



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

What should you include in the recommendation?

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

Answer: A

NEW QUESTION 2

- (Topic 1)

You need to configure Agent1 to meet the security and compliance requirements. What should you use?

- A. prompt shields
- B. Personally Identifiable Information (PII) Detection
- C. self-harm content filtering
- D. violence content filtering

Answer: B

NEW QUESTION 3

- (Topic 1)

You need to configure 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

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:

- Standard
- Global Standard
- Global Provisioned

Version update policy:

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Deployment type:

- Standard
- Global Standard
- Global Provisioned

Version update policy:

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

NEW QUESTION 5

- (Topic 2)

You have a Microsoft Foundry project that contains an agent.

The agent uses a knowledge source built from documents stored in Azure Blob Storage. The documents include digitally scanned PDFs that contain multipage tables.

You have an ingestion job that extracts only plain text, causing loss of table structure, headings, and page-number metadata.

Users frequently ask questions that require the retrieval of specific table rows across the pages.

You need to configure an ingestion job for a Retrieval Augmented Generation (RAG) pipeline that performs optical character recognition (OCR) on scanned PDFs, preserves tables and headings as structure-aware chunks, and stores page-number metadata with each chunk.

How should you configure the ingestion job?

- A. Use basic parsing and fixed-size chunking.
- B. Use advanced data parsing to reingest the documents.
- C. Use OCR and page-level chunking.
- D. Use page-level OCR extraction and store each page as a single chunk.

Answer: B

NEW QUESTION 6

- (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 7

- (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 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 that contains a prompt agent used by a customer support web app. The agent is invoked from a Python service that does NOT run in the Foundry portal. You need to implement end-to-end tracing to capture latency breakdowns and exceptions across agent runs. Which two components can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. a Log Analytics workspace
- B. OpenTelemetry
- C. Application Insights
- D. the Azure Monitor Agent
- E. Microsoft Sentinel

Answer: BC

NEW QUESTION 10

- (Topic 2)

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

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

Answer: A

NEW QUESTION 10

HOTSPOT - (Topic 2)

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

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

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 13

HOTSPOT - (Topic 2)

You have a Microsoft Foundry project that contains an agent named PaymentAgent.

PaymentAgent includes a function tool that issues customer refunds by using an external API.

You are creating a workflow in YAML.

You need to ensure that the workflow pauses for human approval and continues with the refund step only after approval is granted.

How should you complete the workflow definition? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

steps:

- id: propose_refund

type: agent

agent: PaymentAgent

- id: approval

type:

▼

ask_question

basic_chat

data_transformation

- id: execute_refund

type: agent

agent: PaymentAgent

condition:

▼

approval == "approved"

propose_refund.output != null

true

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

steps:

- id: propose_refund

type: agent

agent: PaymentAgent

- id: approval

type:

ask_question

basic_chat

data_transformation

- id: execute_refund

type: agent

agent: PaymentAgent

condition:

approval == "approved"

propose_refund.output != null

true

NEW QUESTION 14

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 16

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 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 22

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 26

- (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 28

- (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 30

- (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 32

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 33

- (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 34

- (Topic 2)

You have a Microsoft Foundry project that contains an agent.

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

What should you configure first?

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

Answer: A

NEW QUESTION 38

- (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 43

- (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 47

- (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 48

- (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 50

- (Topic 2)

You plan to configure an evaluation in Microsoft Foundry for a Retrieval Augmented Generation (RAG) chat app. You need to provide scores for groundedness, relevance, and harmful-content categories. Which two evaluation categories can you use? Each correct answer presents part of the solution.

- A. risk and safety metrics
- B. AI quality (NLP) metrics
- C. AI quality (AI assisted) metrics
- D. fluency evaluator
- E. similarity evaluators

Answer: AC

NEW QUESTION 54

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 58

- (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 62

- (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 65

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 68

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 72

- (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 77

- (Topic 2)

You have a Microsoft Foundry project that contains a customer support agent. The agent calls an internal knowledge API tool before generating responses.

Users report the following issues:

- Some requests take more than 15 seconds to complete.
- Some responses are incorrect, even when the knowledge API returns the expected data. You need to inspect individual agent runs to view the ordered sequence of large language model (LLM) calls, tool invocations, and timing information. Which observability capability should you use?

- A. token usage
- B. safety metrics
- C. tracing
- D. monitoring

Answer: C

NEW QUESTION 80

- (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 83

- (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 87

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 88

.....

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/>