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 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 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
- (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 5
- (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 6
- (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 7
DRAG DROP - (Topic 2)
You have an app that uses Azure AI Custom Vision and a custom trained classifier to identify products in images. You need to add new products to the classifier.
The solution must meet the following requirements:
• Minimize how long it takes to add the products.
• Minimize development effort.
Which five 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.

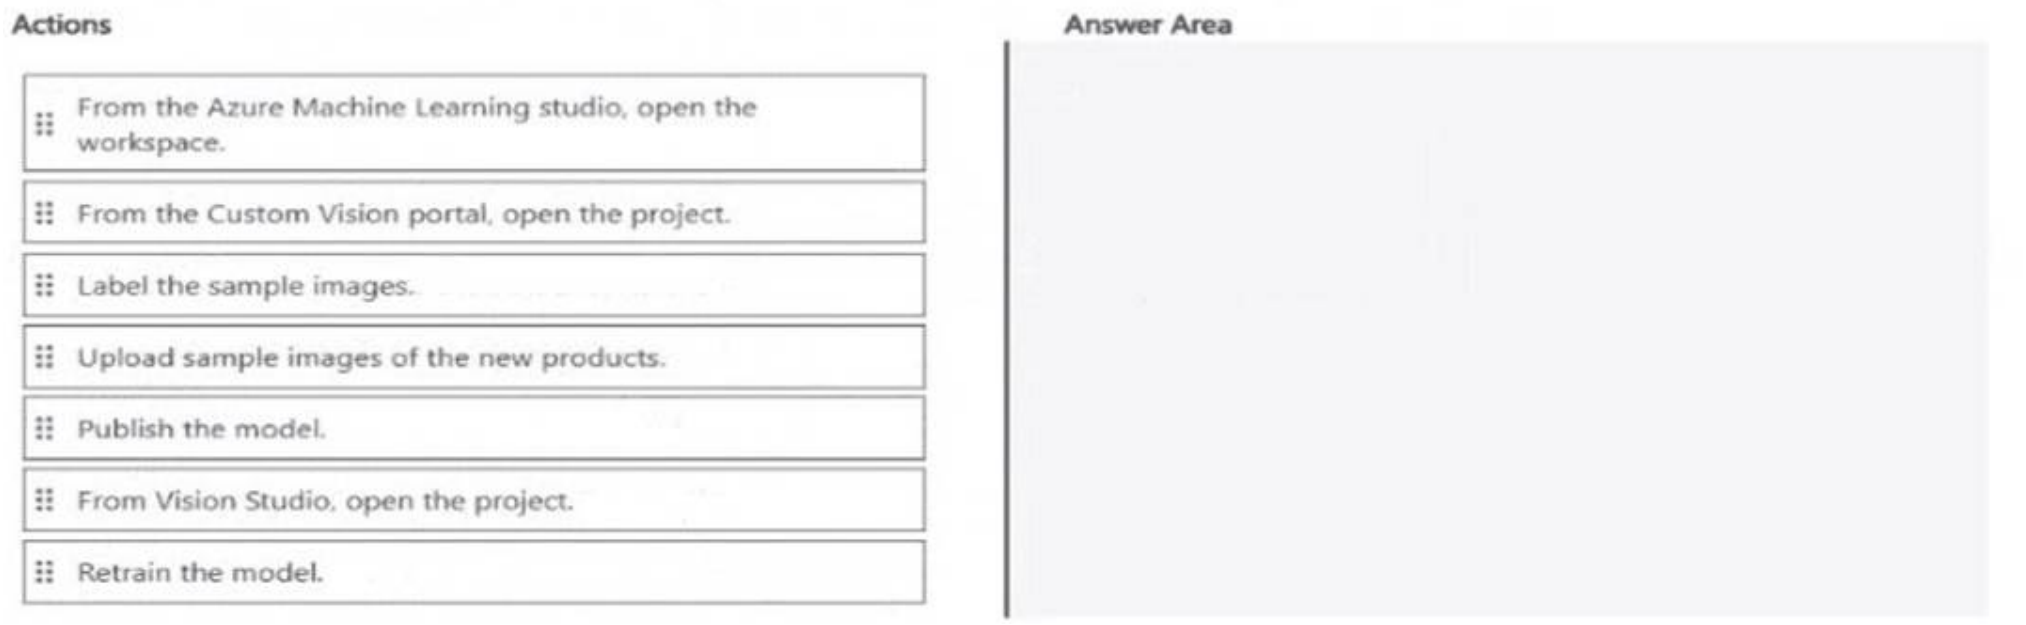

A. Mastered
B. Not Mastered

Answer: A

Explanation:

