

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 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 3

- (Topic 1)

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

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

Answer: A

NEW QUESTION 4

- (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 5

- (Topic 1)

You need to configure personalized user interactions for Agent1 based on the business requirements. What should you include in the solution?

- A. guardrails
- B. memory
- C. tools
- D. instructions

Answer: B

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}")
```

create_version

get

get_version

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
from azure.identity import DefaultAzureCredential
from azure.ai.projects import AIProjectClient
from azure.core.credentials import AzureKeyCredential

myEndpoint = "https://contoso.services.ai.azure.com/api/projects/project1"

project_client = AIProjectClient(
    endpoint=myEndpoint,
    credential=  ,
    

AzureKeyCredential()
        DefaultAzureCredential()
        None


)

myAgent = "Agent1"
agent = project_client.agents.  (agent_name=myAgent)

print(f"Retrieved agent: {agent.name}")
```

create_version
 get
 get_version

NEW QUESTION 7

- (Topic 2)

You have a Microsoft Foundry agent that grounds responses from an Azure AI Search index containing:

? Searchable text fields for product names and product codes.

? A vector field containing embeddings for product descriptions.

You need users to query by exact product names or codes and by natural-language product descriptions.

- A. Semantic search only
- B. Hybrid search
- C. Keyword search only
- D. Vector search only

Answer: B

NEW QUESTION 8

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 9

- (Topic 2)

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

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

Answer: A

NEW QUESTION 10

- (Topic 2)

You 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

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 13

- (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 17

- (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 22

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 27

- (Topic 2)

You have a Microsoft Foundry project that generates product marketing images from text prompts.

After publishing several images, the legal team at your company identifies a competitor's logo on a sign in the background of an image.

You need to remove only the logo, while preserving the rest of the image. What should you do?

- A. Increase the prompt guidance strength.

- B. Modify the original prompt to exclude brand names.
- C. Apply a mask-based inpainting edit to the part of the image that contains the logo.
- D. Rerun the prompt by using a different random seed.

Answer: C

NEW QUESTION 31

- (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 34

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 36

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 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 40

- (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 41

- (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 45

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 <code>OTEL_SERVICE_NAME</code> values separates the services in Application Insights.	☐	☐
When using <code>enable_content_recording=False</code> , prompts and tool data will be captured in the telemetry.	☐	☐

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

NEW QUESTION 46

- (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 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 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 48

- (Topic 2)

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

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

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

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

You need to improve response completeness.

Solution: You increase the value of the max_tokens parameter. Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 50

- (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 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

- (Topic 2)

You are building a web app named App1 that generates responses by using a model deployed to a Microsoft Foundry project named Project1.

Before sending the prompts to the model, App1 must retrieve documents by using Azure AI Search.

You need to integrate Project1 and App1. The solution must meet the following requirements:

- Multiple client applications must use the same search configuration.
- A security policy must prevent key-based authentication.
- Administrative effort must be minimized. What should you do?

- A. Enable a managed identity for each application and call Azure AI Search directly.
- B. Create a custom HTTP connection in Foundry and manually configure Azure AI Search endpoints per application.
- C. Call Azure AI Search directly from each application by using Microsoft Entra authentication.
- D. Configure an Azure AI Search connection in Project1 and reference the connection in each application.

Answer: D

NEW QUESTION 59

- (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 63

- (Topic 2)

You have a Microsoft Foundry project that contains an agent named Agent1. Agent1 runs successfully, but Foundry Control Plane does NOT display values for error rates, runs, and token usage, and the Traces tab is empty. You need to ensure that Foundry Control Plane displays the appropriate values for Agent1.

- A. Restart Agent1 from Foundry Control Plane.
- B. Update Agent1 to a new version.
- C. Assign a Log Analytics workspace to Agent1.
- D. Enable Application Insights for Agent1.

Answer: D

NEW QUESTION 66

- (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 67

- (Topic 2)

You have a Microsoft Foundry project named Project1 that contains the following:

- An OpenAPI tool that calls an external API
- A project connection named Connection1 that stores the API key of the external API When an agent calls the OpenAPI tool, the API returns a 401 unauthorized error, and traces show that the API key header is NOT being sent.

You need to ensure that the OpenAPI tool automatically includes the API key from Connection1 on all requests.

What should you do?

- A. Configure the tool to use the default connection of Project1.
- B. Connect the tool to Connection1.
- C. Enable identity passthrough so that the tool uses the Microsoft Entra token of the caller.
- D. Add the API key header manually to the OpenAPI specification.

Answer: B

NEW QUESTION 71

- (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 72

- (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 75

- (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 79

- (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 83

- (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 84

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)}

NEW QUESTION 86

- (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 90

HOTSPOT - (Topic 2)

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

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

Stakeholders report that small wording differences are causing validation mismatches.

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

How should you complete the Python code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 92

- (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 93

HOTSPOT - (Topic 2)

You have 100,000 images.

You need to build an app that will perform the following actions:

- Identify road signs in the images and extract the text on the signs.
- Analyze the text to identify well-known locations. The solution must minimize development effort.

Answer Area

Extract the text:

Azure Document Intelligence in Foundry Tools

Azure Language in Foundry Tools

Azure AI Search

Azure Vision in Foundry Tools

Identify well-known locations:

Azure Document Intelligence in Foundry Tools

Azure Language in Foundry Tools

Azure AI Search

Azure Vision in Foundry Tools

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Extract the text:

Azure Document Intelligence in Foundry Tools

Azure Language in Foundry Tools

Azure AI Search

Azure Vision in Foundry Tools

Identify well-known locations:

Azure Document Intelligence in Foundry Tools

Azure Language in Foundry Tools

Azure AI Search

Azure Vision in Foundry Tools

NEW QUESTION 97

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=`

`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`

`ApiKey=xxx`

- 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=`

`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`

`ApiKey=xxx`

NEW QUESTION 100

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions 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 First Try
- * AI-103 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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