

Actual exam question from Microsoft's AI-102

Question #: 1

Topic #: 1

[\[All AI-102 Questions\]](#)

DRAG DROP -

You have 100 chatbots that each has its own Language Understanding model.

Frequently, you must add the same phrases to each model.

You need to programmatically update the Language Understanding models to include the new phrases.

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

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Values

AddPhraseListAsync

Phraselist

PhraselistCreateObject

Phrases

SavePhraselistAsync

UploadPhraseListAsync

Answer Area

```
var phraselistId = await client.Features.   
(appId, versionId, new   
{  
    EnabledForAllModels = false,  
    IsExchangeable = true,  
    Name = "PL1",  
    Phrases = "item1,item2,item3,item4,item5"  
});
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 2

Topic #: 1

[\[All AI-102 Questions\]](#)

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.

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.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 3

Topic #: 1

[\[All AI-102 Questions\]](#)

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

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 4

Topic #: 1

[\[All AI-102 Questions\]](#)

Your company wants to reduce how long it takes for employees to log receipts in expense reports. All the receipts are in English.

You need to extract top-level information from the receipts, such as the vendor and the transaction total. The solution must minimize development effort.

Which Azure service should you use?

- A. Custom Vision
- B. Personalizer
- C. Form Recognizer
- D. Computer Vision

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 5

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT -

You need to create a new resource that will be used to perform sentiment analysis and optical character recognition (OCR). The solution must meet the following requirements:

- ⇒ Use a single key and endpoint to access multiple services.
- ⇒ Consolidate billing for future services that you might use.
- ⇒ Support the use of Computer Vision in the future.

How should you complete the HTTP request to create the new resource? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
xxxx-xxxx-  
xxxxxxxxxxxx/resourceGroups/RG1/providers/Microsoft.CognitiveServices/  
accounts/CS1?api-version=2017-04-18  
{  
  "location": "West US",  
  "kind": "",  
  "sku": {  
    "name": "S0"  
  },  
  "properties": {},  
  "identity": {  
    "type": "SystemAssigned"  
  }  
}
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 6

Topic #: 1

[\[All AI-102 Questions\]](#)

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.

NOTE: Each correct selection is worth one point.

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

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 7

Topic #: 1

[\[All AI-102 Questions\]](#)

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.

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

Answer Area

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.

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 8

Topic #: 1

[\[All AI-102 Questions\]](#)

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=xxxxxxxxxxxxxxxxxxxx
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 9

Topic #: 1

[\[All AI-102 Questions\]](#)

You have the following C# method for creating Azure Cognitive Services resources programmatically.

```
static void create_resource(CognitiveServicesManagementClient client, string
resource_name, string kind, string account_tier, string location)
{
    CognitiveServicesAccount parameters =
        new CognitiveServicesAccount(null, null, kind, location, resource_name,
new CognitiveServicesAccountProperties(), new Sku(account_tier));
    var result = client.Accounts.Create(resource_group_name, account_tier,
parameters);
}
```

You need to call the method to create a free Azure resource in the West US Azure region. The resource will be used to generate captions of images automatically.

Which code should you use?

- A. `create_resource(client, "res1", "ComputerVision", "F0", "westus")`
- B. `create_resource(client, "res1", "CustomVision.Prediction", "F0", "westus")`
- C. `create_resource(client, "res1", "ComputerVision", "S0", "westus")`
- D. `create_resource(client, "res1", "CustomVision.Prediction", "S0", "westus")`

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 10

Topic #: 1

[\[All AI-102 Questions\]](#)

You successfully run the following HTTP request.

POST `https://management.azure.com/subscriptions/18c51a87-3a69-47a8-aedc-a54745f708a1/resourceGroups/RG1/providers/`

`Microsoft.CognitiveServices/accounts/contoso1/regenerateKey?api-version=2017-04-18`

`Body{"keyName": "Key2"}`

What is the result of the request?

- A. A key for Azure Cognitive Services was generated in Azure Key Vault.
- B. A new query key was generated.
- C. The primary subscription key and the secondary subscription key were rotated.
- D. The secondary subscription key was reset.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 11

Topic #: 1

[\[All AI-102 Questions\]](#)

You build a custom Form Recognizer model.

You receive sample files to use for training the model as shown in the following table.

Name	Type	Size
File1	PDF	20 MB
File2	MP4	100 MB
File3	JPG	20 MB
File4	PDF	100 MB
File5	GIF	1 MB
File6	JPG	40 MB

Which three files can you use to train the model? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. File1
- B. File2
- C. File3
- D. File4
- E. File5
- F. File6

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 12

Topic #: 1

[\[All AI-102 Questions\]](#)

A customer uses Azure Cognitive Search.

The customer plans to enable a server-side encryption and use customer-managed keys (CMK) stored in Azure.

What are three implications of the planned change? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. The index size will increase.
- B. Query times will increase.
- C. A self-signed X.509 certificate is required.
- D. The index size will decrease.
- E. Query times will decrease.
- F. Azure Key Vault is required.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 13

Topic #: 1

[\[All AI-102 Questions\]](#)

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

You plan to notify users that their data has been processed by the sales system.

Which responsible AI principle does this help meet?

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

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 14

Topic #: 1

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1.

You plan to create a new Azure Cognitive Search service named service1.

You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet.

Solution: You deploy service1 and a public endpoint to a new virtual network, and you configure Azure Private Link.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 15

Topic #: 1

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1.

You plan to create a new Azure Cognitive Search service named service1.

You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet.

Solution: You deploy service1 and a public endpoint, and you configure an IP firewall rule.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 16

Topic #: 1

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1.

You plan to create a new Azure Cognitive Search service named service1.

You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet.

Solution: You deploy service1 and a public endpoint, and you configure a network security group (NSG) for vnet1.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 17

Topic #: 1

[\[All AI-102 Questions\]](#)

You plan to perform predictive maintenance.

You collect IoT sensor data from 100 industrial machines for a year. Each machine has 50 different sensors that generate data at one-minute intervals. In total, you have 5,000 time series datasets.

You need to identify unusual values in each time series to help predict machinery failures.

Which Azure service should you use?

- A. Anomaly Detector
- B. Cognitive Search
- C. Form Recognizer
- D. Custom Vision

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 18

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are developing a streaming Speech to Text solution that will use the Speech SDK and MP3 encoding.

You need to develop a method to convert speech to text for streaming MP3 data.

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

```
var audioFormat =  (AudioStreamContainerFormat.MP3);  
AudioConfig.SetProperty  
AudioStreamFormat.GetCompressedFormat  
AudioStreamFormat.GetWaveFormatPCM  
PullAudioInputStream  
  
var speechConfig = SpeechConfig.FromSubscription("18c51a87-3a69-47a8-aedc-a54745f708a1", "westus");  
var audioConfig = AudioConfig.FromStreamInput(pushStream, audioFormat);  
using (var recognizer = new  (speechConfig, audioConfig))  
{  
    KeywordRecognizer  
    SpeakerRecognizer  
    SpeechRecognizer  
    SpeechSynthesizer  
  
