



AI-102^{Q&As}

Designing and Implementing a Microsoft Azure AI Solution (beta)

Pass Microsoft AI-102 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.lead4pass.com/ai-102.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

DRAG DROP

You plan to use containerized versions of the Anomaly Detector API on local devices for testing and in on-premises datacenters.

You need to ensure that the containerized deployments meet the following requirements:

Prevent billing and API information from being stored in the command-line histories of the devices that run the container.

Control access to the container images by using Azure role-based access control (Azure RBAC).

Which four 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. (Choose four.)

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:

Actions

- Create a custom Dockerfile.
- Pull the Anomaly Detector container image.
- Distribute a docker run script.
- Push the image to an Azure container registry.
- Build the image.
- Push the image to Docker Hub.

Answer Area

Correct Answer:

Actions

-
-
-
-
- Build the image.
- Push the image to Docker Hub.

Answer Area

- Pull the Anomaly Detector container image.
- Create a custom Dockerfile.
- Push the image to an Azure container registry.
- Distribute a docker run script.

Step 1: Pull the Anomaly Detector container image.

Step 2: Create a custom Dockerfile

Step 3: Push the image to an Azure container registry.

To push an image to an Azure Container registry, you must first have an image.

Step 4: Distribute the docker run script

Use the docker run command to run the containers.

Reference:

<https://docs.microsoft.com/en-us/azure/container-registry/container-registry-intro>

QUESTION 2

You are developing the smart e-commerce project.

You need to implement autocomplete as part of the Cognitive Search solution.

Which three actions should you perform? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

- A. Make API queries to the autocomplete endpoint and include suggesterName in the body.
- B. Add a suggester that has the three product name fields as source fields.
- C. Make API queries to the search endpoint and include the product name fields in the searchFields query parameter.
- D. Add a suggester for each of the three product name fields.
- E. Set the searchAnalyzer property for the three product name variants.
- F. Set the analyzer property for the three product name variants.

Correct Answer: ABF

Scenario: Support autocomplete and autosuggestion based on all product name variants.

A: Call a suggester-enabled query, in the form of a Suggestion request or Autocomplete request, using an API. API usage is illustrated in the following call to the Autocomplete REST API.

POST /indexes/myxboxgames/docs/autocomplete?searchandapi-version=2020-06-30

```
{ "search": "minecraf", "suggesterName": "sg"
}
```

B:

In Azure Cognitive Search, typeahead or "search-as-you-type" is enabled through a suggester. A suggester provides a list of fields that undergo additional tokenization, generating prefix sequences to support matches on partial terms. For example, a suggester that includes a City field with a value for "Seattle" will have prefix combinations of "sea", "seat", "seatt", and "seattl" to support typeahead.

F.

Use the default standard Lucene analyzer ("analyzer": null) or a language analyzer (for example, "analyzer":



"en.Microsoft") on the field.

Reference: <https://docs.microsoft.com/en-us/azure/search/index-add-suggesters>

QUESTION 3

DRAG DROP

You are developing the smart e-commerce project.

You need to design the skillset to include the contents of PDFs in searches.

How should you complete the skillset design diagram? To answer, drag the appropriate services to the correct stages. Each service 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.

Select and Place:

Services

Azure Blob storage
Custom Vision API
Azure Files
Language Understanding API
Translator API
Computer Vision API
Azure Cosmos DB

