

➤ **Vendor: Microsoft**

➤ **Exam Code: AI-100**

➤ **Exam Name: Designing and Implementing an Azure AI Solution**

➤ **New Updated Questions from [Braindump2go](#) (Updated in [Dec./2020](#))**

**[Visit Braindump2go and Download Full Version AI-100 Exam Dumps](#)**

**QUESTION 110**

**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, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You have an app named App1 that uses the Face API.

App1 contains several PersonGroup objects.

You discover that a PersonGroup object for an individual named Ben Smith cannot accept additional entries. The PersonGroup object for Ben Smith contains 10,000 entries.

You need to ensure that additional entries can be added to the PersonGroup object for Ben Smith. The solution must ensure that Ben Smith can be identified by all the entries.

Solution: You migrate all the entries to the LargePersonGroup object for Ben Smith.

Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

LargePersonGroup and LargeFaceList are collectively referred to as large-scale operations.

LargePersonGroup can contain up to 1 million persons, each with a maximum of 248 faces. LargeFaceList can contain up to 1 million faces. The large-scale operations are similar to the conventional PersonGroup and FaceList but have some differences because of the new architecture.

References:

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

**QUESTION 111**

Your company plans to develop a mobile app to provide meeting transcripts by using speech-to-text. Audio from the meetings will be streamed to provide real-time transcription.

You need to recommend which task each meeting participant must perform to ensure that the transcripts of the meetings can identify all participants.

Which task should you recommend?

- A. Record the meeting as an MP4.
- B. Create a voice signature.
- C. Sign up for Azure Speech Services.
- D. Sign up as a guest in Azure Active Directory (Azure AD)

**Answer:** B

**Explanation:**

**[AI-100 Exam Dumps](#) [AI-100 Exam Questions](#) [AI-100 PDF Dumps](#) [AI-100 VCE Dumps](#)**

**<https://www.braindump2go.com/ai-100.html>**

The first step is to create voice signatures for the conversation participants. Creating voice signatures is required for efficient speaker identification.

Note: In addition to the standard baseline model used by the Speech Services, you can customize models to your needs with available data, to overcome speech recognition barriers such as speaking style, vocabulary and background noise.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/cognitive-services/speech-service/how-to-use-conversation-transcription-service>

#### **QUESTION 112**

You need to create a prototype of a bot to demonstrate a user performing a task. The demonstration will use the Bot Framework Emulator.

Which botbuilder CLI tool should you use to create the prototype?

- A. Chatdown
- B. QnAMaker
- C. Dispatch
- D. LuDown

**Answer: A**

**Explanation:**

Use Chatdown to produce prototype mock conversations in markdown and convert the markdown to transcripts you can load and view in the new V4 Bot Framework Emulator.

Incorrect Answers:

B: QnA Maker is a cloud-based API service that lets you create a conversational question-and-answer layer over your existing data. Use it to build a knowledge base by extracting questions and answers from your semi-structured content, including FAQs, manuals, and documents. Answer users' questions with the best answers from the QnAs in your knowledge base--automatically. Your knowledge base gets smarter, too, as it continually learns from user behavior.

C: Dispatch lets you build language models that allow you to dispatch between disparate components (such as QnA, LUIS and custom code).

D: LuDown build LUIS language understanding models using markdown files

References:

<https://github.com/microsoft/botframework/blob/master/README.md>

#### **QUESTION 113**

You are designing an AI solution that will provide feedback to teachers who train students over the Internet. The students will be in classrooms located in remote areas. The solution will capture video and audio data of the students in the classrooms.

You need to recommend Azure Cognitive Services for the AI solution to meet the following requirements:

- Alert teachers if a student facial expression indicates the student is angry or scared.
- Identify each student in the classrooms for attendance purposes.
- Allow the teachers to log voice conversations as text.

Which Cognitive Services should you recommend?

- A. Face API and Text Analytics
- B. Computer Vision and Text Analytics
- C. QnA Maker and Computer Vision
- D. Speech to Text and Face API

**Answer: D**

**Explanation:**

Speech-to-text from Azure Speech Services, also known as speech-to-text, enables real-time transcription of audio streams into text that your applications, tools, or devices can consume, display, and take action on as command input.

Face detection: Detect one or more human faces in an image and get back face rectangles for where in the image the faces are, along with face attributes which contain machine learning-based predictions of facial features. The face

attribute features available are: Age, Emotion, Gender, Pose, Smile, and Facial Hair along with 27 landmarks for each face in the image.

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/speech-to-text>

<https://azure.microsoft.com/en-us/services/cognitive-services/face/>

**QUESTION 114**

You need to evaluate trends in fuel prices during a period of 10 years. The solution must identify unusual fluctuations in prices and produce visual representations.

Which Azure Cognitive Services API should you use?

- A. Anomaly Detector
- B. Computer Vision
- C. Text Analytics
- D. Bing Autosuggest

**Answer: A**

**Explanation:**

The Anomaly Detector API enables you to monitor and detect abnormalities in your time series data with machine learning. The Anomaly Detector API adapts by automatically identifying and applying the best-fitting models to your data, regardless of industry, scenario, or data volume. Using your time series data, the API determines boundaries for anomaly detection, expected values, and which data points are anomalies.

