Microsoft

Exam Questions AI-103

Developing AI Apps and Agents on Azure

NEW QUESTION 1
- (Topic 1)
You need to configure Agent1 to meet the security and compliance requirements. What should you use?

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

Answer: B

NEW QUESTION 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 recommend an invoice review solution that resolves the issue reported by the finance department.
What should you include in the recommendation?

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

Answer: A

NEW QUESTION 4
- (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 5
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.                                    (
    model=deployment_name
    prompt="A video of ou
)
while video.status not in ["completed", "failed", "cancelled"]:
    time.sleep(20)
    video = client.videos.                                (video.id)
print(video.status)
```

create
download_content
list
retrieve

A. Mastered
B. Not Mastered

Answer: A

Explanation:

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 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 8
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.

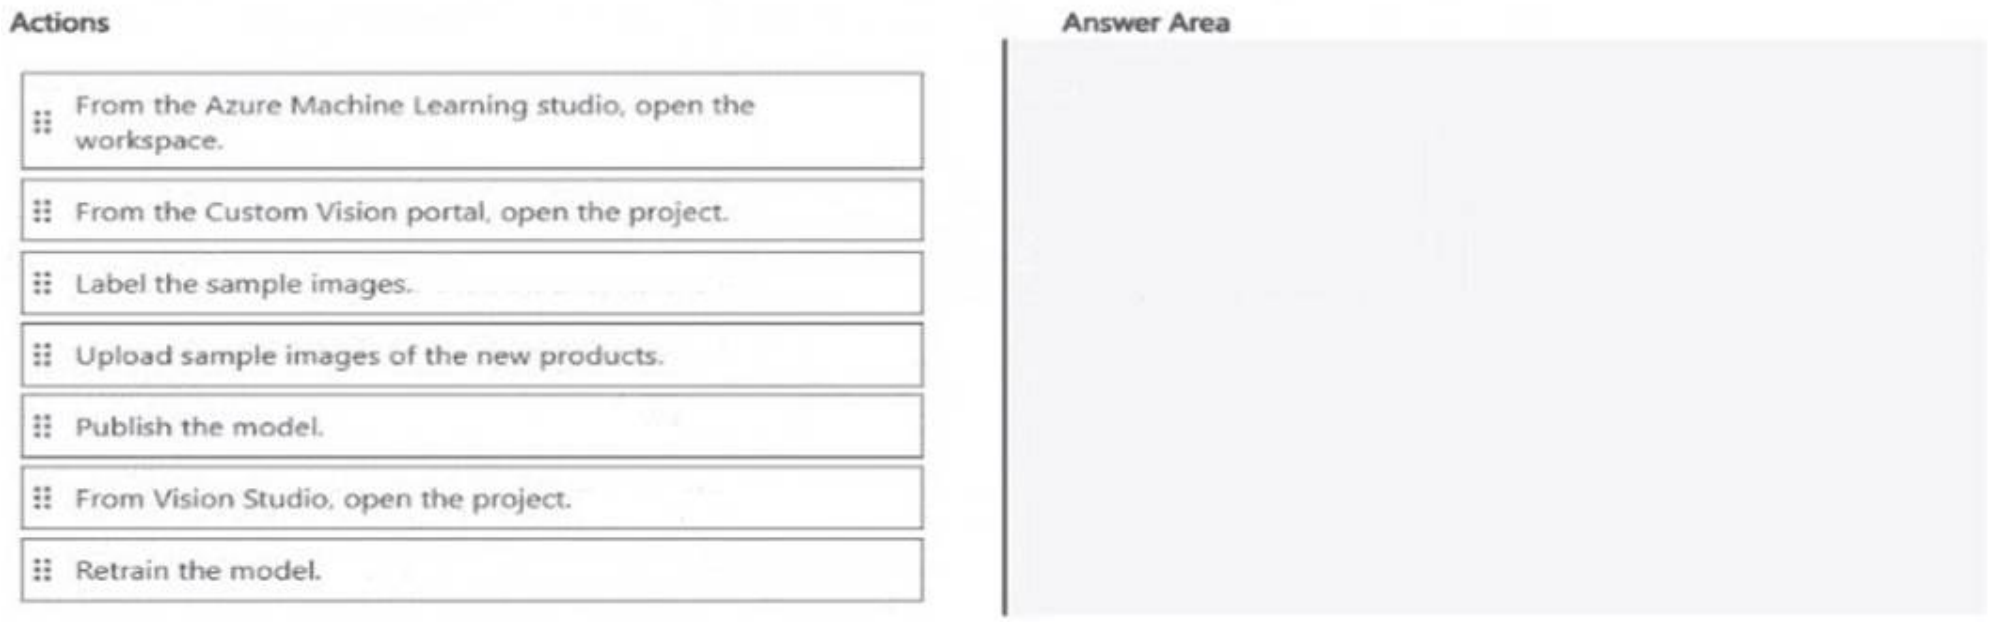

A. Mastered
B. Not Mastered

Answer: A

Explanation:

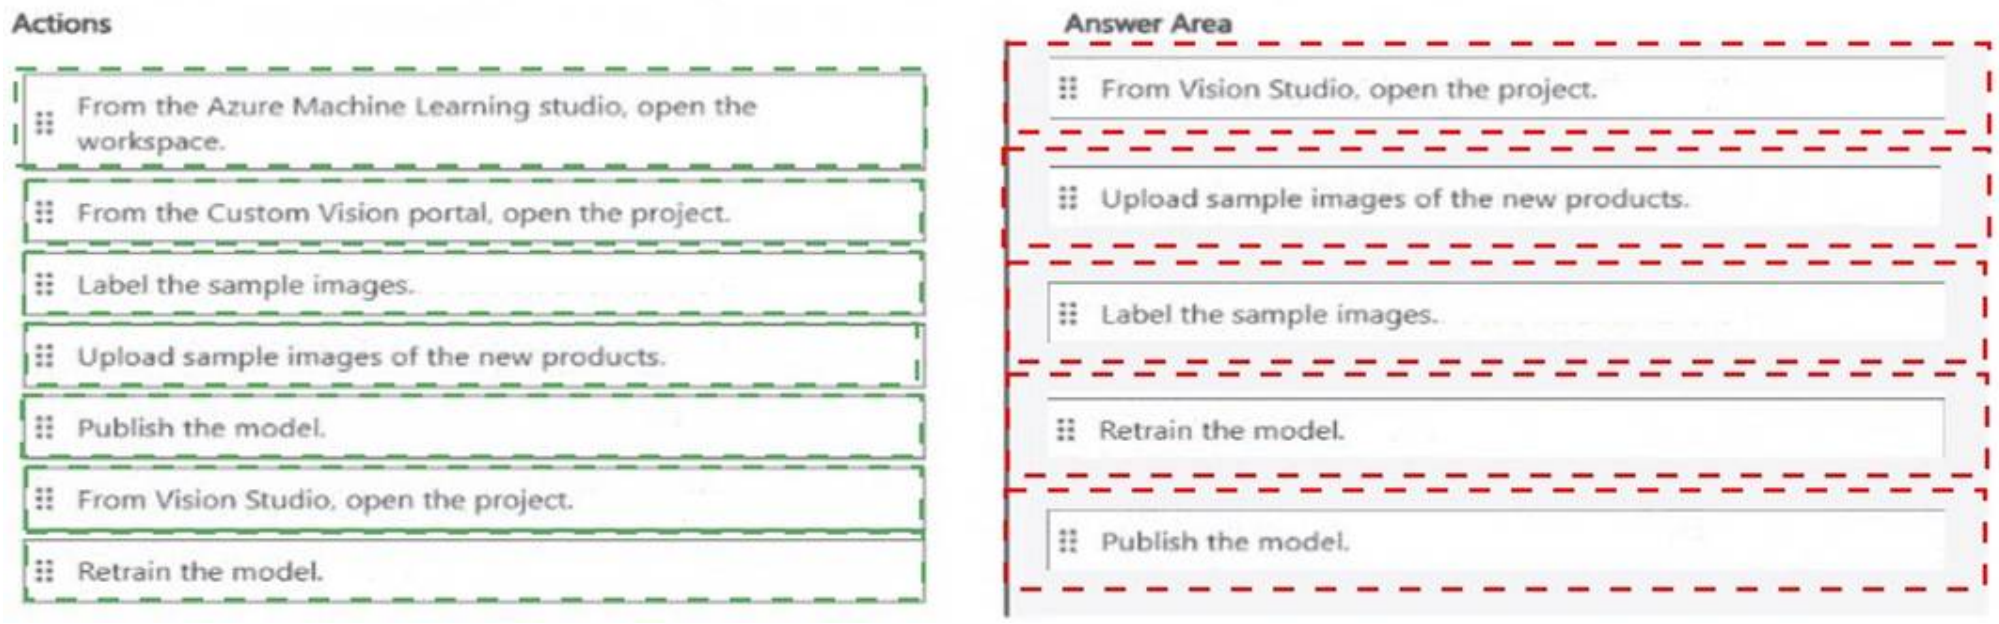

NEW QUESTION 9
- (Topic 2)
You have an Azure Speech in Foundry Tools resource that hosts a custom speech to text model deployed to a custom endpoint. An agent uses the endpoint to perform real-time speech recognition.
