


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 recommend a solution to support the planned changes and technical requirements for Agent1 to use the product information stored in storage1.

What should you include in the recommendation?

- A. Azure AI Search
- B. Azure Translator in Foundry Tools
- C. Azure Document Intelligence in Foundry Tools
- D. Grounding with Bing Search

Answer: A

NEW QUESTION 2

- (Topic 1)

You need to 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 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

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:

Standard
Global Standard
Global Provisioned

Version update policy:

Once the current version expires
Opt out of automatic model version upgrades
Upgrade once a new default version becomes available

NEW QUESTION 5

HOTSPOT - (Topic 1)

You need to ensure that the marketing department can generate videos by using the model deployed to Project2. now should you complete the Python code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
import time

from openai import OpenAI

client = OpenAI(
    base_url="https://Contoso.openai.azure.com/openai/v1/",
    api_key=...,
)

video = client.videos.
    model=deployment_name
    prompt="A video of ou

)

while video.status not in ["completed", "failed", "cancelled"]:
    time.sleep(20)
    video = client.videos.
print(video.status)
```

create
download_content
list
retrieve

create
download_content
list
retrieve

- 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 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 an Azure subscription that contains a Microsoft Foundry resource. You need to build an app that will suggest product names from a given product description. Which Foundry model should you use?

- A. GPT-4
- B. DALL-E
- C. Whisper
- D. embeddings

Answer: A

NEW QUESTION 9

- (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 10

- (Topic 2)
You have a Microsoft Foundry project that contains an agent. The knowledge source for the agent is a set of scanned PDF troubleshooting guides stored in Azure Blob Storage. The guide pages contain two-column layouts and tables. You use Azure Content Understanding in Foundry Tools to process the PDFs. You plan to ingest the processed content into an index for Retrieval Augmented Generation (RAG) and store extracted fields for downstream automation. Stakeholders must be able to verify where each extracted field value came from in the original PDF and route low-reliability extractions for manual review. You need to ensure that the Content Understanding document analyzer output includes a per-field confidence score and source grounding locations within the source document. What should you do?

- A. Enable estimateFieldSourceAndConfidence.
- B. Configure the analyzer to use generative extraction for all fields.
- C. Set enableSegment to true.
- D. Provide labeled samples.

Answer: A

NEW QUESTION 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,  
        "schema": {  
            "type": "object",  
            "properties": {  
                "merchant": {"type": "string"},  
                "order_number": {"type": "string"},  
                "discount_code": {"type": "string"},  
                "line_items": {"type": "array", "items": {"type": "object", "properties": {"description": {"type": "string"}, "price": {"type": "number"}, "quantity": {"type": "number"}}, "required": ["description", "price", "quantity"]}, "required": ["merchant", "order_number", "discount_code", "line_items", "total"]}            }  
        }  
    }
```

Answer Area

Statements	Yes	No
The schema requires that a receipt has a discount_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 merchant, order_number, discount_code, line_items, and total.	☐	☐

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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 11

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 16

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

- A. Mastered
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 18

- (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 22

- (Topic 2)

Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals.
More than one solution in the set might solve the problem. It is also possible that none of the solutions in the set solve the problem. After you answer a question in this section, you will NOT be able to return. As a result, these questions do not appear on the Review Screen.
You have a 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 24

- (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 27

- (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 29

- (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 33

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:

The Leader of IT Certification

visit - <https://www.certleader.com>

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 35

- (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 39

- (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 42

- (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 45

- (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 50

- (Topic 2)
You have an Azure subscription that contains an Azure Language in Foundry Tools service resource. You need to identify the URL of the REST interface for the Language service. Which blade should you use in the Azure portal?

A. Networking
B. Keys and Endpoint
C. Identity
D. Properties

Answer: B

NEW QUESTION 52

- (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 54

- (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 55

DRAG DROP - (Topic 2)
You have a Microsoft Foundry project that contains a customer support agent grounded in internal documentation.
After a recent update, users report the following issues:

- Some answers are unsupported by retrieved documents.
- A small number of responses are flagged for policy violations. You need to evaluate each issue.

Which observability signals should you use for each issue? To answer, drag the appropriate observability signals to the correct issues. Each observability signal may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Observability signals

⋮ Groundedness evaluation metrics

⋮ Latency breakdown traces

⋮ Risk and safety metrics

⋮ Token usage analytics

Answer Area

Unsupported responses:

Observability signal

Policy violations:

Observability signal

A. Mastered
B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 59

- (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 61

- (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 63

- (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 65

- (Topic 2)

You have a Microsoft Foundry project that contains three agents as shown in the following table.

Name	Description
TriageAgent	Classifies incoming customer requests
PolicyAgent	Answers policy questions by searching internal content
ActionAgent	Creates or updates tickets by calling an HTTP API

You need to orchestrate the agents to ensure that the customer requests meet the following requirements:

- Support a deterministic, step-based process that uses conditional branching and shared state across the agents.
- Optionally trigger a ticket action based on the triage result. The solution must minimize development effort.

What should you include in the solution?

- A. separate agent runs coordinated in the application code
- B. a multi-agent group chat session
- C. a workflow
- D. threads and runs without a workflow

Answer: C

NEW QUESTION 70

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains a deployed chat model.


You have a Python service that sends API requests to the model. The service is integrated with an automated validation system that compares generated outputs against approved response patterns.

Stakeholders report that small wording differences are causing validation mismatches.

You need to update the request parameters to improve output stability. The solution must maximize reasoning quality.

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.

```
message = client.messages.create(  
  
    model="deployment-name",  
  
    messages=[  
  
        {"role": "user", "content": "Summarize the release notes in 3 bullet points."}  
  
    ],  
  
    max_tokens=800,  
  
    temperature= ,  
  
    thinking={"type": "enabled"},  
  
    output_config={"effort":  }  
  
)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
message = client.messages.create(
    model="deployment-name",
    messages=[
        {"role": "user", "content": "Summarize the release notes in 3 bullet points."}
    ],
    max_tokens=800,
    temperature=
        0
        1
        2
    thinking={"type": "enabled"},
    output_config={"effort":
        "high"
        "low"
        "medium"
    }
)
```

NEW QUESTION 75

- (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 78

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an agent.

The agent accepts user-uploaded screenshots and uses a multimodal chat model.

Some screenshots contain potentially malicious embedded text.

You need to prevent a prompt injection attack and ensure that third-party content is treated as lower trust.

How should you configure prompt shields for document attacks? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

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 83

.....

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>