

DUMPSBOSS.

Designing and Implementing an Azure AI

Microsoft AI-100

Version Demo

Total Demo Questions: 15

Total Premium Questions: 219

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Case Study 1	2
Topic 2, Case Study 2	3
Topic 3, Mixed Questions	214
Total	219

QUESTION NO: 1 - (HOTSPOT)

HOTSPOT

You are designing a solution that will ingest data from an Azure IoT Edge device, preprocess the data in Azure Machine Learning, and then move the data to Azure HDInsight for further processing.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Machine Learning module to use to move the data into HDInsight:

	▼
Export Data	
Load Trained Model	
Partition and Sample	
Unpack Zipped Datasets	

Query type to use:

	▼
Apache Hive	
Apache Spark	
C#	
Transact-SQL	

Output the data to:

	▼
Azure Cosmos DB	
Azure Data Lake	
Azure Table storage	
HDFS	

ANSWER:

Answer Area

Machine Learning module to use to move the data into HDInsight:

	▼
Export Data	
Load Trained Model	
Partition and Sample	
Unpack Zipped Datasets	

Query type to use:

	▼
Apache Hive	
Apache Spark	
C#	
Transact-SQL	

Output the data to:

	▼
Azure Cosmos DB	
Azure Data Lake	
Azure Table storage	
HDFS	

Explanation:

Box 1: Export Data

The Export data to Hive option in the Export Data module in Azure Machine Learning Studio. This option is useful when you are working with very large datasets, and want to save your machine learning experiment data to a Hadoop cluster or HDInsight distributed storage.

Box 2: Apache Hive

Apache Hive is a data warehouse system for Apache Hadoop. Hive enables data summarization, querying, and analysis of data. Hive queries are written in HiveQL, which is a query language similar to SQL.

Box 3: Azure Data Lake

Default storage for the HDFS file system of HDInsight clusters can be associated with either an Azure Storage account or an Azure Data Lake Storage.

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/export-to-hive-query>
<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/hdinsight-use-hive>

QUESTION NO: 2

Which RBAC role should you assign to the KeyManagers group?

- A. Cognitive Services Contributor
- B. Security Manager
- C. Cognitive Services User
- D. Security Administrator

ANSWER: A

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>

QUESTION NO: 3

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 developing an application that uses an Azure Kubernetes Service (AKS) cluster.

You are troubleshooting a node issue.

You need to connect to an AKS node by using SSH.

Solution: You run the kubectl command, and then you create an SSH connection.

Does this meet the goal?

- A. Yes
- B. No

ANSWER: B

QUESTION NO: 4 - (DRAG DROP)

DRAG DROP

You create an image classification model in Azure Machine Learning Studio.

You need to deploy the model as a containerized web service.

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.

Select and Place:

