



# AI-100<sup>Q&As</sup>

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.passapply.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 AI solution that will use in-memory caching and a columnar storage engine for Apache Hive queries. What HDInsight platform should you use?

- A. Apache Kafka
- B. Apache Spark
- C. Interactive Query
- D. Apache Storm

Correct Answer: C

Interactive Query provides In-memory caching and improved columnar storage engine for Hive queries. Reference: <https://docs.microsoft.com/bs-latn-ba/azure/hdinsight/interactive-query/apache-interactive-query-get-started>

---

### QUESTION 2

You are developing a mobile application that will perform optical character recognition (OCR) from photos.

The application will annotate the photos by using metadata, store the photos in Azure Blob storage, and then score the photos by using an Azure Machine Learning model.

What should you use to process the data?

- A. Azure Event Hubs
- B. Azure Functions
- C. Azure Stream Analytics
- D. Azure Logic Apps

Correct Answer: A

---

### QUESTION 3

You are designing an AI solution in Azure that will perform image classification.

You need to identify which processing platform will provide you with the ability to update the logic over time. The solution must have the lowest latency for inferencing without having to batch.

Which compute target should you identify?

- A. graphics processing units (GPUs)
- B. field-programmable gate arrays (FPGAs)
- C. central processing units (CPUs)



D. application-specific integrated circuits (ASICs)

Correct Answer: B

FPGAs, such as those available on Azure, provide performance close to ASICs. They are also flexible and reconfigurable over time, to implement new logic. Incorrect Answers:

D: ASICs are custom circuits, such as Google's TensorFlow Processor Units (TPU), provide the highest efficiency. They can't be reconfigured as your needs change.

References: <https://docs.microsoft.com/en-us/azure/machine-learning/service/concept-accelerate-with-fpgas>

---

#### QUESTION 4

You have Azure IoT Edge devices that collect measurements every 30 seconds.

You plan to send the measurements to an Azure IoT hub.

You need to ensure that every event is processed as quickly as possible.

What should you use?

- A. Apache Kafka
- B. Azure Stream Analytics record functions
- C. Azure Stream Analytics windowing functions
- D. Azure Machine Learning on the IoT Edge devices

Correct Answer: A

<https://docs.microsoft.com/en-us/azure/iot-edge/tutorial-deploy-stream-analytics>

---

#### 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 need to create an IoT solution that performs the following tasks:

1.  
Identifies hazards
- 2.



Provides a real-time online dashboard

3.

Takes images of an area every minute

4.

Counts the number of people in an area every minute

Solution: You implement Azure Cognitive Services containers on the IoT devices, and then you configure results to be sent to an Azure IoT hub. You configure Microsoft Power BI to connect to the IoT hub by using Azure Stream Analytics.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

There is support for running Azure Cognitive Services containers for Text Analytics and Language Understanding containers on edge devices with Azure IoT Edge. This means that all your workloads can be run locally where your data is

being generated while keeping the simplicity of the cloud to manage them remotely, securely and at scale.

You would have to set up an IoT Edge device and its IoT Hub.

Note: Azure Stream Analytics enables you to take advantage of one of the leading business intelligence tools, Microsoft Power BI.

Get your IoT hub ready for data access by adding a consumer group.

Create, configure, and run a Stream Analytics job for data transfer from your IoT hub to your Power BI account.

Create and publish a Power BI report to visualize the data.

References:

<https://azure.microsoft.com/es-es/blog/running-cognitive-services-on-iot-edge/>

<https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-live-data-visualization-in-power-bi>

---

## QUESTION 6

You are developing a mobile application. You want to implement search functionality in the application.

Your solution must meet the following requirements:

Users must be able to run searches by typing in their search query.

Users must be able to run searches by voice commands.

Which of the following actions should you take?



- A. Make use of Language Understanding (LUIS)
- B. Make use of QnA Maker
- C. Make use of Bing Entity Search
- D. Make use of Azure Cognitive Search

Correct Answer: A

Language Understanding (LUIS) is a natural language processing service provided by Microsoft Azure. It allows you to build applications that can understand and interpret user's intents from their spoken or typed input. By using LUIS, you can train a language model to recognize specific search intents and extract relevant entities from user queries. LUIS supports both text and speech input, making it suitable for implementing search functionality that can be accessed via typing or voice commands.

---

### QUESTION 7

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 create several AI models in Azure Machine Learning Studio.

You deploy the models to a production environment.