You are approaching the expiration date of the custom speech to text model. What is 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)
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 10
- (Topic 2)
You have a Microsoft Foundry project that contains an agent and an image generation model deployment.
The agent generates original images from user-supplied product photos.
You need to ensure that the generated images maintain the product identity and visual characteristics of the provided photo.
What should you do?

A. Set the input_fidelity parameter to high.
B. Apply a groundedness detection filter.
C. Include a prompt and input image in the request.
D. Decrease the value of the temperature parameter.

Answer: A

NEW QUESTION 12
- (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 17
HOTSPOT - (Topic 2)
You have an Azure AI Search resource named Search1 that is used by multiple apps hosted in Azure. You need to secure Search1. The solution must meet the following requirements:
• Prevent access to Search1 from the internet.
• Limit the access of each app to query specific indexes.
What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct answer is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

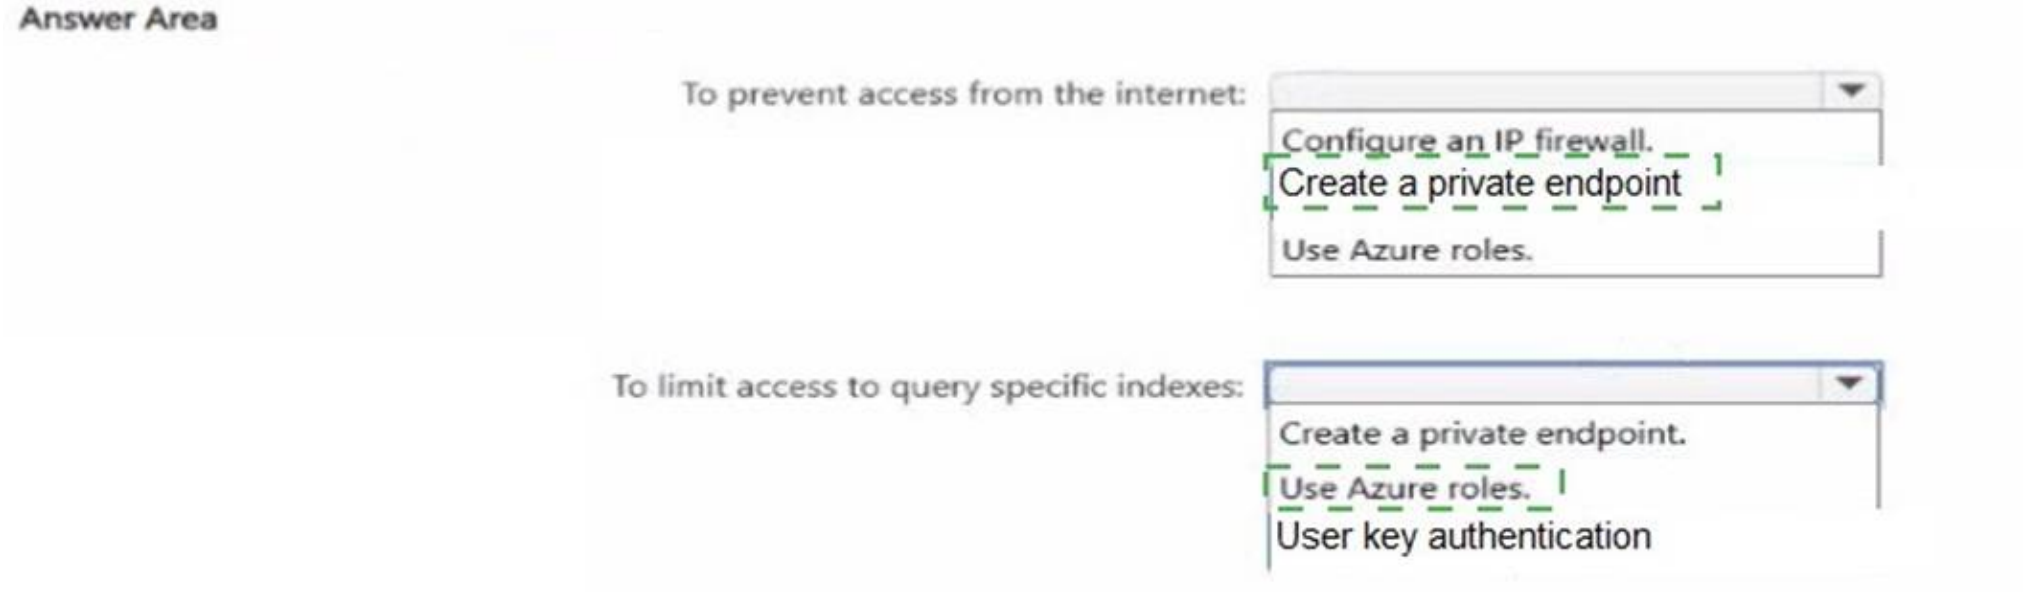

NEW QUESTION 22
- (Topic 2)
Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals.
