



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

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. AudioOutputConfig
- B. SpeechSynthesizer
- C. AudioConfig
- D. SpeechConfig

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

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

Currently there are no comments in this discussion, be the first to comment!

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

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Naesh** 1 week, 2 days ago

Selected Answer: A

SYSTEM INSTRUCTIONS allow you to give high level instructions to the model that governs the behaviour for all the responses.
upvoted 1 times

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.	☐	☐

 **Shifin** 2 days, 6 hours ago

System prompts can be used to authorize users - No (Authorization is handled through authentication and access control mechanisms, not system prompts.)

A system prompt is used to reduce tokens per minute (TPM) - No (TPM is a service quota/rate-limit metric. System prompts do not exist to reduce TPM usage.)

A system prompt guides the behavior of a generative AI model - Yes (System prompts provide instructions that define the model's role, tone, constraints, and behavior, guiding how it responds to users.)

upvoted 1 times

 **Avbud** 1 week, 2 days ago

1. No -> system prompts can't authorize users, either Entra ID or api key does

2. No -> TPM is configured in deployment settings not via system prompts!

3. Yes

upvoted 1 times

 **RRandell** 1 month ago

1. Yes

2. No

3. No

A system prompt guides model behaviour, improves response quality, and can reduce harmful outputs. It's core purpose isn't managing TPM

upvoted 3 times

 **MichaelinX** 1 month, 1 week ago

2. A system prompt is used to reduce tokens per minute (TPM).

-> No. I think the answer is No?

TPM / rate limit / throughput / quota / capacity

upvoted 1 times

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

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

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
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.	☐	☐

Currently there are no comments in this discussion, be the first to comment!

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()

```
...
```

Currently there are no comments in this discussion, be the first to comment!

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

Suggested Answer: *AE*

Currently there are no comments in this discussion, be the first to comment!


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.

Answer Area

```
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
)
...
```

Currently there are no comments in this discussion, be the first to comment!

HOTSPOT -

Select the answer that correctly completes the sentence.

Answer Area

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

- deployment slot
- embedding index
- fine-tuning job
- system prompt

Currently there are no comments in this discussion, be the first to comment!

HOTSPOT -

You are developing an application that analyze invoices by using Azure Content Understanding in Foundry Tools.

You need to ensure that the application retrieves the analysis results after processing completes.

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

```
poller = client.begin_analyze(  
    analyzer_id="invoice",  
    input_url=url  
)  
result = poller.
```

get_results

result

status

wait

Currently there are no comments in this discussion, be the first to comment!

DRAG DROP -

You have a Microsoft Foundry project named project1 that contains an Azure OpenAI resource named Resource1.

To Resource1, you deploy a gpt-4.1-mini model by using a model deployment named my-mini-gpt.

You need to connect to my-mini-gpt from an application.

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.

NOTE: Each correct selection is worth one point.

Values

:: gpt-4.1-mini

:: my-mini-gpt

:: project1

:: resource1

Answer Area

```
client = OpenAI(api_key="...",  
base_url="https:// [ ] .openai.azure.com/openai/v1/",)  
response = client.responses.create(  
    model="[ ]",  
    ...  
)
```

Currently there are no comments in this discussion, be the first to comment!

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
An AI generative model is retrained before performing each user request.	☐	☐
An AI agent responds by copying and pasting answers stored in a database.	☐	☐
An AI agent uses a generative AI model to establish actions based on user input.	☐	☐

Currently there are no comments in this discussion, be the first to comment!

DRAG DROP -

You are reviewing best practices for using AI at your company.

Which Microsoft responsible AI principle is each task an example of? To answer, drag the appropriate principles to the correct tasks. Each task 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 match is worth one point.

Principles☐ Accountability☐ Fairness☐ Inclusiveness☐ Privacy and security☐ Reliability and safety☐ Transparency**Answer Area**

Evaluating model outputs to ensure that decisions are **NOT** biased against specific demographic groups:

Encrypting sensitive customer data and restricting system access to authorized personnel:

Informing users when they are interacting with an AI system and explaining the system's capabilities and limitations:

Testing AI systems under different conditions to reduce unexpected failures:

Currently there are no comments in this discussion, be the first to comment!

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.	☐	☐

Currently there are no comments in this discussion, be the first to comment!

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

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

HOTSPOT -

Select the answer that correctly completes the sentence.

Answer Area

Information extraction solutions that detect and read text
in scanned documents and images rely on

- computer vision
- image generation
- sentiment analysis
- speech synthesis

Currently there are no comments in this discussion, be the first to comment!

You have a Microsoft Foundry project that has a generative AI model deployment.

You need to ensure that responses generated by the model minimize costs and remain within a defined length.

Which parameter should you configure?

- A. Top P
- B. Temperature
- C. Max Completion Tokens
- D. Model version settings

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

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
Human-in-the-loop practices provide accountability for AI-generated decisions.	☐	☐
Deploying an AI system to a production environment eliminates the need for ongoing monitoring.	☐	☐
Disclosing the team that designed and deployed an AI system provides accountability for the system's output.	☐	☐

Currently there are no comments in this discussion, be the first to comment!

Your company processes customer support emails.

You need to implement an AI solution that automatically identifies mentions of people, organizations, and locations in the emails.

Which text analysis technique should you use?

- A. Named Entity Recognition (NER)
- B. key phrase extraction
- C. sentiment analysis
- D. summarization

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

HOTSPOT -

Select the answer that correctly completes the sentence.

Answer Area

Ensuring that human reviewers oversee AI-generated decisions
and remain responsible for the final output
is an example of the Microsoft responsible AI principle of

- accountability
- fairness
- privacy and security
- transparency

Currently there are no comments in this discussion, be the first to comment!

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

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

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.	☐	☐

  **Avbud** 1 week, 2 days ago

1. NO! -> temperature is a per-request inference parameter! you pass it in each API call (or set it in the playground per session)
upvoted 1 times

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

Currently there are no comments in this discussion, be the first to comment!

HOTSPOT -

Select the answer that correctly completes the sentence.

Answer Area

The Microsoft responsible AI principle of transparency requires that AI systems

- be explainable to users
- protect sensitive user data
- reduce bias in decisions
- require human oversight

Currently there are no comments in this discussion, be the first to comment!

You are developing an AI-powered customer support application.

Which task is an example of the Microsoft responsible AI principle of inclusiveness?

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

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!