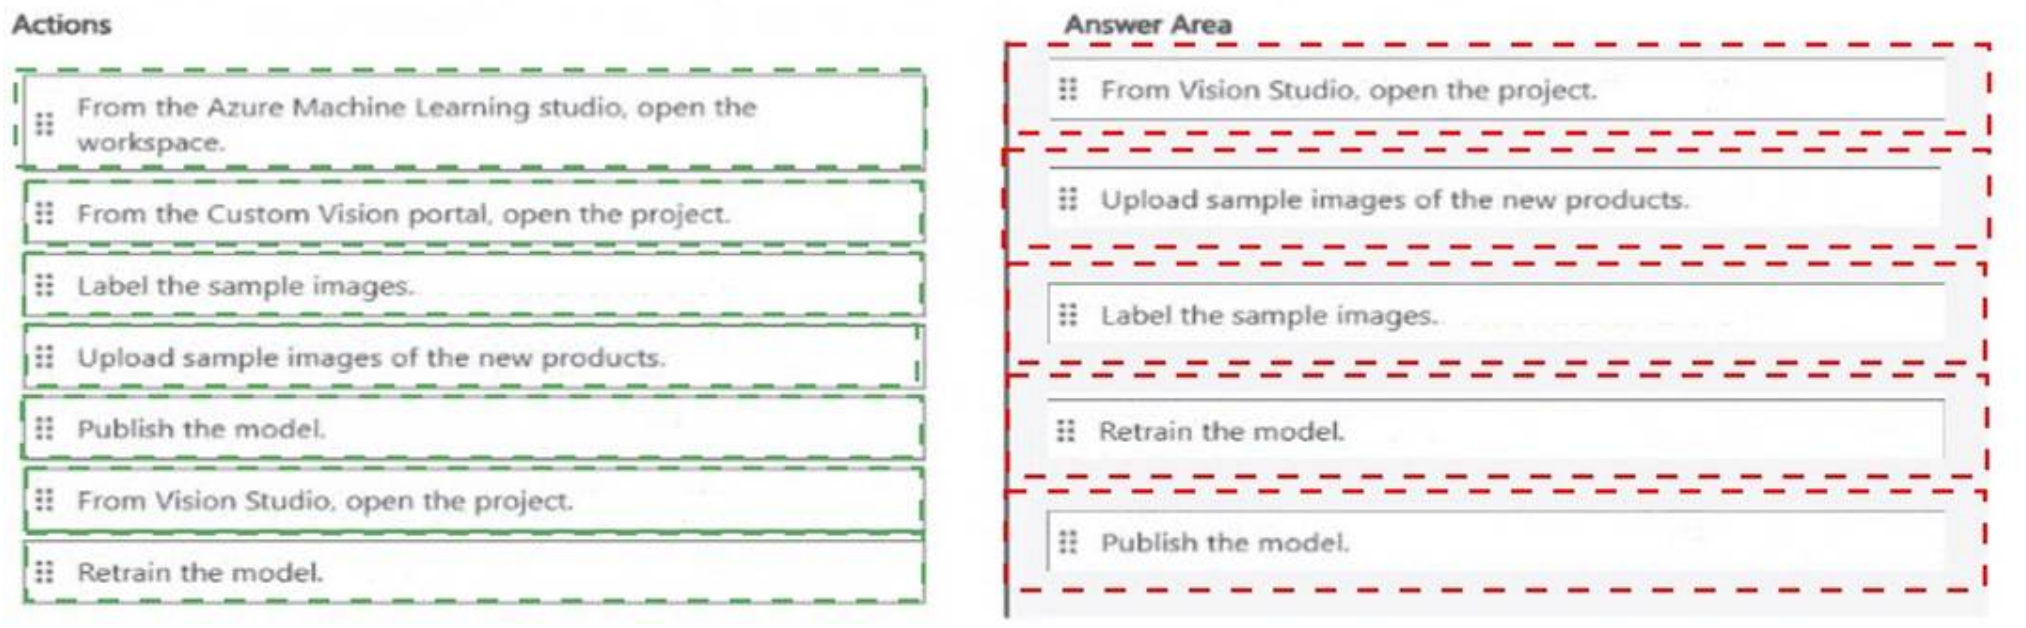

NEW QUESTION 8
- (Topic 2)
You have a Microsoft Foundry project that contains a high-traffic agent. After a recent update, operational costs increase significantly.
Monitoring confirms that the volume of user traffic to the agent remains unchanged.
You suspect that changes to the request or response 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 9
- (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 10
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.

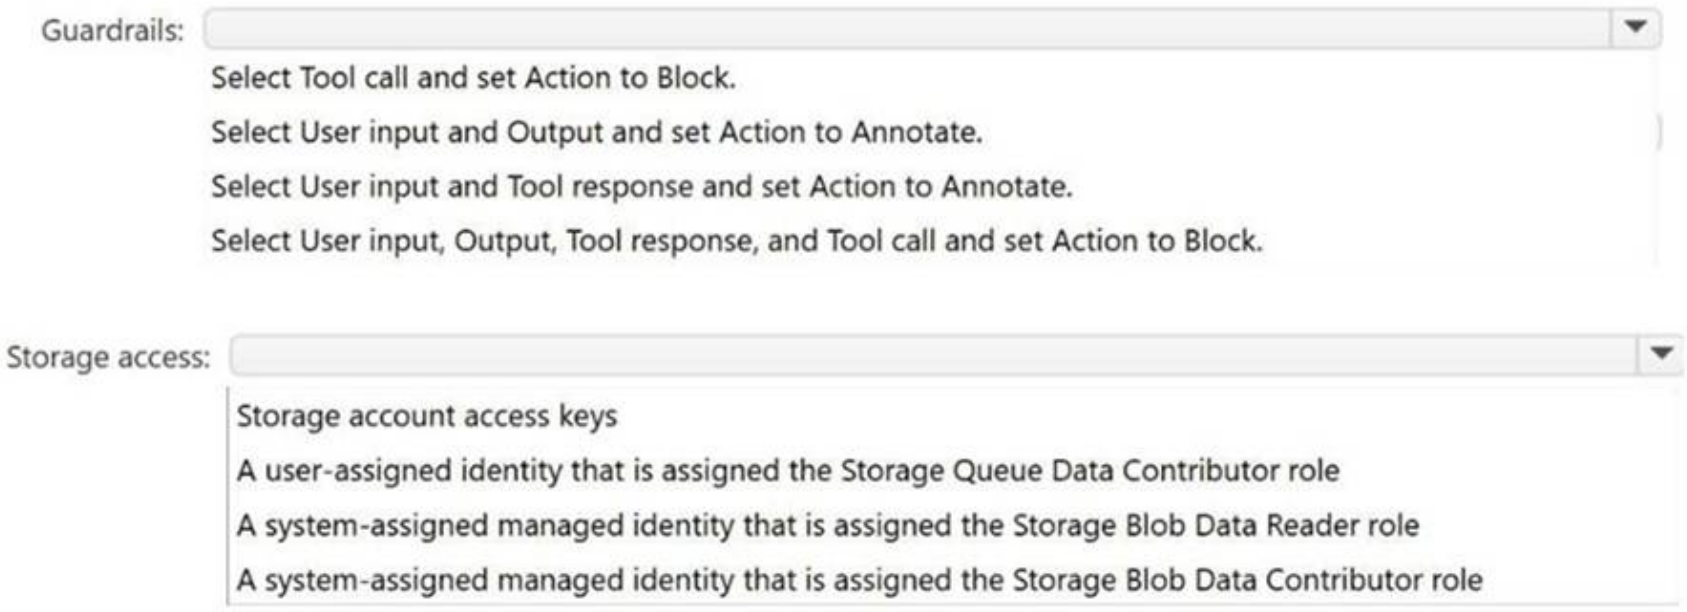

A. Mastered
B. Not Mastered

Answer: A

Explanation:

