

Microsoft

Exam Questions AI-103

Developing AI Apps and Agents on Azure



NEW QUESTION 1

- (Topic 1)
You need to configure Agent1 to answer customer questions about only the Contoso products. The solution must meet the business requirements.
What should you do?

- A. Apply top-p sampling.
- B. Modify the system message instructions.
- C. Add few-shot examples.
- D. Increase the value of the temperature parameter.

Answer: B

NEW QUESTION 2

- (Topic 2)
You have an Azure subscription that contains a Microsoft Foundry resource. You need to build an app that will suggest product names from a given product description. Which Foundry model should you use?

- A. GPT-4
- B. DALL-E
- C. Whisper
- D. embeddings

Answer: A

NEW QUESTION 3

- (Topic 2)
You are building an app that will use Azure AI to monitor workspaces for safety. You need to recommend a service that meets the following requirements:
? Generates alerts when employees enter high-risk areas
? Monitors video feeds in real time
? Minimizes development effort
What should you recommend?

- A. Azure Vision in Foundry Tools Image Analysis
- B. Azure AI Video Indexer
- C. Azure Vision in Foundry Tools Spatial Analysis
- D. Object detection in Azure Custom Vision

Answer: C

NEW QUESTION 4

DRAG DROP - (Topic 2)
You have a Microsoft Foundry project that contains an agent used by the financial analysts at your company.
You need to optimize the agent workflow by providing additional data access and processing capabilities. The solution must meet the following requirements:
• Ensure that the agent can perform calculations during conversations
• Ensure that the agent can access up-to-date information from public websites.
• Ensure that the agent can retrieve information from documents uploaded directly to the agent.
What should you use for each requirement? To answer, drag the appropriate tools to the correct requirements. Each tool 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.

Tools

⋮ Code interpreter

⋮ Computer use

⋮ File search

⋮ Grounding with Bing Search

⋮ Microsoft Fabric

Answer Area

Access up-to-date information from public websites:

Tool

Perform calculations during conversations:

Tool

Retrieve information from documents uploaded directly to the agent:

Tool

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Tools

Code interpreter

Computer use

File search

Grounding with Bing Search

Microsoft Fabric

Answer Area

Access up-to-date information from public websites

Grounding with Bing Search

Perform calculations during conversations

Code interpreter

Retrieve information from documents uploaded directly to the agent

File search

NEW QUESTION 5

- (Topic 2)

You have a Microsoft Foundry project that contains a prompt agent used by a customer support web app.

The agent is invoked from a Python service that does NOT run in the Foundry portal.

You need to implement end-to-end tracing to capture latency breakdowns and exceptions across agent runs.

Which two components can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. a Log Analytics workspace
- B. OpenTelemetry
- C. Application Insights
- D. the Azure Monitor Agent
- E. Microsoft Sentinel

Answer: BC

NEW QUESTION 6

- (Topic 2)

You have a Microsoft Foundry project that generates product marketing images from text prompts.

After publishing several images, the legal team at your company identifies a competitor's logo on a sign in the background of an image.

You need to remove only the logo, while preserving the rest of the image. What should you do?

- A. Increase the prompt guidance strength.
- B. Modify the original prompt to exclude brand names.
- C. Apply a mask-based inpainting edit to the part of the image that contains the logo.
- D. Rerun the prompt by using a different random seed.

Answer: C

NEW QUESTION 7

- (Topic 2)

You have a chat app in a Microsoft Foundry project and an Azure AI Search vectorized index.

You need to connect to the index to meet the following requirements:

- Complex questions must retrieve information from multiple chunks.
- Multi-turn conversations must influence retrieval planning.
- Retrievals must run in parallel to reduce latency. Which retrieval approach should you use?

- A. classic Retrieval Augmented Generation (RAG)
- B. chain of thought
- C. agentic Retrieval Augmented Generation (RAG)
- D. iterative retrieval

Answer: C

NEW QUESTION 8

DRAG DROP - (Topic 2)

You have a Microsoft Foundry project that contains an agent.

You need to enable long-term memory to ensure that the agent can recall user preferences across separate conversations. Stored memories must be isolated per authenticated user without the client application manually generating user IDs.

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

⋮ "session"

⋮ "{{conversationId}}"

⋮ "{{userId}}"

⋮ [mem_store_name]

⋮ [memory_tool]

⋮ MemorySearchTool("support_mem_store")

Answer Area

```
from azure.ai.projects.models import MemorySearchTool, PromptAgentDefinition

mem_store_name = "agent_mem_store"

memory_tool = MemorySearchTool(
    memory_store_name=mem_store_name,
    scope= Value,
)

agent_def = PromptAgentDefinition(
    model="gpt-5.2",
    instructions="You are a customer support assistant.",
    tools= Value
)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Values

⋮ "session"

⋮ "{{conversationId}}"

⋮ "{{userId}}"

⋮ [mem_store_name]

⋮ [memory_tool]

⋮ MemorySearchTool("support_mem_store")

Answer Area

