Encrypt stored customer data and restrict access by using role-based controls.

Answer: B

NEW QUESTION 30

DRAG DROP

You are developing an application that extracts structured information from different types of content by using Azure Content Understanding in Foundry Tools. You need to extract scanned invoices in the PDF format and voicemail recordings in the WAV format. Which type of analyzer should you use for each content type? To answer, drag the appropriate analyzer types to the correct content types. Each analyzer type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Analyzer Types

audio analyzer
document analyzer
image analyzer
video analyzer

Answer Area

Scanned invoices:
Voicemail recordings:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Analyzer Types

audio analyzer

document analyzer

image analyzer

video analyzer

Answer Area

Scanned invoices:

document analyzer

Vocemall recordings:

audio analyzer

NEW QUESTION 33

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

A keyword list

Optical character recognition (OCR)-only processing

A schema

A synchronous API call

defines which fields to extract when analyzing content

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A keyword list

Optical character recognition (OCR)-only processing

A schema

A synchronous API call

defines which fields to extract when analyzing content

NEW QUESTION 37

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

To define an agent's role and behaviors, you must configure a

deployment slot

embedding index

fine tuning job

system prompt

for the agent.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

To define an agent's role and behaviors, you must configure a

deployment slot

embedding index

fine tuning job

system prompt

 for the agent.

NEW QUESTION 38

You have a Microsoft Foundry project that contains a vision-enabled model deployment. You are developing an application that sends images to the model. You need to ensure that the model can analyze the images.
In which two formats can you provide the images? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

- A. a JSON document that describes the image content
- B. a UTF-8 encoded text description of the image
- C. a publicly accessible URL of the image
- D. a base64 encoded image data string

Answer: CD

NEW QUESTION 42

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

Answer Area

Statements	Yes	No
The Temperature parameter can be set before deploying a model.	☐	☐
During inference, the model name is used to route requests to a specific deployment.	☐	☐
After a model is deployed, both code and testing tools can be used to interact with the model.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The Temperature parameter can be set before deploying a model.	☐	☒
During inference, the model name is used to route requests to a specific deployment.	☐	☒
After a model is deployed, both code and testing tools can be used to interact with the model.	☒	☐

NEW QUESTION 45

You are developing an application that processes voicemail recordings by using Azure Content Understanding in Foundry Tools. Which feature does Azure Content Understanding use to convert audio to text?

- A. Voice Live
- B. key phrase extraction
- C. transcription
- D. optical character recognition (OCR)

Answer: C

Explanation:

Azure Content Understanding usestranscriptionto convert audio content, such as voicemail recordings, into text. Microsoft??s Azure Content Understanding audio documentation states that transcription converts conversational audio into searchable and analyzable text-based transcripts.

OptionA. Voice Liveis not the feature used by Content Understanding to convert voicemail recordings into text.

OptionB. key phrase extractionidentifies important phrases after text is available; it is not the audio-to-text conversion step.

OptionD. optical character recognition (OCR)is for extracting text from images or documents, not audio.

Therefore, the correct answer isC. transcription.

NEW QUESTION 49

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

An AI workload that produces new content based on user input is an example of

content understanding

generative AI

information extraction

text analysis

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

An AI workload that produces new content based on user input is an example of

content understanding

generative AI

information extraction

text analysis

NEW QUESTION 50

You have a Microsoft Foundry project that contains a generative AI model deployment. You test the model by using the Foundry playground. You need to develop an application that sends requests to the deployed model. Which information must the application include to call the model?

- A. The model training dataset
- B. The Foundry project display name
- C. The exported playground session history
- D. The model endpoint and authentication credentials

Answer: D

Explanation:

To call a deployed Azure OpenAI model from an application, the app must know the service endpoint and authenticate its request. Microsoft documentation states that Azure OpenAI supports API key authentication or Microsoft Entra ID authentication, and API key authentication requires including the API key in the request. Microsoft quickstart guidance also states that to successfully make a call against Azure OpenAI, you need an endpoint and a key. The application does not need the model training dataset, the Foundry project display name, or exported playground session history to call the deployed model.

NEW QUESTION 51

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

AI-901 Practice Exam Features:

- * AI-901 Questions and Answers Updated Frequently
- * AI-901 Practice Questions Verified by Expert Senior Certified Staff
- * AI-901 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AI-901 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AI-901 Practice Test Here](#)