



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 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 2

- (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 3

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=  ,
    

AzureKeyCredential()
        DefaultAzureCredential()
        None


)

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=  ,
)

myAgent = "Agent1"
agent = project_client.agents.  (agent_name=myAgent)
print(f"Retrieved agent: {agent.name}")
```

create_version

get

get_version

NEW QUESTION 4

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 5

- (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 6

- (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 7

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 8

- (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 9

- (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 10

- (Topic 2)

You have a Microsoft Foundry project that uses Azure AI Search to ground an agent in internal documentation. After a recent content update, users report that the agent's answers have become less accurate. You need to identify whether the retrieved content is negatively influencing the model's generated responses. Which observability signal should you review?

- A. prediction drift metrics
B. groundedness evaluation metrics
C. latency breakdown traces
D. indexer status and failure history

Answer: B

NEW QUESTION 10

DRAG DROP - (Topic 2)

You have a Microsoft Foundry project that contains an agent.

You need to enable long-term memory to ensure that the agent can recall user preferences across separate conversations. Stored memories must be isolated per authenticated user without the client application manually generating user IDs.

How should you complete the Python code? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

NOTE: Each correct selection is worth one point.

Values

⋮ "session"

⋮ "{{\${conversationId}}}"

⋮ "{{\${userId}}}"

⋮ [mem_store_name]

⋮ [memory_tool]

⋮ MemorySearchTool("support_mem_store")

Answer Area

```
from azure.ai.projects.models import MemorySearchTool, PromptAgentDefinition

mem_store_name = "agent_mem_store"

memory_tool = MemorySearchTool(
    memory_store_name=mem_store_name,
    scope= 

Value

 ,
)

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

Value


)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Values

⋮ "session"

⋮ "{{\${conversationId}}}"

⋮ "{{\${userId}}}"

⋮ [mem_store_name]

⋮ [memory_tool]

⋮ MemorySearchTool("support_mem_store")

Answer Area

```
from azure.ai.projects.models import MemorySearchTool, PromptAgentDefinition

mem_store_name = "agent_mem_store"

memory_tool = MemorySearchTool(
    memory_store_name=mem_store_name,
    scope= 

⋮ "{{${userId}}}"

 ,
)

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

⋮ [memory_tool]


)
```

NEW QUESTION 15

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 19

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 AI1 by using the REST API and will use the deployment that was optimized for the app's 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 AI1.

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 21

- (Topic 2)

You have a Microsoft Foundry project that contains an agent.

The agent ingests scanned PDF vendor invoices that contain tables and embedded QR codes.

The agent must preserve the PDF layout in the extracted output to ensure that downstream processing can reference sections and tables.

You plan to call Azure Content Understanding in Foundry Tools.

You need to extract content and layout elements and detect QR codes without requiring a language model deployment.

Which built-in analyzer should you use?

- A. prebuilt-layout
- B. prebuilt-documentFieldSchema
- C. prebuilt-read
- D. prebuilt-documentSearch

Answer: A

NEW QUESTION 25

HOTSPOT - (Topic 2)

You have a Python application named App1 that integrates with a Microsoft Foundry project named Project1.

You need to ensure that App1 meets the following requirements:


- Authenticates by using a Microsoft Entra managed identity
- Sends prompts to a deployed model by using the Azure OpenAI Responses API

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.

```
from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

credential =  ()


project_client = AIProjectClient(

    endpoint="https://contosoai.services.ai.azure.com/api/projects/project1",

    credential=credential,

)

with project_client.get_openai_client() as openai_client:

    response = openai_client.responses.  (

        model="trail-guide-chat",

        input="Create a 3-day hiking itinerary near Seattle.",

    )

    print(response.output_text)
```


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
from azure.identity import DefaultAzureCredential

from azure.ai.projects import AIProjectClient

credential =  ()


project_client = AIProjectClient(

    endpoint="https://contosoai.services.ai.azure.com/api/projects/project1",

    credential=credential,

)

with project_client.get_openai_client() as openai_client:

    response = openai_client.responses.  (

        model="trail-guide-chat",

        input="Create a 3-day hiking itinerary near Seattle.",

    )

    print(response.output_text)
```

NEW QUESTION 27

HOTSPOT - (Topic 2)

You need to recommend a plan to create a customer support agent by using the Microsoft Foundry Agent Service. The agent must meet the following requirements:

- Retain user preferences across multiple conversations.
- Enable users to provide contextual grounding by directly uploading documents during a chat.


Which Foundry capability should you recommend for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

To retain user preferences across conversations, use:




To enable users to provide contextual grounding during chats, use the:



B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 31

- (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 35

- (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 39

- (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 43

- (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 45

- (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 50

- (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 51

- (Topic 2)

You have a Microsoft Foundry project that serves a high-volume chat app. Most requests are simple FAQs, but some require advanced reasoning.

You need to reduce costs and latency for common queries, without degrading the quality of the responses to complex questions.

What should you do?

- A. Increase the value of the max_tokens parameter for all the requests.
- B. Route all the requests to a smaller model.
- C. Route all the requests to the most capable model.
- D. Use a model cascade that routes the requests to different models.

Answer: D

NEW QUESTION 55

- (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 57

- (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 58

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 63

- (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 68

- (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 70

- (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 75

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 78

- (Topic 2)

You need to measure the public perception of your brand on social media by using natural language processing. Which Azure service should you use?

- A. Azure Document Intelligence in Foundry Tools
- B. Content Safety in Foundry Control Plane
- C. Azure Language in Foundry Tools
- D. Azure Vision in Foundry Tools

Answer: C

NEW QUESTION 80

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 84

- (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 88

- (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 90

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=
```

```
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 91

.....

Relate Links

100% Pass Your AI-103 Exam with Exam Bible 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/>