


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 recommend a solution to assess the responses generated by Agent1 when the agent uses the product information stored in storage1. The solution must meet the technical requirements.

- A. a Retrieval Augmented Generation (RAG) evaluator
- B. a coherence evaluator
- C. a custom guardrail
- D. model fine-tuning

Answer: A

NEW QUESTION 3

- (Topic 1)

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

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

Answer: A

NEW QUESTION 4

- (Topic 1)

You need to configure 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 5

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 6

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 7

- (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 8

- (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 9

- (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 10

- (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 10

- (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 12

- (Topic 2)

You have a Microsoft Foundry project that contains a high-traffic agent. After a recent update, operational costs increase significantly.

Monitoring confirms that the volume of user traffic to the agent remains unchanged.

You suspect that changes to the request or response characteristics are causing the increase.

You need to identify whether the additional costs are driven by the model input size, the model output size, or expanded tool usage.

Which observability capability should you use?

- A. evaluation metrics
- B. token usage
- C. latency
- D. run success rate

Answer: B

NEW QUESTION 16

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

⋮ "auto"

⋮ "required"

⋮ "response_format"

⋮ "tool_choice"

⋮ "tools"

⋮ "type"

Answer Area

```
run_payload = {  
    "assistant_id": agent_id,  
    "tool_choice": "required",  
    "metadata": {  
        "scenario": "ticket-triage"  
    }  
}
```

NEW QUESTION 17

DRAG DROP - (Topic 2)

You have a Microsoft Foundry project that processes procurement documents submitted by suppliers.

You need to implement two pipelines by using Azure Content Understanding in Foundry Tools. The solution must meet the following requirements:

- Include a pipeline named Pipeline1 that supports cost-effective, high-volume processing of standalone PDF invoices.
- Include a pipeline named Pipeline2 that supports cross-document validation by using multi-step reasoning and reference data.

How should you configure each pipeline? To answer, drag the appropriate configurations to the correct pipelines. Each configuration 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.

Configurations

⋮ Multi-file task in pro mode

⋮ Multi-file task in standard mode

⋮ Single-file task in pro mode

⋮ Single-file task in standard mode

Answer Area

Pipeline1: Configuration

Pipeline2: Configuration

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Configurations

⋮ Multi-file task in pro mode

⋮ Multi-file task in standard mode

⋮ Single-file task in pro mode

⋮ Single-file task in standard mode

Answer Area

Pipeline1 ⋮ Single-file task in standard mode

Pipeline2 ⋮ Multi-file task in pro mode

NEW QUESTION 20

- (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 24

- (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 29

- (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 34

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 36

- (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 37

- (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 39

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 41

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 42

- (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 46

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 51
HOTSPOT - (Topic 2)
Your company is piloting a customer support agent in a Microsoft Foundry project name Project1. Project1 is connected to an existing Application Insights resource, and the company's support team reviews runs in the Traces tab. The Foundry Agent Service is configured to perform the following actions:

- Retrieve the Application Insights connection string by calling project_client.telemetry.get_application_insights_connection_string().
- Call configure_azure_monitor(connection_string=...) to enable telemetry.

A separate LangChain service configured to use OpenTelemetry and has the following configurations:

- Uses AzureAIOpenTelemetryTracer(connection_string=..., enable_content_recording=False)
- Passes the tracer by using config={"callbacks":[azure_tracer]}

Company policy has the following requirements:

- Telemetry from LangChain and OpenTelemetry must be distinguishable within the same Application Insights resource.
- Secrets and credentials must NOT be stored in prompts, tool arguments, or span attributes.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The LangChain service will appear in Traces without configuring a tracer.	☐	☐
Setting different OTEL_SERVICE_NAME values separates the services in Application Insights.	☐	☐
When using enable_content_recording=False, prompts and tool data will be captured in the telemetry.	☐	☐

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The LangChain service will appear in Traces without configuring a tracer.	☐	☒
Setting different OTEL_SERVICE_NAME values separates the services in Application Insights.	☒	☐
When using enable_content_recording=False, prompts and tool data will be captured in the telemetry.	☐	☒

NEW QUESTION 55
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 59

- (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 60

- (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 64

- (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 67

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 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 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 71

- (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 74

- (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 79

- (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 80

- (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 81

HOTSPOT - (Topic 2)

You have a configurable guardrail in a Microsoft Foundry project.

You have a Python application. A text blocklist named ConfidentialTerms already contains terms that must be detected in user messages. The application stores each incoming user message in a variable named input_text.

You plan to moderate input_text before the text is sent to a language model. You need to analyze input_text by using the existing blocklist.

Answer Area

```
...
blocklist_name = "ConfidentialTerms"
client = ContentSafetyClient(endpoint, credential)
request =
    text=input_
    .. ...
    AddOrUpdateTextBlocklistItemsOptions
    AnalyzeImageOptions
    AnalyzeTextOptions
    TextBlocklist
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
...
blocklist_name = "ConfidentialTerms"
client = ContentSafetyClient(endpoint, credential)
request =
    text=input_
    .. ..
    AddOrUpdateTextBlocklistItemsOptions
    AnalyzeImageOptions
    AnalyzeTextOptions
    TextBlocklist
```

NEW QUESTION 85

- (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 90

- (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 93

- (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 98

- (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 100

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 104

- (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

.....

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>