Answer Area

Source

---->

Cracking

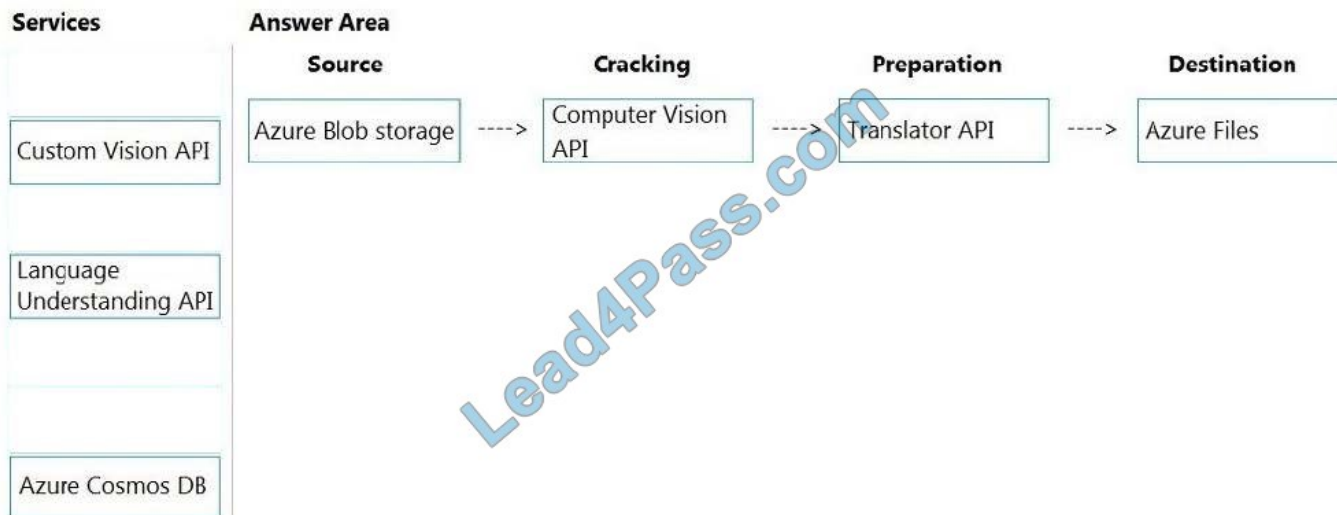
---->

Preparation

---->

Destination

Correct Answer:



Box 1: Azure Blob storage

At the start of the pipeline, you have unstructured text or non-text content (such as images, scanned documents, or JPEG files). Data must exist in an Azure data storage service that can be accessed by an indexer.

Box 2: Computer Vision API

Scenario: Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products.

The Computer Vision Read API is Azure's latest OCR technology (learn what's new) that extracts printed text (in several languages), handwritten text (English only), digits, and currency symbols from images and multi-page PDF documents.

Box 3: Translator API

Scenario: Product descriptions, transcripts, and all text must be available in English, Spanish, and Portuguese.

Box 4: Azure Files

Scenario: Store all raw insight data that was generated, so the data can be processed later.

Incorrect Answers:

The custom vision API from Microsoft Azure learns to recognize specific content in imagery and becomes smarter with training and time.

Reference:

<https://docs.microsoft.com/en-us/azure/search/cognitive-search-concept-intro> <https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/overview-ocr>

QUESTION 4

You are developing a new sales system that will process the video and text from a public-facing website.

You plan to monitor the sales system to ensure that it provides equitable results regardless of the user's location or



background.

Which two responsible AI principles provide guidance to meet the monitoring requirements? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

- A. transparency
- B. fairness
- C. inclusiveness
- D. reliability and safety
- E. privacy and security

Correct Answer: BD

AI systems should treat all people fairly.

AI systems should perform reliably and safely.

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/responsible-ai>

QUESTION 5

You deploy a web app that is used as a management portal for indexing in Azure Cognitive Search. The app is configured to use the primary admin key.

During a security review, you discover unauthorized changes to the search index. You suspect that the primary access key is compromised.

You need to prevent unauthorized access to the index management endpoint. The solution must minimize downtime.

What should you do next?

- A. Regenerate the primary admin key, change the app to use the secondary admin key, and then regenerate the secondary admin key.
- B. Change the app to use a query key, and then regenerate the primary admin key and the secondary admin key.
- C. Regenerate the secondary admin key, change the app to use the secondary admin key, and then regenerate the primary key.
- D. Add a new query key, change the app to use the new query key, and then delete all the unused query keys.

Correct Answer: A

Regenerate admin keys.

Two admin keys are created for each service so that you can rotate a primary key, using the secondary key for business continuity.



