



Microsoft

Exam Questions AI-103

Developing AI Apps and Agents on Azure

NEW QUESTION 1

- (Topic 1)

You need to recommend a solution to support the planned changes and technical requirements for Agent1 to use the product information stored in storage1.

What should you include in the recommendation?

- A. Azure AI Search
- B. Azure Translator in Foundry Tools
- C. Azure Document Intelligence in Foundry Tools
- D. Grounding with Bing Search

Answer: A

NEW QUESTION 2

- (Topic 1)

You need to configure Agent1 to meet the security and compliance requirements. What should you use?

- A. prompt shields
- B. Personally Identifiable Information (PII) Detection
- C. self-harm content filtering
- D. violence content filtering

Answer: B

NEW QUESTION 3

- (Topic 1)

You need to recommend a solution to assess the responses generated by Agent1 when the agent uses the product information stored in storage1. The solution must meet the technical requirements.

- A. a Retrieval Augmented Generation (RAG) evaluator
- B. a coherence evaluator
- C. a custom guardrail
- D. model fine-tuning

Answer: A

NEW QUESTION 4

- (Topic 1)

You need to recommend an invoice review solution that resolves the issue reported by the finance department.

What should you include in the recommendation?

- A. Azure Content Understanding in Foundry Tools
- B. chat completions
- C. Azure Document Intelligence in Foundry Tools
- D. Image Analysis

Answer: A

NEW QUESTION 5

- (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 6

- (Topic 1)

You need to configure personalized user interactions for Agent1 based on the business requirements. What should you include in the solution?

- A. guardrails
- B. memory
- C. tools
- D. instructions

Answer: B

NEW QUESTION 7

HOTSPOT - (Topic 1)

You need to ensure that the marketing department can generate videos by using the model deployed to Project2. now should you complete the Python code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
import time
from openai import OpenAI
client = OpenAI(
    base_url="https://Contoso.openai.azure.com/openai/v1/",
    api_key=...,
)
video = client.videos.create(
    model=deployment_name,
    prompt="A video of our product"
)
while video.status not in ["completed", "failed", "cancelled"]:
    time.sleep(20)
    video = client.videos.retrieve(video.id)
print(video.status)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
import time
from openai import OpenAI
client = OpenAI(
    base_url="https://Contoso.openai.azure.com/openai/v1/",
    api_key=...,
)
video = client.videos.create(
    model=deployment_name,
    prompt="A video of our product"
)
while video.status not in ["completed", "failed", "cancelled"]:
    time.sleep(20)
    video = client.videos.retrieve(video.id)
print(video.status)
```

NEW QUESTION 8

HOTSPOT - (Topic 1)

You need to ensure that Agent1Dev Team can access Agent1. The solution must meet the security and compliance requirements.
How should you complete the Python code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

from azure.core.credentials import AzureKeyCredential

myEndpoint = "https://contoso.services.ai.azure.com/api/projects/project1"

project_client = AIProjectClient(

    endpoint=myEndpoint,

    credential=  ,

    )

myAgent = "Agent1"

agent = project_client.agents.  (agent_name=myAgent)

print(f"Retrieved agent: {agent.name}")
```

create_version

get

get_version

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

from azure.core.credentials import AzureKeyCredential

myEndpoint = "https://contoso.services.ai.azure.com/api/projects/project1"

project_client = AIProjectClient(
    endpoint=myEndpoint,
    credential=  ,
    

AzureKeyCredential()
        DefaultAzureCredential()
        None


)

myAgent = "Agent1"

agent = project_client.agents.  (agent_name=myAgent)

print(f"Retrieved agent: {agent.name}")


