


AI-103 Dumps

Developing AI Apps and Agents on Azure

<https://www.certleader.com/AI-103-dumps.html>



NEW QUESTION 1

- (Topic 1)

You need to configure Agent1 to answer customer questions about only the Contoso products. The solution must meet the business requirements. What should you do?

- A. Apply top-p sampling.
- B. Modify the system message instructions.
- C. Add few-shot examples.
- D. Increase the value of the temperature parameter.

Answer: B

NEW QUESTION 2

- (Topic 1)

You need to configure 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 3

- (Topic 1)

You need to configure an indexing pipeline for Agent1 to retrieve the relevant product information in storage1. The solution must meet the technical requirement. Which two built-in skills should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Language Detection
- B. Entity Recognition
- C. Merge
- D. Azure OpenAI Embedding
- E. Text Split
- F. key phrase extraction

Answer: DE

NEW QUESTION 4

- (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 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.create(
    model=deployment_name,
    prompt="A video of our product",
)
while video.status not in ["completed", "failed", "cancelled"]:
    time.sleep(20)
    video = client.videos.retrieve(video.id)
print(video.status)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
import time
from openai import OpenAI
client = OpenAI(
    base_url="https://Contoso.openai.azure.com/openai/v1/",
    api_key=...,
)
video = client.videos.create(
    model=deployment_name,
    prompt="A video of our product",
)
while video.status not in ["completed", "failed", "cancelled"]:
    time.sleep(20)
    video = client.videos.retrieve(video.id)
print(video.status)
```

NEW QUESTION 6

- (Topic 2)

You are building an app that will use Azure AI to monitor workspaces for safety. You need to recommend a service that meets the following requirements:

- ? Generates alerts when employees enter high-risk areas
 - ? Monitors video feeds in real time
 - ? Minimizes development effort
- What should you recommend?

- A. Azure Vision in Foundry Tools Image Analysis
- B. Azure AI Video Indexer
- C. Azure Vision in Foundry Tools Spatial Analysis
- D. Object detection in Azure Custom Vision

Answer: C

NEW QUESTION 7

DRAG DROP - (Topic 2)

You have a Microsoft Foundry project that contains an agent used by the financial analysts at your company.

You need to optimize the agent workflow by providing additional data access and processing capabilities. The solution must meet the following requirements:

- Ensure that the agent can perform calculations during conversations
- Ensure that the agent can access up-to-date information from public websites.
- Ensure that the agent can retrieve information from documents uploaded directly to the agent.

What should you use for each requirement? To answer, drag the appropriate tools to the correct requirements. Each tool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Tools

⋮ Code interpreter

⋮ Computer use

⋮ File search

⋮ Grounding with Bing Search

⋮ Microsoft Fabric

Answer Area

Access up-to-date information from public websites:

Tool

Perform calculations during conversations:

Tool

Retrieve information from documents uploaded directly to the agent:

Tool

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Tools

⋮ Code interpreter

⋮ Computer use

⋮ File search

⋮ Grounding with Bing Search

⋮ Microsoft Fabric

Answer Area

Access up-to-date information from public websites

⋮ Grounding with Bing Search

Perform calculations during conversations

⋮ Code interpreter

Retrieve information from documents uploaded directly to the agent

⋮ File search

NEW QUESTION 8

- (Topic 2)

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 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 add a reflection pass that regenerates the response if the required clauses are missing.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 10

HOTSPOT - (Topic 2)

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 10

- (Topic 2)

You have an Azure Speech in Foundry Tools resource that hosts a custom speech to text model deployed to a custom endpoint. An agent uses the endpoint to perform real-time speech recognition.

You are approaching the expiration date of the custom speech to text model. What is the expected behavior when the model expires?

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

D. The custom model will be deleted automatically when the model expires.

Answer: A

NEW QUESTION 15

- (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 20

- (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 21

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.

Guardrails:

Select Tool call and set Action to Block.

Select User input and Output and set Action to Annotate.

Select User input and Tool response and set Action to Annotate.

Select User input, Output, Tool response, and Tool call and set Action to Block.

Storage access:

Storage account access keys

A user-assigned identity that is assigned the Storage Queue Data Contributor role

A system-assigned managed identity that is assigned the Storage Blob Data Reader role

A system-assigned managed identity that is assigned the Storage Blob Data Contributor role

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Guardrails:

Select Tool call and set Action to Block.

Select User input and Output and set Action to Annotate.

Select User input and Tool response and set Action to Annotate.