Actions

Start the container

Create a container image

Create an Azure Batch AI account

Get the http endpoint of the web service

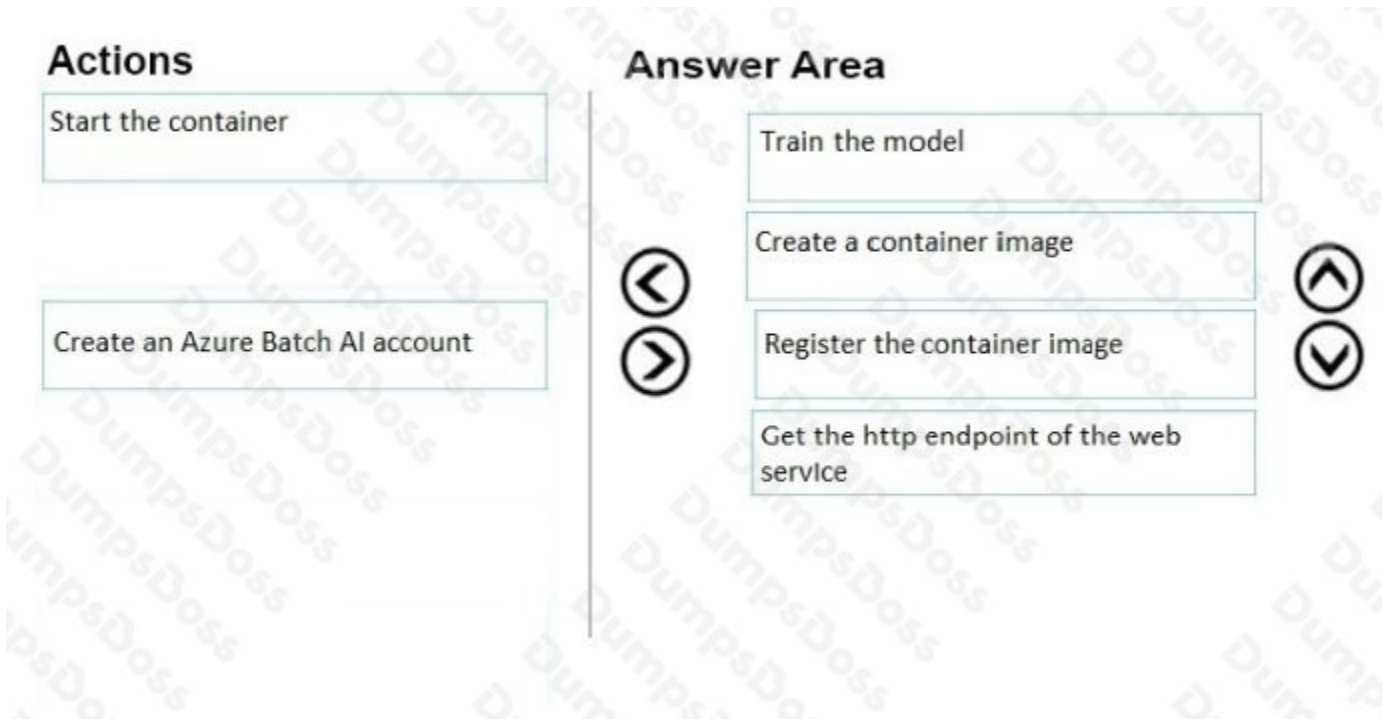
Register the container image

Train the model

Answer Area

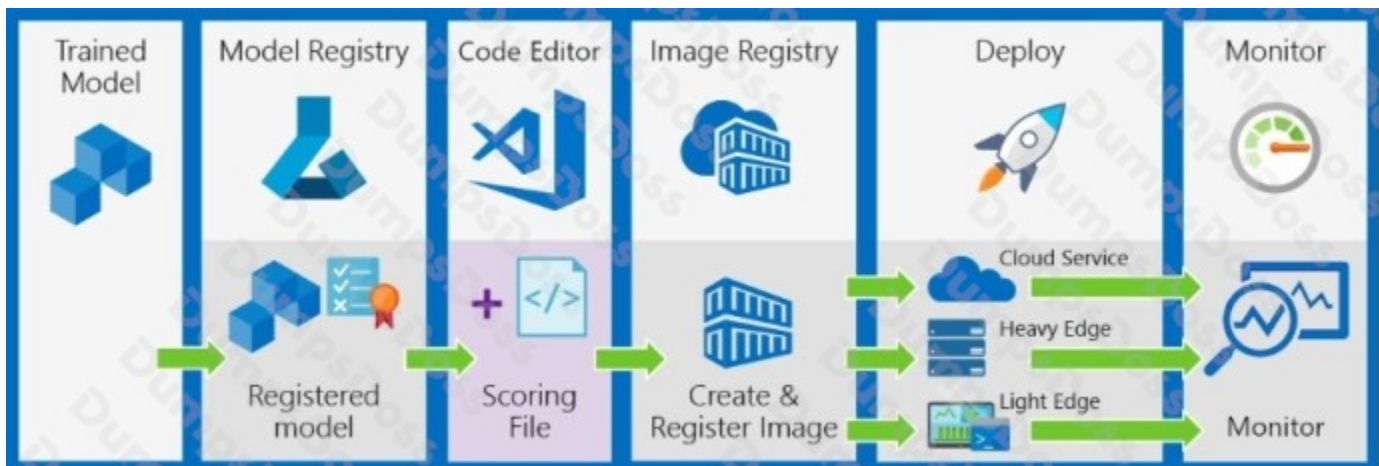


ANSWER:



Explanation:

The following diagram illustrates the complete deployment workflow:



The deployment workflow includes the following steps:

1. Register the model in a registry hosted in your Azure Machine Learning Service workspace
2. Register an image that pairs a model with a scoring script and dependencies in a portable container
3. Deploy the image as a web service in the cloud or to edge devices
4. Monitor and collect data
5. Update a deployment to use a new image.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/machine-learning/service/concept-model-management-and-deployment#step-3-deploy-image>

QUESTION NO: 5

You need to meet the testing requirements for the data scientists.

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

NOTE: Each correct selection is worth one point.

