DUMPSARENA

Designing and Implementing an Azure AI

Microsoft AI-100

Version Demo

Total Demo Questions: 15

Total Premium Questions: 219

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.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

Your company develops an API application that is orchestrated by using Kubernetes. You need to deploy the application.

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

NOTE: Each correct selection is worth one point.

A. Create a Kubernetes cluster.

B. Create an Azure Container Registry instance.

C. Create a container image file.

D. Create a Web App for Containers.

E. Create an Azure container instance.

ANSWER: A B C

Explanation:

References: https://docs.microsoft.com/en-us/azure/aks/tutorial-kubernetes-prepare-app

QUESTION NO: 2

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 modify the custom time interval for the training phase of App1.

Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

Instead, use a LargePersonGroup. 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 NO: 3

You are developing an AI application for a virtual classroom.

Your application must perform the following tasks:

Capture audio and video data of the students in the virtual classroom.

Identify each student in the virtual classroom for attendance purposes.

Alert the teacher when a student's facial expression indicates the student having difficulty.

Which two of the following actions should you take? Each correct selection is part of the solution.

NOTE: Each correct selection is worth one point.

A. Make use of the Face API

B. Make use of the Computer Vision API

C. Make use of the QnA Maker

D. Make use of the Video Indexer API

E. Make use of the Custom Speech Service

ANSWER: A D

Explanation:

Azure Video Indexer is a cloud application built on Azure Media Analytics, Azure Search, Cognitive Services (such as the Face API, Microsoft Translator, the Computer Vision API, and Custom Speech Service). It enables you to extract the insights from your videos using Video Indexer video and audio models.

Face API enables you to search, identify, and match faces in your private repository of up to 1 million people.

The Face API now integrates emotion recognition, returning the confidence across a set of emotions for each face in the image such as anger, contempt, disgust, fear, happiness, neutral, sadness, and surprise. These emotions are understood to be cross-culturally and universally communicated with particular facial expressions.

Reference:

https://docs.microsoft.com/en-us/azure/media-services/video-indexer/video-indexer-overview https://azure.microsoft.com/en-us/services/cognitive-services/face/

QUESTION NO: 4 - (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:

