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.pass4itsure.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

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 Azure Functions as an IoT Edge module.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Instead use Azure Stream Analytics and REST API.

Note. Available in both the cloud and Azure IoT Edge, Azure Stream Analytics offers built-in machine learning based anomaly detection capabilities that can be used to monitor the two most commonly occurring anomalies: temporary and

persistent.

Stream Analytics supports user-defined functions, via REST API, that call out to Azure Machine Learning endpoints.

References:

https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-machine-learning-anomaly-detection

QUESTION 2

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

Correct Answer: CE

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 3

You company\\'s developers have created an Azure Data Factory pipeline that moves data from an on-premises server to Azure Storage. The pipeline consumes

Azure Cognitive Services APIs.

You need to deploy the pipeline. Your solution must minimize custom code.

You use Integration Runtime to move data to the cloud and Azure API Management to consume Cognitive Services APIs.

Does this action accomplish your objective?

A. Yes, it does

B. No, it does not

Correct Answer: B

The given solution does not accomplish the objective of minimizing custom code. While Integration Runtime can be used to move data from an on-premises server to Azure Storage, and Azure API Management can be used to consume

Cognitive Services APIs, it does not eliminate the need for custom code.

To minimize custom code and simplify the deployment of the pipeline, you should consider using Azure Logic Apps. Azure Logic Apps provide a serverless and code-free way to orchestrate workflows and integrate different systems and

services.

Reference:

https://docs.microsoft.com/en-us/azure/data-factory/concepts-integration-runtime https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-examples-and-scenarios

QUESTION 4

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 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 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 configure Azure Monitor for containers.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B