

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 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 3

- (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 4

HOTSPOT - (Topic 1)

You need to configure the model deployment for Agent1 to meet the technical requirements.

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

Deployment type:

Version update policy:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Deployment type:

Standard

Global Standard

Global Provisioned

Version update policy:

Once the current version expires

Opt out of automatic model version upgrades

Upgrade once a new default version becomes available

NEW QUESTION 5

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}")
```

AzureKeyCredential()

DefaultAzureCredential()

None

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=  ,
)

myAgent = "Agent1"

agent = project_client.agents.  (agent_name=myAgent)

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

create_version
get
get_version

NEW QUESTION 6

- (Topic 2)

You have a Microsoft Foundry project that contains an agent.

The agent uses a knowledge source built from documents stored in Azure Blob Storage. The documents include digitally scanned PDFs that contain multipage tables.

You have an ingestion job that extracts only plain text, causing loss of table structure, headings, and page-number metadata.

Users frequently ask questions that require the retrieval of specific table rows across the pages.

You need to configure an ingestion job for a Retrieval Augmented Generation (RAG) pipeline that performs optical character recognition (OCR) on scanned PDFs, preserves tables and headings as structure-aware chunks, and stores page-number metadata with each chunk.

How should you configure the ingestion job?

- A. Use basic parsing and fixed-size chunking.
- B. Use advanced data parsing to reingest the documents.
- C. Use OCR and page-level chunking.
- D. Use page-level OCR extraction and store each page as a single chunk.

Answer: B

NEW QUESTION 7

- (Topic 2)

You have a Microsoft Foundry agent that grounds responses from an Azure AI Search index containing:

? Searchable text fields for product names and product codes.

? A vector field containing embeddings for product descriptions.

You need users to query by exact product names or codes and by natural-language product descriptions.

- A. Semantic search only
- B. Hybrid search
- C. Keyword search only
- D. Vector search only

Answer: B

NEW QUESTION 8

- (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 9

- (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 10

- (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 10

- (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 15

- (Topic 2)

You have an application that processes scanned PDF invoices. The invoices have varied layouts and include multipage tables.

You have a pipeline that uses optical character recognition (OCR) and extracts totals and invoice numbers. The results are often incorrect because the document structure is ignored.

You need to implement a solution that provides OCR, layout analysis, and template- generalizing field extraction. The solution must NOT require training a custom model. The solution must minimize administrative effort. What should you include in the solution?

- A. an Azure Machine Learning model
- B. Azure Content Understanding in Foundry Tools
- C. Azure Language in Foundry Tools

Answer: B

NEW QUESTION 17

- (Topic 2)

You have a Microsoft Foundry project that uses Azure AI Search to ground an agent in internal documentation.

After a recent content update, users report that the agent's answers have become less accurate.

You need to identify whether the retrieved content is negatively influencing the model's generated responses.

Which observability signal should you review?

- A. prediction drift metrics
- B. groundedness evaluation metrics

- C. latency breakdown traces
- D. indexer status and failure history

Answer: B

NEW QUESTION 21

- (Topic 2)

You have a customer support agent built by using the Microsoft Foundry Agent Service.

The agent calls an Azure OpenAI model deployment.

During load testing, calls intermittently fail and return an HTTP 429 rate limit exceeded error.

You need to handle throttling to reduce call failures and improve reliability under load. The solution must remain within the service and model limits. What should you do?

- A. Implement a retry policy that uses exponential backoff and jitter.
- B. Create a new thread and retry the calls immediately.
- C. Reduce the number of registered tools.
- D. Split uploaded content into smaller files.

Answer: A

NEW QUESTION 23

DRAG DROP - (Topic 2)

You have a Docker host named Host1 that contains a container base image.

You have an Azure subscription that contains a custom speech-to-text model named model1.

You need to run model 1 on Host1.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Retrain the model.

Configure disk logging.

Export model1 to Host1.

Run the container.

Request approval to run the container.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Retrain the model.

Configure disk logging.

Export model1 to Host1.

Run the container.

Request approval to run the container.

Answer Area

Request approval to run the container.

Export model1 to Host1.

Run the container.

NEW QUESTION 28

- (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 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

HOTSPOT - (Topic 2)

You have a Python application named App1 that integrates with a Microsoft Foundry project named Project1.

You need to ensure that App1 meets the following requirements:

- Authenticates by using a Microsoft Entra managed identity
- Sends prompts to a deployed model by using the Azure OpenAI Responses API

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.