    {  
        var result = await recognizer.RecognizeOnceAsync();  
        var text = result.Text;  
    }  
}
```

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 19

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are developing an internet-based training solution for remote learners.

Your company identifies that during the training, some learners leave their desk for long periods or become distracted.

You need to use a video and audio feed from each learner's computer to detect whether the learner is present and paying attention. The solution must minimize development effort and identify each learner.

Which Azure Cognitive Services service should you use for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

From a learner's video feed, verify whether the learner is present:

Face
Speech
Text Analytics

From a learner's facial expression in the video feed, verify whether the learner is paying attention:

Face
Speech
Text Analytics

From a learner's audio feed, detect whether the learner is talking:

Face
Speech
Text Analytics

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 20

Topic #: 1

[\[All AI-102 Questions\]](#)

You plan to provision a QnA Maker service in a new resource group named RG1.

In RG1, you create an App Service plan named AP1.

Which two Azure resources are automatically created in RG1 when you provision the QnA Maker service? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Language Understanding
- B. Azure SQL Database
- C. Azure Storage
- D. Azure Cognitive Search
- E. Azure App Service

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 21

Topic #: 1

[\[All AI-102 Questions\]](#)

You are building a language model by using a Language Understanding (classic) service.

You create a new Language Understanding (classic) resource.

You need to add more contributors.

What should you use?

- A. a conditional access policy in Azure Active Directory (Azure AD)
- B. the Access control (IAM) page for the authoring resources in the Azure portal
- C. the Access control (IAM) page for the prediction resources in the Azure portal

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 22

Topic #: 1

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

You discover that some search query requests to the Cognitive Search service are being throttled.

You need to reduce the likelihood that search query requests are throttled.

Solution: You migrate to a Cognitive Search service that uses a higher tier.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 23

Topic #: 1

[\[All AI-102 Questions\]](#)

DRAG DROP -

You need to develop an automated call handling system that can respond to callers in their own language. The system will support only French and English.

Which Azure Cognitive Services service should you use to meet each requirement? To answer, drag the appropriate services to the correct requirements. 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

Answer Area

Speaker Recognition

Speech to Text

Text Analytics

Text to Speech

Translator

Detect the incoming language:

Respond in the callers' own language:

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 24

Topic #: 1

[\[All AI-102 Questions\]](#)

You have receipts that are accessible from a URL.

You need to extract data from the receipts by using Form Recognizer and the SDK. The solution must use a prebuilt model.

Which client and method should you use?

- A. the FormRecognizerClient client and the StartRecognizeContentFromUri method
- B. the FormTrainingClient client and the StartRecognizeContentFromUri method
- C. the FormRecognizerClient client and the StartRecognizeReceiptsFromUri method
- D. the FormTrainingClient client and the StartRecognizeReceiptsFromUri method

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 25

Topic #: 1

[\[All AI-102 Questions\]](#)

You have a collection of 50,000 scanned documents that contain text.

You plan to make the text available through Azure Cognitive Search.

You need to configure an enrichment pipeline to perform optical character recognition (OCR) and text analytics. The solution must minimize costs.

What should you attach to the skillset?

- A. a new Computer Vision resource
- B. a free (Limited enrichments) Cognitive Services resource
- C. an Azure Machine Learning Designer pipeline
- D. a new Cognitive Services resource that uses the S0 pricing tier

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 26

Topic #: 1

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

You discover that some search query requests to the Cognitive Search service are being throttled.

You need to reduce the likelihood that search query requests are throttled.

Solution: You add indexes.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 27

Topic #: 1

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

You discover that some search query requests to the Cognitive Search service are being throttled.

You need to reduce the likelihood that search query requests are throttled.

Solution: You enable customer-managed key (CMK) encryption.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 28

Topic #: 1

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1.

You plan to create a new Azure Cognitive Search service named service1.

You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet.

Solution: You deploy service1 and a private endpoint to vnet1.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 29

Topic #: 1

[\[All AI-102 Questions\]](#)

You have a Language Understanding resource named lu1.

You build and deploy an Azure bot named bot1 that uses lu1.

You need to ensure that bot1 adheres to the Microsoft responsible AI principle of inclusiveness.

How should you extend bot1?

- A. Implement authentication for bot1.
- B. Enable active learning for lu1.
- C. Host lu1 in a container.
- D. Add Direct Line Speech to bot1.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 30

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are building an app that will process incoming email and direct messages to either French or English language support teams.

Which Azure Cognitive Services API should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

https://

<input type="text"/>
api.cognitive.microsofttranslator.com
eastus.api.cognitive.microsoft.com
portal.azure.com

<input type="text"/>
/text/analytics/v3.1/entities/recognition/general
/text/analytics/v3.1/languages
/translator/text/v3.0/translate?to=en
/translator/text/v3.0/translate?to=fr

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 31

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an Azure Cognitive Search instance that indexes purchase orders by using Form Recognizer.

You need to analyze the extracted information by using Microsoft Power BI. The solution must minimize development effort.

What should you add to the indexer?

- A. a projection group
- B. a table projection
- C. a file projection
- D. an object projection

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 32

Topic #: 1

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

You discover that some search query requests to the Cognitive Search service are being throttled.

You need to reduce the likelihood that search query requests are throttled.

Solution: You add replicas.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 33

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

You need to create a Text Analytics service named Text12345678, and then enable logging for Text12345678. The solution must ensure that any changes to Text12345678 will be stored in a Log Analytics workspace.

To complete this task, sign in to the Azure portal.

[Show Suggested Answer](#)





Actual exam question from Microsoft's AI-102

Question #: 34

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

You need to create a search service named search12345678 that will index a sample Azure Cosmos DB database named hotels-sample. The solution must ensure that only English language fields are retrievable.

To complete this task, sign in to the Azure portal.

[Show Suggested Answer](#)





Actual exam question from Microsoft's AI-102

Question #: 35

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

You plan to create a solution to generate captions for images that will be read from Azure Blob Storage.

You need to create a service in Azure Cognitive Services for the solution. The service must be named captions12345678 and must use the Free pricing tier.

To complete this task, sign in to the Azure portal.

[Show Suggested Answer](#)





Actual exam question from Microsoft's AI-102

Question #: 36

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

You need to create a Form Recognizer resource named fr12345678.

Use the Form Recognizer sample labeling tool at <https://fott-2-1.azurewebsites.net/> to analyze the invoice located in the C:\Resources\Invoices folder.

Save the results as C:\Resources\Invoices\Results.json.

To complete this task, sign in to the Azure portal and open the Form Recognizer sample labeling tool.

[Show Suggested Answer](#)





Actual exam question from Microsoft's AI-102

Question #: 37

Topic #: 1

[\[All AI-102 Questions\]](#)

You have a factory that produces food products.

You need to build a monitoring solution for staff compliance with personal protective equipment (PPE) requirements. The solution must meet the following requirements:

- * Identify staff who have removed masks or safety glasses.
- * Perform a compliance check every 15 minutes.
- * Minimize development effort.
- * Minimize costs.

Which service should you use?

- A. Face
- B. Computer Vision
- C. Azure Video Analyzer for Media (formerly Video Indexer)

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 38

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an Azure Cognitive Search solution and a collection of blog posts that include a category field.

You need to index the posts. The solution must meet the following requirements:

- * Include the category field in the search results.
- * Ensure that users can search for words in the category field.
- * Ensure that users can perform drill down filtering based on category.

Which index attributes should you configure for the category field?

- A. searchable, sortable, and retrievable
- B. searchable, facetable, and retrievable
- C. retrievable, filterable, and sortable
- D. retrievable, facetable, and key

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 39

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

Use the following login credentials as needed:

To enter your username, place your cursor in the Sign in box and click on the username below.

To enter your password, place your cursor in the Enter password box and click on the password below.

Azure Username: admin@abc.com -

Azure Password: XXXXXXXXXXXX -

The following information is for technical support purposes only:

Lab Instance: 12345678 -

Task -

You plan to build an API that will identify whether an image includes a Microsoft Surface Pro or Surface Studio.

You need to deploy a service in Azure Cognitive Services for the API. The service must be named AAA12345678 and must be in the East US Azure region. The solution must use the Free pricing tier.

To complete this task, sign in to the Azure portal.

[Show Suggested Answer](#)



Actual exam question from Microsoft's AI-102

Question #: 40

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

Use the following login credentials as needed:

To enter your username, place your cursor in the Sign in box and click on the username below.

To enter your password, place your cursor in the Enter password box and click on the password below.

Azure Username: admin@abc.com -

Azure Password: XXXXXXXXXXXX -

The following information is for technical support purposes only:

Lab Instance: 12345678 -

Task -

You need to build an API that uses the service in Azure Cognitive Services named AAA12345678 to identify whether an image includes a Microsoft Surface Pro or Surface Studio.

To achieve this goal, you must use the sample images in the C:\Resources\Images folder.

To complete this task, sign in to the Azure portal.

[Show Suggested Answer](#)



Actual exam question from Microsoft's AI-102

Question #: 41

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

Use the following login credentials as needed:

To enter your username, place your cursor in the Sign in box and click on the username below.

To enter your password, place your cursor in the Enter password box and click on the password below.

Azure Username: admin@abc.com -

Azure Password: XXXXXXXXXXXX -

The following information is for technical support purposes only:

Lab Instance: 12345678 -

Task -

You need to get insights from a video file located in the C:\Resources\Video\Media.mp4 folder.

Save the insights to the C:\Resources\Video\Insights.json folder.

To complete this task, sign in to the Azure Video Analyzer for Media at <https://www.videoindexer.ai/> by using admin@abc.com

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 42

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

Use the following login credentials as needed:

To enter your username, place your cursor in the Sign in box and click on the username below.

To enter your password, place your cursor in the Enter password box and click on the password below.

Azure Username: admin@abc.com -

Azure Password: XXXXXXXXXXXX -

The following information is for technical support purposes only:

Lab Instance: 12345678 -

Task -

You plan to analyze stock photography and automatically generate captions for the images.

You need to create a service in Azure to analyze the images. The service must be named caption12345678 and must be in the East US Azure region. The solution must use the Free pricing tier.

In the C:\Resources\Caption\Params.json folder, enter the value for Key 1 and the endpoint for the new service.

To complete this task, sign in to the Azure portal.

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 43

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

Use the following login credentials as needed:

To enter your username, place your cursor in the Sign in box and click on the username below.

To enter your password, place your cursor in the Enter password box and click on the password below.

Azure Username: admin@abc.com -

Azure Password: XXXXXXXXXXXX -

The following information is for technical support purposes only:

Lab Instance: 12345678 -

Task -

You plan to build an application that will use caption12345678. The application will be deployed to a virtual network named VNet1.

You need to ensure that only virtual machines on VNet1 can access caption12345678.

To complete this task, sign in to the Azure portal.

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 44

Topic #: 1

[\[All AI-102 Questions\]](#)

SIMULATION -

Use the following login credentials as needed:

To enter your username, place your cursor in the Sign in box and click on the username below.

To enter your password, place your cursor in the Enter password box and click on the password below.

Azure Username: admin@abc.com -

Azure Password: XXXXXXXXXXXX -

The following information is for technical support purposes only:

Lab Instance: 12345678 -

Task -

You need to ensure that a user named admin@abc.com can regenerate the subscription keys of AAA12345678. The solution must use the principle of least privilege.

To complete this task, sign in to the Azure portal.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 45

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an Azure IoT hub that receives sensor data from machinery.

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

- Perform anomaly detection across multiple correlated sensors.
- Identify the root cause of process stops.
- Send incident alerts.

The solution must minimize development time.

Which Azure service should you use?

- A. Azure Metrics Advisor
- B. Form Recognizer
- C. Azure Machine Learning
- D. Anomaly Detector

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 46

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an app that analyzes images by using the Computer Vision API.

You need to configure the app to provide an output for users who are vision impaired. The solution must provide the output in complete sentences.

Which API call should you perform?

- A. `readInStreamAsync`
- B. `analyzeImagesByDomainInStreamAsync`
- C. `tagImageInStreamAsync`
- D. `describeImageInStreamAsync`

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 47

Topic #: 1

[\[All AI-102 Questions\]](#)

DRAG DROP

You have a Custom Vision service project that performs object detection. The project uses the General domain for classification and contains a trained model.

You need to export the model for use on a network that is disconnected from the internet.

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

Change the classification type.

Export the model.

Retrain the model.

Change Domains to **General (compact)**.

Create a new classification model.



Answer Area



Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 48

Topic #: 1

[\[All AI-102 Questions\]](#)

You are building an AI solution that will use Sentiment Analysis results from surveys to calculate bonuses for customer service staff.

You need to ensure that the solution meets the Microsoft responsible AI principles.

What should you do?

- A. Add a human review and approval step before making decisions that affect the staff's financial situation.
- B. Include the Sentiment Analysis results when surveys return a low confidence score.
- C. Use all the surveys, including surveys by customers who requested that their account be deleted and their data be removed.
- D. Publish the raw survey data to a central location and provide the staff with access to the location.

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 49

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an Azure subscription that contains a Language service resource named ta1 and a virtual network named vnet1.

You need to ensure that only resources in vnet1 can access ta1.

What should you configure?

- A. a network security group (NSG) for vnet1
- B. Azure Firewall for vnet1
- C. the virtual network settings for ta1
- D. a Language service container for ta1

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 50

Topic #: 1

[\[All AI-102 Questions\]](#)

You are developing a monitoring system that will analyze engine sensor data, such as rotation speed, angle, temperature, and pressure. The system must generate an alert in response to atypical values.

What should you include in the solution?

- A. Application Insights in Azure Monitor
- B. metric alerts in Azure Monitor
- C. Multivariate Anomaly Detection
- D. Univariate Anomaly Detection

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 51

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an app named App1 that uses an Azure Cognitive Services model to identify anomalies in a time series data stream.

You need to run App1 in a location that has limited connectivity. The solution must minimize costs.

What should you use to host the model?

- A. Azure Kubernetes Service (AKS)
- B. Azure Container Instances
- C. a Kubernetes cluster hosted in an Azure Stack Hub integrated system
- D. the Docker Engine

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 52

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You have an Azure Cognitive Search resource named Search1 that is used by multiple apps.

You need to secure Search1. The solution must meet the following requirements:

- Prevent access to Search1 from the internet.
- Limit the access of each app to specific queries.

What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

To prevent access from the internet:

- Configure an IP firewall.
- Create a private endpoint.
- Use Azure roles.

To limit access to queries:

- Create a private endpoint.
- Use Azure roles.
- Use key authentication.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 53

Topic #: 1

[\[All AI-102 Questions\]](#)

You are building a solution that will detect anomalies in sensor data from the previous 24 hours.

You need to ensure that the solution scans the entire dataset, at the same time, for anomalies.

Which type of detection should you use?

- A. batch
- B. streaming
- C. change points

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 54

Topic #: 1

[\[All AI-102 Questions\]](#)

DRAG DROP

You are building an app that will scan confidential documents and use the Language service to analyze the contents.

You provision an Azure Cognitive Services resource.

You need to ensure that the app can make requests to the Language service endpoint. The solution must ensure that confidential documents remain on-premises.

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

Run the container and specify an App ID and Client Secret.

Provision an on-premises Kubernetes cluster that is isolated from the internet.

Pull an image from the Microsoft Container Registry (MCR).

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Provision an on-premises Kubernetes cluster that has internet connectivity.

Pull an image from Docker Hub.

Provision an Azure Kubernetes Service (AKS) resource.



Answer Area



Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 55

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT

You have an Azure subscription that has the following configurations:

- Subscription ID: 8d3591aa-96b8-4737-ad09-00f9b1ed35ad
- Tenant ID: 3edfe572-cb54-3ced-ae12-c5c177f39a12

You plan to create a resource that will perform sentiment analysis and optical character recognition (OCR).

You need to use an HTTP request to create the resource in the subscription. The solution must use a single key and endpoint.

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

NOTE: Each correct selection is worth one point.

Answer Area

`https://management.azure.com/`

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 56

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an Azure subscription that contains an Anomaly Detector resource.

You deploy a Docker host server named Server1 to the on-premises network.

You need to host an instance of the Anomaly Detector service on Server1.

Which parameter should you include in the docker run command?

- A. Fluentd
- B. Billing
- C. Http Proxy
- D. Mounts

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 57

Topic #: 1

[\[All AI-102 Questions\]](#)

You are building an app that will use the Speech service.

You need to ensure that the app can authenticate to the service by using a Microsoft Azure Active Directory (Azure AD), part of Microsoft Entra, token.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable a virtual network service endpoint.
- B. Configure a custom subdomain.
- C. Request an X.509 certificate.
- D. Create a private endpoint.
- E. Create a Conditional Access policy.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 58

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You plan to deploy an Azure OpenAI resource by using an Azure Resource Manager (ARM) template.

You need to ensure that the resource can respond to 600 requests per minute.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity":  : 
    "count": 
    "maxValue": 
    "size": 
  },
  "properties": {
    "model": {
      "format": "OpenAI",
    }
  }
}
```

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 59

Topic #: 1

[\[All AI-102 Questions\]](#)

DRAG DROP

You have an app that manages feedback.

