

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 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 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 3

- (Topic 2)

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

? Searchable text fields for product names and product codes.

? A vector field containing embeddings for product descriptions.

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

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

Answer: B

NEW QUESTION 4

- (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 5

- (Topic 2)

You have a Microsoft Foundry project that contains a prompt agent used by a customer support web app.

The agent is invoked from a Python service that does NOT run in the Foundry portal.

You need to implement end-to-end tracing to capture latency breakdowns and exceptions across agent runs.

Which two components can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. a Log Analytics workspace
- B. OpenTelemetry
- C. Application Insights
- D. the Azure Monitor Agent
- E. Microsoft Sentinel

Answer: BC

NEW QUESTION 6

- (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 7

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 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.

Actions

From the Azure Machine Learning studio, open the workspace.

From the Custom Vision portal, open the project.

Label the sample images.

Upload sample images of the new products.

Publish the model.

From Vision Studio, open the project.

Retrain the model.

Answer Area

B. Not Mastered

Answer: A

Explanation:

Actions

- From the Azure Machine Learning studio, open the workspace.
- From the Custom Vision portal, open the project.
- Label the sample images.
- Upload sample images of the new products.
- Publish the model.
- From Vision Studio, open the project.
- Retrain the model.

Answer Area

- From Vision Studio, open the project.
- Upload sample images of the new products.
- Label the sample images.
- Retrain the model.
- Publish the model.

NEW QUESTION 9

- (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 10

- (Topic 2)

You have a Microsoft Foundry project that generates product marketing images from text prompts.

After publishing several images, the legal team at your company identifies a competitor's logo on a sign in the background of an image.

You need to remove only the logo, while preserving the rest of the image. What should you do?

- A. Increase the prompt guidance strength.
- B. Modify the original prompt to exclude brand names.
- C. Apply a mask-based inpainting edit to the part of the image that contains the logo.
- D. Rerun the prompt by using a different random seed.

Answer: C

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 uses Azure AI Search to ground an agent in internal documentation.

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

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

Which observability signal should you review?

- A. prediction drift metrics
- B. groundedness evaluation metrics
- C. latency breakdown traces
- D. indexer status and failure history

Answer: B

NEW QUESTION 14

DRAG DROP - (Topic 2)

You have a Microsoft Foundry project that contains an agent.

You need to enable long-term memory to ensure that the agent can recall user preferences across separate conversations. Stored memories must be isolated per authenticated user without the client application manually generating user IDs.

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.

NOTE: Each correct selection is worth one point.

Values

⋮ "session"

⋮ "{{\${conversationId}}}"

⋮ "{{\${userId}}}"

⋮ [mem_store_name]

⋮ [memory_tool]

⋮ MemorySearchTool("support_mem_store")

Answer Area

```
from azure.ai.projects.models import MemorySearchTool, PromptAgentDefinition

mem_store_name = "agent_mem_store"

memory_tool = MemorySearchTool(

    memory_store_name=mem_store_name,

    scope= 

Value

 ,

)

agent_def = PromptAgentDefinition(

    model="gpt-5.2",

    instructions="You are a customer support assistant.",

    tools= 

Value



)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Values

⋮ "session"

⋮ "{{\${conversationId}}}"

⋮ "{{\${userId}}}"

⋮ [mem_store_name]

⋮ [memory_tool]

⋮ MemorySearchTool("support_mem_store")

Answer Area

```
from azure.ai.projects.models import MemorySearchTool, PromptAgentDefinition

mem_store_name = "agent_mem_store"

memory_tool = MemorySearchTool(

    memory_store_name=mem_store_name,

    scope= 

⋮ "{{${userId}}}"

 ,

)

agent_def = PromptAgentDefinition(

    model="gpt-5.2",

    instructions="You are a customer support assistant.",

    tools= 

⋮ [memory_tool]



)
```

NEW QUESTION 19

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.

Answer Area

To prevent access from the internet:

Configure an IP firewall.
Create a private endpoint
Use Azure roles.

To limit access to query specific indexes:

Create a private endpoint.
Use Azure roles.
User key authentication

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Answer Area

To prevent access from the internet:

Configure an IP firewall.
Create a private endpoint
Use Azure roles.

To limit access to query specific indexes:

Create a private endpoint.
Use Azure roles.
User key authentication

NEW QUESTION 23

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 27

- (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 30

- (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 35

HOTSPOT - (Topic 2)

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

You need to ensure that App1 meets the following requirements:


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

How should you complete the Python code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

credential =  ()


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)
```