```
from azure.ai.projects.models import MemorySearchTool, PromptAgentDefinition

mem_store_name = "agent_mem_store"

memory_tool = MemorySearchTool(
    memory_store_name=mem_store_name,
    scope= ⋮ "{{userId}}",
)

agent_def = PromptAgentDefinition(
    model="gpt-5.2",
    instructions="You are a customer support assistant.",
    tools= ⋮ [memory_tool]
)
```

NEW QUESTION 9

HOTSPOT - (Topic 2)

You have an Azure AI Search resource named Search1 that is used by multiple apps hosted in Azure. You need to secure Search1. The solution must meet the following requirements:

- Prevent access to Search1 from the internet.
- Limit the access of each app to query specific indexes.

What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct answer is worth one point.

Answer Area

To prevent access from the internet:

Configure an IP firewall.
Create a private endpoint
Use Azure roles.

To limit access to query specific indexes:

Create a private endpoint.
Use Azure roles.
User key authentication

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 10

- (Topic 2)

You have a Microsoft Foundry project that contains an agent. The agent uses two tools to perform the following actions:

? Use Azure AI Search to retrieve answers from a private product documentation index.

? Use the web search tool to retrieve public information on the internet.

You need to ensure that for a specific run, the agent deterministically retrieves information only from the internet. To what should you set tool_choice?

- A. "required"
- B. {"type": "azure_ai_search"}
- C. "auto"
- D. {"type": "bing_grounding"}

Answer: D

NEW QUESTION 10

- (Topic 2)

You are creating an image-editing workflow in a Microsoft Foundry project. The workflow must meet the following requirements:

- Ensure that background objects can be removed by applying a mask-based inpainting edit.
- Preserve the original lighting and style of the edited images.
- Use the built-in image editing controls, NOT a custom model.

You need to ensure that image edits apply exclusively inside the masked area.

How should you configure the workflow?

- A. Enable text_to_image mode and a prompt describing the desired background removal.
- B. Set generation mode to image_variation and provide the original image as a reference.
- C. Enable image_to_image mode and a high-strength value to regenerate the full image based on the prompt.
- D. Enable mask_inpainting and supply both the input image and a mask indicating which part of the image to modify.

Answer: D

NEW QUESTION 12

- (Topic 2)

You have a Microsoft Foundry project that contains an agent.

You need to process mixed-format documents that contain scanned text, tables, and multicolumn layouts. The extracted content must preserve the document structure and be converted into the Markdown format for downstream reasoning.

What should you configure first?

- A. an Azure Content Understanding in Foundry Tools analyzer
- B. a generative chat completion request
- C. an Azure Language in Foundry Tools text analysis model deployment
- D. an Azure OpenAI Responses API call that uses a multimodal model

Answer: A

NEW QUESTION 17

- (Topic 2)

You have a Microsoft Foundry project that contains an agent.

The agent ingests scanned PDF vendor invoices that contain tables and embedded QR codes.

The agent must preserve the PDF layout in the extracted output to ensure that downstream processing can reference sections and tables.

You plan to call Azure Content Understanding in Foundry Tools.

You need to extract content and layout elements and detect QR codes without requiring a language model deployment. Which built-in analyzer should you use?

- A. prebuilt-layout
- B. prebuilt-documentFieldSchema
- C. prebuilt-read
- D. prebuilt-documentSearch

Answer: A

NEW QUESTION 19

HOTSPOT - (Topic 2)

You need to recommend a plan to create a customer support agent by using the Microsoft Foundry Agent Service. The agent must meet the following requirements:

- Retain user preferences across multiple conversations.
- Enable users to provide contextual grounding by directly uploading documents during a chat.

Which Foundry capability should you recommend for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

To retain user preferences across conversations, use:

Agent memory that uses persistent storage

Conversation history

Orchestration-managed session context

To enable users to provide contextual grounding during chats, use the:

Azure AI Search tool

Code interpreter tool

File search tool

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

To retain user preferences across conversations, use:

Agent memory that uses persistent storage

Conversation history

Orchestration-managed session context

To enable users to provide contextual grounding during chats, use the:

Azure AI Search tool

Code interpreter tool

File search tool

NEW QUESTION 22

- (Topic 2)

Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals.

More than one solution in the set might solve the problem. It is also possible that none of the solutions in the set solve the problem. After you answer a question in this section, you will NOT be able to return. As a result, these questions do not appear on the Review Screen.

You have a multimodal AI generative model that accepts image uploads and uses extracted image text to generate responses.

You discover that users can upload unsafe images and embed hidden instructions into images to manipulate the model.

You need to implement controls to mitigate the risk. Solution: You configure a prompt shield for documents. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 24

- (Topic 2)

You have a Microsoft Foundry project named Project1.

Project1 contains an application that processes PDF vendor invoices.

You need to configure Azure Document Intelligence in Foundry Tools to generate a Markdown output that preserves the sections and table structure of the PDFs. The solution must minimize development effort. What should you do?

- A. Increase the confidence threshold.
- B. Configure output=figures when you analyze the PDF.
- C. Configure content=markdown when you analyze the document.
- D. Set the output_content_format=ContentForma
- E. MARKDOWN value.

Answer: D

NEW QUESTION 26

HOTSPOT - (Topic 2)


You have a configurable guardrail in a Microsoft Foundry project.

You have a Python application. A text blocklist named ConfidentialTerms already contains terms that must be detected in user messages. The application stores each incoming user message in a variable named input_text.

You plan to moderate input_text before the text is sent to a language model. You need to analyze input_text by using the existing blocklist.

Answer Area