You need to ensure that the app can detect negative comments by using the Sentiment Analysis API in Azure AI Language. The solution must ensure that the managed feedback remains on your company's internal network.

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.

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

Actions

Identify the Language service endpoint URL and query the prediction endpoint.

Provision the Language service resource in Azure.

Run the container and query the prediction endpoint.

Deploy a Docker container to an on-premises server.

Deploy a Docker container to an Azure container instance.

Answer Area

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 60

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You have an Azure OpenAI 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

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 61

Topic #: 1

[\[All AI-102 Questions\]](#)

You build a bot by using the Microsoft Bot Framework SDK.

You start the bot on a local computer.

You need to validate the functionality of the bot.

What should you do before you connect to the bot?

- A. Run the Bot Framework Emulator.
- B. Run the Bot Framework Composer.
- C. Register the bot with Azure Bot Service.
- D. Run Windows Terminal.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 62

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an Azure OpenAI model named AI1.

You are building a web app named App1 by using the Azure OpenAI SDK.

You need to configure App1 to connect to AI1.

What information must you provide?

- A. the endpoint, key, and model name
- B. the deployment name, key, and model name
- C. the deployment name, endpoint, and key
- D. the endpoint, key, and model type

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 63

Topic #: 1

[\[All AI-102 Questions\]](#)

You are building a solution in Azure that will use Azure Cognitive Service for Language to process sensitive customer data.

You need to ensure that only specific Azure processes can access the Language service. The solution must minimize administrative effort.

What should you include in the solution?

- A. IPsec rules
- B. Azure Application Gateway
- C. a virtual network gateway
- D. virtual network rules

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 64

Topic #: 1

[\[All AI-102 Questions\]](#)

You plan to perform predictive maintenance.

You collect IoT sensor data from 100 industrial machines for a year. Each machine has 50 different sensors that generate data at one-minute intervals. In total, you have 5,000 time series datasets.

You need to identify unusual values in each time series to help predict machinery failures.

Which Azure service should you use?

- A. Azure AI Computer Vision
- B. Cognitive Search
- C. Azure AI Document Intelligence
- D. Azure AI Anomaly Detector

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 65

Topic #: 1

[\[All AI-102 Questions\]](#)

HOTSPOT

You plan to deploy a containerized version of an Azure Cognitive Services 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-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
```

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 66

Topic #: 1

[\[All AI-102 Questions\]](#)

You are developing a system that will monitor temperature data from a data stream. The system must generate an alert in response to atypical values. The solution must minimize development effort.

What should you include in the solution?

- A. Multivariate Anomaly Detection
- B. Azure Stream Analytics
- C. metric alerts in Azure Monitor
- D. Univariate Anomaly Detection

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 67

Topic #: 1

[\[All AI-102 Questions\]](#)

You have a Microsoft OneDrive folder that contains a 20-GB video file named File1.avi.

You need to index File1.avi by using the Azure Video Indexer website.

What should you do?

- A. Upload File1.avi to the www.youtube.com webpage, and then copy the URL of the video to the Azure AI Video Indexer website.
- B. Download File1.avi to a local computer, and then upload the file to the Azure AI Video Indexer website.
- C. From OneDrive, create a download link, and then copy the link to the Azure AI Video Indexer website.
- D. From OneDrive, create a sharing link for File1.avi, and then copy the link to the Azure AI Video Indexer website.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 68

Topic #: 1

[\[All AI-102 Questions\]](#)

You have an Azure subscription that contains an Azure AI Service resource named CSAccount1 and a virtual network named VNet1. CSAccount1 is connected to VNet1.

You need to ensure that only specific resources can access CSAccount1. The solution must meet the following requirements:

- Prevent external access to CSAccount1.
- Minimize administrative effort.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct answer is worth one point.

- A. In VNet1, enable a service endpoint for CSAccount1.
- B. In CSAccount1, configure the Access control (IAM) settings.
- C. In VNet1, modify the virtual network settings.
- D. In VNet1, create a virtual subnet.
- E. In CSAccount1, modify the virtual network settings.

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 69

Topic #: 1

[\[All AI-102 Questions\]](#)

You are building an internet-based training solution. The solution requires that a user's camera and microphone remain enabled.

You need to monitor a video stream of the user and detect when the user asks an instructor a question. The solution must minimize development effort.

What should you include in the solution?

- A. speech-to-text in the Azure AI Speech service
- B. language detection in Azure AI Language Service
- C. the Face service in Azure AI Vision
- D. object detection in Azure AI Custom Vision

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 1

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are developing an application that will use the Computer Vision client library. The application has the following code.

```
public async Task AnalyzeImage(ComputerVisionClient client, string localImage)
{
    List<VisualFeatureTypes> features = new List<VisualFeatureTypes>()
    {
        VisualFeatureTypes.Description,
        VisualFeatureTypes.Tags,
    };
    using (Stream imageStream = File.OpenRead(localImage))
    {
        try
        {
            ImageAnalysis results = await client.AnalyzeImageInStreamAsync(imageStream, features);

            foreach (var caption in results.Description.Captions)
            {
                Console.WriteLine($"{caption.Text} with confidence {caption.Confidence}");
            }

            foreach (var tag in results.Tags)
            {
                Console.WriteLine($"{tag.Name} {tag.Confidence}");
            }
        }
        catch (Exception ex)
        {
            Console.WriteLine(ex.Message);
        }
    }
}
```

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 perform face recognition.	<input type="radio"/>	<input type="radio"/>
The code will list tags and their associated confidence.	<input type="radio"/>	<input type="radio"/>
The code will read a file from the local file system.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 2

Topic #: 2

[\[All AI-102 Questions\]](#)

You are developing a method that uses the Computer Vision client library. The method will perform optical character recognition (OCR) in images. The method has the following code.

```
public static async Task ReadFileUrl(ComputerVisionClient client, string urlFile)
{
    const int numberOfCharsInOperationId = 36;

    var txtHeaders = await client.ReadAsync(urlFile, language: "en");

    string opLocation = txtHeaders.OperationLocation;
    string operationId = opLocation.Substring(opLocation.Length -
numberOfCharsInOperationId);

    ReadOperationResult results;

    results = await client.GetReadResultAsync(Guid.Parse(operationId));

    var textUrlFileResults = results.AnalyzeResult.ReadResults;
    foreach (ReadResult page in textUrlFileResults)
    {
        foreach (Line line in page.Lines)
        {
            Console.WriteLine(line.Text);
        }
    }
}
```

During testing, you discover that the call to the `GetReadResultAsync` method occurs before the read operation is complete.

You need to prevent the `GetReadResultAsync` method from proceeding until the read operation is complete.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Remove the `Guid.Parse(operationId)` parameter.
- B. Add code to verify the `results.Status` value.
- C. Add code to verify the status of the `txtHeaders.Status` value.
- D. Wrap the call to `GetReadResultAsync` within a loop that contains a delay.

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 3

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT -

You have a Computer Vision resource named contoso1 that is hosted in the West US Azure region.

You need to use contoso1 to make a different size of a product photo by using the smart cropping feature.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

```
curl -H "Ocp-Apim-Subscription-Key: xxx" /
```

```
-o "sample.png" -H "Content-Type: application/json" /
```

/vision/v3.1/

?width=100&height=100&smartCropping=true" /

```
-d "{\"url\":\"https://upload.litwareinc.org/litware/bicycle.jpg\"}"
```

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 4

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP -

You are developing a webpage that will use the Azure Video Analyzer for Media (previously Video Indexer) service to display videos of internal company meetings.

You embed the Player widget and the Cognitive Insights widget into the page.

You need to configure the widgets to meet the following requirements:

- ⇒ Ensure that users can search for keywords.
- ⇒ Display the names and faces of people in the video.
- ⇒ Show captions in the video in English (United States).

How should you complete the URL for each widget? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Values

en-US

false

people,keywords

people,search

search

true

Answer Area

Cognitive Insights Widget

`https://www.videoindexer.ai/embed/insights/<accountId>/<videoId>/?widgets=`

Value

controls=

Value

Player Widget

`https://www.videoindexer.ai/embed/player/<accountId>/<videoId>/? showcaptions=`

Value

captions=

Value

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 5

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP -

You train a Custom Vision model to identify a company's products by using the Retail domain.

You plan to deploy the model as part of an app for Android phones.

You need to prepare the model for deployment.

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.

Select and Place:

Actions

Change the model domain.

Retrain the model.

Test the model.

Export the model.

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 6

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are developing an application to recognize employees' faces by using the Face Recognition API. Images of the faces will be accessible from a URI endpoint.

The application has the following code.

```
def add_face(subscription_key, person_group_id, person_id, image_uri):
    headers = {
        'Content-Type': 'application/json',
        'Ocp-Apim-Subscription-Key': subscription_key
    }
    body = {
        'url': image_uri }
    conn = httplib.HTTPSConnection('westus.api.cognitive.microsoft.com')
    conn.request('POST',
f'/face/v1.0/persongroups/{person_group_id}/persons/{person_id}/persistedFaces', f'{body}', headers)
    response = conn.getresponse()
    response_data = response.read()
```

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 add a face image to a person object in a person group.	<input type="radio"/>	<input type="radio"/>
The code will work for up to 10,000 people.	<input type="radio"/>	<input type="radio"/>
add_face can be called multiple times to add multiple face images to a person object.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 7

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP -

You have a Custom Vision resource named acvdev in a development environment.

You have a Custom Vision resource named acvprod in a production environment.

In acvdev, you build an object detection model named obj1 in a project named proj1.

You need to move obj1 to acvprod.

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.

Select and Place:

Actions

Use the `ExportProject` endpoint on acvdev.

Use the `GetProjects` endpoint on acvdev.

Use the `ImportProject` endpoint on acvprod.

Use the `ExportIteration` endpoint on acvdev.

Use the `GetIterations` endpoint on acvdev.

Use the `UpdateProject` endpoint on acvprod.

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 8

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP -

You are developing an application that will recognize faults in components produced on a factory production line. The components are specific to your business.

You need to use the Custom Vision API to help detect common faults.

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.

Select and Place:

Actions

Train the classifier model.

Upload and tag images.

Initialize the training dataset.

Train the object detection model.

Create a project.



Answer Area



Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 9

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are building a model that will be used in an iOS app.

You have images of cats and dogs. Each image contains either a cat or a dog.

You need to use the Custom Vision service to detect whether the images is of a cat or a dog.

How should you configure the project in the Custom Vision portal? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Project Types: ▼

Classification
Object Detection

Classification Types: ▼

Multiclass (Single tag per image)
Multilabel (Multiple tags per image)

Domains: ▼

Audit
Food
General
General (compact)
Landmarks
Landmarks (compact)
Retail
Retail (compact)

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 10

Topic #: 2

[\[All AI-102 Questions\]](#)

You have an Azure Video Analyzer for Media (previously 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.

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 11

Topic #: 2

[\[All AI-102 Questions\]](#)

You use the Custom Vision service to build a classifier.

After training is complete, you need to evaluate the classifier.

Which two metrics are available for review? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. recall
- B. F-score
- C. weighted accuracy
- D. precision
- E. area under the curve (AUC)

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 12

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP -

You are developing a call to the Face API. The call must find similar faces from an existing list named `employeefaces`. The `employeefaces` list contains 60,000 images.

How should you complete the body of the HTTP request? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Values

"faceListId"

"LargeFaceListId"

"matchFace"

"matchPerson"

Answer Area

```
{  
  "faceId": "18c51a87-3a69-47a8-aedc-a54745f708a1",  
  [ ]: "employeefaces",  
  "maxNumOfCandidatesReturned": 1,  
  "mode": [ ]  
}
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 13

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP -

You are developing a photo application that will find photos of a person based on a sample image by using the Face API.