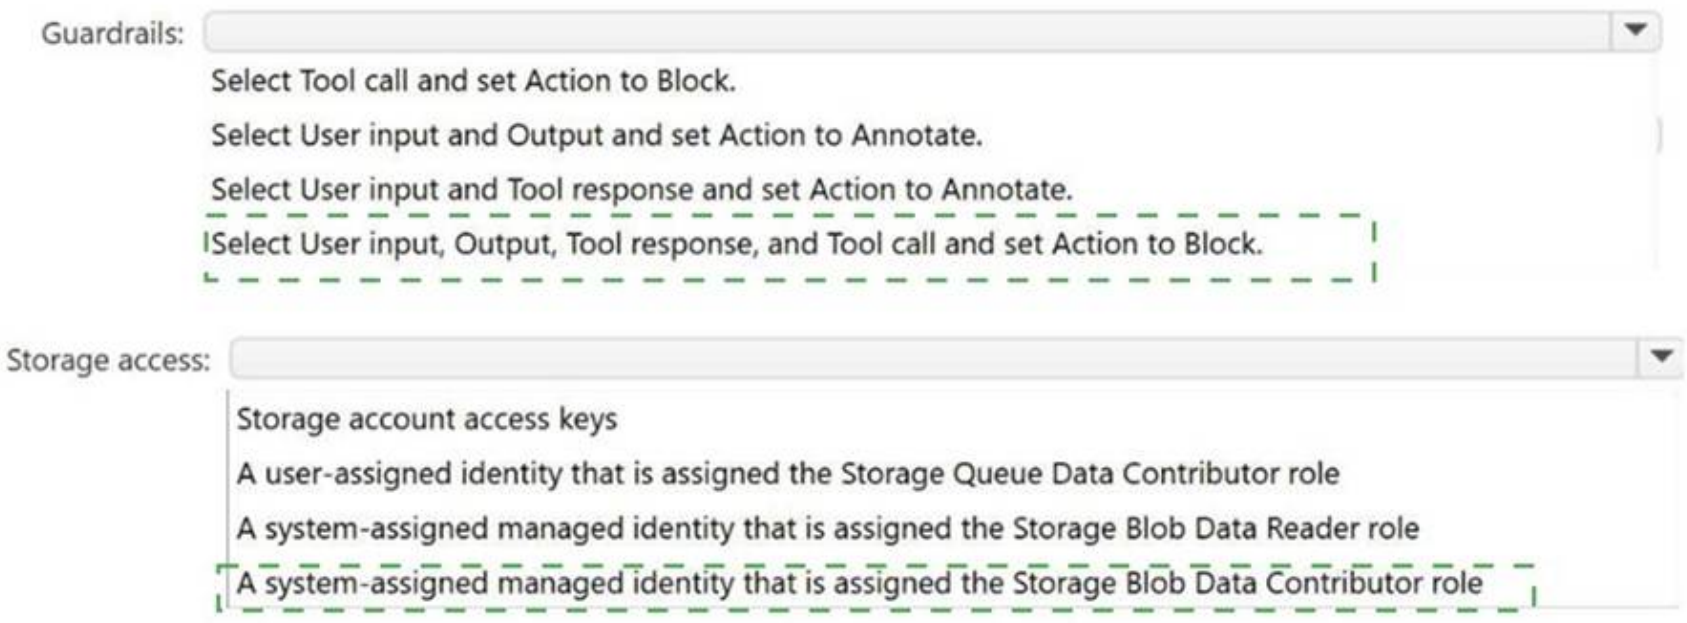

NEW QUESTION 10
- (Topic 2)
You have a chat app in a Microsoft Foundry project and an Azure AI Search vectorized index.

You need to connect to the index to meet the following requirements:
• Complex questions must retrieve information from multiple chunks.
• Multi-turn conversations must influence retrieval planning.
• Retrievals must run in parallel to reduce latency. Which retrieval approach should you use?

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

Answer: C

NEW QUESTION 11
- (Topic 2)
You have a Microsoft Foundry project that contains an agent. The agent generates summaries from retrieved policy documents.
You need to improve response completeness. The solution must be implemented in the logic of the application code before responses are returned.
What should you do?

A. Decrease the value of the temperature parameter.
B. Add a reflection pass before the responses are returned.
C. Switch to Retrieval Augmented Generation (RAG).
D. Enable response streaming.

Answer: B

NEW QUESTION 14
- (Topic 2)
You are creating an agent workflow in a Microsoft Foundry project to support natural voice interactions.
The agent must receive continuous audio input, convert the input into text for reasoning, and then return spoken responses to a user. The workflow must meet the following requirements:
. Support turn-taking dynamics, where the agent begins to generate the speech output before the user finishes speaking.
. Operate with low latency to maintain a conversational experience.
You need to enable both speech to text and text to speech in a real-time agent interaction. What should you do?

A. Use an embeddings model to encode the audio, and then decode the audio into text and speech.
B. Use batch transcription to convert the audio input and return text responses from the agent.
C. Use speech translation to convert the audio into another language and return the translated text.
D. Use real-time speech to text for incoming audio and text to speech for agent responses.

Answer: D

NEW QUESTION 18
- (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 23
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 24
DRAG DROP - (Topic 2)
You have a Docker host named Host! 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.

