



VirtuLearn
Trusted, Accurate & Reliable
vp

TRUSTED, ACCURATE AND RELIABLE

The most comprehensive IT certification
preparation materials in the industry!



✓ Expert-Verified

✓ Exam-Ready

✓ Regularly Updated

All rights reserved. No part of this document may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law. Unauthorized copying, reselling, or distribution of this document is strictly prohibited and may result in legal action.

<https://www.virtulearner.com>
support@virtulearner.com



Microsoft

AI-901

Microsoft Azure AI

Fundamentals

Sections:

1. Identify AI Concepts and Responsibilities
2. Implement AI Solutions Using Microsoft Foundry

Exam Topic: Identify AI Concepts and Responsibilities**QUESTION: 1**

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`

Answer(s): D

Explanation:

To configure the connection to the Azure Speech service, you should use the `azure.cognitiveservices.speech.SpeechConfig` class.

Implementation Details Initialization: This class holds authentication and endpoint details, such as your Azure resource key and region.

Usage: You pass this configuration object into your real-time processing classes, like `SpeechRecognizer` (for speech-to-text) or `SpeechSynthesizer` (for text-to-speech).

Reference:

<https://docs.azure.cn/en-us/ai-services/speech-service/how-to-recognize-speech>

QUESTION: 2

HOTSPOT (Drag and Drop is not supported)

Select the answer that correctly completes the sentence.

Hot Area:**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. See Explanation section for answer.

Answer(s): 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	

Box: a base64-encoded image data When using the OpenAI Responses API and a vision-enabled model, you can include an image in a request by providing the image as _____.

When using the OpenAI Responses API with a vision-enabled model, you can include an image in a request by providing it in three different formats: A fully qualified image URL: Passing the web link where the image is hosted online. *-> A Base64-encoded data URL: Encoding a local image file directly into a text string (formatted as data:image/jpeg;base64,...). A file ID: Uploading the image via the Files API first and referencing its unique ID.

Reference:

<https://developers.openai.com/api/docs/guides/image-generation>

QUESTION: 3

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

Answer(s): A

Explanation:

In Microsoft Foundry, instructions (formerly referred to as system messages) are the primary way to define the role, behavior, tone, and constraints of your AI agent.

Reference:

<https://learn.microsoft.com/en-us/azure/foundry-classic/agents/quickstart>

QUESTION: 4

HOTSPOT (Drag and Drop is not supported)

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

Note: Each correct selection is worth one point.

Hot Area: