

Exam Questions AI-103

Developing AI Apps and Agents on Azure

<https://www.2passeasy.com/dumps/AI-103/>



NEW QUESTION 1

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:

Version update policy:

NEW QUESTION 2

- (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 user prompts. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 3

- (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 4

- (Topic 2)

You have a Microsoft Foundry project that contains an agent.

The knowledge source for the agent is a set of scanned PDF troubleshooting guides stored in Azure Blob Storage. The guide pages contain two-column layouts and tables.

You use Azure Content Understanding in Foundry Tools to process the PDFs.

You plan to ingest the processed content into an index for Retrieval Augmented Generation (RAG) and store extracted fields for downstream automation.

Stakeholders must be able to verify where each extracted field value came from in the original PDF and route low-reliability extractions for manual review.

You need to ensure that the Content Understanding document analyzer output includes a per-field confidence score and source grounding locations within the source document.

What should you do?

- A. Enable estimateFieldSourceAndConfidence.
- B. Configure the analyzer to use generative extraction for all fields.
- C. Set enableSegment to true.
- D. Provide labeled samples.

Answer: A

NEW QUESTION 5

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an agent named PaymentAgent.

PaymentAgent includes a function tool that issues customer refunds by using an external API.

You are creating a workflow in YAML.

You need to ensure that the workflow pauses for human approval and continues with the refund step only after approval is granted.

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

NOTE: Each correct selection is worth one point.

steps:

- id: propose_refund

type: agent

agent: PaymentAgent

- id: approval

type:

ask_question
basic_chat
data_transformation

- id: execute_refund

type: agent

agent: PaymentAgent

condition:

approval == "approved"
propose_refund.output != null
true

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

steps:

- id: propose_refund

type: agent

agent: PaymentAgent

- id: approval

type:

ask_question

basic_chat

data_transformation

- id: execute_refund

type: agent

agent: PaymentAgent

condition:

approval == "approved"

propose_refund.output != null

true

NEW QUESTION 6

- (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 temperature parameter. Does this meet the goal?

A. Yes

B. No

Answer: B

NEW QUESTION 7

HOTSPOT - (Topic 2)

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

The agent uploads user-provided screenshots to Azure Storage through a ticketing tool and receives a blob URL for additional reasoning.

You need to use image moderation during agent runs and prevent harmful content from being returned during runs. Azure AI Content Safety must access the images by using the blob URL. The solution must follow the principle of least privilege.

What should you configure for Content Safety? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Guardrails:

Select Tool call and set Action to Block.

Select User input and Output and set Action to Annotate.

Select User input and Tool response and set Action to Annotate.

Select User input, Output, Tool response, and Tool call and set Action to Block.

Storage access:

Storage account access keys

A user-assigned identity that is assigned the Storage Queue Data Contributor role

A system-assigned managed identity that is assigned the Storage Blob Data Reader role

A system-assigned managed identity that is assigned the Storage Blob Data Contributor role

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Guardrails:

Select Tool call and set Action to Block.

Select User input and Output and set Action to Annotate.

Select User input and Tool response and set Action to Annotate.

Select User input, Output, Tool response, and Tool call and set Action to Block.

Storage access:

Storage account access keys

A user-assigned identity that is assigned the Storage Queue Data Contributor role

A system-assigned managed identity that is assigned the Storage Blob Data Reader role

A system-assigned managed identity that is assigned the Storage Blob Data Contributor role

NEW QUESTION 8

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 9

HOTSPOT - (Topic 2)

You have a Microsoft Foundry resource named AI1 that hosts three deployments of the GPT 3.5 model. Each deployment is optimized for a unique workload. You plan to deploy three apps. Each app will access AM by using the REST API and will use the deployment that was optimized for the apps intended workload. You need to provide each app with access to AI1 and the appropriate deployment. The solution must ensure that only the apps can access AH. What should you use to provide access to AI1. and what should each app use to connect to its appropriate deployment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Provide access to AI1 by using:

An API key
A bearer token
A shared access signature (SAS) token

Connect to the deployment by using:

An API key
A deployment endpoint
A deployment name
A deployment type

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Provide access to AI1 by using:

An API key
A bearer token
A shared access signature (SAS) token

Connect to the deployment by using:

An API key
A deployment endpoint
A deployment name
A deployment type

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 have a customer support agent that uses the Microsoft Foundry Agent Service.

Sometimes, customers return to a session days later to continue the same support case, and the agent must resume with the full historical context. The agent must provide the following:

- Multi-turn continuity within the session
- Cross-session continuity for the same case
- Access to the full interaction history, including user messages, agent messages, tool calls, and tool outputs

You need to ensure that the agent automatically reloads the complete history on each new turn.

What should you do?

- A. Persist only the final model response stored in the client application and prepend the response to future prompts.
- B. Enable memory summarization on the agent definition to persist the context automatically.
- C. Create and reuse a conversation by storing the conversation??s ID and supplying the ID on subsequent requests.

Answer: C

NEW QUESTION 11

HOTSPOT - (Topic 2)

Your company is piloting a customer support agent in a Microsoft Foundry project name Project1. Project1 is connected to an existing Application Insights resource, and the company's support team reviews runs in the Traces tab.

The Foundry Agent Service is configured to perform the following actions:

- Retrieve the Application Insights connection string by calling project_client.telemetry.get_application_insights_connection_string().
- Call configure_azure_monitor(connection_string=...) to enable telemetry.