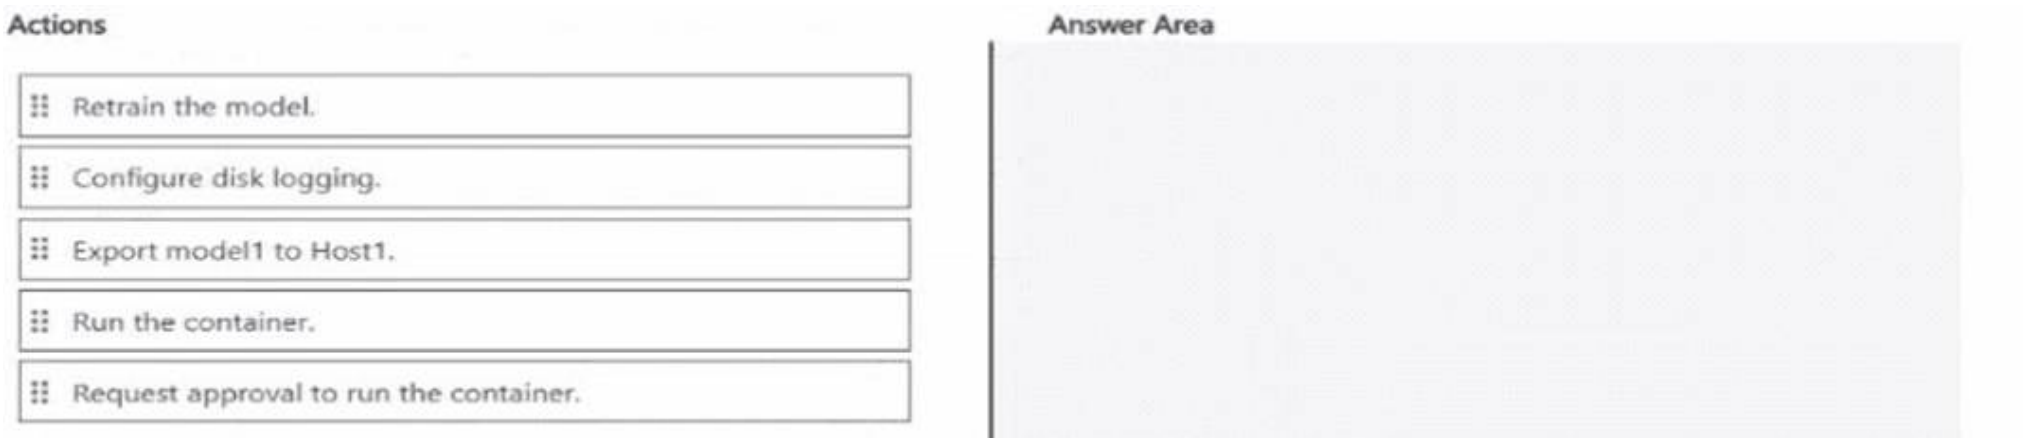

A. Mastered
B. Not Mastered

Answer: A

Explanation:

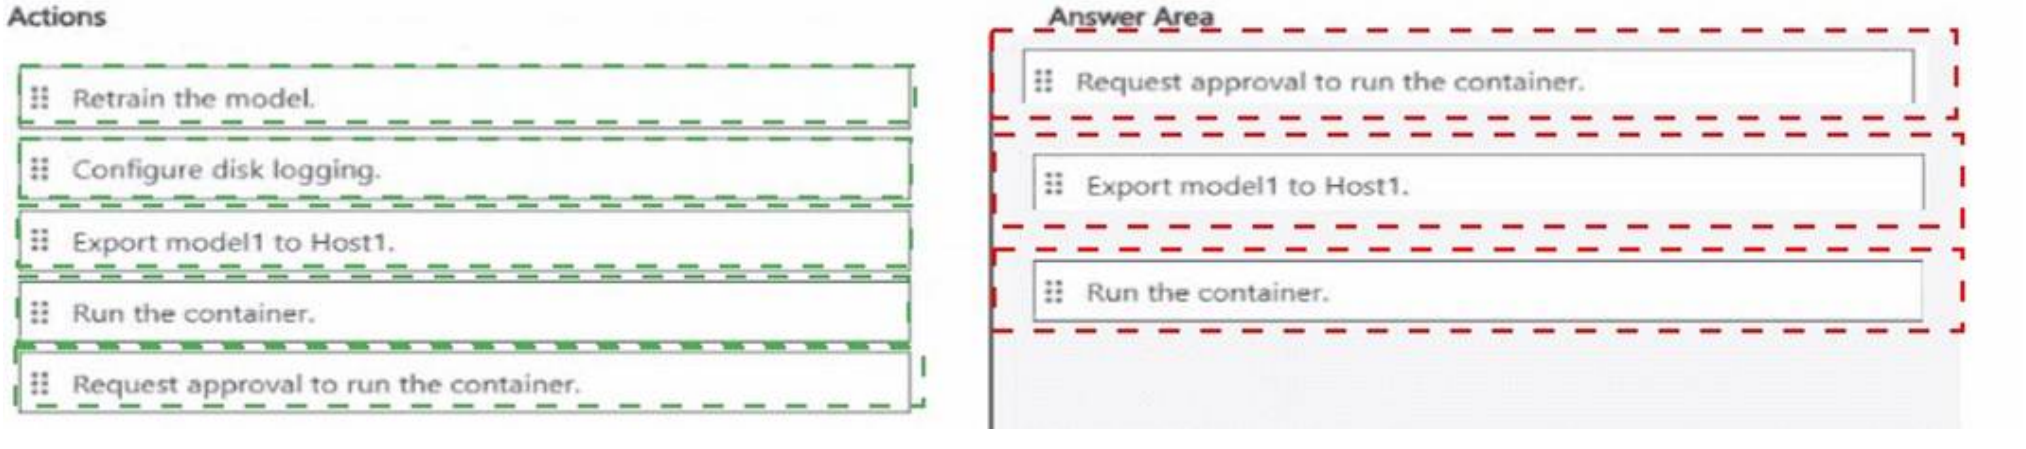

NEW QUESTION 25

- (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 26
HOTSPOT - (Topic 2)
You need to recommend a plan to create a customer support agent by using the Microsoft Foundry Agent Service. The agent must meet the following requirements:
• Retain user preferences across multiple conversations.
• Enable users to provide contextual grounding by directly uploading documents during a chat.
Which Foundry capability should you recommend for each requirement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

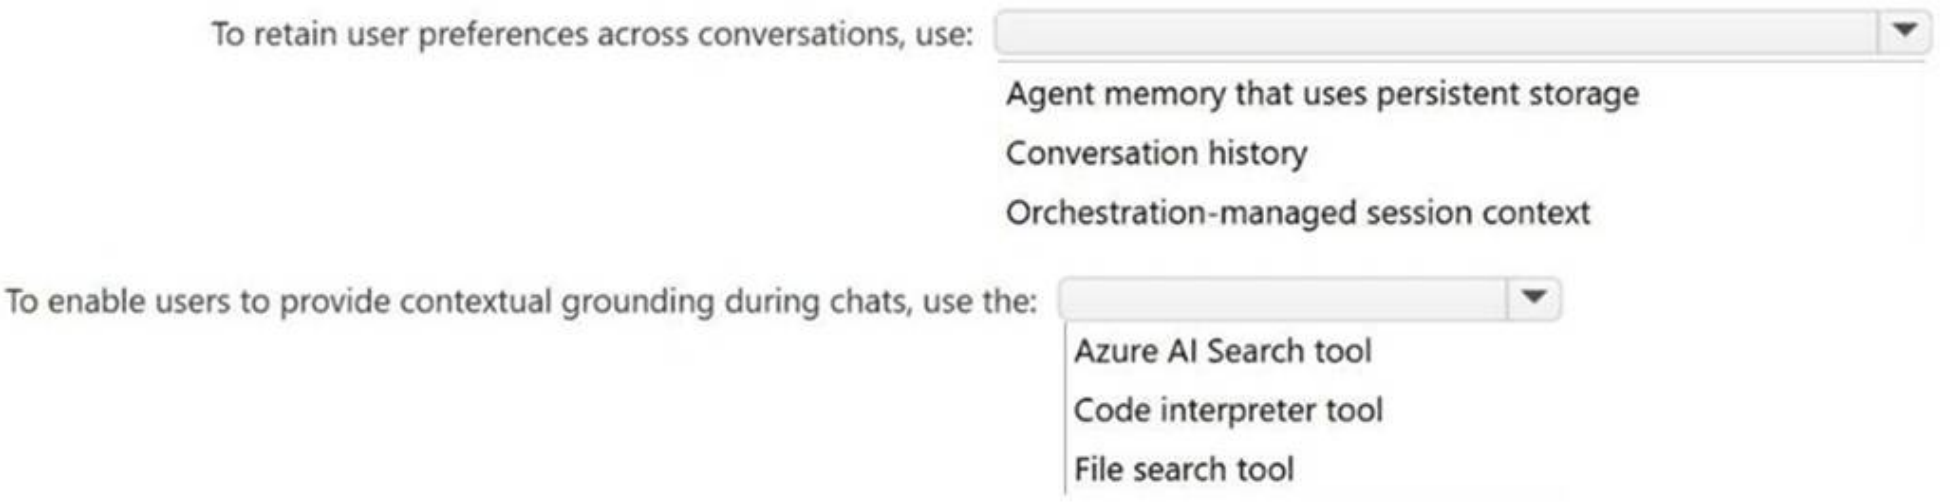

A. Mastered
B. Not Mastered

Answer: A

Explanation:

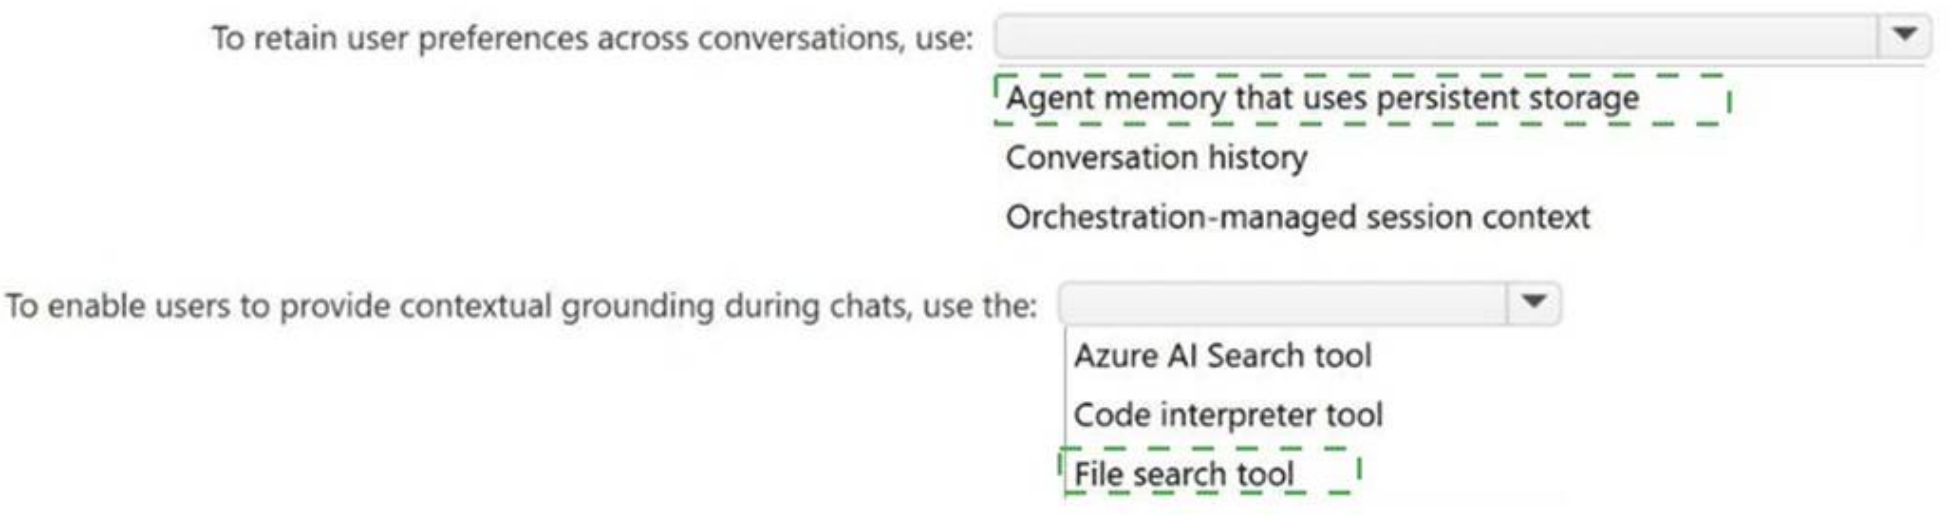

NEW QUESTION 30
- (Topic 2)
Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals.
More than one solution in the set might solve the problem. It is also possible that none of the solutions in the set solve the problem. After you answer a question in this section, you will NOT be able to return. As a result, these questions do not appear on the Review Screen.
You have a multimodal AI generative model that accepts image uploads and uses extracted image text to generate responses.
You discover that users can upload unsafe images and embed hidden instructions into images to manipulate the model.
You need to implement controls to mitigate the risk. Solution: You configure a prompt shield for documents. Does this meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 33
- (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 38
- (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 39
- (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 43
- (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 45
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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