1.

In the Settings >Keys page, copy the secondary key.

2.

For all applications, update the API key settings to use the secondary key.

3.

Regenerate the primary key.

4.

Update all applications to use the new primary key.

Note: Two admin api-keys, referred to as primary and secondary keys in the portal, are automatically generated when the service is created and can be individually regenerated on demand. Having two keys allows you to roll over one key while using the second key for continued access to the service.

Reference: <https://docs.microsoft.com/en-us/azure/search/search-security-api-keys#regenerate-admin-keys>

QUESTION 6

You need to build a chatbot that meets the following requirements:

Supports chit-chat, knowledge base, and multilingual models
Performs sentiment analysis on user messages
Selects the best language model automatically

What should you integrate into the chatbot?

- A. QnA Maker, Language Understanding, and Dispatch
- B. Translator, Speech, and Dispatch
- C. Language Understanding, Text Analytics, and QnA Maker
- D. Text Analytics, Translator, and Dispatch

Correct Answer: C

Language Understanding: An AI service that allows users to interact with your applications, bots, and IoT devices by using natural language.

QnA Maker is a cloud-based Natural Language Processing (NLP) service that allows you to create a natural conversational layer over your data. It is used to find the most appropriate answer for any input from your custom knowledge base

(KB) of information.

Text Analytics: Mine insights in unstructured text using natural language processing (NLP)--no machine learning expertise required. Gain a deeper understanding of customer opinions with sentiment analysis. The Language Detection feature

of the Azure Text Analytics REST API evaluates text input



Incorrect Answers:

A, B, D: Dispatch uses sample utterances for each of your bot's different tasks (LUIS, QnA Maker, or custom), and builds a model that can be used to properly route your user's request to the right task, even across multiple bots.

Reference:

<https://azure.microsoft.com/en-us/services/cognitive-services/text-analytics/> <https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/overview/overview>

QUESTION 7

HOTSPOT

You develop a test method to verify the results retrieved from a call to the Computer Vision API. The call is used to analyze the existence of company logos in images. The call returns a collection of brands named brands.