- A. Deploy an Azure Kubernetes Service (AKS) cluster to the East US 2 region
- B. Get the docker image from mcr.microsoft.com/azure-cognitive-services/sentiment:latest
- C. Deploy an Azure an Azure Container Service cluster to the West Europe region
- D. Export the production version of the Language Understanding (LUIS) app
- E. Deploy a Kubernetes cluster to Azure Stack
- F. Get the docker image from mcr.microsoft.com/azure-cognitive-services/luis:latest
- G. Export the staging version of the Language and Understanding (LUIS) app

ANSWER: E F G

Explanation:

Scenario: Data scientists must test Butler by using ASDK.

Note: Contoso wants to provide a new version of the Bookings app that will provide a highly available, reliable service for booking travel packages by interacting with a chatbot named Butler.

E: The ASDK (Azure Stack Development Kit) is meant to provide an environment in which you can evaluate Azure Stack and develop modern applications using APIs and tooling consistent with Azure in a non-production environment.

Microsoft Azure Stack integrated systems range in size from 4-16 nodes, and are jointly supported by a hardware partner and Microsoft.

F: The Language Understanding (LUIS) container loads your trained or published Language Understanding model, also known as a LUIS app, into a docker container and provides access to the query predictions from the container's API endpoints.

Use the docker pull command to download a container image from the mcr.microsoft.com/azure-cognitive-services/luis repository:

```
docker pull mcr.microsoft.com/azure-cognitive-services/luis:latest
```

G: You can test using the endpoint with a maximum of two versions of your app. With your main or live version of your app set as the production endpoint, add a second version to the staging endpoint.

Reference: <https://docs.microsoft.com/en-us/azure-stack/asdk/asdk-what-is> <https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-container-howto> <https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-test>

QUESTION NO: 6

Your company has factories in 10 countries. Each factory contains several thousand IoT devices.

The devices present status and trending data on a dashboard.

You need to ingest the data from the IoT devices into a data warehouse.

Which two Microsoft Azure technologies should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Azure Stream Analytics
- B. Azure Data Factory
- C. an Azure HDInsight cluster
- D. Azure Batch
- E. Azure Data Lake

ANSWER: C E

Explanation:

With Azure Data Lake Store (ADLS) serving as the hyper-scale storage layer and HDInsight serving as the Hadoop-based compute engine services. It can be used for prepping large amounts of data for insertion into a Data Warehouse

References: <https://www.blue-granite.com/blog/azure-data-lake-analytics-holds-a-unique-spot-in-the-modern-data-architecture>

QUESTION NO: 7

Your company has a data team of Transact-SQL experts.

You plan to ingest data from multiple sources into Azure Event Hubs.

You need to recommend which technology the data team should use to move and query data from Event Hubs to Azure Storage. The solution must leverage the data team's existing skills.

What is the best recommendation to achieve the goal? More than one answer choice may achieve the goal.

- A. Azure Notification Hubs
- B. Azure Event Grid
- C. Apache Kafka streams
- D. Azure Stream Analytics

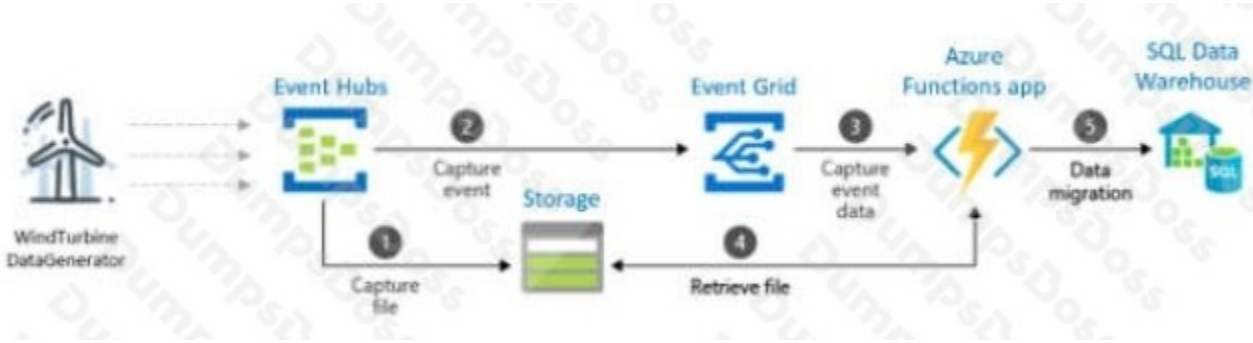
ANSWER: B

Explanation:

Event Hubs Capture is the easiest way to automatically deliver streamed data in Event Hubs to an Azure Blob storage or Azure Data Lake store. You can subsequently process and deliver the data to any other storage destinations of your choice, such as SQL Data Warehouse or Cosmos DB.