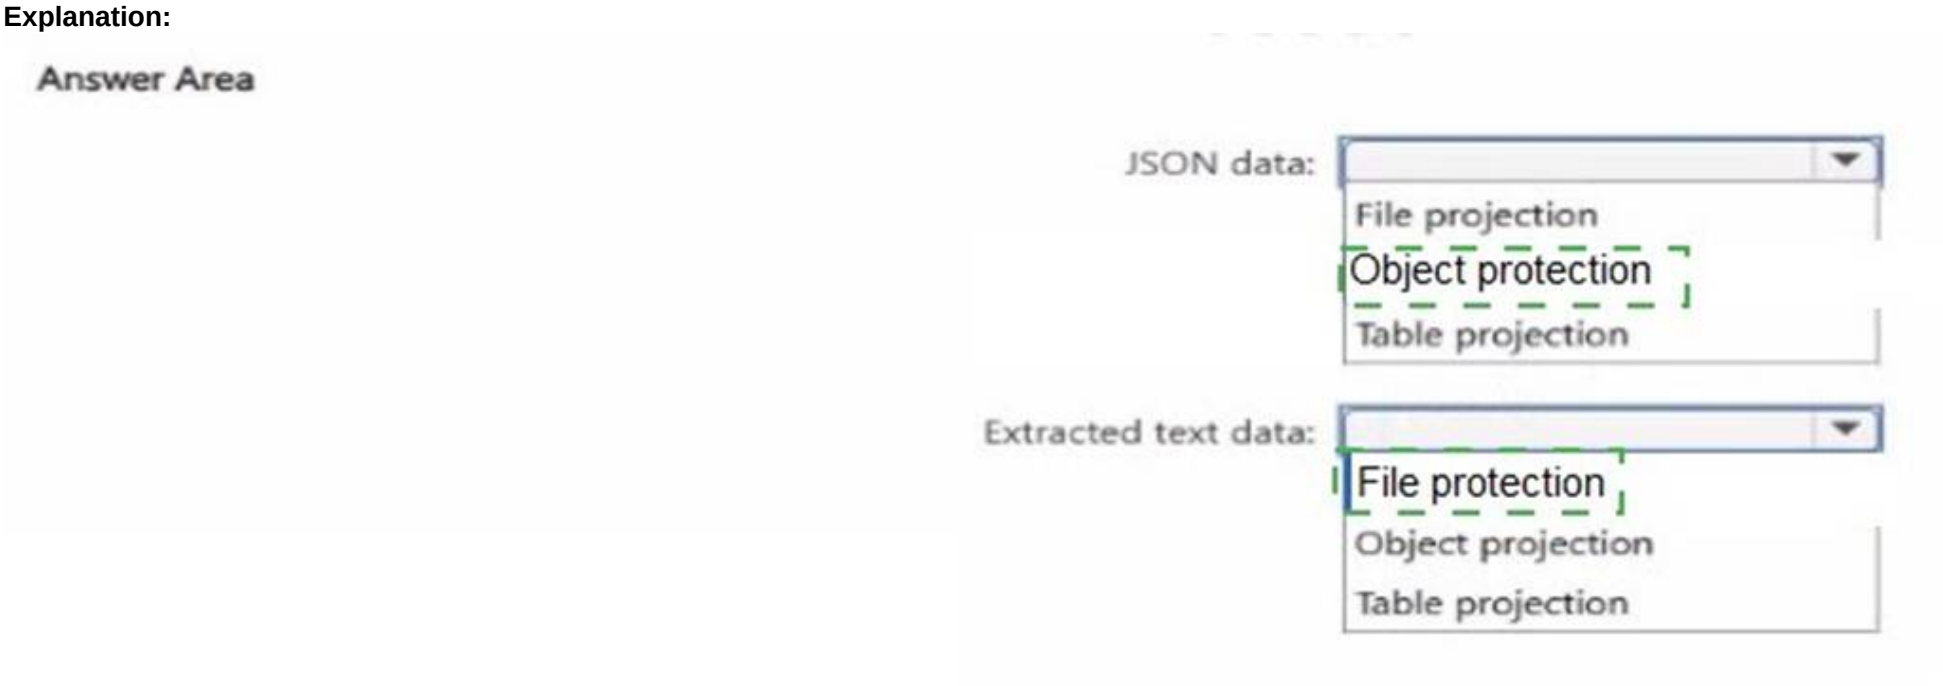

NEW QUESTION 50
- (Topic 2)
You have a Microsoft Foundry project. You need to deploy a model from the model catalog to support real-time inference. The solution must meet the following requirements:
? Use key-based authentication
? Support real-time REST API access
? Not consume the vCPU quota of the virtual machines in the Azure subscription
Which type of deployment should you use?

A. serverless API
B. self-hosted container
C. standard
D. batch

Answer: A

NEW QUESTION 53
HOTSPOT - (Topic 2)
You have a configurable guardrail in a Microsoft Foundry project.
You have a Python application. A text blocklist named Conf identialTerns already contains terms that must be detected in user messages. The application stores each incoming user message in a variable named inpot_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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 54
- (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 58
- (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 63
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.

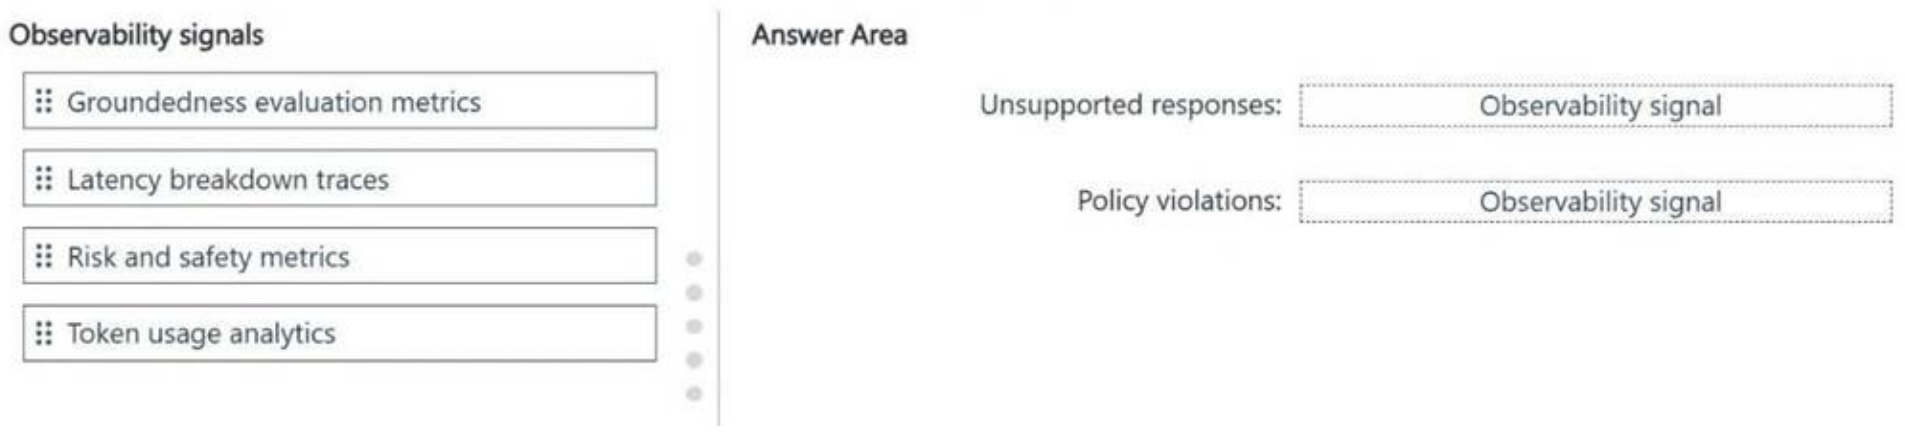

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 66
- (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 70
- (Topic 2)
You have a Microsoft Foundry project that contains an agent. The agent has a Model Context Protocol (MCP) tool that queries a knowledge base stored in Azure AI Search.
Some agent runs return answers from the base model without invoking the knowledge base, which results in responses without grounded citations.
You are provided with the following code snippet that runs the agent. run = project_client.agents.runs.create_and_process(thread_id=thread.id, agent_id=agent.id,
)
You need to add the correct tool_choice parameter to the code to deterministically force the agent to invoke the MCP tool on each run.
What should you add?

A. tool_choice ={"type":"mcp"}
B. tool_choice={"auto"}
C. tool_choice={"type":"knowledge_base"}

D. tool_choice={"required"}

Answer: D

NEW QUESTION 75
- (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 76
- (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 77
HOTSPOT - (Topic 2)
You plan to deploy a containerized version of an Azure Al Language service that will be used for sentiment analysis.
You configure https: //contoso.cognitiveservices.azure, cow 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.