You need to create a POST request to find the photos.

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

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Values

Answer Area

POST {Endpoint}/face/v1.0/

Request Body

```
{
  "faceId": "c5c24a82-6845-4031-9d5d-978df9175426",
  "largeFaceListId": "sample_list",
  "maxNumOfCandidatesReturned": 10,
  "mode": " "
}
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 14

Topic #: 2

[\[All AI-102 Questions\]](#)

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.

```
for brand in image_analysis.brands:
    if brand_confidence >= 0.75:
        print(f"\nLogo 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="radio"/>	<input type="radio"/>
The code will return coordinates for the top-left corner of the rectangle that contains the brand logo of the displayed brands.	<input type="radio"/>	<input type="radio"/>
The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 15

Topic #: 2

[\[All AI-102 Questions\]](#)

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);
        }
    }
});
```

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 16

Topic #: 2

[\[All AI-102 Questions\]](#)

Your company uses an Azure Cognitive Services solution to detect faces in uploaded images. The method to detect the faces uses the following code.

```
static async Task DetectFaces(string imagePath)
{
    HttpClient client = new HttpClient();
    DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", subscriptionKey);
    string requestParameter = "detectionModel=detection_01&returnFaceId=true&returnFaceLandmarks=false";
    string uri = endpoint + "/face/v1.0/detect?" + requestParameters;
    HttpResponseMessage response;
    byte[] byteData = GetImagesAsByteArray(imageFilePath);
    using (ByteArrayContent content = new ByteArrayContent(byteData))
    {
        Headers.ContentType = new MediaTypeHeaderValue("application/octet-stream");
        response = await PostAsync(uri, content);
        string contentString = await Content.ReadAsStringAsync();
        ProcessDetection(contentString);
    }
}
```

You discover that the solution frequently fails to detect faces in blurred images and in images that contain sideways faces.

You need to increase the likelihood that the solution can detect faces in blurred images and images that contain sideways faces.

What should you do?

- A. Use a different version of the Face API.
- B. Use the Computer Vision service instead of the Face service.
- C. Use the Identify method instead of the Detect method.
- D. Change the detection model.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 17

Topic #: 2

[\[All AI-102 Questions\]](#)

You have the following Python function for creating Azure Cognitive Services resources programmatically. `def create_resource (resource_name, kind, account_tier, location) : parameters = CognitiveServicesAccount(sku=Skus(name=account_tier), kind=kind, location=location, properties={}) result = client.accounts.create(resource_group_name, resource_name, parameters)`

You need to call the function to create a free Azure resource in the West US Azure region. The resource will be used to generate captions of images automatically. Which code should you use?

- A. `create_resource("res1", "ComputerVision", "F0", "westus")`
- B. `create_resource("res1", "CustomVision.Prediction", "F0", "westus")`
- C. `create_resource("res1", "ComputerVision", "S0", "westus")`
- D. `create_resource("res1", "CustomVision.Prediction", "S0", "westus")`

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 18

Topic #: 2

[\[All AI-102 Questions\]](#)

You are developing a method that uses the Computer Vision client library. The method will perform optical character recognition (OCR) in images. The method has the following code.

```
def read_file_url(computervision_client, url_file):
    read_response = computervision_client.read(url_file, raw=True)
    read_operation_location = read_response.headers["Operation-Location"]
    operation_id = read_operation_location.split("/")[-1]
    read_result = computervision_client.get_read_result(operation_id)

    for page in read_result.analyze_result.read_results:
        for line in page.lines:
            print(line.text)
```

During testing, you discover that the call to the GetReadResultAsync method occurs before the read operation is complete.

You need to prevent the GetReadResultAsync method from proceeding until the read operation is complete.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Remove the operation_id parameter.
- B. Add code to verify the read_results.status value.
- C. Add code to verify the status of the read_operation_location value.
- D. Wrap the call to get_read_result within a loop that contains a delay.

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 19

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are building an app that will enable users to upload images. The solution must meet the following requirements:

- * Automatically suggest alt text for the images.
- * Detect inappropriate images and block them.
- * Minimize development effort.

You need to recommend a computer vision endpoint for each requirement.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Generate alt text:

<input type="checkbox"/>	https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate
<input type="checkbox"/>	https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image
<input type="checkbox"/>	https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Descnption

Detect inappropriate content:

<input type="checkbox"/>	https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate
<input type="checkbox"/>	https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image
<input type="checkbox"/>	https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description
<input type="checkbox"/>	https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 20

Topic #: 2

[\[All AI-102 Questions\]](#)

You need to build a solution that will use optical character recognition (OCR) to scan sensitive documents by using the Computer Vision API. The solution must NOT be deployed to the public cloud.

What should you do?

- A. Build an on-premises web app to query the Computer Vision endpoint.
- B. Host the Computer Vision endpoint in a container on an on-premises server.
- C. Host an exported Open Neural Network Exchange (ONNX) model on an on-premises server.
- D. Build an Azure web app to query the Computer Vision endpoint.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 21

Topic #: 2

[\[All AI-102 Questions\]](#)

You have an Azure Cognitive Search solution and a collection of handwritten letters stored as JPEG files.

You plan to index the collection. The solution must ensure that queries can be performed on the contents of the letters.

You need to create an indexer that has a skillset.

Which skill should you include?

- A. image analysis
- B. optical character recognition (OCR)
- C. key phrase extraction
- D. document extraction

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 22

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You have a library that contains thousands of images.

You need to tag the images as photographs, drawings, or clipart.

Which service endpoint and response property should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Service endpoint:

- Computer Vision analyze images
- Computer Vision object detection
- Custom Vision image classification
- Custom Vision object detection

Property:

- categories
- description
- imageType
- metadata
- objects

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 23

Topic #: 2

[\[All AI-102 Questions\]](#)

You have an app that captures live video of exam candidates.

You need to use the Face service to validate that the subjects of the videos are real people.

What should you do?

- A. Call the face detection API and retrieve the face rectangle by using the FaceRectangle attribute.
- B. Call the face detection API repeatedly and check for changes to the FaceAttributes.HeadPose attribute.
- C. Call the face detection API and use the FaceLandmarks attribute to calculate the distance between pupils.
- D. Call the face detection API repeatedly and check for changes to the FaceAttributes.Accessories attribute.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 24

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT

You make an API request and receive the results shown in the following exhibits.

HTTP request

```
POST https://facetesting.cognitiveservices.azure.com/face/v1.0/detect?returnFaceId=true&returnFaceLandmarks=false&returnFaceAttributes=qualityforrecognition&recognitionModel=recognition_04&returnRecognitionModel=false&detectioModel=detection_03&faceIdTimeToLive=86400 HTTP/1.1
Host: facetesting.cognitiveservices.azure.com
Content-Type: application/json
Ocp-Apim-Subscription-Key: *****

{
  "url": "https://news.microsoft.com/wp-content/uploads/prod/sites/68/2021/11/EDU19_HigherEdStudentsOnCampus_002-1536x1024.jpg"
}
```

Send

Response status

200 OK

Response content

```
x-envoy-upstream-service-time: 1292
apim-request-id: 8a3aa72f-5bad-45d0-b8a4-584312258f06
Strict-Transport-Security: max-age=31536000; includeSubDomains; preload
x-content-type-options: nosniff
CSP-Billing-Usage: CognitiveServices.Face.Transaction=1
Date: Sat, 04 Dec 2021 11:15:33 GMT
Content-Length: 655
Content-Type: application/json; charset=utf-8
```