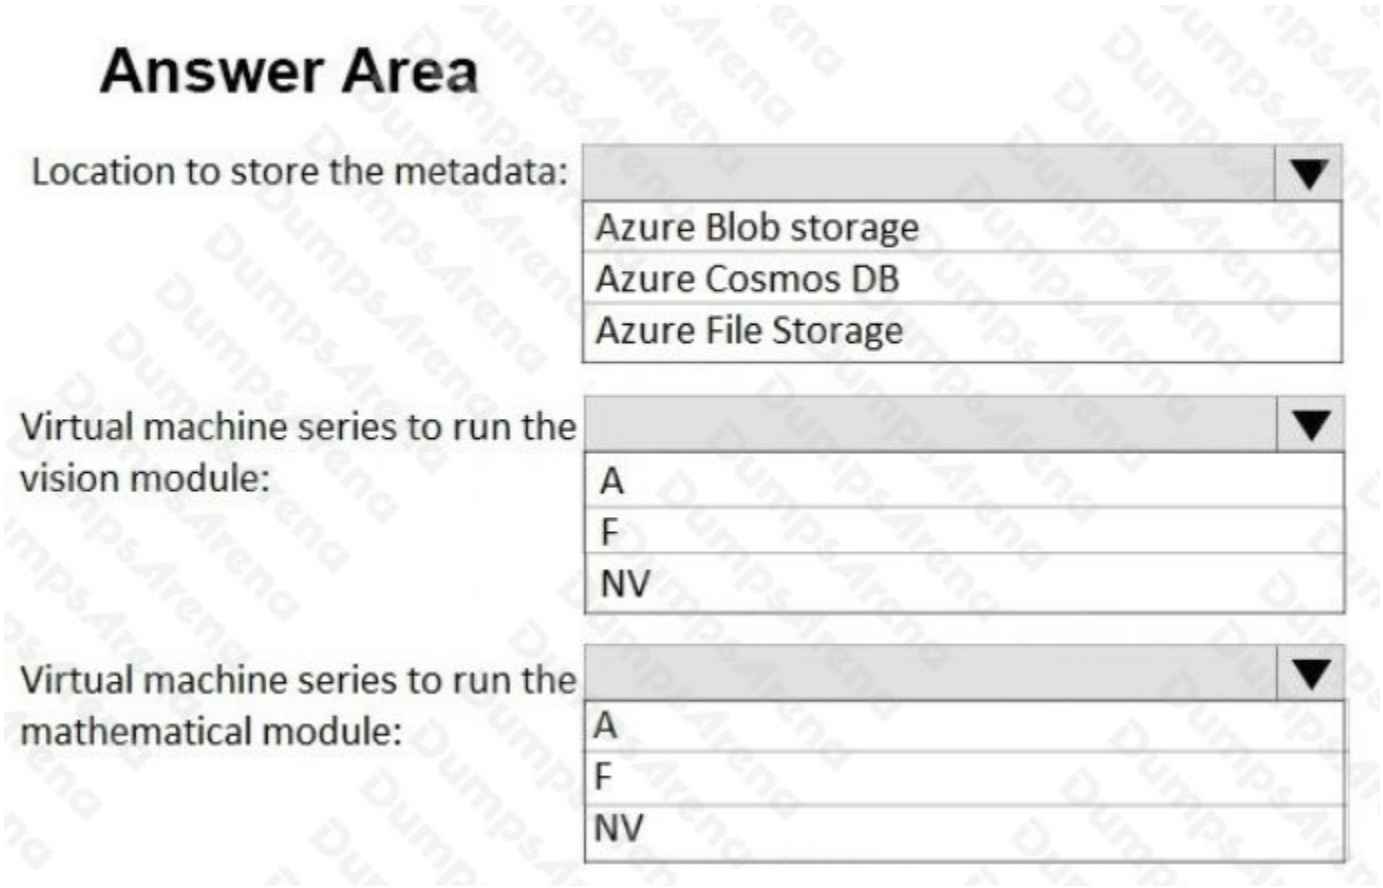
Answer Area
Location to store the metadata:
Azure Blob storage
Azure Cosmos DB
Azure File Storage
Virtual machine series to run the vision module:
A
F
NV
Virtual machine series to run the mathematical module:
A
F
NV

ANSWER:

