

# Microsoft AI-901

Microsoft Azure AI Fundamentals (Updated Version)

For More Information – Visit link below:

<https://www.examsempire.com/>

**Product Version**

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/ai-901>

# Latest Version: 6.0

**Answer: C**

You have a Microsoft Foundry project that contains an agent named Agent1. You need to ensure that Agent1 always calls an Azure function when the agent responds to user input. To what should you set tool\_choice for Agent1?

- A. auto
- B. none
- C. required

**Answer: C**

Explanation:

Microsoft's Foundry Agent Service documentation states that tool\_choice provides deterministic control over tool calling:

auto means the model decides whether to call tools.

required means the model must call one or more tools.

none means the model does not call tools.

Therefore:

A . auto = Incorrect, because the model may or may not call the Azure function.

B . none = Incorrect, because this prevents tool/function calls.

C . required = Correct, because it forces the agent to call a tool.

The Azure OpenAI function-calling documentation also confirms that tool\_choice="auto" lets the model decide whether to call a function, while tool\_choice="none" forces a user-facing response without a tool call.

## Question: 2

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



Explanation:

### Question: 3

DRAGDROP

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:// [project1] .openai.azure.com/openai/v1/",)
response = client.responses.create(
model=" [my-mini-gpt] ",
...
)
```

**Answer:**

Values

- gpt-4.1-mini
- my-mini-gpt
- project1
- resource1

Answer Area

```
client = OpenAI(api_key="...",
base_url="https:// [project1] .openai.azure.com/openai/v1/",)
response = client.responses.create(
model=" [my-mini-gpt] ",
...
)
```

Explanation:

### Question: 4

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: A, B**

Explanation:

Microsoft Foundry Agent Service documentation states that instructions define goals, constraints, and behavior for an agent. Therefore, instructions are used to guide how the generative AI model or agent should respond and behave.

Option A is correct because instructions can define constraints the model must follow.

Option B is correct because instructions can define the agent's role and behavior.

Options C, D, and E are incorrect because Azure region, model selection, and TPM allocation are configuration or deployment/resource settings, not purposes of prompt instructions.

## Question: 5

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**

Explanation:

Voicemail recordings are audio content. Azure Content Understanding analyzers define what type of content to process, including documents, images, audio, or video, and what elements to extract, including transcripts and structured fields.

Microsoft's custom analyzer documentation also shows an audio example based on prebuilt-audio for processing customer support call recordings, which is the same content type as voicemail recordings.

Therefore, to extract a transcript and structured information from voicemail recordings, you should use an audio analyzer.

**Thank You for Trying Our Product**  
**Special 16 USD Discount Coupon: NSZUBG3X**  
**Email: [support@examsempire.com](mailto:support@examsempire.com)**

**Check our Customer Testimonials and ratings  
available on every product page.**

**Visit our website.**

**<https://examsempire.com/>**