```
[{
  "faceId": "d14d131c-76ba-43e9-9e3d-dcf6466e5022",
  "faceRectangle": {
    "top": 201,
    "left": 797,
    "width": 121,
    "height": 160
  },
  "faceAttributes": {
    "qualityForRecognition": "high"
  }
}, {
  "faceId": "a3a0f2ff-b015-464c-b87c-0dd09d0698da",
  "faceRectangle": {
    "top": 249,
    "left": 1167,
    "width": 103,
    "height": 159
  },
  "faceAttributes": {
    "qualityForRecognition": "medium"
  }
}, {
  "faceId": "45481ce8-dcc4-4564-a21c-3c15cdc9c4fa",
  "faceRectangle": {
    "top": 191,
    "left": 497,
    "width": 85,
    "height": 178
  },
  "faceAttributes": {
    "qualityForRecognition": "low"
  }
}, {
  "faceId": "eac17649-effd-42c9-9093-4dd60fd4cfc7",
  "faceRectangle": {
    "top": 754,
    "left": 118,
    "width": 30,
    "height": 44
  },
  "faceAttributes": {
    "qualityForRecognition": "low"
  }
}]
```

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer Area

The API **[answer choice]** faces.

- detects
- finds similar
- recognizes
- verifies

A face that can be used in person enrollment is at position **[answer choice]** within the photo.

- 118, 754
- 497, 191
- 797, 201
- 1167, 249

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 25

Topic #: 2

[\[All AI-102 Questions\]](#)

You have an Azure subscription that contains an AI enrichment pipeline in Azure Cognitive Search and an Azure Storage account that has 10 GB of scanned documents and images.

You need to index the documents and images in the storage account. The solution must minimize how long it takes to build the index.

What should you do?

- A. From the Azure portal, configure parallel indexing.
- B. From the Azure portal, configure scheduled indexing.
- C. Configure field mappings by using the REST API.
- D. Create a text-based indexer by using the REST API.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 26

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP

You need to analyze video content to identify any mentions of specific company names.

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

Add the specific company names to the exclude list.

Add the specific company names to the include list.

From Content model customization, select **Language**.

Sign in to the Custom Vision website.

Sign in to the Azure Video Analyzer for Media website.

From Content model customization, select **Brands**.



Answer Area



Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 27

Topic #: 2

[\[All AI-102 Questions\]](#)

You have a mobile app that manages printed forms.

You need the app to send images of the forms directly to Forms Recognizer to extract relevant information. For compliance reasons, the image files must not be stored in the cloud.

In which format should you send the images to the Form Recognizer API endpoint?

- A. raw image binary
- B. form URL encoded
- C. JSON

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 28

Topic #: 2

[\[All AI-102 Questions\]](#)

You plan to build an app that will generate a list of tags for uploaded images. The app must meet the following requirements:

- Generate tags in a user's preferred language.
- Support English, French, and Spanish.
- Minimize development effort.

You need to build a function that will generate the tags for the app.

Which Azure service endpoint should you use?

- A. Content Moderator Image Moderation
- B. Custom Vision image classification
- C. Computer Vision Image Analysis
- D. Custom Translator

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 29

Topic #: 2

[\[All AI-102 Questions\]](#)

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.

Answer Area

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

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 30

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP

-

You have a factory that produces cardboard packaging for food products. The factory has intermittent internet connectivity.

The packages are required to include four samples of each product.

You need to build a Custom Vision model that will identify defects in packaging and provide the location of the defects to an operator. The model must ensure that each package contains the four products.

Which project type and domain should you use? To answer, drag the appropriate options to the correct targets. Each option 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.

Options

Food

General

General (compact)

Image classification

Logo

Object detection

Answer Area

Project type:

Domain:

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 31

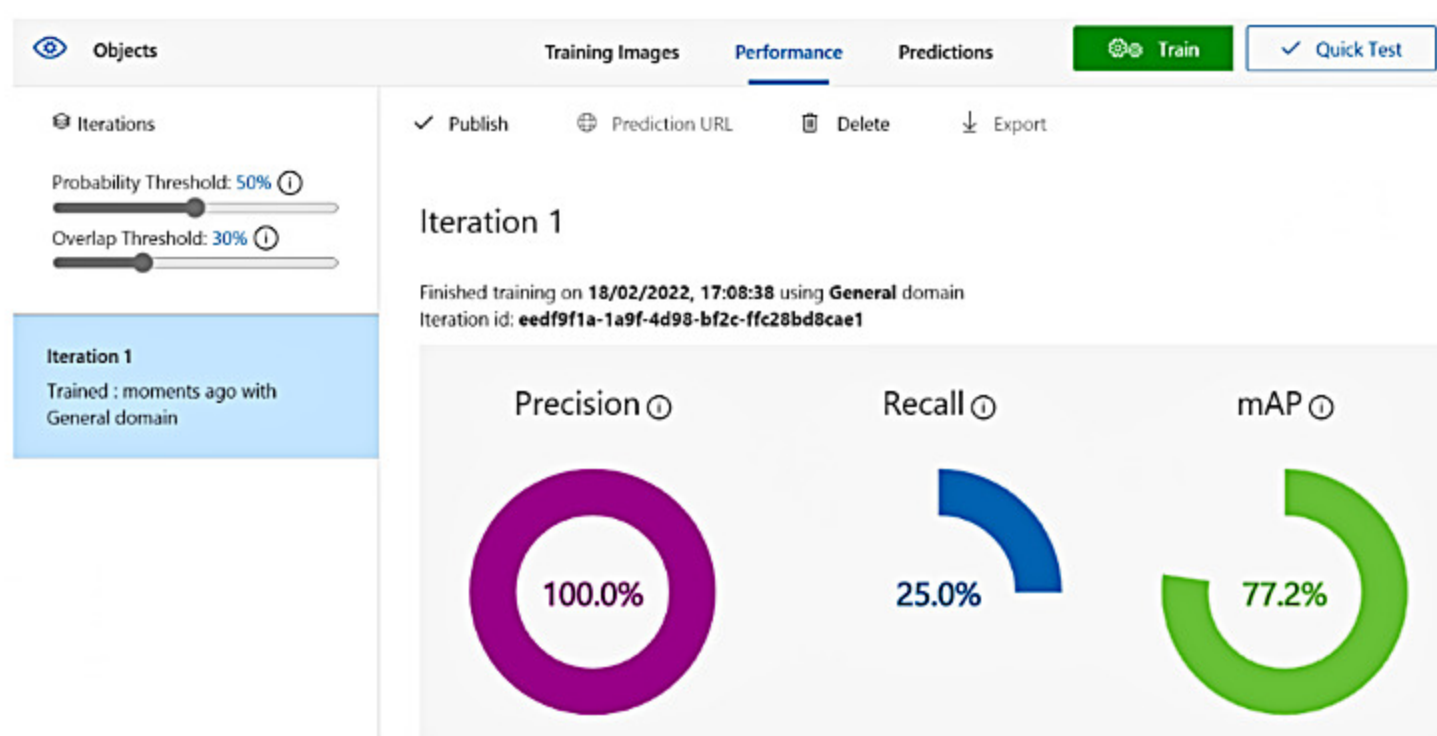
Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT

You are building a model to detect objects in images.

The performance of the model based on training data is shown in the following exhibit.



Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer Area

The percentage of false positives is [answer choice].

The value for the number of true positives divided by the total number of true positives and false negatives is [answer choice]%.

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 32

Topic #: 2

[\[All AI-102 Questions\]](#)

You are building an app that will include one million scanned magazine articles. Each article will be stored as an image file.

You need to configure the app to extract text from the images. The solution must minimize development effort.

What should you include in the solution?

- A. Computer Vision Image Analysis
- B. the Read API in Computer Vision
- C. Form Recognizer
- D. Azure Cognitive Service for Language

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 33

Topic #: 2

[\[All AI-102 Questions\]](#)

You have a 20-GB video file named File1.avi that is stored on a local drive.

You need to index File1.avi by using the Azure Video Indexer website.

What should you do first?

- A. Upload File1.avi to an Azure Storage queue.
- B. Upload File1.avi to the Azure Video Indexer website.
- C. Upload File1.avi to Microsoft OneDrive.
- D. Upload File1.avi to the www.youtube.com webpage.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 34

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are building an app that will share user images.

You need to configure the app to meet the following requirements:

- Uploaded images must be scanned and any text must be extracted from the images.
- Extracted text must be analyzed for the presence of profane language.
- The solution must minimize development effort.

What should you use for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Text extraction:

- Azure AI Language
- Azure AI Computer Vision
- Content Moderator
- Azure AI Custom Vision
- Azure AI Document Intelligence

Profane language detection:

- Azure AI Language
- Azure AI Computer Vision
- Content Moderator
- Azure AI Custom Vision
- Azure AI Document Intelligence

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 35

Topic #: 2

[\[All AI-102 Questions\]](#)

You are building an app that will share user images.

You need to configure the app to perform the following actions when a user uploads an image:

- Categorize the image as either a photograph or a drawing.
- Generate a caption for the image.

The solution must minimize development effort.

Which two services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. object detection in Azure AI Computer Vision
- B. content tags in Azure AI Computer Vision
- C. image descriptions in Azure AI Computer Vision
- D. image type detection in Azure AI Computer Vision
- E. image classification in Azure AI Custom Vision

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 36

Topic #: 2

[\[All AI-102 Questions\]](#)

You are building an app that will use the Azure AI Video Indexer service.

You plan to train a language model to recognize industry-specific terms.

You need to upload a file that contains the industry-specific terms.

Which file format should you use?

- A. XML
- B. TXT
- C. XLS
- D. PDF

[Show Suggested Answer](#)



Actual exam question from Microsoft's AI-102

Question #: 37

Topic #: 2

[\[All AI-102 Questions\]](#)

DRAG DROP

You have an app that uses Azure AI 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

Label the sample images.

From Vision Studio, open the project.

Publish the model.

From the Custom Vision portal, open the project.

Retrain the model.

Upload sample images of the new products.

From the Azure Machine Learning studio, open the workspace.

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 38

Topic #: 2

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are developing an application that will use the Azure AI Vision client library. The application has the following code.

```
def analyze_image(local_image):
    with open(local_image, "rb") as image_stream:
        image_analysis = client.analyze_image_in_stream(
            image=image_stream,
            visual_features=[
                VisualFeatureTypes.tags,
                VisualFeatureTypes.description
            ]
        )
    for caption in image_analysis.description.captions:
        print(f"\n{caption.text} with confidence {caption.confidence}")
    for tag in image_analysis.tags:
        print(f"\n{tag.name} with confidence {tag.confidence}")
```

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 code will perform face recognition.	<input type="radio"/>	<input type="radio"/>
The code will list tags and their associated confidence.	<input type="radio"/>	<input type="radio"/>
The code will read an image file from the local file system.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 1

Topic #: 3

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training.

- ⇒ Find contacts in London.
- ⇒ Who do I know in Seattle?
- ⇒ Search for contacts in Ukraine.

You need to implement the phrase list in Language Understanding.

Solution: You create a new pattern in the FindContact intent.

Does this meet the goal?

- A. Yes
- B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 2

Topic #: 3

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You develop an application to identify species of flowers by training a Custom Vision model.

You receive images of new flower species.

You need to add the new images to the classifier.

Solution: You add the new images, and then use the Smart Labeler tool.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 3

Topic #: 3

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You develop an application to identify species of flowers by training a Custom Vision model.

You receive images of new flower species.

You need to add the new images to the classifier.

Solution: You add the new images and labels to the existing model. You retrain the model, and then publish the model.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 4

Topic #: 3

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You develop an application to identify species of flowers by training a Custom Vision model.

You receive images of new flower species.

You need to add the new images to the classifier.

Solution: You create a new model, and then upload the new images and labels.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 5

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are developing a service that records lectures given in English (United Kingdom).

You have a method named `AppendToTranscriptFile` that takes translated text and a language identifier.

You need to develop code that will provide transcripts of the lectures to attendees in their respective language. The supported languages are English, French, Spanish, and German.

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

```
static async Task TranslateSpeechAsync()
{
    var config = SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");

    var lang = new List<string>
    {
        {"en-GB"},
        {"fr", "de", "es"},
        {"French", "Spanish", "German"}
    };

    config.SpeechRecognitionLanguage = "en-GB";
    lang.ForEach(config.AddTargetLanguage);

    using var audioConfig = AudioConfig.FromDefaultMicrophoneInput();
    using var recognizer = new
    {
        IntentRecognizer,
        SpeakerRecognizer,
        SpeechSynthesizer,
        TranslationRecognizer
    } (config, audioConfig);

    var result = await recognizer.RecognizeOnceAsync();
    if (result.Reason == ResultReason.TranslatedSpeech)
```

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 6

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP -

You train a Custom Vision model used in a mobile app.

You receive 1,000 new images that do not have any associated data.

You need to use the images to retrain the model. The solution must minimize how long it takes to retrain the model.

Which three actions should you perform in the Custom Vision portal? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

Actions

Upload the images by category.

Get suggested tags.

Upload all the images.

Group the images locally into category folders.

Review the suggestions and confirm the tags.

Tag the images manually.

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 7

Topic #: 3

[\[All AI-102 Questions\]](#)

You are building a Conversational Language Understanding model for an e-commerce chatbot. Users can speak or type their billing address when prompted by the chatbot.

You need to construct an entity to capture billing addresses.

Which entity type should you use?

- A. machine learned
- B. Regex
- C. list
- D. Pattern.any

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 8

Topic #: 3

[\[All AI-102 Questions\]](#)

You are building an Azure WebJob that will create knowledge bases from an array of URLs.

You instantiate a QnAMakerClient object that has the relevant API keys and assign the object to a variable named client.

You need to develop a method to create the knowledge bases.

Which two actions should you include in the method? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Create a list of FileDTO objects that represents data from the WebJob.
- B. Call the client.Knowledgebase.CreateAsync method.
- C. Create a list of QnADTO objects that represents data from the WebJob.
- D. Create a CreateKbDTO object.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 9

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are developing an application that includes language translation.

The application will translate text retrieved by using a function named `getTextToBeTranslated`. The text can be in one of many languages. The content of the text must remain within the Americas Azure geography.

You need to develop code to translate the text to a single language.

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

...

```
var endpoint =
```

<code>"https://api.cognitive.microsofttranslator.com/translate"</code>
<code>"https://api.cognitive.microsofttranslator.com/transliterate"</code>
<code>"https://api-apc.cognitive.microsofttranslator.com/detect"</code>
<code>"https://api-nam.cognitive.microsofttranslator.com/detect"</code>
<code>"https://api-nam.cognitive.microsofttranslator.com/translate"</code>

```
;
```

```
var apiKey = "FF956C68B83B21B38691ABD200A4C606";  
var text = getTextToBeTranslated();  
var body = '[{"Text":"' + text + '"}]';  
var client = new HttpClient();  
client.DefaultRequestHeaders.Add("Ocp-Api-Subscription-Key", apiKey);
```

<code>var uri = endpoint + "?from=en";</code>
<code>var uri = endpoint + "?suggestedFrom=en";</code>
<code>var uri = endpoint + "?to=en";</code>

```
HttpResponseMessage response;  
var content = new StringContent(body, Encoding.UTF8, "application/json");  
var response = await client.PutAsync(uri, content);  
...
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 10

Topic #: 3

[\[All AI-102 Questions\]](#)

You are building a conversational language understanding model.

