

Microsoft

Exam Questions AI-103

Developing AI Apps and Agents on Azure



NEW QUESTION 1

- (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 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 Agent1Dev Team can access Agent1. The solution must meet the security and compliance requirements.

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.

Answer Area

```

from azure.identity import DefaultAzureCredential
from azure.ai.projects import AIProjectClient
from azure.core.credentials import AzureKeyCredential
myEndpoint = "https://contoso.services.ai.azure.com/api/projects/project1"
project_client = AIProjectClient(
    endpoint=myEndpoint,
    credential=  ,
)
myAgent = "Agent1"
agent = project_client.agents.  (agent_name=myAgent)
print(f"Retrieved agent: {agent.name}")

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

from azure.core.credentials import AzureKeyCredential

myEndpoint = "https://contoso.services.ai.azure.com/api/projects/project1"

project_client = AIProjectClient(
    endpoint=myEndpoint,
    credential=  ,
    

AzureKeyCredential()
        DefaultAzureCredential()
        None


)

myAgent = "Agent1"

agent = project_client.agents.  (agent_name=myAgent)

print(f"Retrieved agent: {agent.name}")
```

create_version
 get
 get_version

NEW QUESTION 5

- (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 6

- (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 7

- (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 8

- (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 9

- (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

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	Answer Area
⋮ "auto"	<pre>run_payload = { "assistant_id": agent_id, "tool_choice": "required", "metadata": { "scenario": "ticket-triage" } }</pre>
⋮ "required"	
⋮ "response_format"	
⋮ "tool_choice"	
⋮ "tools"	
⋮ "type"	

NEW QUESTION 10

- (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 12

- (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 15

- (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 20

- (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 22

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 27

- (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 32

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 33

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 34

- (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 39

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:

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

NEW QUESTION 41

- (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 44

- (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 45

- (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 49

- (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 51

- (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 54

- (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 58

- (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 59

HOTSPOT - (Topic 2)

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

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

NOTE: Each correct selection is worth one point.

Answer Area

JSON data:

File projectionObject protectionTable projection

Extracted text data:

File protectionObject protectionTable projection

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

JSON data:

File projectionObject protectionTable projection

Extracted text data:

File protectionObject protectionTable projection

NEW QUESTION 64

- (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 65

- (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 70

- (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 73

- (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 76

- (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 79

- (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 81

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 85

- (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 90

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.

Answer Area

```
docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \
```

<http://contoso.blob.core.windows.net>
<https://contoso.cognitiveservices.azure.com>
<mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase>
<mcr.microsoft.com/azure-cognitive-service/textanalytics/sentiment>

Eula=accept \

Billing=

ApiKey=xxx

<http://contoso.blob.core.windows.net>
<https://contoso.cognitiveservices.azure.com>
<mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase>
<mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \
```

<http://contoso.blob.core.windows.net>
<https://contoso.cognitiveservices.azure.com>
<mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase>
<mcr.microsoft.com/azure-cognitive-service/textanalytics/sentiment>

Eula=accept \

Billing=

ApiKey=xxx

<http://contoso.blob.core.windows.net>
<https://contoso.cognitiveservices.azure.com>
<mcr.microsoft.com/azure-cognitive-services/textanalytics/keyphrase>
<mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment>

NEW QUESTION 95

- (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 100

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions 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 First Try
- * AI-103 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AI-103 Practice Test Here](#)