You need to monitor the compute performance of the models.

Solution: You enable AppInsights diagnostics.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

You need to enable Model data collection.

References: <https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-enable-data-collection>

---

### QUESTION 8

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 add an SSH key to the node, and then you create an SSH connection.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

By default, SSH keys are generated when you create an AKS cluster. If you did not specify your own SSH keys when you created your AKS cluster, add your public SSH keys to the AKS nodes. You also need to create an SSH connection to the AKS node.

References: <https://docs.microsoft.com/en-us/azure/aks/ssh>

---

## QUESTION 9

You are developing an AI application for your company. The application will use Microsoft Azure Stream Analytics.

You save the outputs from the Stream Analytics workflows to the cloud.

Which of the following actions should you take?

A. Make use of a Hive table in Azure HDInsight

B. Make use of Azure Cosmos DB

C. Make use of Azure File storage

D. Make use of Azure Table storage

Correct Answer: B

While options like Azure HDInsight with a Hive table, Azure File storage, and Azure Table storage have their own use cases, they are not specifically designed for efficiently storing and querying real-time streaming data from Azure Stream Analytics. Azure Cosmos DB is the more appropriate choice for this scenario.

---

## QUESTION 10

You have Azure IoT Edge devices that collect measurements every 30 seconds.



You plan to send the measurements to an Azure IoT hub.

You need to process events in the cloud.

What should you use?

- A. Apache Kafka
- B. Azure Stream Analytics record functions
- C. Azure Stream Analytics windowing functions
- D. Azure Machine Learning on the IoT Edge devices

Correct Answer: D

Use Azure Notebooks to develop a machine learning module and deploy it to a Linux device running Azure IoT Edge. You can use IoT Edge modules to deploy code that implements your business logic directly to your IoT Edge devices.

References: <https://docs.microsoft.com/en-us/azure/iot-edge/tutorial-deploy-machine-learning>

---

## QUESTION 11

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 an Azure Data Factory pipeline that contains the Machine Learning Batch Execution activity. Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Using the Batch Execution Activity in an Azure Data Factory pipeline, you can invoke an Azure Machine Learning Studio (classic) web service to make predictions on the data in batch

Reference: <https://docs.microsoft.com/en-us/azure/data-factory/transform-data-using-machine-learning>

---

## QUESTION 12



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.

Correct Answer: B

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>

---

### QUESTION 13

You have an Azure Machine Learning experiment.

You need to validate that the experiment meets GDPR regulation requirements and stores documentation about the experiment.

What should you use?

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

Correct Answer: A

Compliance Manager for Azure helps you assess and manage GDPR compliance. Compliance Manager is a free, Microsoft cloud services solution designed to help organizations meet complex compliance obligations, including the GDPR, ISO 27001, ISO 27018, and NIST 800-53. Generally available today for Azure customers, the Compliance Manager GDPR dashboard enables you to assign, track, and record your GDPR compliance activities so you can collaborate across teams and manage your documents for creating audit reports more easily.





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

---

#### QUESTION 14

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

Correct Answer: B

---

#### QUESTION 15

You plan to perform analytics of the medical records of patients located around the world. You need to recommend a solution that avoids storing and processing data in the cloud. What should you include in the recommendation?

A. Azure Machine Learning Studio

B. the Text Analytics API that has container support

C. Azure Machine Learning services

D. an Apache Spark cluster that uses MMLSpark

Correct Answer: D

The Microsoft Machine Learning Library for Apache Spark (MMLSpark) assists in provisioning scalable machine learning models for large datasets, especially for building deep learning problems. MMLSpark works with SparkML pipelines, including Microsoft CNTK and the OpenCV library, which provide end-to-end support for the ingress and processing of image input data, categorization of images, and text analytics using pre-trained deep learning algorithms.

References: [https://subscription.packtpub.com/book/big\\_data\\_and\\_business\\_intelligence/9781789131956/10/ch10lv1s/ec61/an-overview-of-the-microsoft-machine-learning-library-for-apache-spark-mmlspark](https://subscription.packtpub.com/book/big_data_and_business_intelligence/9781789131956/10/ch10lv1s/ec61/an-overview-of-the-microsoft-machine-learning-library-for-apache-spark-mmlspark)



VCE & PDF

PassApply.com

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

2024 Latest passapply AI-100 PDF and VCE dumps Download

---

[AI-100 PDF Dumps](#)

[AI-100 Practice Test](#)

[AI-100 Study Guide](#)