You need to enable active learning.

What should you do?

- A. Add show-all-intents=true to the prediction endpoint query.
- B. Enable speech priming.
- C. Add log=true to the prediction endpoint query.
- D. Enable sentiment analysis.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 11

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT -

You run the following command.

```
docker run --rm -it -p 5000:5000 --memory 10g --cpus 2 \  
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment\  
Eula=accept \  
Billing={ENDPOINT_URI} \  
ApiKey={API_KEY}
```

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
Going to <code>http://localhost:5000/status</code> will query the Azure endpoint to verify whether the API key used to start the container is valid.	<input type="radio"/>	<input type="radio"/>
The container logging provider will write log data.	<input type="radio"/>	<input type="radio"/>
Going to <code>http://localhost:5000/swagger</code> will provide the details to access the documentation for the available endpoints.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 12

Topic #: 3

[\[All AI-102 Questions\]](#)

You are building a Language Understanding model for an e-commerce platform.

You need to construct an entity to capture billing addresses.

Which entity type should you use for the billing address?

- A. machine learned
- B. Regex
- C. geographyV2
- D. Pattern.any
- E. list

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 13

Topic #: 3

[\[All AI-102 Questions\]](#)

You need to upload speech samples to a Speech Studio project for use in training.

How should you upload the samples?

- A. Combine the speech samples into a single audio file in the .wma format and upload the file.
- B. Upload a .zip file that contains a collection of audio files in the .wav format and a corresponding text transcript file.
- C. Upload individual audio files in the FLAC format and manually upload a corresponding transcript in Microsoft Word format.
- D. Upload individual audio files in the .wma format.

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 14

Topic #: 3

[\[All AI-102 Questions\]](#)

You are developing a method for an application that uses the Translator API.

The method will receive the content of a webpage, and then translate the content into Greek (el). The result will also contain a transliteration that uses the Roman alphabet.

You need to create the URI for the call to the Translator API.

You have the following URI.

`https://api.cognitive.microsofttranslator.com/translate?api-version=3.0`

Which three additional query parameters should you include in the URI? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. toScript=Cyrl

B. from=el

C. textType=html

D. to=el

E. textType=plain

F. toScript=Latn

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 15

Topic #: 3

[\[All AI-102 Questions\]](#)

You have a chatbot that was built by using the Microsoft Bot Framework.

You need to debug the chatbot endpoint remotely.

Which two tools should you install on a local computer? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Fiddler
- B. Bot Framework Composer
- C. Bot Framework Emulator
- D. Bot Framework CLI
- E. ngrok
- F. nginx

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 16

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP -

You are building a retail chatbot that will use a QnA Maker service.

You upload an internal support document to train the model. The document contains the following question: "What is your warranty period?"

Users report that the chatbot returns the default QnA Maker answer when they ask the following question: "How long is the warranty coverage?"

The chatbot returns the correct answer when the users ask the following question: "What is your warranty period?"

Both questions should return the same answer.

You need to increase the accuracy of the chatbot responses.

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.

Select and Place:

Actions

Answer Area

Add a new question and answer (QnA) pair.

Retrain the model.

Add additional questions to the document.

Republish the model.

Add alternative phrasing to the question and answer (QnA) pair.

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 17

Topic #: 3

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training.

- ⇒ Find contacts in London.
- ⇒ Who do I know in Seattle?
- ⇒ Search for contacts in Ukraine.

You need to implement the phrase list in Language Understanding.

Solution: You create a new intent for location.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 18

Topic #: 3

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training.

⇒ Find contacts in London.

⇒ Who do I know in Seattle?

Search for contacts in Ukraine.

▪

You need to implement the phrase list in Language Understanding.

Solution: You create a new entity for the domain.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 19

Topic #: 3

[\[All AI-102 Questions\]](#)

You are training a Language Understanding model for a user support system.

You create the first intent named GetContactDetails and add 200 examples.

You need to decrease the likelihood of a false positive.

What should you do?

- A. Enable active learning.
- B. Add a machine learned entity.
- C. Add additional examples to the GetContactDetails intent.
- D. Add examples to the None intent.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 20

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP -

You are building a Language Understanding model for purchasing tickets.

You have the following utterance for an intent named PurchaseAndSendTickets.

Purchase [2 audit business] tickets to [Paris] [next Monday] and send tickets to [email@domain.com]

You need to select the entity types. The solution must use built-in entity types to minimize training data whenever possible.

Which entity type should you use for each label? To answer, drag the appropriate entity types to the correct labels. Each entity type 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.

Select and Place:

Entity Types

Answer Area

Email

List

Regex

GeographyV2

Machine learned

Paris:

email@domain.com:

2 audit business:

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 21

Topic #: 3

[\[All AI-102 Questions\]](#)

You have the following C# method.

```
static void create_resource(string resource_name, string kind, string account_tier, string location)
{
    CognitiveServicesAccount parameters =
        new CognitiveServicesAccount(null, null, kind, location, resource_name, new CognitiveServicesAccountProperties(), new Sku(account_tier));
    var result = cog_svc_client.Accounts.Create(resource_group_name, account_tier, parameters);
}
```

You need to deploy an Azure resource to the East US Azure region. The resource will be used to perform sentiment analysis.

How should you call the method?

- A. `create_resource("res1", "ContentModerator", "S0", "eastus")`
- B. `create_resource("res1", "TextAnalytics", "S0", "eastus")`
- C. `create_resource("res1", "ContentModerator", "Standard", "East US")`
- D. `create_resource("res1", "TextAnalytics", "Standard", "East US")`

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 22

Topic #: 3

[\[All AI-102 Questions\]](#)

You build a Conversational Language Understanding model by using the Language Services portal.

You export the model as a JSON file as shown in the following sample.

```
{
  "text": "average amount of rain by month at chicago last year",
  "intent": "Weather.CheckWeatherValue",
  "entities": [
    {
      "entity": "Weather.WeatherRange",
      "startPos": 0,
      "endPos": 6,
      "children": []
    },
    {
      "entity": "Weather.WeatherCondition",
      "startPos": 18,
      "endPos": 21,
      "children": []
    },
    {
      "entity": "Weather.Historic",
      "startPos": 23,
      "endPos": 30,
      "children": []
    }
  ]
}
```

To what does the Weather.Historic entity correspond in the utterance?

- A. by month
- B. chicago
- C. rain
- D. location

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 23

Topic #: 3

[\[All AI-102 Questions\]](#)

You are examining the Text Analytics output of an application.

The text analyzed is: `Our tour guide took us up the Space Needle during our trip to Seattle last week.`

The response contains the data shown in the following table.

Text	Category	ConfidenceScore
Tour guide	PersonType	0.45
Space Needle	Location	0.38
Trip	Event	0.78
Seattle	Location	0.78
Last week	DateTime	0.80

Which Text Analytics API is used to analyze the text?

- A. Entity Linking
- B. Named Entity Recognition
- C. Sentiment Analysis
- D. Key Phrase Extraction

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 25

Topic #: 3

[\[All AI-102 Questions\]](#)

SIMULATION -

You need to configure bot12345678 support the French (FR-FR) language.

Export the bot to C:\Resources\Bot\Bot1.zip.

To complete this task, use the Microsoft Bot Framework Composer.

[Show Suggested Answer](#)





Actual exam question from Microsoft's AI-102

Question #: 27

Topic #: 3

[\[All AI-102 Questions\]](#)

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. Language service
- B. Content Moderator
- C. Computer Vision
- D. Form Recognizer

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 28

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are developing an application that includes language translation.

The application will translate text retrieved by using a function named `get_text_to_be_translated`. The text can be in one of many languages. The content of the text must remain within the Americas Azure geography.

You need to develop code to translate the text to a single language.

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

```
...
api_key = "FF956C68B83B21B38691ABD200A4C606"
text = get_text_to_be_translated()
headers = {
    'Content-Type': 'application/json',
    'Ocp-Apim-Subscription-Key': api_key
}
body = {
    'Text': text
}
conn = httplib.HTTPSConnection
```

("api.cogninve.microsofttranslator.com")
("api-apc.cognitive.microsofttranslator.com")
("api-nam.cognitive.microsofttranslator.com")

```
conn.request("POST",
```

"/translate?from=en"
"/translate?suggestedFrom=en"
"/translate?to=en"
"/detect?to=en"
"/detect?from=en"

```
str(body), headers)

response = conn.getresponse()
response_data = response.read()
...
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 30

Topic #: 3

[\[All AI-102 Questions\]](#)

SIMULATION -

Use the following login credentials as needed:

To enter your username, place your cursor in the Sign in box and click on the username below.

To enter your password, place your cursor in the Enter password box and click on the password below.

Azure Username: admin@abc.com -

Azure Password: XXXXXXXXXXXX -

The following information is for technical support purposes only:

Lab Instance: 12345678 -

Task -

You need to create and publish a Language Understanding (classic) model named 1u12345678. The model will contain an intent of Travel that has an utterance of Boat.

To complete this task, sign in to the Language Understanding portal at <http://www.luis-ai/>.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 32

Topic #: 3

[\[All AI-102 Questions\]](#)

You have a Language service resource that performs the following:

- Sentiment analysis
- Named Entity Recognition (NER)
- Personally Identifiable Information (PII) identification

You need to prevent the resource from persisting input data once the data is analyzed.

Which query parameter in the Language service API should you configure?

- A. model-version
- B. piiCategories
- C. showStats
- D. loggingOptOut

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 33

Topic #: 3

[\[All AI-102 Questions\]](#)

You have an Azure Cognitive Services model named Model1 that identifies the intent of text input.

You develop an app in C# named App1.

You need to configure App1 to use Model1.

Which package should you add to App1?

- A. `Universal.Microsoft.CognitiveServices.Speech`
- B. `SpeechServicesToolkit`
- C. `Azure.AI.Language.Conversations`
- D. `Xamarin.Cognitive.Speech`

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 34

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are building content for a video training solution.

You need to create narration to accompany the video content. The solution must use Custom Neural Voice.

What should you use to create a custom neural voice, and which service should you use to generate the narration? To answer, select the appropriate options in the answer area.

NOTE: Each correct answer is worth one point.

Answer Area

Custom neural voice:

- Microsoft Bot Framework Composer
- The Azure portal
- The Language Understanding portal
- The Speech Studio portal

Narration:

- Language Understanding
- Speaker Recognition
- Speech-to-text
- Text-to-speech

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 35

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

You are building a call handling system that will receive calls from French-speaking and German-speaking callers. The system must perform the following tasks:

- Capture inbound voice messages as text.
- Replay messages in English on demand.

Which Azure Cognitive Services services should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

To capture messages:

Speaker Recognition
Speech-to-text
Text-to-speech
Translator

To replay messages:

Speech-to-text only
Speech-to-text and Language
Speaker Recognition and Language
Text-to-speech and Language
Text-to-speech and Translator

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 36

Topic #: 3

[\[All AI-102 Questions\]](#)

You are building a social media extension that will convert text to speech. The solution must meet the following requirements:

- Support messages of up to 400 characters.
- Provide users with multiple voice options.
- Minimize costs.

You create an Azure Cognitive Services resource.

Which Speech API endpoint provides users with the available voice options?

- A. `https://uksouth.api.cognitive.microsoft.com/speechtotext/v3.0/models/base`
- B. `https://uksouth.customvoice.api.speech.microsoft.com/api/texttospeech/v3.0/longaudiosynthesis/voices`
- C. `https://uksouth.tts.speech.microsoft.com/cognitiveservices/voices/list`
- D. `https://uksouth.voice.speech.microsoft.com/cognitiveservices/v1?deploymentId={deploymentId}`

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 37

Topic #: 3

[\[All AI-102 Questions\]](#)

You develop a custom question answering project in Azure Cognitive Service for Language. The project will be used by a chatbot.

You need to configure the project to engage in multi-turn conversations.

What should you do?

- A. Add follow-up prompts.
- B. Enable active learning.
- C. Add alternate questions.
- D. Enable chit-chat.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 38

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are building a solution that students will use to find references for essays.

You use the following code to start building the solution.

