

AI-100^{Q&As}

Designing and Implementing an Azure AI Solution

Pass Microsoft AI-100 Exam with 100% Guarantee

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

<https://www.leads4pass.com/ai-100.html>

100% Passing Guarantee
100% Money Back Assurance

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

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



QUESTION 1

You have an Azure Machine Learning experiment that must comply with GDPR regulations. You need to track compliance of the experiment and store documentation about the experiment.

What should you use?

- A. Azure Table storage
- B. Azure Security Center
- C. An Azure Log Analytics workspace
- D. Compliance Manager

Correct Answer: D

References: <https://azure.microsoft.com/en-us/blog/new-capabilities-to-enable-robust-gdpr-compliance/>

QUESTION 2

You have created an AI solution that uses several PersonGroup objects.

One of the PersonGroup objects contains thousands of entries and cannot accept any new entries.

You want to be able to add new entries to the PersonGroup object. The PersonGroup object must be identifiable by all the entries.

Which of the following actions should you take?

- A. Compress the entries from the PersonGroup object.
- B. Create another PersonGroup object with the same name.
- C. Migrate the PersonGroup to a LargePersonGroup object.
- D. Archive some of the entries from the PersonGroup object.

Correct Answer: C

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.

Reference:

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

QUESTION 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 have an Azure SQL database, an Azure Data Lake Storage Gen 2 account, and an API developed by using Azure Machine Learning Studio.

You need to ingest data once daily from the database, score each row by using the API, and write the data to the storage account.

Solution: You create a scheduled Jupyter Notebook in Azure Databricks.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

We need to schedule the job in Azure Data Factory.

QUESTION 4

You have deployed 1,000 sensors for an AI application that you are developing. The sensors generate large amounts data that is ingested on an hourly basis.

You want your application to analyze the data generated by the sensors in real-time.

Which of the following actions should you take?

- A. Make use of Azure Kubernetes Service (AKS)
- B. Make use of Azure Cosmos DB
- C. Make use of an Azure HDInsight Hadoop cluster
- D. Make use of Azure Data Factory

Correct Answer: A

To analyze the data generated by the sensors in real-time, you can make use of Azure Kubernetes Service (AKS). AKS is a container orchestration service that allows you to deploy, manage, and scale containerized applications. By deploying your AI application on AKS, you can process the data in real-time using containers and take advantage of the scalability and flexibility offered by AKS.

QUESTION 5

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

1.

Send an email message to the marketing department when negative Twitter messages are detected.

2.

Run sentiment analysis on Twitter messages that mention specific tags.

3.

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

Correct Answer: BE

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>

[Latest AI-100 Dumps](#)

[AI-100 PDF Dumps](#)

[AI-100 VCE Dumps](#)