

Exam Questions AI-103

Developing AI Apps and Agents on Azure

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



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 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 3

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 4

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 5

- (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 6

- (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 7

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 8

- (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 9

HOTSPOT - (Topic 2)

You have a Python application that uses Azure OpenAt structured outputs to extract fields from unstructured receipt text and support ticket text. The application uses the following schema.

```
receipt_response_format = {
    "type": "json_schema",
    "json_schema": {
        "name": "ReceiptExtraction",
        "strict": True,
```

Answer Area

Statements	Yes	No
The schema requires that a receipt has a discount_code value.	☐	☐
The schema can be used as is for Azure OpenAI structured outputs.	☐	☐
The generated JSON object will preserve the schema property order of merchant, order_number, discount_code, line_items, and total.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The schema requires that a receipt has a discount_code value.	☐	☒
The schema can be used as is for Azure OpenAI structured outputs.	☐	☒
The generated JSON object will preserve the schema property order of merchant, order_number, discount_code, line_items, and total.	☒	☐

NEW QUESTION 10

- (Topic 2)

You have an Azure Speech in Foundry Tools resource that hosts a custom speech to text model deployed to a custom endpoint. An agent uses the endpoint to perform real-time speech recognition.
You are approaching the expiration date of the custom speech to text model. What is the expected behavior when the model expires?

- A. Speech recognition requests will fall back to the most recent base model for the same locale.
- B. Speech recognition requests will continue to use the expired custom model until the model is removed manually.
- C. Speech recognition requests will return a 4xx error until a new custom model is deployed.
- D. The custom model will be deleted automatically when the model expires.

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

DRAG DROP - (Topic 2)
You have a Microsoft Foundry project that contains a deployed ticket-triage agent.
You discover that sometimes the agent responds without calling any tools, even when a tool is required.
You need to ensure that the agent calls a tool during execution.
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. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Values

"auto"

"required"

"response_format"

"tool_choice"

"tools"

"type"

Answer Area

```
run_payload = {  
    "assistant_id": agent_id,  
    Value : Value ,  
    "metadata": {  
        "scenario": "ticket-triage"  
    }  
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Values

"auto"

"required"

"response_format"

"tool_choice"

"tools"

"type"

Answer Area

```
run_payload = {  
    "assistant_id": agent_id,  
    "tool_choice" : "required" ,  
    "metadata": {  
        "scenario": "ticket-triage"  
    }  
}
```

NEW QUESTION 19

- (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 24

- (Topic 2)

You have an invoice-processing application named App1 that uses Azure Content Understanding in Foundry Tools. You are building a new Content Understanding pipeline named Pipeline1 that must meet the following requirements:

? Compare an invoice to its related purchase order.

? Validate the invoice against static vendor contract documents.

? Return a single structured output that includes discrepancy findings.

You need to configure Pipeline1 and expose the pipeline as a single analyzer endpoint. What should you configure?

A. A single-file task in standard mode that uses confidence scores enabled for the extracted fields

B. A multi-file task in pro mode that uses the vendor contract files as reference data

C. A multi-file task in standard mode that uses the invoice and purchase order as input to the analyzer

D. A single-file task in standard mode that uses the vendor contract provided as an additional document during analysis

Answer: B

NEW QUESTION 29

- (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 33

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 AI1.

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 35

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 36

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 38

- (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 43

- (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 run an evaluation flow that scores responses for completeness and blocks responses that fall below a defined threshold.

Does this meet the goal?

- A. Yes
B. No

Answer: B

NEW QUESTION 47

- (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 50

- (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 52

- (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 57

- (Topic 2)

You are building a web app named App1 that generates responses by using a model deployed to a Microsoft Foundry project named Project1.

Before sending the prompts to the model, App1 must retrieve documents by using Azure AI Search.

You need to integrate Project1 and App1. The solution must meet the following requirements:

- Multiple client applications must use the same search configuration.
- A security policy must prevent key-based authentication.
- Administrative effort must be minimized. What should you do?

- A. Enable a managed identity for each application and call Azure AI Search directly.
- B. Create a custom HTTP connection in Foundry and manually configure Azure AI Search endpoints per application.
- C. Call Azure AI Search directly from each application by using Microsoft Entra authentication.
- D. Configure an Azure AI Search connection in Project1 and reference the connection in each application.

Answer: D

NEW QUESTION 62

- (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 66

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
  ...
  AddOrUpdateTextBlocklistItemsOptions
  AnalyzeImageOptions
  AnalyzeTextOptions
  TextBlocklist
```

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Answer Area

```
...
blocklist_name = "ConfidentialTerms"
client = ContentSafetyClient(endpoint, credential)
request =
  text=Input
  ...
  AddOrUpdateTextBlocklistItemsOptions
  AnalyzeImageOptions
  AnalyzeTextOptions
  TextBlocklist
```

NEW QUESTION 71

- (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 73

- (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 75

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 78

- (Topic 2)

You need to measure the public perception of your brand on social media by using natural language processing. Which Azure service should you use?

- A. Azure Document Intelligence in Foundry Tools
- B. Content Safety in Foundry Control Plane
- C. Azure Language in Foundry Tools
- D. Azure Vision in Foundry Tools

Answer: C

NEW QUESTION 79

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 83

- (Topic 2)

You have a Microsoft Foundry project that contains a customer support agent. The agent calls an internal knowledge API tool before generating responses.

Users report the following issues:

- Some requests take more than 15 seconds to complete.
- Some responses are incorrect, even when the knowledge API returns the expected data. You need to inspect individual agent runs to view the ordered sequence

of large language model (LLM) calls, tool invocations, and timing information. Which observability capability should you use?

- A. token usage
- B. safety metrics
- C. tracing
- D. monitoring

Answer: C

NEW QUESTION 87

- (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 92

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":  }

)
```

- 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 94

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 98

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 102

- (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 105

.....

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