

Microsoft

Exam Questions AI-103

Developing AI Apps and Agents on Azure



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 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 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 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

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an internal Q&A agent. Users report the following issues when they ask the agent questions:

- An increase in the following response: ??No relevant information found??
- Periodic HTTP 429 rate limit exceeded errors during peak hours

You need to identify whether each issue is caused by model unavailability, resource limits, or inference failures.

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

Metrics to enable:

Model Availability Rate and Provisioned Utilization
 Only Tokens Cache Match Rate
 Only Total Requests filtered to status code 200
 Time To Response and Total Tokens

Diagnostic log to collect:

AllMetrics
 audit
 RequestResponse
 trace

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Metrics to enable:

- Model Availability Rate and Provisioned Utilization
- Only Tokens Cache Match Rate
- Only Total Requests filtered to status code 200
- Time To Response and Total Tokens

Diagnostic log to collect:

- AllMetrics
- audit
- RequestResponse
- trace

NEW QUESTION 8

- (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 9

- (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 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

- (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 13

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 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:

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.	☒	☐

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 21

- (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 26

- (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 27

- (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 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 31

- (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 33

- (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 35

- (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 38

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 41

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 42

- (Topic 2)

You have a Microsoft Foundry project that contains an agent. The agent uses two tools to perform the following actions:

? Use Azure AI Search to retrieve answers from a private product documentation index.

? Use the web search tool to retrieve public information on the internet.

You need to ensure that for a specific run, the agent deterministically retrieves information only from the internet. To what should you set tool_choice?

- A. "required"
- B. {"type": "azure_ai_search"}
- C. "auto"
- D. {"type": "bing_grounding"}

Answer: D

NEW QUESTION 47

- (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 48

- (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 49

- (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 50

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 51

- (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 53

- (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 55

- (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 56

- (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 58

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 projection

Object protection

Table projection

Extracted text data:

File protection

Object projection

Table projection

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

JSON data:

File projection
Object protection
Table projection

Extracted text data:

File protection
Object projection
Table projection

NEW QUESTION 61

- (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 63

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_text  
    ...
```

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 67

- (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 68

- (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 69

- (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 74

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an agent.

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

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

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

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

Set tool_choice to:

auto

none

required

Configure the tool to authenticate by:

Storing API keys in prompts

Using the shared project agent identity

Using a distinct agent identity bound to the client application

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Set tool_choice to:

auto

none

required

Configure the tool to authenticate by:

Storing API keys in prompts

Using the shared project agent identity

Using a distinct agent identity bound to the client application

NEW QUESTION 76

- (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 78

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:

Observability signals

⋮ Groundedness evaluation metrics

⋮ Latency breakdown traces

⋮ Risk and safety metrics

⋮ Token usage analytics

Answer Area

Unsupported responses:

⋮ Groundedness evaluation metrics

Policy violations:

⋮ Risk and safety metrics

NEW QUESTION 83

- (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 87

- (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 88

- (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 92

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=  


0



1



2

  
    thinking={"type": "enabled"},  
    output_config={"effort":  


"high"



"low"




"medium"

  
    }  
)
```

- 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= ,  
    thinking={"type": "enabled"},  
    output_config={"effort":  }  
)
```

NEW QUESTION 96

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 100

- (Topic 2)
You have a Microsoft Foundry project named Project1 that contains an agent. The agent uses an OpenAPI 3.0 specification to call an external weather service. The weather service requires a key to be passed in an HTTP header. The key value is stored as a connection in Project1. You need to ensure that the key value from the connection is included automatically whenever the OpenAPI tool is invoked. What should you configure in the OpenAPI specification?

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

Answer: C

NEW QUESTION 104

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 105

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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 FirstTry
- * AI-103 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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