More than one solution in the set might solve the problem. 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 image moderation to block unsafe content before processing the images.
Does this meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 23
- (Topic 2)
You have a Microsoft Foundry project that contains a model deployment.
You have an application that calls the deployment by using the Azure OpenAI v1 API and DefaultAzureCredential.
The developers at your company receive HTTP 403 errors when they send inference requests, even after running az login.
You need to ensure that the developers can perform model inference. The solution must follow the principle of least privilege.

Which role-based access control (RBAC) role should you assign to the developers?

A. Cognitive Services OpenAI User
B. Cognitive Services Data Reader
C. Cognitive Services User
D. Contributor

Answer: A

NEW QUESTION 24
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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

```
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. [          ] (
        model="trail-guide-chat",
        input="Create a 3-day hiking itinerary near Seattle.",
    )
    print(response.output_text)
```

- compact
- create
- retrieve

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.

To retain user preferences across conversations, use: [ ]
- Agent memory that uses persistent storage
- Conversation history
- Orchestration-managed session context

To enable users to provide contextual grounding during chats, use the: [ ]
- Azure AI Search tool
- Code interpreter tool
- File search tool

A. Mastered

B. Not Mastered

Answer: A

Explanation:

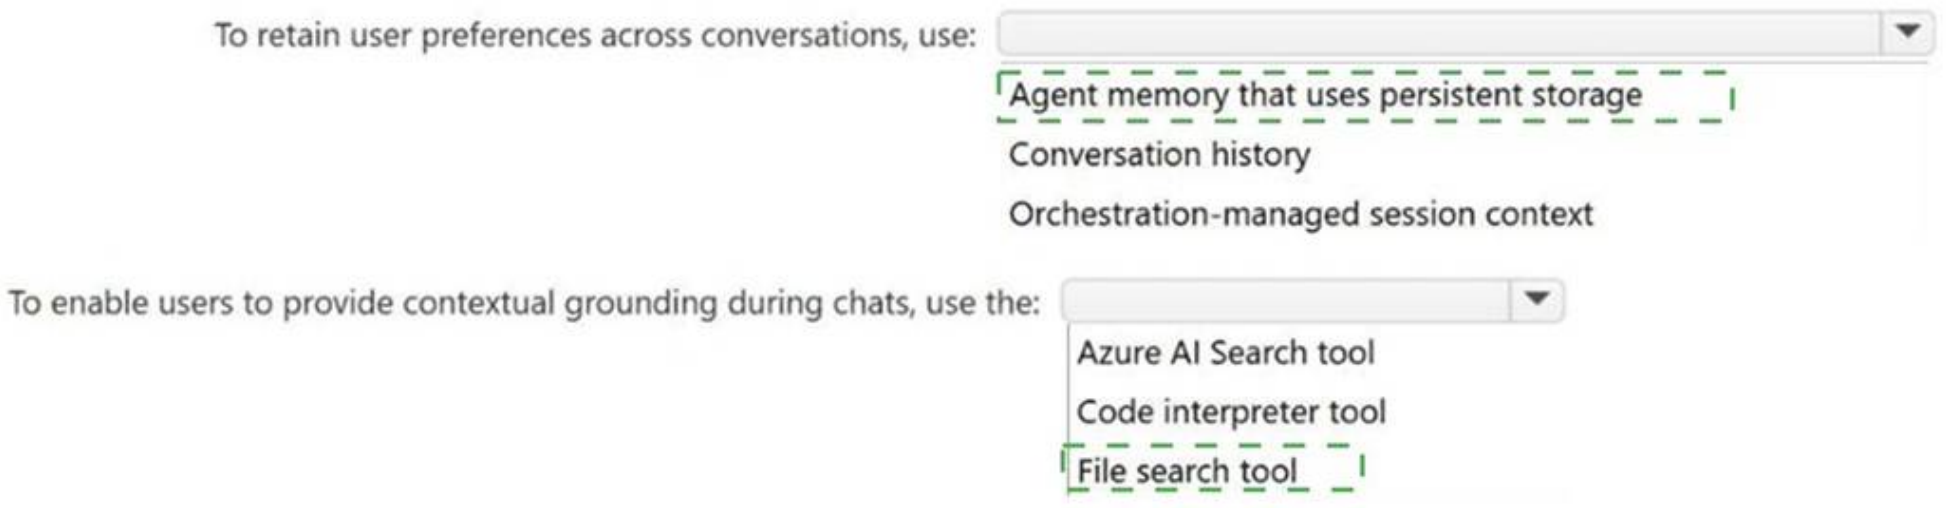

NEW QUESTION 30
- (Topic 2)
Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals.
More than one solution in the set might solve the problem. 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)