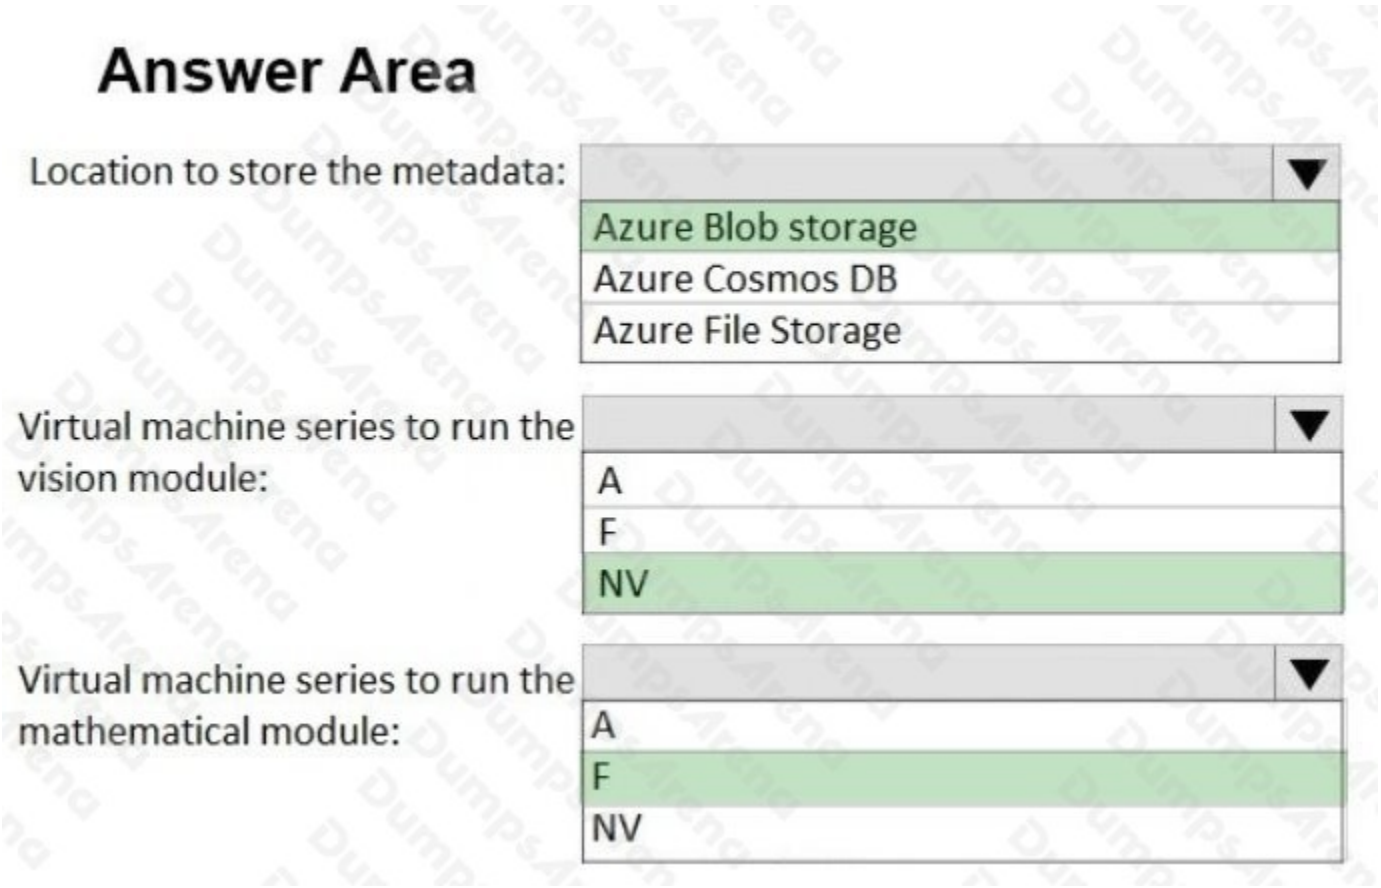

Explanation:

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: 5

You are designing an Azure Batch AI solution that will perform image recognition. The solution will be used to train several Azure Machine Learning models.

You need to recommend a compute infrastructure for the solution that minimizes the processing time.

What should you recommend?

A. Compute optimized virtual machines.

B. Memory optimized virtual machines.

C. GPU optimized virtual machines.

D. General purpose virtual machines.

ANSWER: C

Explanation:

GPU optimized virtual machines are specialized virtual machines targeted for heavy graphic rendering and video editing, as well as model training and inferencing (ND) with deep learning.

Reference:

https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sizes-gpu

QUESTION NO: 6

You are developing a bot for an ecommerce application. The bot will support five languages.

The bot will use Language Understanding (LUIS) to detect the language of the customer, and QnA Maker to answer common customer questions. LUIS supports all the languages.

You need to determine the minimum number of Azure resources that you must create for the bot.

You create five instances of QnA Maker and one instance Language Understanding (LUIS).

Does this action accomplish your objective?

A. Yes, it does

B. No, it does not

ANSWER: B

Explanation:

You need to have a new QnA Maker resource for each language.

If LUIS supports all the languages, you develop a LUIS app for each language. Each LUIS app has a unique app ID, and endpoint log. If you need to provide language understanding for a language LUIS does not support, you can use Microsoft Translator API to translate the utterance into a supported language, submit the utterance to the LUIS endpoint, and receive the resulting scores.

Reference:

https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/overview/language-support
https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-language-support

QUESTION NO: 7

You create an Azure Machine Learning Studio experiment.

You plan to publish the experiment as a Machine Learning Web service.

You need to ensure that you can consume the web service from Microsoft Excel spreadsheets.

