


AI-901 Dumps

Microsoft Azure AI Fundamentals (Updated Version)

<https://www.certleader.com/AI-901-dumps.html>



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 recognition identifies or verifies speakers by voice.
- * C. speech recognition converts spoken audio into text.
- * D. speech translation translates 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=key, region=region)
audio_config = speechsdk.audio.

synthesizer = speechsdk.SpeechSynthesizer(
    speech_config=speech_config,
    audio_config=audio_config
)

...
```

AudioOutputConfig(filename="output.wav")

AudioOutputConfig(stream)

AudioStreamFormat(wave_stream_format=AudioStreamWaveFormat.PCM)

- 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

You need to create an AI agent in Microsoft Foundry that follows a specific role and behavior when responding to users. What should you configure?

- A. tokens per minute (TPM)
- B. system instructions
- C. temperature
- D. max completion tokens

Answer: B

Explanation:

To create an AI agent that follows a specific role and behavior, you configuresystem instructions. Microsoft Foundry Agent Service documentation states that agentinstructions define goals, constraints, and behavior.

OptionA. tokens per minute (TPM)controls throughput quota, not behavior.

OptionC. temperaturecontrols response randomness/creativity, not the agent??s role.

OptionD. max completion tokenscontrols response length, not the agent??s role or behavioral rules.

Therefore, the correct answer isB. system instructions.

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
Fairness can be achieved by focusing solely on improving the overall accuracy of an AI model.	☐	☐
Evaluating AI system outputs to identify and reduce bias across demographic groups supports fairness.	☐	☐
Ensuring fairness for an AI system means that all users always receive the same output from the system.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Fairness can be achieved by focusing solely on improving the overall accuracy of an AI model.	☐	☒
Evaluating AI system outputs to identify and reduce bias across demographic groups supports fairness.	☒	☐
Ensuring fairness for an AI system means that all users always receive the same output from the system.	☐	☒

NEW QUESTION 7

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 8

You have a Microsoft Foundry project that contains a vision-enabled model deployment. You use the Azure OpenAI Responses API to send a prompt to the model. You need to provide an image for analysis. Which content item should you include in the request?

- A. image_base64
- B. image_generation
- C. output_image
- D. input_image

Answer: D

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

endpoint of the model

evaluation pipeline of model

Foundry playground

training dataset of the model

NEW QUESTION 10

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

NEW QUESTION 10

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
In the new Microsoft Foundry portal, you must fine-tune a model before you can deploy the model.	☐	☐
In the new Microsoft Foundry portal, you can test a model from the model catalog only after you deploy the model.	☐	☐
In the new Microsoft Foundry portal, you can deploy a model from the model catalog only after retraining the model.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Statement 1:In the new Microsoft Foundry portal, you must fine-tune a model before you can deploy the model.=No
Fine-tuning is optional. Microsoft's Foundry model deployment documentation describes deploying Foundry Models directly from the model catalog for inference. It does not require fine-tuning first.

Statement 2:In the new Microsoft Foundry portal, you can test a model from the model catalog only after you deploy the model.=No
Microsoft documentation states thatsome Foundry Tools are available to try via the model catalog without a project, and Foundry playgrounds are used for prototyping and validation before production. Therefore, the statement using ??only after you deploy?? is too restrictive.

Statement 3:In the new Microsoft Foundry portal, you can deploy a model from the model catalog only after retraining the model.=No

Retraining/fine-tuning is not required before deployment. Microsoft states that after you deploy a Foundry Model, you can interact with it in the Foundry Playground and use it from code, and the deployment workflow starts by selecting a model from the model catalog and choosing Deploy.

NEW QUESTION 13

HOTSPOT

Select the answer that correctly completes the sentence.

Answer Area

Evaluating model outcomes across demographic groups to reduce bias

is an example of the Microsoft responsible AI principle of

accountability

fairness

privacy and security

transparency

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Evaluating model outcomes across demographic groups to reduce bias

is an example of the Microsoft responsible AI principle of

accountability

fairness

privacy and security

transparency

NEW QUESTION 16

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



NEW QUESTION 19

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.

NEW QUESTION 24

Based on the image provided, here is the transcribed text:

You need to build an AI solution that produces new product images based on written descriptions provided by users. Which AI workload should you use?

- A. image generation
- B. image analysis
- C. object detection
- D. optical character recognition (OCR)

Answer: A

NEW QUESTION 29

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 34

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

Voice Live returns only transcribed text.

Voice Live requires you to separately implement speech to text and text to speech services.

Voice Live combines speech to text, reasoning, and text to speech into a single conversational experience.

Yes

No

☐

☐

☐

☐

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements

Voice Live returns only transcribed text.

Voice Live requires you to separately implement speech to text and text to speech services.

Voice Live combines speech to text, reasoning, and text to speech into a single conversational experience.

Yes

No

☐

☒

☐

☒

NEW QUESTION 38

You need to compare the costs of large language models (LLMs) for a generative AI solution. What should you use in the Microsoft Foundry portal?

- A. Evaluator catalog
- B. Model leaderboard
- C. Compliance
- D. Tools

Answer: B

Explanation:

To compare the costs of large language models in Microsoft Foundry portal, use theModel leaderboard. Microsoft documentation states that themodel leaderboardhelps compare models acrossquality, safety, estimated cost, and throughput. It also supports trade-off charts and side-by-side model comparison for features, performance, and estimated cost. Why the other options are incorrect:
* A. Evaluator catalogis for selecting evaluators to measure model or application outputs, not comparing LLM costs.
* C. Compliancerelates to governance and compliance, not model cost comparison.
* D. Toolsprovides Foundry tools, not benchmarked cost comparison across models.

NEW QUESTION 41

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 46

Your company has thousands of recorded customer support calls in multiple languages stored as audio files in Azure Storage. You need to generate text transcripts of all the recordings. Which Azure Speech in Foundry Tools capability should you use?

- A. speech to text batch transcription
- B. speech to text real-time transcription
- C. text to speech
- D. speech translation

Answer: A

NEW QUESTION 49

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
Generating a response to a user prompt occurs during the inference stage.	☐	☐
A generative AI model generates responses by copying stored documents directly from the model's training data.	☐	☐
A generative AI model produces output by predicting the next token based on patterns learned from the model's training data.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Generating a response to a user prompt occurs during the inference stage.	☒	☐
A generative AI model generates responses by copying stored documents directly from the model's training data.	☐	☒
A generative AI model produces output by predicting the next token based on patterns learned from the model's training data.	☒	☐

NEW QUESTION 52

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AI-901 Exam with Our Prep Materials Via below:

<https://www.certleader.com/AI-901-dumps.html>