```
from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

credential =  ()
    AzureKeyCredential
    ClientSecretCredential
    DefaultAzureCredential

project_client = AIProjectClient(
    endpoint="https://contosoai.services.ai.azure.com/api/projects/project1",
    credential=credential,
)

with project_client.get_openai_client() as openai_client:
    response = openai_client.responses.  (
        compact
        create
        retrieve

    model="trail-guide-chat",
    input="Create a 3-day hiking itinerary near Seattle.",
)

print(response.output_text)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```

from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

credential =  ()


project_client = AIProjectClient(

    endpoint="https://contosoai.services.ai.azure.com/api/projects/project1",

    credential=credential,

)

with project_client.get_openai_client() as openai_client:

    response = openai_client.responses.  (

        model="trail-guide-chat",

        input="Create a 3-day hiking itinerary near Seattle.",

    )

    print(response.output_text)

```

NEW QUESTION 42

- (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 47

- (Topic 2)

You have a Microsoft Foundry project that contains a customer support agent built on a deployed chat model. The agent responses are validated by using an automated testing system that compares generated answers to stored expected outputs. Identical prompts must return consistent responses to prevent automated test failures. You need to reduce response variability, without modifying the prompt or reducing factual accuracy.

- A. Increase the max_tokens parameter.
- B. Remove stop sequences from the requests.
- C. Increase the temperature parameter.
- D. Decrease the temperature parameter.

Answer: D

NEW QUESTION 49

- (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 51

- (Topic 2)

You have a Microsoft Foundry project that serves a high-volume chat app. Most requests are simple FAQs, but some require advanced reasoning.

You need to reduce costs and latency for common queries, without degrading the quality of the responses to complex questions.

What should you do?

- A. Increase the value of the max_tokens parameter for all the requests.
- B. Route all the requests to a smaller model.
- C. Route all the requests to the most capable model.
- D. Use a model cascade that routes the requests to different models.

Answer: D

NEW QUESTION 55

- (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 59

- (Topic 2)

You have an Azure subscription that contains an Azure Language in Foundry Tools service resource. You need to identify the URL of the REST interface for the Language service. Which blade should you use in the Azure portal?

- A. Networking
- B. Keys and Endpoint
- C. Identity
- D. Properties

Answer: B

NEW QUESTION 62

- (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 protected material detection. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 63

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains a workflow for a customer support triage process.

You have an Ask a question node that stores user responses in a local variable named Var01.

You need to create the following Power Fx expressions:

? An if/else condition expression that ensures that Var01 contains a value

? A Send message expression that returns the stored user response in uppercase

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

If/else condition expression:

▼

IsBlank(Local.Var01)

IsEmpty(Local.Var01)

Not(IsBlank(Local.Var01))

Send message expression:

▼

{Local.Var01}

{Upper(Local.Var01)}

{Upper(Var01)}

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

If/else condition expression:

▼

IsBlank(Local.Var01)

IsEmpty(Local.Var01)

Not(IsBlank(Local.Var01))

Send message expression:

▼

{Local.Var01}

{Upper(Local.Var01)}

{Upper(Var01)}

NEW QUESTION 64
- (Topic 2)
You have a Microsoft Foundry project that contains three agents as shown in the following table.
Name Description TriageAgent

Classifies incoming customer requests PolicyAgent
Answers policy questions by searching internal content ActionAgent
Creates or updates tickets by calling an HTTP API
You need to orchestrate the agents to ensure that the customer requests meet the following requirements:

- Support a deterministic, step-based process that uses conditional branching and shared state across the agents.
- Optionally trigger a ticket action based on the triage result. The solution must minimize development effort.

What should you include in the solution?

- A. separate agent runs coordinated in the application code
- B. a multi-agent group chat session
- C. a workflow
- D. threads and runs without a workflow

Answer: C

NEW QUESTION 66

- (Topic 2)
You are deploying a support agent that enables users to upload photos.
You need to automatically classify uploaded images for harmful content. The solution must block content based on severity levels.
What should you do?

- A. Implement image moderation.
- B. Enable prompt shields.
- C. Apply keyword scanning to optical character recognition (OCR) output by using Azure Vision in Foundry Tools.
- D. Use blocklists.

Answer: A

NEW QUESTION 67

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 72

- (Topic 2)

You have a Microsoft Foundry project that contains an agent. The agent uses Azure AI Search as the retriever.

You plan to ingest PDFs into an Azure AI Search index to ensure that the agent can ground responses in texts in both documents and embedded images.

Users require citations that link to the source files.

You need to ensure that during indexing, the images are extracted into a structure that can be used as input for the built-in optical character recognition (OCR) skill.

Which indexing approach should you use?

- A. a skillset to run the OCR skill directly against the content field of the index
- B. the outputFieldMappings parameter to write image data to a searchable field
- C. an indexer to extract image data into a normalized_images collection
- D. a Shaper skill to restructure the OCR input

Answer: C

NEW QUESTION 76

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions 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 First Try
- * AI-103 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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