



Microsoft

Exam Questions AI-103

Developing AI Apps and Agents on Azure

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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 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 3

- (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 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 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:

Diagnostic log to collect:

- 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 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

DRAG DROP - (Topic 2)

You have a Microsoft Foundry project that contains an agent used by the financial analysts at your company. You need to optimize the agent workflow by providing additional data access and processing capabilities. The solution must meet the following requirements:

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

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

NOTE: Each correct selection is worth one point.

Tools

Code interpreter

Computer use

File search

Grounding with Bing Search

Microsoft Fabric

Answer Area

Access up-to-date information from public websites:

Tool

Perform calculations during conversations:

Tool

Retrieve information from documents uploaded directly to the agent:

Tool

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Tools

⋮ Code interpreter

⋮ Computer use

⋮ File search

⋮ Grounding with Bing Search

⋮ Microsoft Fabric

Answer Area

Access up-to-date information from public websites ⋮ Grounding with Bing Search

Perform calculations during conversations ⋮ Code interpreter

Retrieve information from documents uploaded directly to the agent ⋮ File search

NEW QUESTION 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)

Note: This section contains one or more sets of questions with the same scenario and problem. Each question presents a unique solution to the problem. You must determine whether the solution meets the stated goals.

More than one solution in the set might solve the problem. It is also possible that none of the solutions in the set solve the problem. After you answer a question in this section, you will NOT be able to return. As a result, these questions do not appear on the Review Screen.

You have a Microsoft Foundry project that contains an agent. The agent generates summaries from retrieved policy documents.

Users report that some responses omit required regulatory clauses, even when the clauses are present in the retrieved content.

You need to improve response completeness.

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

Does this meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 10

HOTSPOT - (Topic 2)

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 14

- (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 18

- (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 23

- (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 28

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 32

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 | Multi-file task in standard mode

Pipeline2 | Multi-file task in pro mode

NEW QUESTION 37

- (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 39

- (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 42

- (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 47

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= ,
)

agent_def = PromptAgentDefinition(
    model="gpt-5.2",
    instructions="You are a customer support assistant.",
    tools= ,
)
```

- 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= ,
)

agent_def = PromptAgentDefinition(
    model="gpt-5.2",
    instructions="You are a customer support assistant.",
    tools= ,
)
```

NEW QUESTION 48

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 51

- (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 52

HOTSPOT - (Topic 2)

You have a Microsoft Foundry resource named AI1 that hosts three deployments of the GPT 3.5 model. Each deployment is optimized for a unique workload. You plan to deploy three apps. Each app will access AM by using the REST API and will use the deployment that was optimized for the apps intended workload. You need to provide each app with access to AI1 and the appropriate deployment. The solution must ensure that only the apps can access AH. What should you use to provide access to AI1. and what should each app use to connect to its appropriate deployment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Provide access to AI1 by using:

An API key

A bearer token

A shared access signature (SAS) token

Connect to the deployment by using:

An API key

A deployment endpoint

A deployment name

A deployment type

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Provide access to AI1 by using:

An API key
A bearer token
A shared access signature (SAS) token

Connect to the deployment by using:

An API key
A deployment endpoint
A deployment name
A deployment type

NEW QUESTION 55

- (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 59

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

The LangChain service will appear in Traces without configuring a tracer.

Yes

☐

No

☒

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 63

- (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 documents. Does this meet the goal?

A. Yes

B. No

Answer: B

NEW QUESTION 64

- (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 65

- (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 run an evaluation flow that scores responses for completeness and blocks responses that fall below a defined threshold.

Does this meet the goal?

A. Yes

B. No

Answer: B

NEW QUESTION 68

- (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 72

- (Topic 2)

You have a web app named App1 that sends requests to a multimodal chat model deployment in a Microsoft Foundry project. User messages can contain both text and images. Currently, App1 includes image URLs as plain text inside the message content, so the model cannot recognize them as images. You need to send the message as a structured array that includes both the text portion and the image reference.

- A. Add the image URL to request metadata.
- B. Set the user message content array to include items with type: text and type: image_url.
- C. Place the image URL inside the system message.
- D. Include base64 data directly inside the content string.

Answer: B

NEW QUESTION 73

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 74

- (Topic 2)

You have a Microsoft Foundry project. You need to deploy a model from the model catalog to support real-time inference. The solution must meet the following requirements:

- ? Use key-based authentication
- ? Support real-time REST API access
- ? Not consume the vCPU quota of the virtual machines in the Azure subscription

Which type of deployment should you use?

- A. serverless API
- B. self-hosted container
- C. standard

D. batch

Answer: A

NEW QUESTION 76

- (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 80

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 are building a speech processing solution in Microsoft Foundry for a customer support platform.

The platform will transcribe live phone calls, so that supervisors at your company can view call transcripts and detect issues while the calls are in progress. The call audio will arrive as a continuous stream from the telephony system.

You need to ensure that the call transcripts appear within only a few seconds of the audio stream.

What should you do?

- A. Run a batch transcription job on recorded audio files.
- B. Use real-time speech to text to process streaming audio input.
- C. Use speech translation to generate the transcripts into multiple languages.
- D. Use text to speech by using a custom neural voice.

Answer: B

NEW QUESTION 90

- (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 92

- (Topic 2)

You have a Microsoft Foundry project that ingests scanned PDF invoices stored in Azure Blob Storage. Each invoice contains printed line items and has a table-based layout.

Extracted results are stored as structured JSON and used as grounding data for an agent in a Retrieval Augmented Generation (RAG) solution.

You need to create a single analyzer that meets the following requirements:

- Extracts the invoice number, invoice date, vendor name, and total amount across varying templates
- Returns confidence scores so that results with confidence below 0.80 can be routed for supervisor review

What should you use?

- A. the Azure Content Understanding in Foundry Tools prebuilt-layout analyzer
- B. a Foundry agent that has groundedness guardrails enabled to extract invoice fields and confidence scores
- C. a custom Azure Content Understanding in Foundry Tools analyzer that defines the required fields as the extracted fields and the returned confidence scores for routing
- D. the Azure Content Understanding in Foundry Tools prebuilt-documentSearch analyzer and search.score from the Azure AI Search results for routing

Answer: C

NEW QUESTION 93

- (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 98

- (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 99

- (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 101

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 103

- (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 104

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains a workflow for a customer support triage process.

You have an Ask a question node that stores user responses in a local variable named Var01.

You need to create the following Power Fx expressions:

? An if/else condition expression that ensures that Var01 contains a value
? A Send message expression that returns the stored user response in uppercase
How should you configure the expressions? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

If/else condition expression:

▼

IsBlank(Local.Var01)

IsEmpty(Local.Var01)

Not(IsBlank(Local.Var01))

Send message expression:

▼

{Local.Var01}

{Upper(Local.Var01)}

{Upper(Var01)}

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

If/else condition expression:

▼

IsBlank(Local.Var01)

IsEmpty(Local.Var01)

Not(IsBlank(Local.Var01))

Send message expression:

▼

{Local.Var01}

{Upper(Local.Var01)}

{Upper(Var01)}

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 110

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 115

- (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 116

- (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 117

- (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 122

.....

Relate Links

100% Pass Your AI-103 Exam with ExamBible Prep Materials

<https://www.exambible.com/AI-103-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>