```
...
blocklist_name = "ConfidentialTerms"
client = ContentSafetyClient(endpoint, credential)
request =
    text=input_text
    ...
```




- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
...
blocklist_name = "ConfidentialTerms"
client = ContentSafetyClient(endpoint, credential)
request =
    text=input_text
    ...
```



NEW QUESTION 30

- (Topic 2)

You are building a speech processing solution in Microsoft Foundry for a customer support platform.

The platform will transcribe live phone calls, so that supervisors at your company can view call transcripts and detect issues while the calls are in progress. The call audio will arrive as a continuous stream from the telephony system.

You need to ensure that the call transcripts appear within only a few seconds of the audio stream.

What should you do?

- A. Run a batch transcription job on recorded audio files.
- B. Use real-time speech to text to process streaming audio input.
- C. Use speech translation to generate the transcripts into multiple languages.
- D. Use text to speech by using a custom neural voice.

Answer: B

NEW QUESTION 33

- (Topic 2)

You have a Microsoft Foundry project named Project1 that contains the following:

- An OpenAPI tool that calls an external API
 - A project connection named Connection1 that stores the API key of the external API When an agent calls the OpenAPI tool, the API returns a 401 unauthorized error, and traces show that the API key header is NOT being sent.
- You need to ensure that the OpenAPI tool automatically includes the API key from Connection1 on all requests. What should you do?
- A. Configure the tool to use the default connection of Project1.
 - B. Connect the tool to Connection1.
 - C. Enable identity passthrough so that the tool uses the Microsoft Entra token of the caller.
 - D. Add the API key header manually to the OpenAPI specification.

Answer: B

NEW QUESTION 37

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an agent. The agent uses tools to retrieve internal content and call external APIs. The agent is configured to let the model decide when to call the tools. You need to publish the agent for a compliance workflow. The solution must meet the following requirements:

- Each workflow run must include a retrieval step before generating a response.
- Tool calls must authenticate by using the published agent??s own identity.
- Tool access must use an identity isolated from other project resources.
- Tool access must support audit tracing.

What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Set tool_choice to:

auto

none

required

Configure the tool to authenticate by:

Storing API keys in prompts

Using the shared project agent identity

Using a distinct agent identity bound to the client application

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Set tool_choice to:

auto

none

required

Configure the tool to authenticate by:

Storing API keys in prompts

Using the shared project agent identity

Using a distinct agent identity bound to the client application

NEW QUESTION 42

HOTSPOT - (Topic 2)

You plan to deploy a containerized version of an Azure AI Language service that will be used for sentiment analysis. You configure `https://contoso.cognitiveservices.azure.com` as the endpoint URI for the service. You need to run the container on an Azure virtual machine by using Docker. How should you complete the command? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

`docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \`

http://contoso.blob.core.windows.net
https://contoso.cognitiveservices.azure.com
mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
mcr.microsoft.com/azure-cognitive-service/textanalytics/sentiment

Eula=accept \

Billing=

ApiKey=xxx

http://contoso.blob.core.windows.net
https://contoso.cognitiveservices.azure.com
mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

`docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \`

http://contoso.blob.core.windows.net
https://contoso.cognitiveservices.azure.com
mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
mcr.microsoft.com/azure-cognitive-service/textanalytics/sentiment

Eula=accept \

Billing=

ApiKey=xxx

http://contoso.blob.core.windows.net
https://contoso.cognitiveservices.azure.com
mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment

NEW QUESTION 44

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an agent.

The agent accepts user-uploaded screenshots and uses a multimodal chat model.

Some screenshots contain potentially malicious embedded text.

You need to prevent a prompt injection attack and ensure that third-party content is treated as lower trust.

How should you configure prompt shields for document attacks? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Prompt shields action:

Disable the shield.

Set action to block.

Set action to annotate.

Additional mitigation:

Enable Spotlighting.

Create a custom blocklist.

Use optical character recognition (OCR) to extract the text from the images first.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Prompt shields action:

Disable the shield.

Set action to block.

Set action to annotate.

Additional mitigation:

Enable Spotlighting.

Create a custom blocklist.

Use optical character recognition (OCR) to extract the text from the images first.

NEW QUESTION 45
.....

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-103 Practice Exam Features:

- * AI-103 Questions and Answers Updated Frequently
- * AI-103 Practice Questions Verified by Expert Senior Certified Staff
- * AI-103 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AI-103 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AI-103 Practice Test Here](#)