



# 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.geekcert.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 plan to implement a new data warehouse for a planned AI solution. You have the following information regarding the data warehouse:

1.

The data files will be available in one week.

2.

Most queries that will be executed against the data warehouse will be ad-hoc queries.

3.

The schemas of data files that will be loaded to the data warehouse will change often.

4.

One month after the planned implementation, the data warehouse will contain 15 TB of data.

You need to recommend a database solution to support the planned implementation.

What two solutions should you include in the recommendation? Each correct answer is a complete solution.

NOTE: Each correct selection is worth one point.

A. Apache Hadoop

B. Apache Spark

C. A Microsoft Azure SQL database

D. An Azure virtual machine that runs Microsoft SQL Server

Correct Answer: C

References: <https://docs.microsoft.com/en-us/azure/sql-database/saas-multitenantdb-adhoc-reporting>

---

### QUESTION 2

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 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 an Azure Data Factory pipeline that contains a Jupyter Notebook activity.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

---

### QUESTION 4

You are developing a Microsoft Bot Framework application. The application consumes structured NoSQL data that must be stored in the cloud.

You implement Azure Blob storage for the application. You want access to the blob store to be controlled by using a role.

You implement On-premises Active Directory Domain Services (AD DS).

Does this action accomplish your objective?

A. Yes, it does

B. No, it does not



Correct Answer: B

Implementing On-premises Active Directory Domain Services (AD DS) alone does not accomplish the objective of controlling access to Azure Blob storage based on a role.

On-premises AD DS is designed for managing user accounts, groups, and authentication within an on-premises network. It does not directly integrate with Azure Blob storage to control access based on roles.

To achieve role-based access control (RBAC) for Azure Blob storage, you should consider integrating Azure Active Directory (Azure AD) with your application. Azure AD provides a cloud-based identity and access management solution that can be used to define roles, assign permissions, and control access to Azure resources, including Blob storage.

By implementing Azure AD integration, you can manage access to Blob storage based on roles defined in Azure AD, granting or restricting access to users or groups based on their assigned roles.

---

#### QUESTION 5

You are designing an AI application that will perform real-time processing by using Microsoft Azure Stream Analytics.

You need to identify the valid outputs of a Stream Analytics job.

What are three possible outputs? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. A Hive table in Azure HDInsight
- B. Azure SQL Database
- C. Azure Cosmos DB
- D. Azure Blob storage
- E. Azure Redis Cache

Correct Answer: BCD

References: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-define-outputs>