What should you use?

A. a Batch Execution Service (BES) and an API key

B. a Batch Execution Service (BES) and an Azure managed identity

C. a Request-Response Service (RRS) and an Azure managed identity

D. a Request-Response Service (RRS) and an API key

ANSWER: D

Explanation:

Steps to Add a New web service

1. Deploy a web service or use an existing Web service.

2. Click Consume.

3. Look for the Basic consumption info section. Copy and save the Primary Key and the Request-Response URL.

4. In Excel, go to the Web Services section (if you are in the Predict section, click the back arrow to go to the list of web services).

5. Click Add Web Service.

6. Paste the URL into the Excel add-in text box labeled URL.

7. Paste the API/Primary key into the text box labeled API key.

8. Click Add.

References: https://docs.microsoft.com/en-us/azure/machine-learning/studio/excel-add-in-for-web-services

QUESTION NO: 8

You have a database that contains sales data.

You plan to process the sales data by using two data streams named Stream1 and Stream2. Stream1 will be used for purchase order data. Stream2 will be used for reference data.

The reference data is stored in CSV files.

You need to recommend an ingestion solution for each data stream.

What two solutions should you recommend? Each correct answer is a complete solution.

NOTE: Each correct selection is worth one point.

A. an Azure event hub for Stream1 and Azure Blob storage for Stream2

B. Azure Blob storage for Stream1 and Stream2

C. an Azure event hub for Stream1 and Stream2

D. Azure Blob storage for Stream1 and Azure Cosmos DB for Stream2

E. Azure Cosmos DB for Stream1 and an Azure event hub for Stream2

ANSWER: A B

Explanation:

Stream1 - Azure Event

Stream2 - Blob Storage

Azure Event Hubs is a highly scalable data streaming platform and event ingestion service, capable of receiving and processing millions of events per second. Event Hubs can process and store events, data, or telemetry produced by distributed software and devices. Data sent to an event hub can be transformed and stored using any real-time analytics provider or batching/storage adapters. Event Hubs provides publish-subscribe capabilities with low latency at massive scale, which makes it appropriate for big data scenarios.

Stream1, Stream2 - Blob Storage

Stream Analytics has first-class integration with Azure data streams as inputs from three kinds of resources: Azure Event Hubs

Azure IoT Hub

Azure Blob storage

These input resources can live in the same Azure subscription as your Stream Analytics job or a different subscription.

References: https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/real-time-ingestion

QUESTION NO: 9

You create an Azure Cognitive Services resource.

A data scientist needs to make API calls to the Cognitive Services resource.

Which two values should you provide to the data scientist? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Endpoint URL

B. Resource name

C. Access key

D. Resource group name

E. Subscription ID

ANSWER: A C

Explanation:

Reference:

https://www.c-sharpcorner.com/article/using-cognitive-service-face-api-with-azure-logic-app/

QUESTION NO: 10 - (DRAG DROP)

DRAG DROP

You need to build a pipeline for an Azure Machine Learning experiment.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