You have the following code segment.

```
foreach (var brand in brands)
{
    if (brand.Confidence >= .75)
        Console.WriteLine($"Logo of {brand.Name} between {brand.Rectangle.X},
{brand.Rectangle.Y} and {brand.Rectangle.W}, {brand.Rectangle.H}");
}
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
The code will return the name of each detected brand with a confidence equal to or higher than 75 percent.	<input type="checkbox"/>	<input type="checkbox"/>
The code will return coordinates for the bottom-left corner of the rectangle that contains the brand logo of the displayed brands.	<input type="checkbox"/>	<input type="checkbox"/>
The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.	<input type="checkbox"/>	<input type="checkbox"/>

Correct Answer:



Answer Area

Statements

Yes

No

The code will return the name of each detected brand with a confidence equal to or higher than 75 percent.

☒☐

The code will return coordinates for the bottom-left corner of the rectangle that contains the brand logo of the displayed brands.

☒☐

The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.

☐☒

Box 1: Yes

Box 2: Yes

If several logs are detected, or the logo image and the stylized brand name are detected as two separate logos, it starts numbering them from the bottom-left corner.

Box 3: No

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/concept-brand-detection>

QUESTION 8

DRAG DROP

You are using a Language Understanding service to handle natural language input from the users of a web-based customer agent.

The users report that the agent frequently responds with the following generic response: "Sorry, I don't understand that."

You need to improve the ability of the agent to respond to requests.

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. (Choose three.)

Select and Place:



Actions

- Add prebuilt domain models as required.
- Validate the utterances logged for review and modify the model.
- Migrate authoring to an Azure resource authoring key.
- Enable active learning.
- Enable log collection by using Log Analytics.
- Train and republish the Language Understanding model.

Answer Area

Correct Answer:

Actions

- Validate the utterances logged for review and modify the model.
- Migrate authoring to an Azure resource authoring key.
- Enable log collection by using Log Analytics.

Answer Area

- Add prebuilt domain models as required.
- Enable active learning.
- Train and republish the Language Understanding model.

Step 1: Add prebuilt domain models as required.

Prebuilt models provide domains, intents, utterances, and entities. You can start your app with a prebuilt model or add a relevant model to your app later.

Note: Language Understanding (LUIS) provides prebuilt domains, which are pre-trained models of intents and entities that work together for domains or common categories of client applications.

The prebuilt domains are trained and ready to add to your LUIS app. The intents and entities of a prebuilt domain are

fully customizable once you've added them to your app.

Step 2: Enable active learning

To enable active learning, you must log user queries. This is accomplished by calling the endpoint query with the log=true querystring parameter and value.

Step 3: Train and republish the Language Understanding model

The process of reviewing endpoint utterances for correct predictions is called Active learning. Active learning captures endpoint queries and selects user's endpoint utterances that it is unsure of. You review these utterances to select the intent

and mark entities for these real-world utterances. Accept these changes into your example utterances then train and publish. LUIS then identifies utterances more accurately.

Incorrect Answers:

Enable log collection by using Log Analytics

Application authors can choose to enable logging on the utterances that are sent to a published application. This is not done through Log Analytics.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-how-to-review-endpoint-utterances#log-user-queries-to-enable-active-learning>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-prebuilt-model>

QUESTION 9

HOTSPOT

You plan to deploy a containerized version of an Azure Cognitive Services service that will be used for text analysis.

You configure <https://contoso.cognitiveservices.azure.com> as the endpoint URI for the service, and you pull the latest version of the Text Analytics Sentiment Analysis container.

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.

Hot Area:



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-services/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=xxxxxxxxxxxxxxxxxxxxxx

Correct Answer:

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-services/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=xxxxxxxxxxxxxxxxxxxxxx

Box 1: mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment

To run the Sentiment Analysis v3 container, execute the following docker run command.

```
docker run --rm -it -p 5000:5000 --memory 8g --cpus 1 \
```

```
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment \
```

Eula=accept \

Billing={ENDPOINT_URI} \

ApiKey={API_KEY} is the endpoint for accessing the Text Analytics API. <https://cognitiveservices.azure.com>

Box 2: <https://contoso.cognitiveservices.azure.com>

{ENDPOINT_URI} is the endpoint for accessing the Text Analytics API: <https://cognitiveservices.azure.com> The endpoint for

accessing the Text Analytics API. zure.com

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-install-containers?tabs=sentiment>

QUESTION 10

You have a Video Indexer service that is used to provide a search interface over company videos on your company's website.

You need to be able to search for videos based on who is present in the video.

What should you do?

- A. Create a person model and associate the model to the videos.
- B. Create person objects and provide face images for each object.
- C. Invite the entire staff of the company to Video Indexer.
- D. Edit the faces in the videos.
- E. Upload names to a language model.

Correct Answer: A

Video Indexer supports multiple Person models per account. Once a model is created, you can use it by providing the model ID of a specific Person model when uploading/indexing or reindexing a video. Training a new face for a video updates the specific custom model that the video was associated with.

Note: Video Indexer supports face detection and celebrity recognition for video content. The celebrity recognition feature covers about one million faces based on commonly requested data source such as IMDB, Wikipedia, and top LinkedIn influencers. Faces that aren't recognized by the celebrity recognition feature are detected but left unnamed. Once you label a face with a name, the face and name get added to your account's Person model. Video Indexer will then recognize this face in your future videos and past videos.

Reference: <https://docs.microsoft.com/en-us/azure/media-services/video-indexer/customize-person-model-with-api>

QUESTION 11

DRAG DROP

You plan to use a Language Understanding application named app1 that is deployed to a container.

App1 was developed by using a Language Understanding authoring resource named lu1.

App1 has the versions shown in the following table.



Version	Trained date	Published date
V1.2	None	None
V1.1	2020-10-01	None
V1.0	2020-09-01	2020-09-15

You need to create a container that uses the latest deployable version of app1.

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. (Choose three.)

Select and Place:

Actions

Answer Area

Run a container that has version set as an environment variable.

Export the model by using the Export as JSON option.

Select v1.1 of app1.

Run a container and mount the model file.

Select v1.0 of app1.

Export the model by using the Export for containers (GZIP) option.

Select v1.2 of app1.

Correct Answer:



Actions

Run a container that has version set as an environment variable.

Export the model by using the Export as JSON option.

Select v1.0 of app1.

Select v1.2 of app1.

Answer Area

Export the model by using the Export for containers (GZIP) option.

Select v1.1 of app1.

Run a container and mount the model file.

QUESTION 12

HOTSPOT

You are planning the product creation project.

You need to build the REST endpoint to create the multilingual product descriptions.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

▼	▼	?api-version=3.0&to=es&to=pt
api.cognitive.microsofttranslator.com	/detect	
api-nam.cognitive.microsofttranslator.com	/languages	
westus.tts.speech.microsoft.com	/text-to-speech	
wwics.cognitiveservices.azure.com/translator	/translate	



Correct Answer:

Answer Area

api.cognitive.microsofttranslator.com	/detect	?api-version=3.0&to=es&to=pt
api-nam.cognitive.microsofttranslator.com	/languages	
westus.tts.speech.microsoft.com	/text-to-speech	
wwics.cognitiveservices.azure.com/translator	/translate	

Box 1: api.cognitive.microsofttranslator.com Translator 3.0: Translate. Send a POST request to:
<https://api.cognitive.microsofttranslator.com/translate?api-version=3.0>

Box 2: /translate Reference: <https://docs.microsoft.com/en-us/azure/cognitive-services/translator/reference/v3-0-translate>

QUESTION 13

HOTSPOT

You develop an application that uses the Face API.

You need to add multiple images to a person group.

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

