

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.pass2lead.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 are developing an app that will analyze sensitive data from global users.

Your app must adhere the following compliance policies:

The app must not store data in the cloud.

The app not use services in the cloud to process the data.

Which of the following actions should you take?

- A. Make use of Azure Machine Learning Studio
- B. Make use of Docker containers for the Text Analytics
- C. Make use of a Text Analytics container deployed to Azure Kubernetes Service
- D. Make use of Microsoft Machine Learning (MML) for Apache Spark

Correct Answer: D

<https://github.com/MicrosoftDocs/azure-docs/blob/ccf49761e4aefed30d723805f4f09e753615fb09/articles/cognitive-services/cognitive-services-container-support.md>

QUESTION 2

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

Correct Answer: BC

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

QUESTION 3

You are developing an app that consumes data from several Azure IoT Edge devices.

You need to implement a storage solution for the app. Your solution must allow data to be queried in real-time as it streams into the solution. You need to ensure that your solution provides the least amount of latency for loading data.

You want the data files to persist on the devices for at least 14 days.

What storage solution should you implement?

- A. Azure Data Lake Analytics
- B. Azure Data Factory Edge
- C. Azure HDInsight Hadoop cluster
- D. Azure SQL database with In-Memory OLTP

Correct Answer: B

To implement a storage solution that allows real-time querying of data with low latency and ensures data persistence on the devices for at least 14 days, you should use Azure Data Factory Edge.

Azure Data Factory Edge is designed to collect, transform, and analyze data from edge devices. It provides a scalable and reliable solution for ingesting, processing, and storing data at the edge. With Azure Data Factory Edge, you can configure data flows to ingest data from your Azure IoT Edge devices and store it in the desired storage format (e.g., Azure Blob storage, Azure Data Lake Storage).

By using Azure Data Factory Edge, you can set up data pipelines to stream and process the data in real-time, allowing you to query the data as it streams into the solution with minimal latency. Additionally, you can define retention policies to ensure that the data files persist on the devices for at least 14 days.

QUESTION 4

Von have an Azure SQL database w 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

QUESTION 5

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 create a second PersonGroup object for Ben Smith. Does this meet the goal?

A. Yes

B. No

Correct Answer: B

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>

[AI-100 Practice Test](#)

[AI-100 Exam Questions](#)

[AI-100 Braindumps](#)