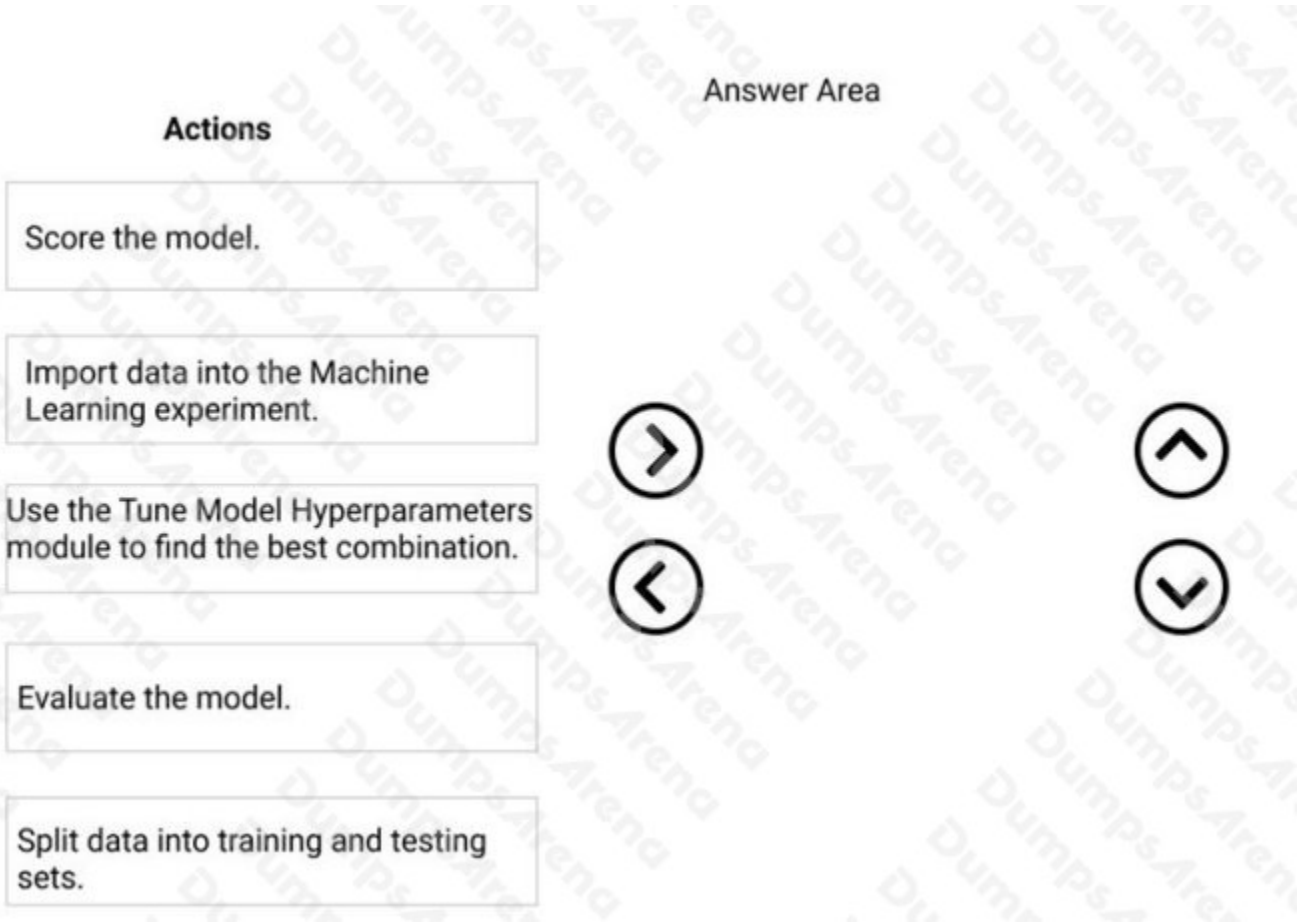
Actions
Score the model.
Import data into the Machine Learning experiment.
Use the Tune Model Hyperparameters module to find the best combination.
Evaluate the model.
Split data into training and testing sets.
Answer Area

ANSWER:

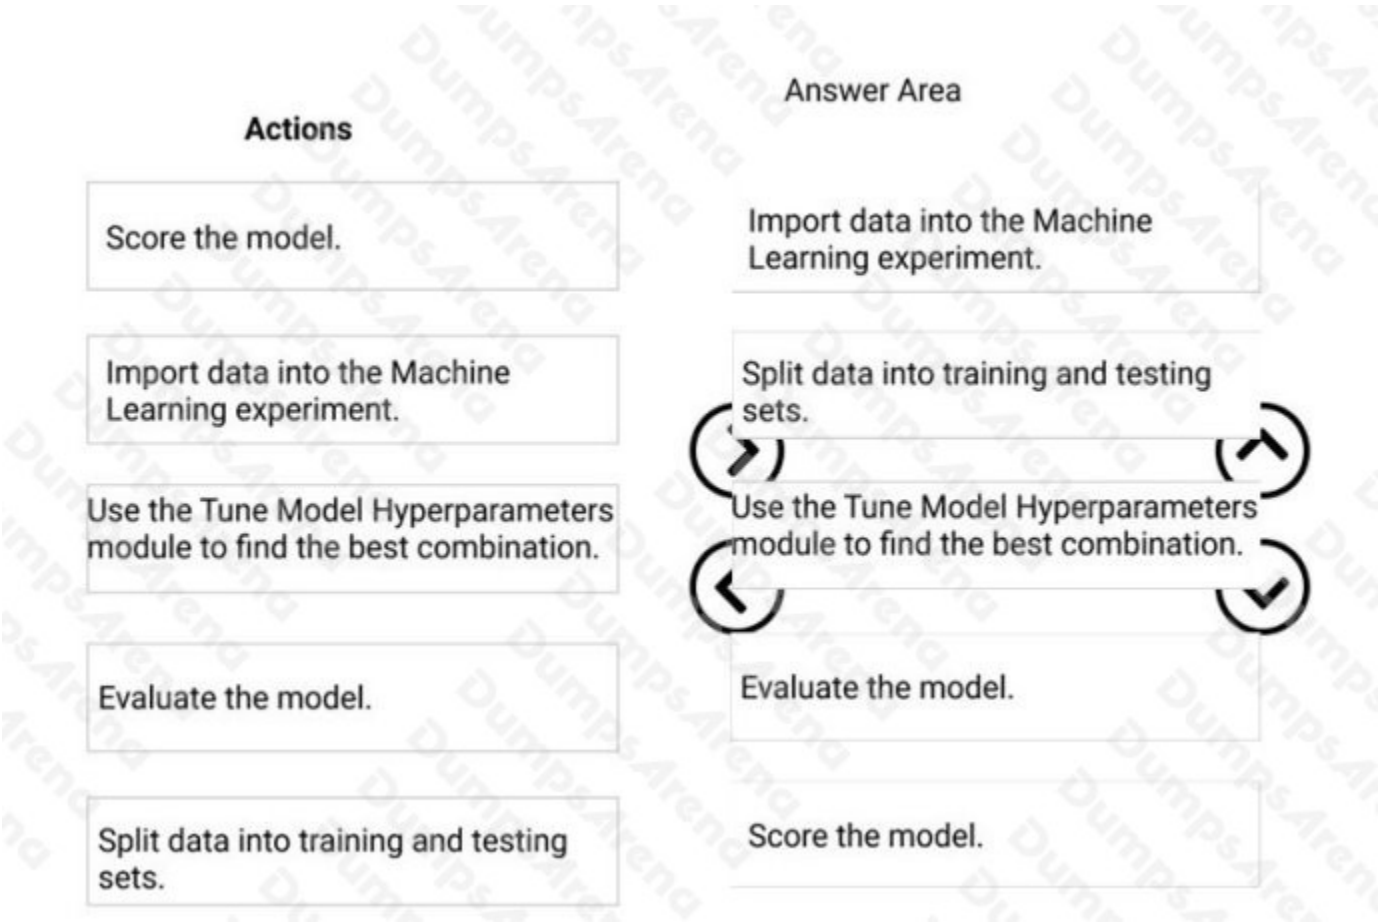

Explanation:

References: https://azure.microsoft.com/en-in/blog/experimentation-using-azure-machine-learning/
https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/machine-learning-modules

QUESTION NO: 11 - (HOTSPOT)

HOTSPOT

You are creating an Azure Cognitive Search capability for the internal documents of your company.

You plan to use Azure Cognitive Services to improve the knowledge gained from the documents.

You configure data import into Azure Cognitive Search as shown in the following exhibit.