```
using Azure;
using System;
using Azure.AI.TextAnalytics;

private static readonly AzureKeyCredential credentials = new AzureKeyCredential("<key>");
private static readonly Uri endpoint = new Uri("<endpoint>");

static void EntityLinker(TextAnalyticsClient client)
{
    var response = client.RecognizeLinkedEntities(
        "Our tour guide took us up the Space Needle during our trip to Seattle last week.");
}
```

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 code will detect the language of documents.

The `url` attribute returned for each linked entity will be a Bing search link.

The `matches` attribute returned for each linked entity will provide the location in a document where the entity is referenced.

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 39

Topic #: 3

[\[All AI-102 Questions\]](#)

You train a Conversational Language Understanding model to understand the natural language input of users.

You need to evaluate the accuracy of the model before deploying it.

What are two methods you can use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. From the language authoring REST endpoint, retrieve the model evaluation summary.
- B. From Language Studio, enable Active Learning, and then validate the utterances logged for review.
- C. From Language Studio, select Model performance.
- D. From the Azure portal, enable log collection in Log Analytics, and then analyze the logs.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 40

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

You develop an app in C# named App1 that performs speech-to-speech translation.

You need to configure App1 to translate English to German.

How should you complete the `SpeechTranslationConfig` object? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

addTargetLanguage

speechSynthesisLanguage

speechRecognitionLanguage

voiceName

Answer Area

```
var translationConfig = SpeechTranslationConfig.FromSubscription(SPEECH__SUBSCRIPTION__KEY, SPEECH__SERVICE__REGION);  
translationConfig.  = "en-US";  
translationConfig.  ("de");
```

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 41

Topic #: 3

[\[All AI-102 Questions\]](#)

You have an Azure subscription that contains an Azure Cognitive Service for Language 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. Identity
- B. Keys and Endpoint
- C. Networking
- D. Properties

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 42

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

-

You are building a transcription service for technical podcasts.

Testing reveals that the service fails to transcribe technical terms accurately.

You need to improve the accuracy of the service.

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

Deploy the model.

Create a Custom Speech project.

Upload training datasets.

Create a speech-to-text model.

Create a Speaker Recognition model.

Train the model.

Create a Conversational Language Understanding model.

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 43

Topic #: 3

[\[All AI-102 Questions\]](#)

You are building a retail kiosk system that will use a custom neural voice.

You acquire audio samples and consent from the voice talent.

You need to create a voice talent profile.

What should you upload to the profile?

- A. a .zip file that contains 10-second .wav files and the associated transcripts as .txt files
- B. a five-minute .flac audio file and the associated transcript as a .txt file
- C. a .wav or .mp3 file of the voice talent consenting to the creation of a synthetic version of their voice
- D. a five-minute .wav or .mp3 file of the voice talent describing the kiosk system

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 44

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

You have a Language Understanding solution that runs in a Docker container.

You download the Language Understanding container image from the Microsoft Container Registry (MCR).

You need to deploy the container image to a host computer.

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

From the host computer, move the package file to the Docker input directory.

From the Language Understanding portal, export the solution as a package file.

From the host computer, build the container and specify the output directory.

From the host computer, run the container and specify the input directory.

From the Language Understanding portal, retrain the model.

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 45

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

You are building a text-to-speech app that will use a custom neural voice.

You need to create an SSML file for the app. The solution must ensure that the voice profile meets the following requirements:

- Expresses a calm tone
- Imitates the voice of a young adult female

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

NOTE: Each correct selection is worth one point.

Answer Area

...

```
<mstts:express-as
```

```
  How can I assist you?
```

```
</mstts:express-as>
```

...

- role
- style
- styledegree
- type
- voice

```
="YoungAdultFemale"
```


- role
- style
- styledegree
- type
- voice

```
="gentle">
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 46

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

You have a collection of press releases stored as PDF files.

You need to extract text from the files and perform sentiment analysis.

Which service should you use for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Extract text:

- Azure Cognitive Search
- Computer Vision
- Form Recognizer

Perform sentiment analysis:

- Azure Cognitive Search
- Computer Vision
- Form Recognizer
- Language

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 47

Topic #: 3

[\[All AI-102 Questions\]](#)

You have a text-based chatbot.

You need to enable content moderation by using the Text Moderation API of Content Moderator.

Which two service responses should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. personal data
- B. the adult classification score
- C. text classification
- D. optical character recognition (OCR)
- E. the racy classification score

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 48

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are developing a text processing solution.

You have the function shown below.

```
static void GetKeywords(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.RecognizeEntities (text);
    Console.WriteLine("Key words:");

    foreach (CategorizedEntity entity in response.Value)
    {
        Console.WriteLine($"{entity.Text}");
    }
}
```

For the second argument, you call the function and specify the following string.

Our tour of Paris included a visit to the Eiffel Tower

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer Area

Statements	Yes	No
The output will include the following words: our and included.	<input type="radio"/>	<input type="radio"/>
The output will include the following words: Paris, Eiffel, and Tower.	<input type="radio"/>	<input type="radio"/>
The function will output all the key phrases from the input string to the console.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 49

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

You are building an Azure web app named App1 that will translate text from English to Spanish.

You need to use the Text Translation REST API to perform the translation. The solution must ensure that you have data sovereignty in the United States.

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

NOTE: Each correct selection is worth one point.

Answer Area

https:// / ?api-version=3.0&to=es

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 50

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

-

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 model1 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.

Request approval to run the container.

Export model1 to Host1.

Run the container.

Configure disk logging.

Answer Area



Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 51

Topic #: 3

[\[All AI-102 Questions\]](#)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You build a language model by using a Conversational Language Understanding. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training.

- Find contacts in London.
- Who do I know in Seattle?
- Search for contacts in Ukraine.

You need to implement the phrase list in Conversational Language Understanding.

Solution: You create a new utterance for each phrase in the FindContact intent.

Does this meet the goal?

A. Yes

B. No

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 52

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

You have a question answering project in Azure Cognitive Service for Language.

You need to move the project to a Language service instance in a different Azure region.

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

From the new Language service instance, train and publish the project.

From the new Language service instance, import the project file.

From the new Language service instance, enable custom text classification.

From the original Language service instance, export the existing project.

From the new Language service instance, regenerate the keys.

From the original Language service instance, train and publish the model.



Answer Area



Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 53

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

-

You are building a customer support chatbot.

You need to configure the bot to identify the following:

- Code names for internal product development
- Messages that include credit card numbers

The solution must minimize development effort.

Which Azure Cognitive Service for Language feature should you use for each requirement? To answer, drag the appropriate features to the correct requirements. Each feature 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.

Features

Custom named entity recognition (NER)

Key phrase extraction

Language detection

Named Entity Recognition (NER)

Personally Identifiable Information (PII) detection

Sentiment analysis

Answer Area

Identify code names for internal product development:

Identify messages that include credit card numbers:

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 54

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are building an app by using the Speech SDK. The app will translate speech from French to German by using natural language processing.

You need to define the source language and the output language.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
var speechTranslationConfig =
```

```
SpeechTranslationConfig.FromSubscription(speechKey, speechRegion);
```

```
speechTranslationConfig.  = "fr"
```

- AddTargetLanguage
- SpeechRecognitionLanguage
- SpeechSynthesisLanguage
- TargetLanguages
- VoiceName

```
speech_translation_config.  ("de")
```

- AddTargetLanguage
- SpeechRecognitionLanguage
- SpeechSynthesisLanguage
- TargetLanguages
- VoiceName

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 55

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

-

You have a collection of Microsoft Word documents and PowerPoint presentations in German.

You need to create a solution to translate the files to French. The solution must meet the following requirements:

- Preserve the original formatting of the files.
- Support the use of a custom glossary.

You create a blob container for German files and a blob container for French files. You upload the original files to the container for German files.

Which three actions should you perform in sequence to complete the solution? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Perform an asynchronous translation by using the list of files to be translated.

Perform an asynchronous translation by using the document translation specification.

Generate a list of files to be translated.

Upload a glossary file to the container for German files.

Upload a glossary file to the container for French files.

Define a document translation specification that has a French target.

Answer Area



Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 56

Topic #: 3

[\[All AI-102 Questions\]](#)

You have the following C# function.

```
static void MyFunction(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.ExtractKeyPhrases(text);
    Console.WriteLine("Key phrases:");

    foreach (string keyphrase in response.Value)
    {
        Console.WriteLine($"{keyphrase}");
    }
}
```

You call the function by using the following code.

```
MyFunction(textAnalyticsClient, "the quick brown fox jumps over the lazy dog");
```

Which output will you receive?

- A. The quick -
The lazy
- B. the quick brown fox jumps over the lazy dog
- C. jumps over the
- D. quick brown fox
lazy dog

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 57

Topic #: 3

[\[All AI-102 Questions\]](#)

You have the following Python method.

```
def create_resource (resource_name, kind, account_tier, location) :  
    parameters = CognitiveServicesAccount(sku=sku(name=account_tier), kind=kind, location=location, properties={})  
    result = cogSvcClient.accounts.create(resource_group_name, resource_name, parameters)
```

You need to deploy an Azure resource to the East US Azure region. The resource will be used to perform sentiment analysis.

How should you call the method?

- A. create_resource("res1", "TextAnalytics", "Standard", "East US")
- B. create_resource("res1", "ContentModerator", "S0", "eastus")
- C. create_resource("res1", "ContentModerator", "Standard", "East US")
- D. create_resource("res1", "TextAnalytics", "S0", "eastus")

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 58

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

You develop a Python app named App1 that performs speech-to-speech translation.

You need to configure App1 to translate English to German.

How should you complete the `SpeechTranslationConfig` object? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

add_target_language

speech_synthesis_language

speech_recognition_language

voice_name

Answer Area

```
def translate_speech_to_text():  
    translation_config = speechsdk.translation.SpeechTranslationConfig(subscription=speech_key, region=service_region)  
    translation_config.  = "en-US";  
    translation_config.  ("de");
```

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 59

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

You are developing a streaming Speech to Text solution that will use the Speech SDK and MP3 encoding.

You need to develop a method to convert speech to text for streaming MP3 data.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
audio_format = speechsdk.audio.  (compressed_stream_format=speechsdk.AudioStreamContainerFormat.MP3)
```

- AudioConfig.SetProperty
- AudioStreamFormat
- GetWaveFormatPCM
- PullAudioInputStream

```
stream = speechsdk.audio.PullAudioInputStream(stream_format=audio_format, pull_stream_callback=callback)
```

```
speech_config = speechsdk.SpeechConfig("18c51a87-3a69-47a8-aedc-a54745f708a1", "westus")
```

```
audio_config = speechsdk.audio.AudioConfig(stream=stream)
```

```
recognizer = speechsdk.  (speech_config=speech_config, audio_config=audio_config)
```

- KeywordRecognizer
- SpeakerRecognizer
- SpeechRecognizer
- SpeechSynthesizer

```
result = recognizer.recognize_once()
```

```
text = result.text
```

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 60

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are building a chatbot.