You to capture data from your event hub into a SQL data warehouse by using an Azure function triggered by an event grid.

Example:



First, you create an event hub with the Capture feature enabled and set an Azure blob storage as the destination. Data generated by WindTurbineGenerator is streamed into the event hub and is automatically captured into Azure Storage as Avro files.

Next, you create an Azure Event Grid subscription with the Event Hubs namespace as its source and the Azure Function endpoint as its destination.

Whenever a new Avro file is delivered to the Azure Storage blob by the Event Hubs Capture feature, Event Grid notifies the Azure Function with the blob URI. The Function then migrates data from the blob to a SQL data warehouse.

References: <https://docs.microsoft.com/en-us/azure/event-hubs/store-captured-data-data-warehouse>

QUESTION NO: 8 - (HOTSPOT)

HOTSPOT You are designing an AI solution that will be used to find buildings in aerial pictures.

Users will upload the pictures to an Azure Storage account. A separate JSON document will contain for the pictures.

The solution must meet the following requirements:

- Store metadata for the pictures in a data store.
- Run a custom vision Azure Machine Learning module to identify the buildings in a picture and the position of the buildings' edges.
- Run a custom mathematical module to calculate the dimensions of the buildings in a picture based on the metadata and data from the vision module.

You need to identify which Azure infrastructure services are used for each component of the AI workflow. The solution must execute as quickly as possible.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Location to store the metadata:

	▼
Azure Blob storage	
Azure Cosmos DB	
Azure File Storage	

Virtual machine series to run the vision module:

	▼
A	
F	
NV	

Virtual machine series to run the mathematical module:

	▼
A	
F	
NV	

ANSWER:

Answer Area

Location to store the metadata:

	▼
Azure Blob storage	
Azure Cosmos DB	
Azure File Storage	

Virtual machine series to run the vision module:

	▼
A	
F	
NV	

Virtual machine series to run the mathematical module:

	▼
A	
F	
NV	

Explanation:

Box 1: Azure Blob Storage

Containers and blobs support custom metadata, represented as HTTP headers.

Box 2: NV

The NV-series enables powerful remote visualisation workloads and other graphics-intensive applications backed by the NVIDIA Tesla M60 GPU.

Note: The N-series is a family of Azure Virtual Machines with GPU capabilities. GPUs are ideal for compute and graphics-intensive workloads, helping customers to fuel innovation through scenarios like high-end remote visualisation, deep learning and predictive analytics.

Box 3: F

F-series VMs feature a higher CPU-to-memory ratio. Example use cases include batch processing, web servers, analytics and gaming.

Incorrect:

A-series VMs have CPU performance and memory configurations best suited for entry level workloads like development and test.

References:

<https://azure.microsoft.com/en-in/pricing/details/virtual-machines/series/>

QUESTION NO: 9

You have deployed several Azure IoT Edge devices for an AI solution. The Azure IoT Edge devices generate measurement data from temperature sensors.

You need a solution to process the sensor data. Your solution must be able to write configuration changes back to the devices.

You make use of Microsoft Azure Event Hub.

Does this action accomplish your objective?

- A. Yes, it does
- B. No, it does not

ANSWER: B

Explanation:

Use Microsoft Azure IoT Hub instead.

Reference:

<https://azure.microsoft.com/en-us/resources/samples/functions-js-iot-hub-processing/>

QUESTION NO: 10

You have an AI application that uses keys in Azure Key Vault.

Recently, a key used by the application was deleted accidentally and was unrecoverable.

You need to ensure that if a key is deleted, it is retained in the key vault for 90 days.

Which two features should you configure? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. The expiration date on the keys
- B. Soft delete
- C. Purge protection
- D. Auditors
- E. The activation date on the keys

ANSWER: B C

Explanation:

References: <https://docs.microsoft.com/en-us/azure/architecture/best-practices/data-partitioning>

QUESTION NO: 11

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 NO: 12

You deploy an application that performs sentiment analysis on the data stored in Azure Cosmos DB.

Recently, you loaded a large amount of data to the database. The data was for a customer named Contoso, Ltd. You discover that queries for the Contoso data are slow to complete, and the queries slow the entire application.

You need to reduce the amount of time it takes for the queries to complete. The solution must minimize costs. What should you do? More than one answer choice may achieve the goal. (Choose two.)

- A. Change the request units.
- B. Change the partitioning strategy.

- C. Change the transaction isolation level.
- D. Migrate the data to the Cosmos DB database.