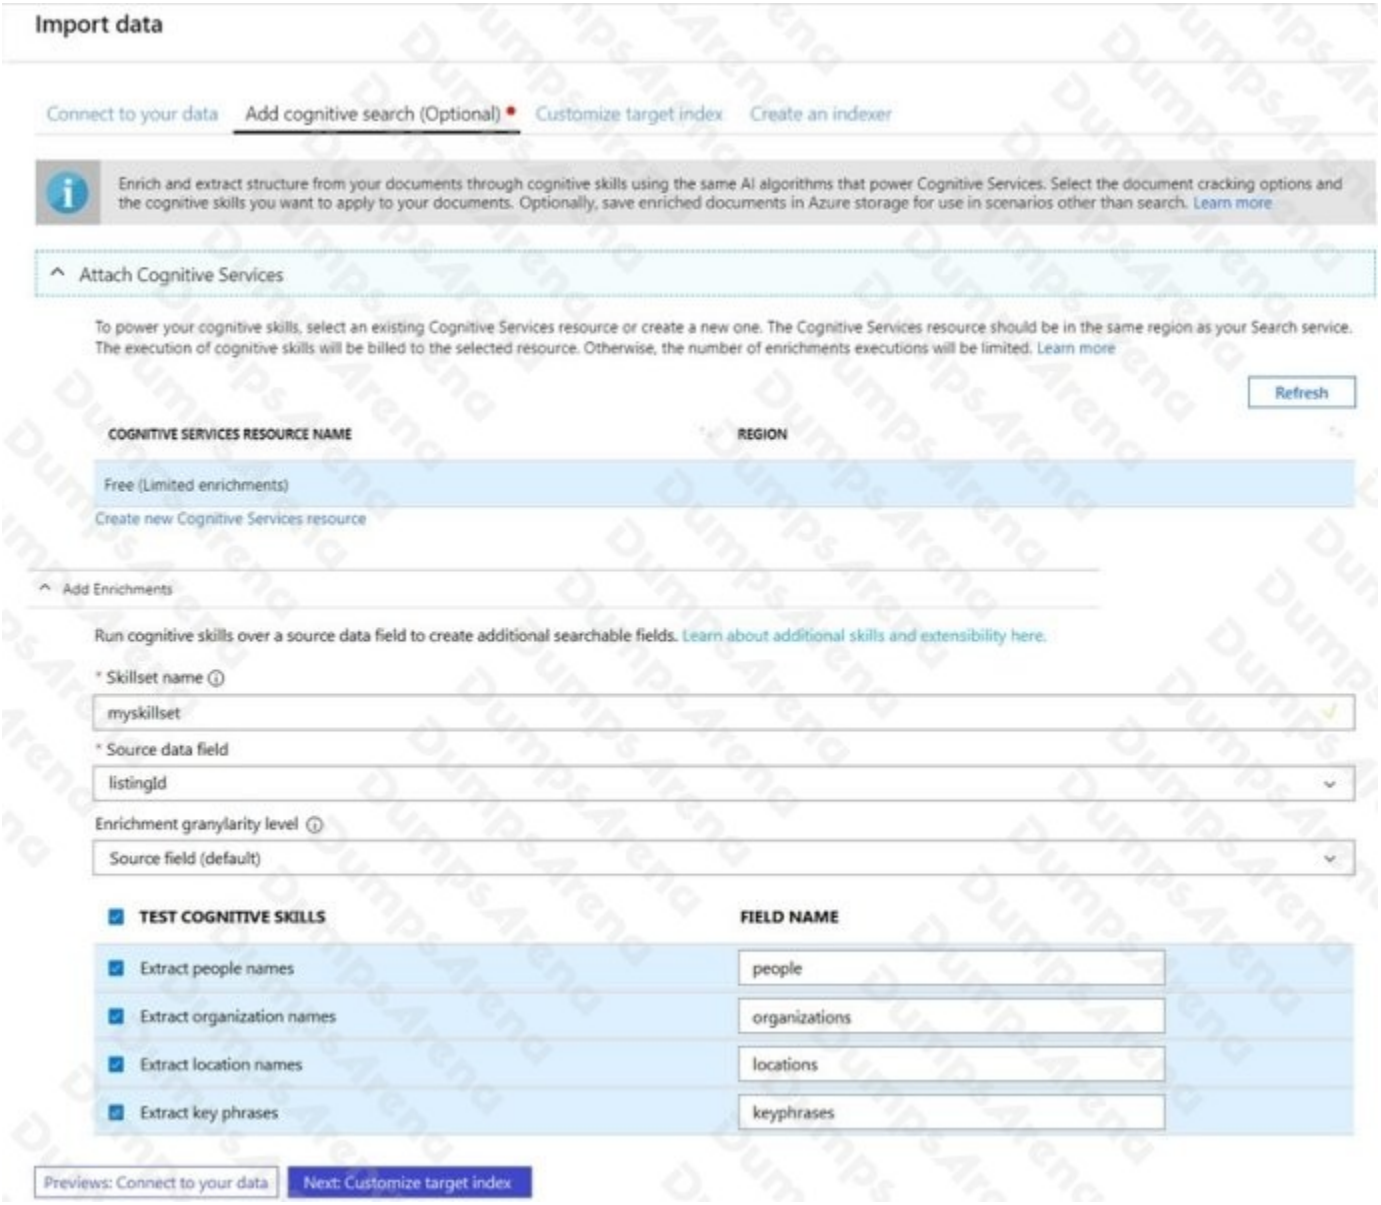

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

To extract relevant information from the documents, the source data field must **[answer choice]**.

- be left blank
- stay the same
- change to a different field

To support GDPR compliance, you can **[answer choice]**.

- clear the Extract people names check box
- clear all the Text Cognitive Skills check boxes
- clear the Extract organization names check box

ANSWER:

Answer Area

To extract relevant information from the documents, the source data field must **[answer choice]**.

- be left blank
- stay the same
- change to a different field

To support GDPR compliance, you can **[answer choice]**.

- clear the Extract people names check box
- clear all the Text Cognitive Skills check boxes
- clear the Extract organization names check box

Explanation:

Reference:

https://docs.microsoft.com/en-us/azure/search/cognitive-search-quickstart-blob

Deploy and manage solutions

QUESTION NO: 12

You plan to build an application that will perform predictive analytics. Users will be able to consume the application data by using Microsoft Power BI or a custom website.

You need to ensure that you can audit application usage.

Which auditing solution should you use?

A. Azure Storage Analytics

B. Azure Application Insights

C. Azure diagnostics logs

D. Azure Active Directory (Azure AD) reporting

ANSWER: D

Explanation:

References: https://docs.microsoft.com/en-us/azure/active-directory/reports-monitoring/concept-audit-logs

QUESTION NO: 13

You plan to deploy two AI applications named AI1 and AI2. The data for the applications will be stored in a relational database.