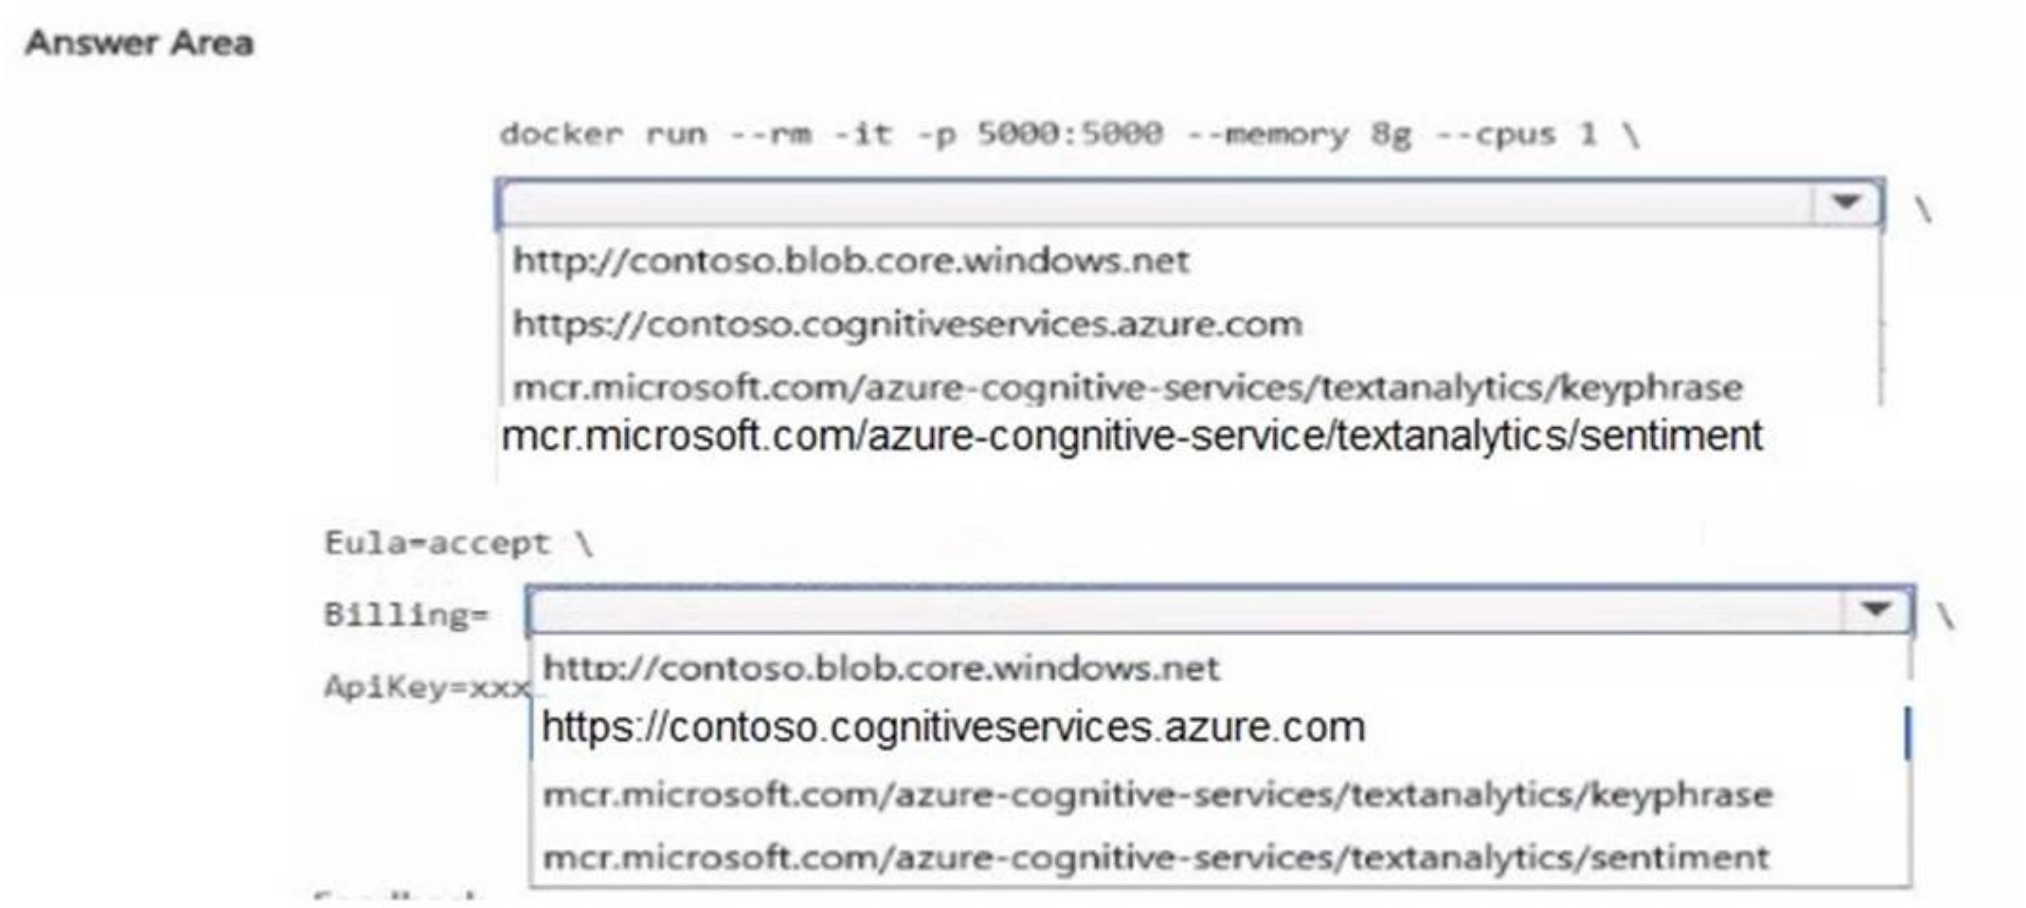

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

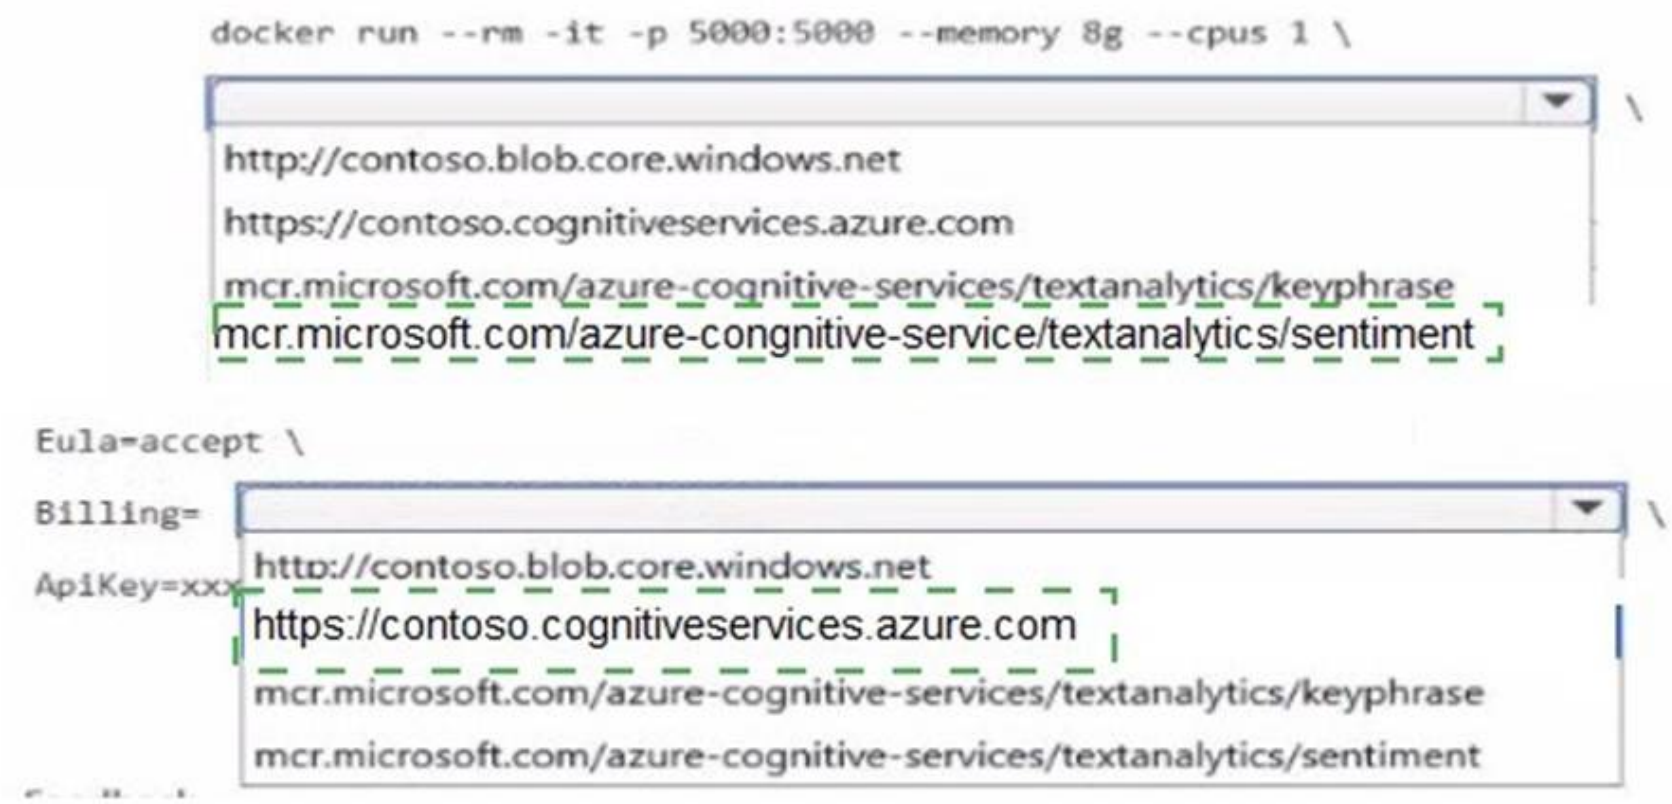

NEW QUESTION 82

.....

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