You need to use the Content Moderator API to identify aggressive and sexually explicit language.

Which three settings should you configure? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Content Moderator - Moderate Text - Screen

The operation detects profanity in more than 100 languages and match against custom and shared blacklists.

Host

Name

Resource Name

Query parameters

autocorrect

[✖ Remove parameter](#)

PII

[✖ Remove parameter](#)

listid

[✖ Remove parameter](#)

classify

[✖ Remove parameter](#)

language

[✖ Remove parameter](#)

[+ Add parameter](#)

Headers

Content-Type

[✖ Remove header](#)

Ocp-Apim-Subscription-Key

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 61

Topic #: 3

[\[All AI-102 Questions\]](#)

You are developing an app that will use the Decision and Language APIs.

You need to provision resources for the app. The solution must ensure that each service is accessed by using a single endpoint and credential.

Which type of resource should you create?

- A. Language
- B. Speech
- C. Azure Cognitive Services
- D. Content Moderator

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 62

Topic #: 3

[\[All AI-102 Questions\]](#)

You are building a chatbot.

You need to ensure that the bot will recognize the names of your company's products and codenames. The solution must minimize development effort.

Which Azure Cognitive Service for Language service should you include in the solution?

- A. custom text classification
- B. entity linking
- C. custom Named Entity Recognition (NER)
- D. key phrase extraction

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 63

Topic #: 3

[\[All AI-102 Questions\]](#)

You have an Azure subscription that contains an Azure App Service app named App1.

You provision a multi-service Azure Cognitive Services resource named CSAccount1.

You need to configure App1 to access CSAccount1. The solution must minimize administrative effort.

What should you use to configure App1?

- A. a system-assigned managed identity and an X.509 certificate
- B. the endpoint URI and an OAuth token
- C. the endpoint URI and a shared access signature (SAS) token
- D. the endpoint URI and subscription key

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 64

Topic #: 3

[\[All AI-102 Questions\]](#)

You have an Azure subscription that contains a multi-service Azure Cognitive Services Translator resource named Translator1.

You are building an app that will translate text and documents by using Translator1.

You need to create the REST API request for the app.

Which headers should you include in the request?

- A. the access control request, the content type, and the content length
- B. the subscription key and the client trace ID
- C. the resource ID and the content language
- D. the subscription key, the subscription region, and the content type

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 65

Topic #: 3

[\[All AI-102 Questions\]](#)

You have a file share that contains 5,000 images of scanned invoices.

You need to analyze the images. The solution must extract the following data:

- Invoice items
- Sales amounts
- Customer details

What should you use?

- A. Custom Vision
- B. Azure AI Computer Vision
- C. Azure AI Immersive Reader
- D. Azure AI Document Intelligence

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 66

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are developing a text processing solution.

You have the function shown below.

```
def get_key_words(textAnalyticsClient, text):  
  
    response = textAnalyticsClient.recognize_entities (documents = [text])[0]  
    print("Key Words:")  
    for entity in response.entities:  
        print("\t\t", entity.text)
```

For the second argument, you call the function and specify the following string.

Our tour of Paris included a visit to the Eiffel Tower

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer Area

Statements

Yes **No**

The output will include the following words: our and included.

The output will include the following words: Paris, Eiffel, and Tower.

The function will output all the key phrases from the input string to the console.

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 67

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are developing a text processing solution.

You develop the following method.

```
def get_key_phrases(text_analytics_client, text):  
    response = text_analytics_client.extract_key_phrases(text, language="en")  
    print('Key phrases:')  
    for keyphrase in response.key_phrases:  
        print(f'\t{keyphrase}')
```

You call the method by using the following code.

```
get_key_phrases(text_analytics_client, "the cat sat on the mat")
```

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

The call will output key phrases from the input string to the console.

The output will contain the following words: the, cat, sat, on, and mat.

The output will contain the confidence level for key phrases.

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 68

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are developing a service that records lectures given in English (United Kingdom).

You have a method named `append_to_transcript_file` that takes translated text and a language identifier.

You need to develop code that will provide transcripts of the lectures to attendees in their respective language. The supported languages are English, French, Spanish, and German.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
speech_key = os.environ['SPEECH__SUBSCRIPTION__KEY']
service_region = os.environ['SPEECH__SERVICE__REGION']
def translate_speech():
    translation_config = speechsdk.translation.SpeechTranslationConfig(
        subscription=speech_key, region=service_region)
    translation_config.speech_recognition_language = "en-GB"
    languages = 
    for language in languages: translation_config.add_target_language(language)
    audio_config = speechsdk.audio.AudioConfig(use_default_microphone=True)
    recognizer = speechsdk.translation.
        translation_config=translation_config, audio_config=audio_config)
    result = recognizer.recognize_once()
    if result.reason == speechsdk.ResultReason.TranslatedSpeech:
        append_to_transcript_file(result.text, "en")
        for language in result.translations:
            append_to_transcript_file(result.translations[language], language)
```

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 69

Topic #: 3

[\[All AI-102 Questions\]](#)

You are developing an app that will use the text-to-speech capability of the Azure AI Speech service. The app will be used in motor vehicles.

You need to optimize the quality of the synthesized voice output.

Which Speech Synthesis Markup Language (SSML) attribute should you configure?

- A. the style attribute of the mstts:express-as element
- B. the effect attribute of the voice element
- C. the pitch attribute of the prosody element
- D. the level attribute of the emphasis element

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 70

Topic #: 3

[\[All AI-102 Questions\]](#)

You are designing a content management system.

You need to ensure that the reading experience is optimized for users who have reduced comprehension and learning differences, such as dyslexia. The solution must minimize development effort.

Which Azure service should you include in the solution?

- A. Azure AI Immersive Reader
- B. Azure AI Translator
- C. Azure AI Document Intelligence
- D. Azure AI Language

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 71

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are building an app that will answer customer calls about the status of an order. The app will query a database for the order details and provide the customers with a spoken response.

You need to identify which Azure AI service APIs to use. The solution must minimize development effort.

Which object should you use for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Convert customer calls into text queries:

- SpeechRecognizer
- SpeechSynthesizer
- TranslationRecognizer
- VoiceProfileClient

Provide customers with the order details:

- SpeechRecognizer
- SpeechSynthesizer
- TranslationRecognizer
- VoiceProfileClient

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 72

Topic #: 3

[\[All AI-102 Questions\]](#)

You have an Azure AI service model named Model1 that identifies the intent of text input.

You develop a Python app named App1.

You need to configure App1 to use Model1.

Which package should you add to App1?

- A. azure-cognitiveservices-language-textanalytics
- B. azure-ai-language-conversations
- C. azure-mgmt-cognitiveservices
- D. azure-cognitiveservices-speech

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 73

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You are building an app that will automatically translate speech from English to French, German, and Spanish by using Azure AI service.

You need to define the output languages and configure the Azure AI Speech service.

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

NOTE: Each correct selection is worth one point.

Answer Area

```
speech_key, service_region = os.environ['SPEECH__SERVICE__KEY'], os.environ['SPEECH__SERVICE__REGION']  
languages =   
def translate_speech_to_text():  
    translation_config = speechsdk.translation.SpeechTranslationConfig(subscription=speech_key, region=service_region)  
    for lang in languages:  
        translation_config.add_target_language(lang)  
    for lang in languages:  
        translation_config.add_target_language(lang)  
    recognizer = speechsdk.translation. (translation_config=translation_config)  
    ...
```

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 74

Topic #: 3

[\[All AI-102 Questions\]](#)

DRAG DROP

-

You plan to implement an Azure AI Search resource that will use custom skill based on sentiment analysis.

You need to create a custom model and configure Azure AI Search use the model.

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

Create an endpoint for the model.

Rerun the indexer to enrich the index.

Create an Azure Machine Learning workspace.

Create and train the model in the Azure Machine Learning studio.

Provision an Azure AI Services resource and obtain the endpoint.

Connect the custom skill the endpoint.



Answer Area



Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 75

Topic #: 3

[\[All AI-102 Questions\]](#)

HOTSPOT

-

You have a collection of press releases stored as PDF files.

You need to extract text from the files and perform sentiment analysis.

Which service should you use for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Extract text:

- Azure AI Search
- Azure AI Vision
- Azure AI Document Intelligence

Perform sentiment analysis:

- Azure Cognitive Search
- Azure AI Computer Vision
- Azure AI Document Intelligence
- Azure AI Language

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 1

Topic #: 4

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are developing a text processing solution.

You develop the following method.

```
static void GetKeyPhrases(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.ExtractKeyPhrases(text);
    Console.WriteLine("Key phrases:");

    foreach (string keyphrase in response.Value)
    {
        Console.WriteLine($"{keyphrase}");
    }
}
```

You call the method by using the following code.

```
GetKeyPhrases(textAnalyticsClient, "the cat sat on the mat");
```

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 call will output key phrases from the input string to the console.	<input type="radio"/>	<input type="radio"/>
The output will contain the following words: the, cat, sat, on, and mat.	<input type="radio"/>	<input type="radio"/>
The output will contain the confidence level for key phrases.	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 2

Topic #: 4

[\[All AI-102 Questions\]](#)

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.

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 3

Topic #: 4

[\[All AI-102 Questions\]](#)

You have an existing Azure Cognitive Search service.

You have an Azure Blob storage account that contains millions of scanned documents stored as images and PDFs.

You need to make the scanned documents available to search as quickly as possible.

What should you do?

- A. Split the data into multiple blob containers. Create a Cognitive Search service for each container. Within each indexer definition, schedule the same runtime execution pattern.
- B. Split the data into multiple blob containers. Create an indexer for each container. Increase the search units. Within each indexer definition, schedule a sequential execution pattern.
- C. Create a Cognitive Search service for each type of document.
- D. Split the data into multiple virtual folders. Create an indexer for each folder. Increase the search units. Within each indexer definition, schedule the same runtime execution pattern.

Show Suggested Answer





Actual exam question from Microsoft's AI-102

Question #: 4

Topic #: 4

[\[All AI-102 Questions\]](#)

You need to implement a table projection to generate a physical expression of an Azure Cognitive Search index.

Which three properties should you specify in the skillset definition JSON configuration table node? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. tableName
- B. generatedKeyName
- C. dataSource
- D. dataSourceConnection
- E. source

Show Suggested Answer



Actual exam question from Microsoft's AI-102

Question #: 5

Topic #: 4

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are creating an enrichment pipeline that will use Azure Cognitive Search. The knowledge store contains unstructured JSON data and scanned PDF documents that contain text.

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.

Hot Area:

Answer Area

JSON data:

	▼
File projection	
Object projection	
Table projection	

Scanned data:

	▼
File projection	
Object projection	
Table projection	

Show Suggested Answer

Actual exam question from Microsoft's AI-102

Question #: 6

Topic #: 4

[\[All AI-102 Questions\]](#)

HOTSPOT -

You are building an Azure Cognitive Search custom skill.

You have the following custom skill schema definition.

```
{
  "@odata.type": "#Microsoft.Skills.Custom.WebApiSkill",
  "description": "My custom skill description",
  "uri": "https://contoso-webskill.azurewebsites.net/api/process",
  "context": "/document/organizations/*",
  "inputs": [
    {
      "name": "companyName",
      "source": "/document/organizations/*"
    }
  ],
  "outputs": [
    {
      "name": "companyDescription",
    }
  ]
}
```

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
CompanyDescription is available for indexing.	<input type="radio"/>	<input type="radio"/>
The definition calls a web API as part of the enrichment process.	<input type="radio"/>	<input type="radio"/>
The enrichment step is called only for the first organization under "/document/organizations".	<input type="radio"/>	<input type="radio"/>

Show Suggested Answer