You need to ensure that the users of AI1 and AI2 can see only data in each user's respective geographic region. The solution must be enforced at the database level by using row-level security.

Which database solution should you use to store the application data?

A. Microsoft SQL Server on a Microsoft Azure virtual machine

B. Microsoft Azure Database for MySQL

C. Microsoft Azure Data Lake Store

D. Microsoft Azure Cosmos DB

ANSWER: A

Explanation:

Row-level security is supported by SQL Server, Azure SQL Database, and Azure SQL Data Warehouse.

References: https://docs.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-2017

QUESTION NO: 14

You need to build a solution to monitor Twitter. The solution must meet the following requirements:

▪ Send an email message to the marketing department when negative Twitter messages are detected. ▪ Run sentiment analysis on Twitter messages that mention specific tags. ▪ Use the least amount of custom code possible.

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. Azure Databricks

B. Azure Stream Analytics

C. Azure Functions

D. Azure Cognitive Services

E. Azure Logic Apps

ANSWER: B E

Explanation:

References: https://docs.microsoft.com/en-us/azure/stream-analytics/streaming-technologies https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-twitter-sentiment-analysis-trends

QUESTION NO: 15 - (DRAG DROP)

You have developed an AI application for your company.

You want to prepare the application for deployment to Kubernetes.

Which three of the following actions should you perform? To answer, move the selected actions from the list of actions to the answer area and rearrange them in the right order.

NOTE: Each correct selection is worth one point.

Select and Place:

ANSWER:

Explanation:

Reference: https://docs.microsoft.com/en-us/azure/aks/tutorial-kubernetes-prepare-app