- A. Mastered
- B. Not Mastered


Answer: A

Explanation:

```

from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

credential =  ()


project_client = AIProjectClient(

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

    credential=credential,

)

with project_client.get_openai_client() as openai_client:

    response = openai_client.responses.  (

        model="trail-guide-chat",

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

    )

    print(response.output_text)

```

NEW QUESTION 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 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 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 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 43

- (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 46

- (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 47

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:

Extracted text data:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 52

- (Topic 2)

You have a Microsoft Foundry project that serves a high-volume chat app. Most requests are simple FAQs, but some require advanced reasoning. You need to reduce costs and latency for common queries, without degrading the quality of the responses to complex questions. What should you do?

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

Answer: D

NEW QUESTION 55

- (Topic 2)

You have a Microsoft Foundry project that contains an agent. The agent uses Azure AI Search for Retrieval Augmented Generation (RAG). You plan to ingest and index PDF product manuals. You need to build a solution that supports semantic similarity matching. The solution must ensure that the agent retrieves relevant data when user questions use different wording than the product manuals.

- A. vector search
- B. semantic ranking
- C. suggesters
- D. analyzers

Answer: A

NEW QUESTION 57

- (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 60

- (Topic 2)

You are building an app that will share user images. You need to configure the app to:

? Categorize each image as a photograph or drawing.

? Generate a caption for the image.

? Minimize development effort.

Which two services should you include?

- A. Image type detection in Azure Vision in Foundry Tools
- B. Object detection in Azure Vision in Foundry Tools
- C. Content tags in Azure Vision in Foundry Tools
- D. Image classification in Azure Custom Vision
- E. Image descriptions in Azure Vision in Foundry Tools

Answer: AE

NEW QUESTION 63

- (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 68

- (Topic 2)

You have an agent named Agent1 that uses Model Context Protocol (MCP) calls to retrieve external data. You need to implement guardrails to ensure that Agent1 cannot send any content tagged as Violence to the MCP server. Which intervention point should you use?

- A. output
- B. user input
- C. tool call
- D. tool response

Answer: C

NEW QUESTION 69

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an agent.


The agent uses tools to retrieve internal content and call external APIs. The agent is configured to let the model decide when to call the tools.

You need to publish the agent for a compliance workflow. The solution must meet the following requirements:


- Each workflow run must include a retrieval step before generating a response.
- Tool calls must authenticate by using the published agent's own identity.
- Tool access must use an identity isolated from other project resources.
- Tool access must support audit tracing.

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

Set tool_choice to:



Configure the tool to authenticate by:




- A. Mastered
- B. Not Mastered


Answer: A

Explanation:

Set tool_choice to:



Configure the tool to authenticate by:



NEW QUESTION 70

- (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 75

- (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 77

- (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 82

- (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

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains two agents named PolicyWriter and RiskReviewer. PolicyWriter generates draft updates for customer policies, and RiskReviewer reviews the drafts. In the visual builder, you need to create a workflow that meets the following requirements:

- Finalizes low-risk updates without manual intervention
- Ensures predictable execution across the agents

Answer Area

Orchestration pattern:

The sequential template that passes outputs node-by-node

The group chat template to dynamically route control between the agents

The human-in-the-loop template that pauses execution of the workflow for input

Approval checkpoints:

Add a Basic chat node.

Add a condition statement.

Add an Ask a question node.

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Orchestration pattern:

The sequential template that passes outputs node-by-node

The group chat template to dynamically route control between the agents

The human-in-the-loop template that pauses execution of the workflow for input

Approval checkpoints:

Add a Basic chat node.

Add a condition statement.

Add an Ask a question node.

NEW QUESTION 91

- (Topic 2)

You have a Microsoft Foundry project named Project1 that contains an agent. The agent uses an OpenAPI 3.0 specification to call an external weather service. The weather service requires a key to be passed in an HTTP header. The key value is stored as a connection in Project1.

You need to ensure that the key value from the connection is included automatically whenever the OpenAPI tool is invoked.

What should you configure in the OpenAPI specification?

- A. an Azure Key Vault connection
B. a header parameter defined for each operation
C. an API key security scheme
D. a Bearer token security scheme

Answer: C

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 95

HOTSPOT - (Topic 2)

You plan to deploy a containerized version of an Azure AI Language service that will be used for sentiment analysis. You configure `https://contoso.cognitiveservices.azure.com` 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.

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 106

.....

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