create_version
    get
    get_version


```

NEW QUESTION 9

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an internal Q&A agent. Users report the following issues when they ask the agent questions:

- An increase in the following response: ??No relevant information found??
- Periodic HTTP 429 rate limit exceeded errors during peak hours

You need to identify whether each issue is caused by model unavailability, resource limits, or inference failures.

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

Metrics to enable:

Model Availability Rate and Provisioned Utilization
 Only Tokens Cache Match Rate
 Only Total Requests filtered to status code 200
 Time To Response and Total Tokens

Diagnostic log to collect:

AllMetrics
 audit
 RequestResponse
 trace

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Metrics to enable:

- Model Availability Rate and Provisioned Utilization
- Only Tokens Cache Match Rate
- Only Total Requests filtered to status code 200
- Time To Response and Total Tokens

Diagnostic log to collect:

- AllMetrics
- audit
- RequestResponse
- trace

NEW QUESTION 10

- (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 10

- (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 14

- (Topic 2)

You have a Microsoft Foundry project.

You plan to build a customer support solution that contains an agent. The solution must meet the following requirements:

- Provide accurate, context-aware responses grounded in internal product documentation stored in Azure AI Search.
- Require deep, multi-step reasoning across long contexts.
- Generate detailed natural language responses.

Which type of model should you use to power the agent?

- A. a multimodal model
- B. a key phrase extraction model
- C. a small language model (SLM)
- D. a large language model (LLM)

Answer: D

NEW QUESTION 18

- (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 Microsoft Foundry project that contains an agent. The agent generates summaries from retrieved policy documents.

Users report that some responses omit required regulatory clauses, even when the clauses are present in the retrieved content.

You need to improve response completeness.

Solution: You add a reflection pass that regenerates the response if the required clauses are missing.

Does this meet the goal?

A. Yes

B. No

Answer: A

NEW QUESTION 23

- (Topic 2)

You have a Microsoft Foundry project that contains a high-traffic agent. After a recent update, operational costs increase significantly.

Monitoring confirms that the volume of user traffic to the agent remains unchanged.

You suspect that changes to the request or response characteristics are causing the increase.

You need to identify whether the additional costs are driven by the model input size, the model output size, or expanded tool usage.

Which observability capability should you use?

A. evaluation metrics

B. token usage

C. latency

D. run success rate

Answer: B

NEW QUESTION 25

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an agent. You use a GitHub Actions workflow for CI/CD.

You need to configure the workflow to automatically evaluate the agent when a pull request (PR) is created and prevent branches from merging if the evaluation results do NOT meet the defined thresholds.

How should you configure the workflow? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Authentication method:

- A personal access token (PAT)
- A user-assigned managed identity
- An Azure Login action that uses OpenID Connect (OIDC)

If the evaluation results are NOT met, configure the workflow to:

- Fail
- Lock the target branch
- Send an alert

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Authentication method:

- A personal access token (PAT)
- A user-assigned managed identity
- An Azure Login action that uses OpenID Connect (OIDC)

If the evaluation results are NOT met, configure the workflow to:

- Fail
- Lock the target branch
- Send an alert

NEW QUESTION 27

- (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 32

- (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 37

- (Topic 2)

You have an app named App1 that uses a Microsoft Foundry multimodal model deployment.

App1 runs optical character recognition (OCR) on uploaded images and appends the OCR output to the prompt as additional context.

Some uploaded images contain embedded text.

You need to prevent potentially malicious instructions from being processed by the model. What should you use?

- A. protected material text
- B. prompt shields for user prompts
- C. image moderation
- D. prompt shields for documents

Answer: D

NEW QUESTION 40

- (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 Microsoft Foundry project that contains an agent. The agent generates summaries from retrieved policy documents.

Users report that some responses omit required regulatory clauses, even when the clauses are present in the retrieved content.

You need to improve response completeness.

Solution: You increase the value of the max_tokens parameter. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 45

- (Topic 2)

You have a Microsoft Foundry project that contains an agent. The agent uses Azure Speech in Foundry Tools.

You fine-tune a baseline speech to text model for the en-us locale and publish the model. The agent calls the Speech to text REST API and returns an error message indicating that the project ID is invalid.