```
Parallel.For(0, PersonCount, async i =>
{
    Guid personId = persons[i].PersonId;
    string personImageDir = $"/path/to/person/{i}/images";
    foreach (string imagePath in Directory.GetFiles(personImageDir, "*.jpg"))
    {
        using ( 

|        |
|--------|
| ▼      |
| File   |
| Stream |
| Uri    |
| Url    |

 t = File.OpenRead(imagePath))

        {
            await faceClient.PersonGroupPerson. 

|                        |
|------------------------|
| ▼                      |
| AddFaceFromStreamAsync |
| AddFaceFromUrlAsync    |
| CreateAsync            |
| GetAsync               |



            (personGroupId, personId, t);
        }
    }
});
```

Correct Answer:

Answer Area

```
Parallel.For(0, PersonCount, async i =>
{
    Guid personId = persons[i].PersonId;
    string personImageDir = $"/path/to/person/{i}/images";
    foreach (string imagePath in Directory.GetFiles(personImageDir, "*.jpg"))
    {
        using ( 

|        |
|--------|
| ▼      |
| File   |
| Stream |
| Uri    |
| Url    |

 t = File.OpenRead(imagePath))

        {
            await faceClient.PersonGroupPerson. 

|                        |
|------------------------|
| ▼                      |
| AddFaceFromStreamAsync |
| AddFaceFromUrlAsync    |
| CreateAsync            |
| GetAsync               |



            (personGroupId, personId, t);
        }
    }
});
```



Box 1: Stream

The File.OpenRead(String) method opens an existing file for reading.

Example: Open the stream and read it back. using (FileStream fs = File.OpenRead(path))

Box 2: CreateAsync

Create the persons for the PersonGroup. Persons are created concurrently.

Example:

```
await faceClient.PersonGroupPerson.CreateAsync(personGroupId, personName);
```

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-add-faces>

[AI-102 Practice Test](#)

[AI-102 Study Guide](#)

[AI-102 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications.
You can view Vendor list of All Certification Exams offered:

<https://www.lead4pass.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.	 Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.	 Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © lead4pass, All Rights Reserved.