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/anomaly-detector/overview>

**QUESTION 115**

You plan to perform analytics of the medical records of patients located around the world.

You need to recommend a solution that avoids storing and processing data in the cloud.

What should you include in the recommendation?

- A. Azure Machine Learning Studio
- B. the Text Analytics API that has container support
- C. Azure Machine Learning services
- D. an Apache Spark cluster that uses MMLSpark

**Answer: D**

**Explanation:**

The Microsoft Machine Learning Library for Apache Spark (MMLSpark) assists in provisioning scalable machine learning models for large datasets, especially for building deep learning problems. MMLSpark works with SparkML pipelines, including Microsoft CNTK and the OpenCV library, which provide end-to-end support for the ingress and processing of image input data, categorization of images, and text analytics using pre-trained deep learning algorithms.

References:

[https://subscription.packtpub.com/book/big\\_data\\_and\\_business\\_intelligence/9781789131956/10/ch10lv1sec61/an-overview-of-the-microsoft-machine-learning-library-for-apache-spark-mmlspark](https://subscription.packtpub.com/book/big_data_and_business_intelligence/9781789131956/10/ch10lv1sec61/an-overview-of-the-microsoft-machine-learning-library-for-apache-spark-mmlspark)

**QUESTION 116**

Your company has an on-premises datacenter.

You plan to publish an app that will recognize a set of individuals by using the Face API. The model is trained.

You need to ensure that all images are processed in the on-premises datacenter.

What should you deploy to host the Face API?

- A. a Docker container
- B. Azure File Sync
- C. Azure Application Gateway
- D. Azure Data Box Edge

**Answer:** A

**Explanation:**

A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another. A Docker container image is a lightweight, standalone, executable package of software that includes everything needed to run an application: code, runtime, system tools, system libraries and settings.

Incorrect Answers:

D: Azure Data Box Edge is an AI-enabled edge computing device with network data transfer capabilities.

This article provides you an overview of the Data Box Edge solution, benefits, key capabilities, and the scenarios where you can deploy this device.

Data Box Edge is a Hardware-as-a-service solution. Microsoft ships you a cloud-managed device with a built-in Field Programmable Gate Array (FPGA) that enables accelerated AI-inferencing and has all the capabilities of a storage gateway.

References:

<https://www.docker.com/resources/what-container>

#### **QUESTION 117**

You have a Bing Search service that is used to query a product catalog.

You need to identify the following information:

- The locale of the query
- The top 50 query strings
- The number of calls to the service
- The top geographical regions of the service

What should you implement?

- A. Bing Statistics
- B. Azure API Management (APIM)
- C. Azure Monitor
- D. Azure Application Insights

**Answer:** A

**Explanation:**

The Bing Statistics add-in provides metrics such as call volume, top queries, API response, code distribution, and market distribution. The rich slicing-and-dicing capability lets you gather deeper understanding of your users and their usage to inform your business strategy.

References:

<https://www.bingapistatistics.com/>

#### **QUESTION 118**

You have a Face API solution that updates in real time. A pilot of the solution runs successfully on a small dataset.

When you attempt to use the solution on a larger dataset that continually changes, the performance degrades, slowing how long it takes to recognize existing faces.

You need to recommend changes to reduce the time it takes to recognize existing faces without increasing costs.

What should you recommend?

- A. Change the solution to use the Computer Vision API instead of the Face API.
- B. Separate training into an independent pipeline and schedule the pipeline to run daily.
- C. Change the solution to use the Bing Image Search API instead of the Face API.
- D. Distribute the face recognition inference process across many Azure Cognitive Services instances.

**Answer:** B

**Explanation:**

Incorrect Answers:

A: The purpose of Computer Vision is to inspect each image associated with an incoming article to (1) scrape out written words from the image and (2) determine what types of objects are present in the image.

C: The Bing API provides an experience similar to Bing.com/search by returning search results that Bing determines

are relevant to a user's query. The results include Web pages and may also include images, videos, and more.

D: That would increase cost.

References:

<https://github.com/Azure/cognitive-services>

#### **QUESTION 119**

**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, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You are deploying an Azure Machine Learning model to an Azure Kubernetes Service (AKS) container.

You need to monitor the scoring accuracy of each run of the model.

Solution: You modify the Config.json file.

Does this meet the goal?

A. Yes

B. No

**Answer: B**

**Explanation:**

Instead update the manifest file.

Reference:

[https://azure.github.io/learnAnalytics-UsingAzureMachineLearningforAIWorkloads/lab07-deploying\\_a\\_scoring\\_service\\_to\\_aks/0\\_README.html](https://azure.github.io/learnAnalytics-UsingAzureMachineLearningforAIWorkloads/lab07-deploying_a_scoring_service_to_aks/0_README.html)

#### **QUESTION 120**

**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, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.**

You have Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.

Solution: You deploy an Azure Machine Learning model as an IoT Edge module.

Does this meet the goal?

A. Yes

B. No

**Answer: A**

**Explanation:**

You can use IoT Edge modules to deploy code that implements your business logic directly to your IoT Edge devices.

For example, you can deploy an Azure Machine Learning module that predicts when a device fails based on simulated machine temperature data.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/iot-edge/tutorial-deploy-machine-learning>