A separate LangChain service configured to use OpenTelemetry and has the following configurations:

- Uses AzureAIOpenTelemetryTracer(connection_string=..., enable_content_recording=False)
- Passes the tracer by using config={"callbacks":[azure_tracer]} Company policy has the following requirements:
- Telemetry from LangChain and OpenTelemetry must be distinguishable within the same Application Insights resource.
- Secrets and credentials must NOT be stored in prompts, tool arguments, or span attributes.

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

NOTE: Each correct selection is worth one point.

Answer Area

Statements

Yes

No

The LangChain service will appear in Traces without configuring a tracer.

☐☐

Setting different OTEL_SERVICE_NAME values separates the services in Application Insights.

☐☐

When using enable_content_recording=False, prompts and tool data will be captured in the telemetry.

☐☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements

Yes

No

The LangChain service will appear in Traces without configuring a tracer.

☐☒

Setting different OTEL_SERVICE_NAME values separates the services in Application Insights.

☒☐

When using enable_content_recording=False, prompts and tool data will be captured in the telemetry.

☐☒

NEW QUESTION 16

- (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 18

- (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 22

- (Topic 2)

You plan to configure an evaluation in Microsoft Foundry for a Retrieval Augmented Generation (RAG) chat app. You need to provide scores for groundedness, relevance, and harmful-content categories. Which two evaluation categories can you use? Each correct answer presents part of the solution.

- A. risk and safety metrics
- B. AI quality (NLP) metrics
- C. AI quality (AI assisted) metrics
- D. fluency evaluator
- E. similarity evaluators

Answer: AC

NEW QUESTION 26

- (Topic 2)

You have a web app named App1 that sends requests to a multimodal chat model deployment in a Microsoft Foundry project. User messages can contain both text and images. Currently, App1 includes image URLs as plain text inside the message content, so the model cannot recognize them as images. You need to send the message as a structured array that includes both the text portion and the image reference.

- A. Add the image URL to request metadata.
- B. Set the user message content array to include items with type: text and type: image_url.
- C. Place the image URL inside the system message.
- D. Include base64 data directly inside the content string.

Answer: B

NEW QUESTION 29

HOTSPOT - (Topic 2)

You are creating an enrichment pipeline that will use Azure AI Search. The knowledge store contains unstructured JSON data and the text from scanned PDF documents.

Which projection type should you use for each data type? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

JSON data:	File projection Object protection Table projection
Extracted text data:	File protection Object projection Table projection

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 32

- (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 37

- (Topic 2)

In Microsoft Foundry, you use the Chat playground with the GPT-35 Turbo model. You have a prompt that contains the following code.

```
function F(n)
{
    var f = [0, 1];
    for (var i = 2; i < n; i++) f[i] = f[i-1] + f[i-2];
    return f;
}
```

You need the model to create an explanation of the code. The solution must minimize costs. What should you do?

- A. Add function F (explanation) to the prompt.
- B. Change the model to GPT-4-32k.
- C. Add// what does function F do? to the prompt.
- D. Set the temperature parameter to 1.

Answer: C

NEW QUESTION 39

- (Topic 2)

You have a Microsoft Foundry project that ingests scanned PDF invoices stored in Azure Blob Storage. Each invoice contains printed line items and has a table-based layout.

Extracted results are stored as structured JSON and used as grounding data for an agent in a Retrieval Augmented Generation (RAG) solution.

You need to create a single analyzer that meets the following requirements:

- Extracts the invoice number, invoice date, vendor name, and total amount across varying templates
- Returns confidence scores so that results with confidence below 0.80 can be routed for supervisor review

What should you use?

- A. the Azure Content Understanding in Foundry Tools prebuilt-layout analyzer
- B. a Foundry agent that has groundedness guardrails enabled to extract invoice fields and confidence scores
- C. a custom Azure Content Understanding in Foundry Tools analyzer that defines the required fields as the extracted fields and the returned confidence scores for

routing
D. the Azure Content Understanding in Foundry Tools prebuilt-documentSearch analyzer and search.score from the Azure AI Search results for routing

Answer: C

NEW QUESTION 41

- (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 44

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains a deployed chat model.

You have a Python service that sends API requests to the model. The service is integrated with an automated validation system that compares generated outputs against approved response patterns.

Stakeholders report that small wording differences are causing validation mismatches.

You need to update the request parameters to improve output stability. The solution must maximize reasoning quality.

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.

```
message = client.messages.create(

    model="deployment-name",

    messages=[

        {"role": "user", "content": "Summarize the release notes in 3 bullet points."}

    ],

    max_tokens=800,

    temperature=  ,

    thinking={"type": "enabled"},

    output_config={"effort":  }

)
```

0

1

2

"high"

"low"



"medium"

A. Mastered

B. Not Mastered

Answer: A

Explanation:

```
message = client.messages.create(  
    model="deployment-name",  
    messages=[  
        {"role": "user", "content": "Summarize the release notes in 3 bullet points."}  
    ],  
    max_tokens=800,  
    temperature=  
          
    ,  
    thinking={"type": "enabled"},  
    output_config={"effort":  
          
    }  
)
```

NEW QUESTION 46

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AI-103 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AI-103 Product From:

<https://www.2passeasy.com/dumps/AI-103/>

Money Back Guarantee

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