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 35
- (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 38
- (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 43
- (Topic 2)
You have a Microsoft Foundry project that contains an agent named Agent1. Agent1 runs successfully, but Foundry Control Plane does NOT display values for error rates, runs, and token usage, and the Traces tab is empty. You need to ensure that Foundry Control Plane displays the appropriate values for Agent1.

A. Restart Agent1 from Foundry Control Plane.
B. Update Agent1 to a new version.
C. Assign a Log Analytics workspace to Agent1.
D. Enable Application Insights for Agent1.

Answer: D

NEW QUESTION 44
- (Topic 2)
You have a Microsoft Foundry project named Project1 that contains the following:
• An OpenAPI tool that calls an external API
• A project connection named Connection1 that stores the API key of the external API When an agent calls the OpenAPI tool, the API returns a 401 unauthorized error, and traces show that the API key header is NOT being sent.
You need to ensure that the OpenAPI tool automatically includes the API key from Connection1 on all requests.
What should you do?

A. Configure the tool to use the default connection of Project1.
B. Connect the tool to Connection1.
C. Enable identity passthrough so that the tool uses the Microsoft Entra token of the caller.
D. Add the API key header manually to the OpenAPI specification.

Answer: B

NEW QUESTION 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 multimodal Al 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 52
- (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 56
- (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 60
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.

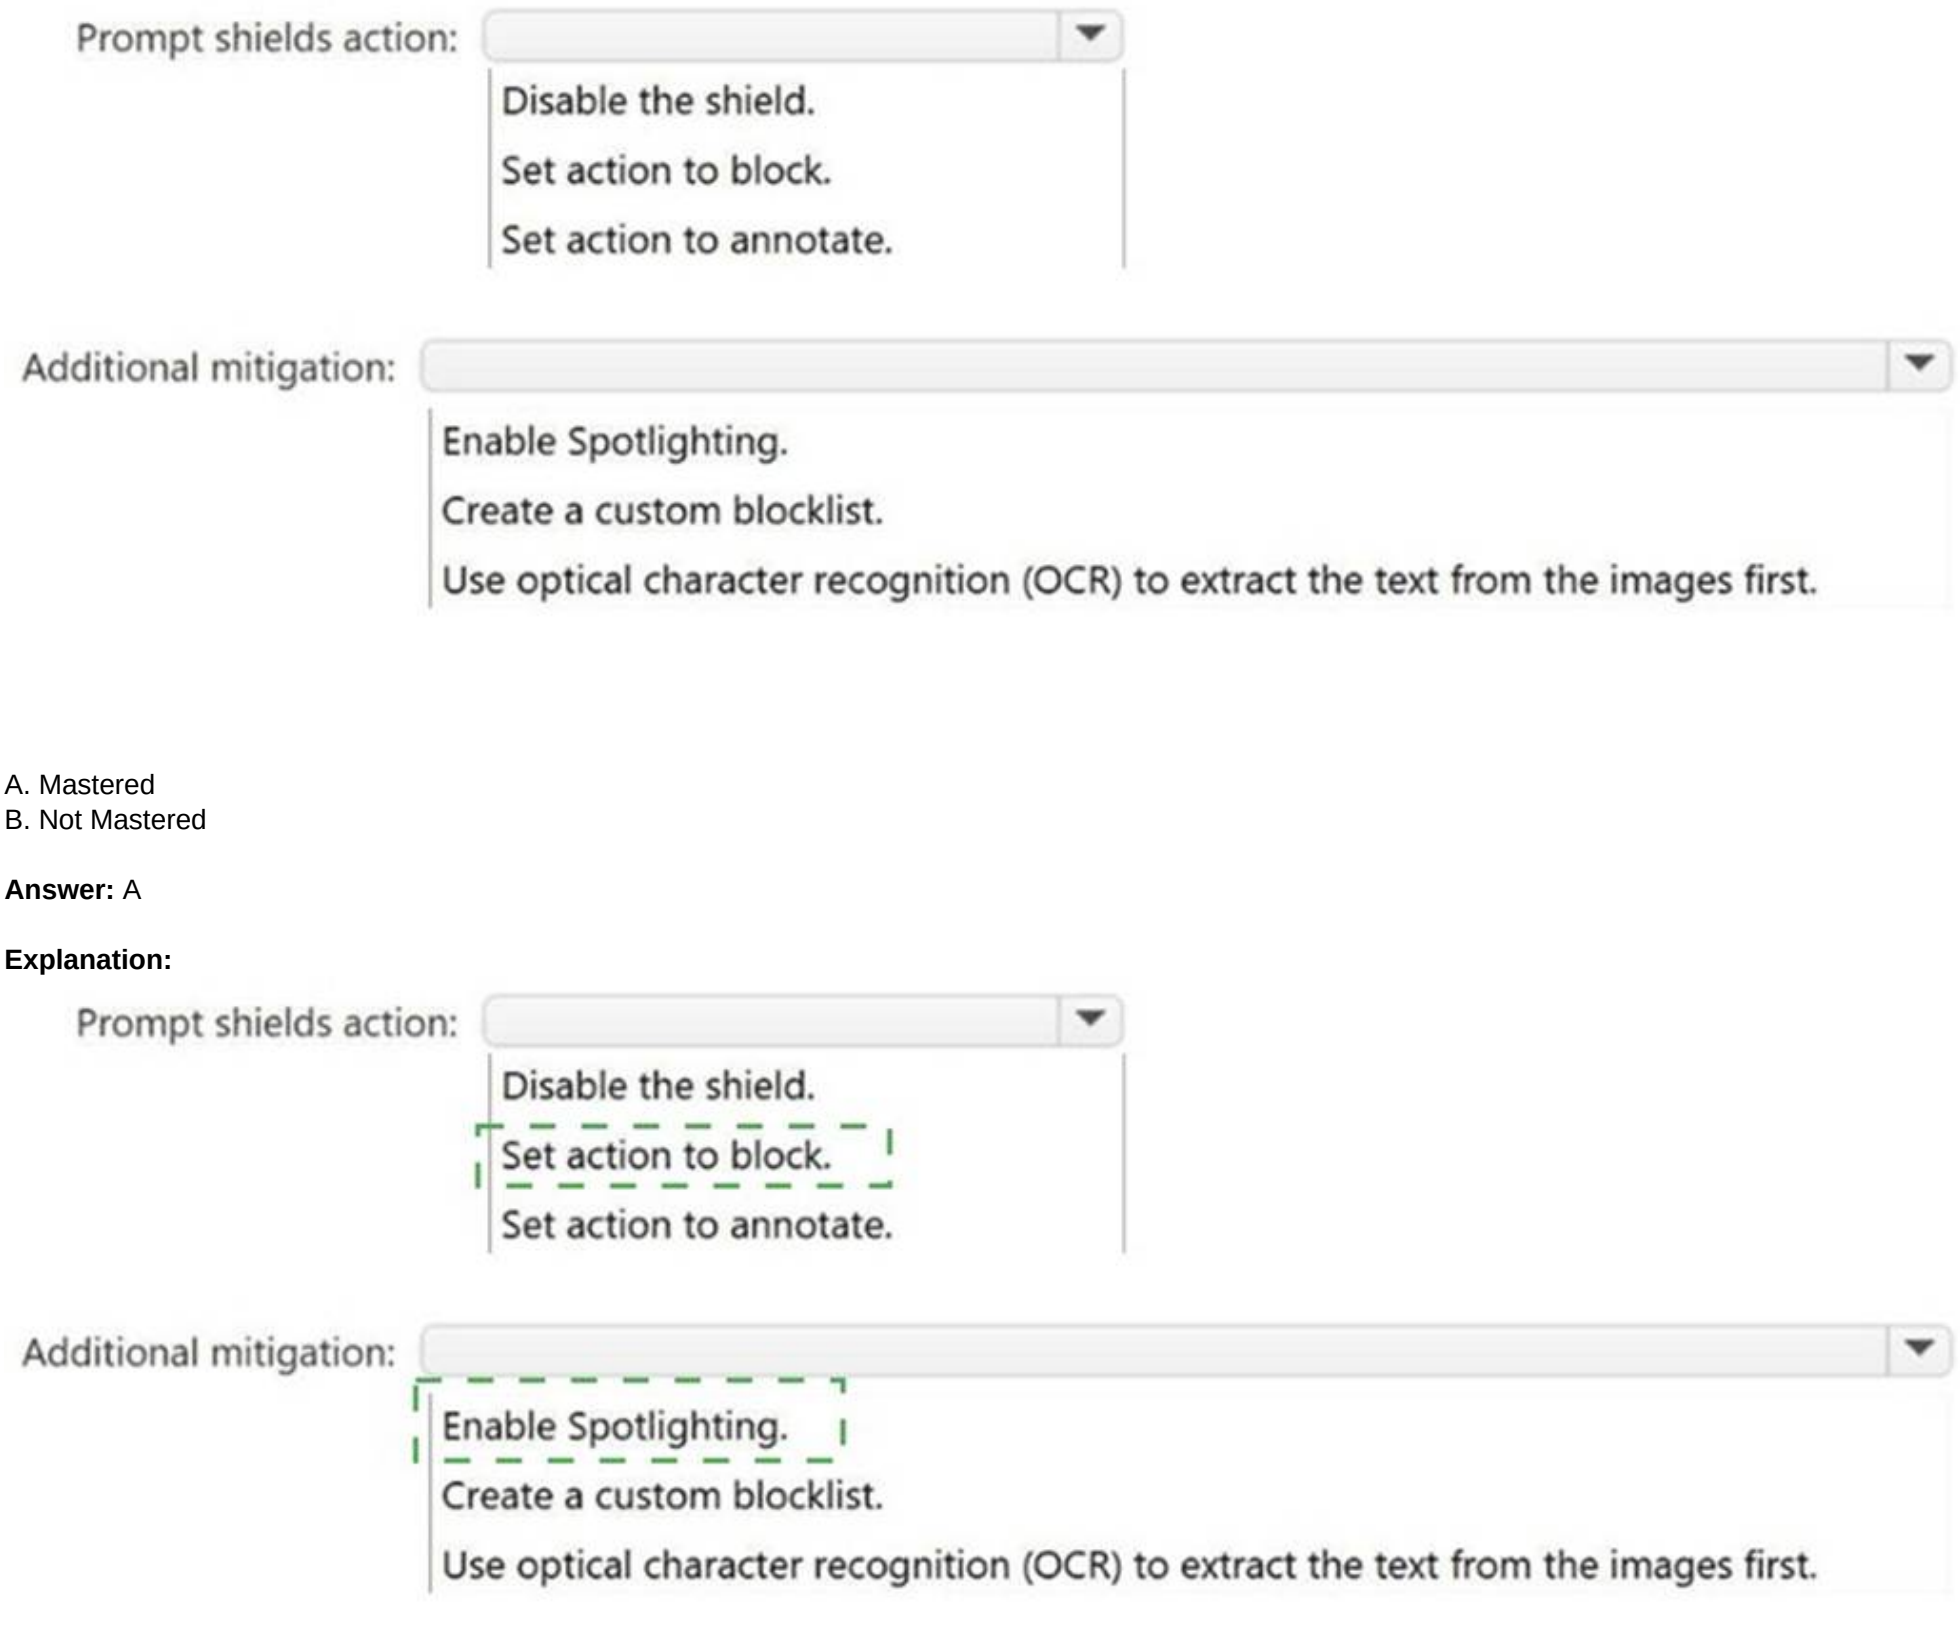

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 63
.....

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