ANSWER: A B

Explanation:

Increasing request units would improve throughput, but at a cost.

Throughput provisioned for a container is divided evenly among physical partitions.

References: <https://docs.microsoft.com/en-us/azure/architecture/best-practices/data-partitioning>

QUESTION NO: 13

You are designing a solution that will use the Azure Content Moderator service to moderate user-generated content.

You need to moderate content containing certain phrases without repeatedly scanning the collected content.

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

NOTE: Each correct selection is worth one point.)

- A. Term List API
- B. Text Moderation API
- C. Image Moderation API
- D. Workflow API

ANSWER: A C

Explanation:

The default global list of terms in Azure Content Moderator is sufficient for most content moderation needs. However, you might need to screen for terms that are specific to your organization. For example, you might want to tag competitor names for further review.

Use the List Management API to create custom lists of terms to use with the Text Moderation API. The Text - Screen operation scans your text for profanity, and also compares text against custom and shared blacklists.

C: Use Content Moderator's machine-assisted image moderation and human-in-the-loop Review tool to moderate images for adult and racy content.

Instead of moderating the same image multiple times, you add the offensive images to your custom list of blocked content. That way, your content moderation system compares incoming images against your custom lists and stops any further processing.

Incorrect Answers:

B: Use the Text Moderation API in Azure Content Moderator to scan your text content. The operation scans your content for profanity, and compares the content against custom and shared blacklists.

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/content-moderator/try-terms-list-api> <https://docs.microsoft.com/en-us/azure/cognitive-services/content-moderator/image-moderation-api>

QUESTION NO: 14 - (DRAG DROP)

DRAG DROP

You have an intelligent edge solution that processes data and outputs the data to an Azure Cosmos DB account that uses the SQL API.

You need to ensure that you can perform full text searches of the data.

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

Create a cognitive pipeline.

Create an Azure Search index.

Create an Azure Service Bus.

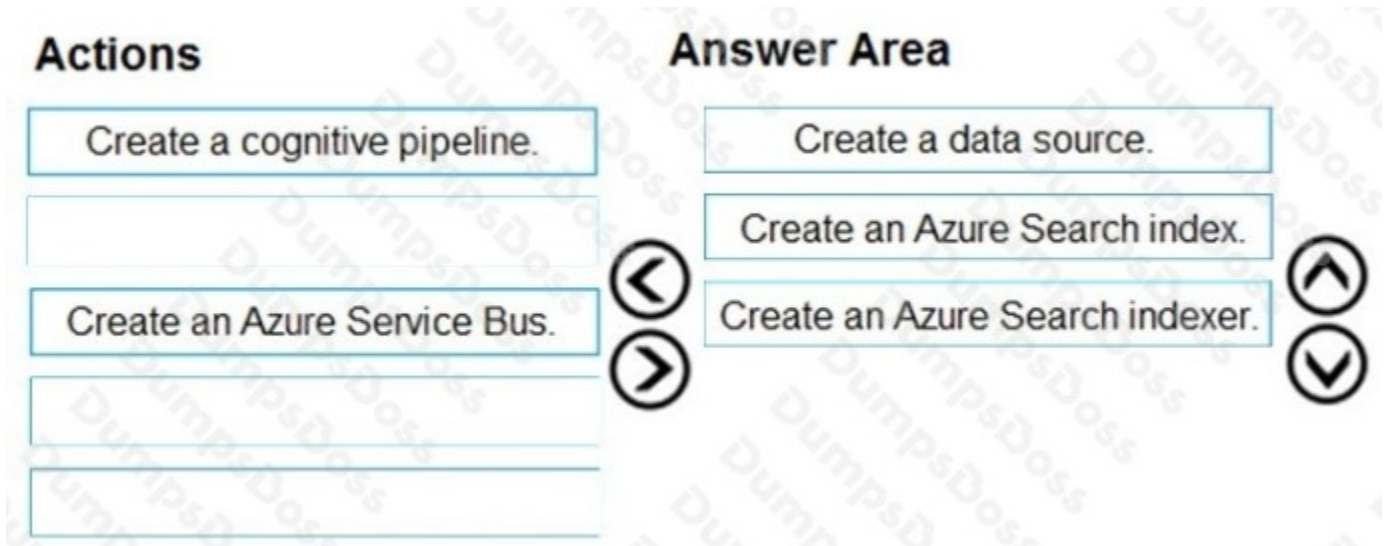
Create a data source.

Create an Azure Search indexer.

Answer Area



ANSWER:



Explanation:

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

Integrate AI models into solutions

QUESTION NO: 15

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>