Select User input, Output, Tool response, and Tool call and set Action to Block.

Storage access:

Storage account access keys

A user-assigned identity that is assigned the Storage Queue Data Contributor role

A system-assigned managed identity that is assigned the Storage Blob Data Reader role

A system-assigned managed identity that is assigned the Storage Blob Data Contributor role

NEW QUESTION 24

- (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 26

- (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 29

HOTSPOT - (Topic 2)

You have a Microsoft Foundry resource named AI1 that hosts three deployments of the GPT 3.5 model. Each deployment is optimized for a unique workload.

You plan to deploy three apps. Each app will access AM by using the REST API and will use the deployment that was optimized for the apps intended workload.

You need to provide each app with access to AI1 and the appropriate deployment. The solution must ensure that only the apps can access AI1.

What should you use to provide access to AI1. and what should each app use to connect to its appropriate deployment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Provide access to AI1 by using:

An API key

A bearer token

A shared access signature (SAS) token

Connect to the deployment by using:

An API key

A deployment endpoint

A deployment name

A deployment type

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 33

- (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 38

- (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 43

- (Topic 2)

You are creating an image-editing workflow in a Microsoft Foundry project. The workflow must meet the following requirements:

- Ensure that background objects can be removed by applying a mask-based inpainting edit.
- Preserve the original lighting and style of the edited images.
- Use the built-in image editing controls, NOT a custom model.

You need to ensure that image edits apply exclusively inside the masked area.

How should you configure the workflow?

- A. Enable text_to_image mode and a prompt describing the desired background removal.
- B. Set generation mode to image_variation and provide the original image as a reference.
- C. Enable image_to_image mode and a high-strength value to regenerate the full image based on the prompt.
- D. Enable mask_inpainting and supply both the input image and a mask indicating which part of the image to modify.

Answer: D

NEW QUESTION 44

- (Topic 2)

You have a Microsoft Foundry project that contains an agent.

You need to process mixed-format documents that contain scanned text, tables, and multicolumn layouts. The extracted content must preserve the document structure and be converted into the Markdown format for downstream reasoning.

What should you configure first?

- A. an Azure Content Understanding in Foundry Tools analyzer
- B. a generative chat completion request
- C. an Azure Language in Foundry Tools text analysis model deployment
- D. an Azure OpenAI Responses API call that uses a multimodal model

Answer: A

NEW QUESTION 47

- (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 51

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 55

- (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 56

- (Topic 2)

You are planning a Microsoft Foundry project named Project1 that will contain multiple agents. Each agent will access the same Azure AI Search resource. You need to recommend a solution to centrally manage the Azure AI Search credentials within Project1. The solution must be implemented across all the agents. What should you recommend?

- A. Enable role-based access control (RBAC) for the Azure AI Search resource.
- B. Add a connection to the Azure AI Search resource.
- C. Disable key-based access control on the Azure AI Search resource.
- D. Create a managed private endpoint that connects to the Azure AI Search resource.

Answer: B

NEW QUESTION 60

- (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 64

- (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 68

- (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 69

- (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 72

- (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 77

- (Topic 2)

You have an application named App1 that uses Azure Speech in Foundry Tools to transcribe live calls.

Transcript segments often contain both English and Spanish. App1 sends each segment to Azure Translator in Foundry Tools to translate to another language.

Sometimes, mixed-language segments result in incomplete or incorrect translations.

You need to reduce translation errors. The solution must ensure that the entire transcript is translated successfully.

What should you do before sending the segments to Translator?

- A. Specify English as the source language in the translation request for all the segments.
- B. Enable automatic language detection for the translation request.
- C. Split the mixed-language segments into single-language segments and translate each segment separately.
- D. Use document translation to translate the entire transcript as a single document.

Answer: C

NEW QUESTION 80

- (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 83

- (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 86

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 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 93

- (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 95

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AI-103 Exam with Our Prep Materials Via below:

<https://www.certleader.com/AI-103-dumps.html>