You need to set the project property to the correct ID. To what should you set the project property?

- A. the custom speech endpoint URL
- B. the project URL
- C. the project ID
- D. the custom speech project ID

Answer: D

NEW QUESTION 46

- (Topic 2)

You are planning a Microsoft Foundry project named Project1 that will contain multiple agents. Each agent will access the same Azure AI Search resource. You need to recommend a solution to centrally manage the Azure AI Search credentials within Project1. The solution must be implemented across all the agents. What should you recommend?

- A. Enable role-based access control (RBAC) for the Azure AI Search resource.
- B. Add a connection to the Azure AI Search resource.
- C. Disable key-based access control on the Azure AI Search resource.
- D. Create a managed private endpoint that connects to the Azure AI Search resource.

Answer: B

NEW QUESTION 51

- (Topic 2)

You have a Microsoft Foundry project that contains an agent. The agent uses Azure AI Search for Retrieval Augmented Generation (RAG). You plan to ingest and index PDF product manuals. You need to build a solution that supports semantic similarity matching. The solution must ensure that the agent retrieves relevant data when user questions use different wording than the product manuals.

- A. vector search
- B. semantic ranking
- C. suggesters
- D. analyzers

Answer: A

NEW QUESTION 54

- (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 58

- (Topic 2)

You are building an app that will share user images. You need to configure the app to:

? Categorize each image as a photograph or drawing.

? Generate a caption for the image.

? Minimize development effort.

Which two services should you include?

- A. Image type detection in Azure Vision in Foundry Tools
- B. Object detection in Azure Vision in Foundry Tools
- C. Content tags in Azure Vision in Foundry Tools
- D. Image classification in Azure Custom Vision
- E. Image descriptions in Azure Vision in Foundry Tools

Answer: AE

NEW QUESTION 63

- (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 68

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 72

- (Topic 2)
- You have an application named App1 that uses Azure Speech in Foundry Tools to transcribe live calls. Transcript segments often contain both English and Spanish. App1 sends each segment to Azure Translator in Foundry Tools to translate to another language. Sometimes, mixed-language segments result in incomplete or incorrect translations. You need to reduce translation errors. The solution must ensure that the entire transcript is translated successfully. What should you do before sending the segments to Translator?
- A. Specify English as the source language in the translation request for all the segments.
B. Enable automatic language detection for the translation request.
C. Split the mixed-language segments into single-language segments and translate each segment separately.
D. Use document translation to translate the entire transcript as a single document.

Answer: C

NEW QUESTION 74

- DRAG DROP - (Topic 2)
- You have a Microsoft Foundry project that contains a customer support agent grounded in internal documentation. After a recent update, users report the following issues:
- Some answers are unsupported by retrieved documents.
 - A small number of responses are flagged for policy violations. You need to evaluate each issue.
- Which observability signals should you use for each issue? To answer, drag the appropriate observability signals to the correct issues. Each observability signal 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.

Observability signals

⋮ Groundedness evaluation metrics

⋮ Latency breakdown traces

⋮ Risk and safety metrics

⋮ Token usage analytics

Answer Area

Unsupported responses:

Observability signal

Policy violations:

Observability signal

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Observability signals

⋮ Groundedness evaluation metrics

⋮ Latency breakdown traces

⋮ Risk and safety metrics

⋮ Token usage analytics

Answer Area

Unsupported responses

⋮ Groundedness evaluation metrics

Policy violations:

⋮ Risk and safety metrics

NEW QUESTION 79

HOTSPOT - (Topic 2)

You have 100,000 images.

You need to build an app that will perform the following actions:

- Identify road signs in the images and extract the text on the signs.
- Analyze the text to identify well-known locations. The solution must minimize development effort.

Answer Area

Extract the text:

Azure Document Intelligence in Foundry Tools

Azure Language in Foundry Tools

Azure AI Search

Azure Vision in Foundry Tools

Identify well-known locations:

Azure Document Intelligence in Foundry Tools

Azure Language in Foundry Tools

Azure AI Search

Azure Vision in Foundry Tools

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Extract the text:

Azure Document Intelligence in Foundry Tools
Azure Language in Foundry Tools
Azure AI Search
Azure Vision in Foundry Tools

Identify well-known locations:

Azure Document Intelligence in Foundry Tools
Azure Language in Foundry Tools
Azure AI Search
Azure Vision in Foundry Tools

NEW QUESTION 84

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 \
```

`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 89

.....

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](#)