

**100%** Money Back  
**Guarantee**

**Vendor:**Microsoft

**Exam Code:**AI-100

**Exam Name:**Designing and Implementing an Azure AI Solution

**Version:**Demo

**QUESTION 1**

DRAG DROP

You need to build an AI solution that will be shared between several developers and customers.

You plan to write code, host code, and document the runtime all within a single user experience.

You build the environment to host the solution.

Which three actions should you perform in sequence next? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

**Actions**

Create stream Inputs and outputs

Create a new experiment

Build an experiment

Build a notebook

Create an Azure Machine Learning Studio workspace

Create a new notebook

Implement Transact-SQL for the stream query

Create an Azure Stream Analytics service

**Answer Area**



Pass4Lead.com

Correct Answer:

## Actions

Create stream Inputs and outputs

Build an experiment

Build a notebook

Implement Transact-SQL for the stream query

Create an Azure Stream Analytics service

## Answer Area

Create an Azure Machine Learning Studio workspace

Create a new notebook

Create a new experiment

Step 1: Create an Azure Machine Learning Studio workspace

Step 2: Create a notebook You can manage notebooks using the UI, the CLI, and by invoking the Workspace API. To create a notebook

1.

Click the Workspace button Workspace Icon or the Home button Home Icon in the sidebar. Do one of the following: Next to any folder, click the Menu Dropdown on the right side of the text and select Create > Notebook. Create Notebook In the Workspace or a user folder, click Down Caret and select Create > Notebook.

2.

In the Create Notebook dialog, enter a name and select the notebook's primary language.

3.

If there are running clusters, the Cluster drop-down displays. Select the cluster to attach the notebook to.

4.

Click Create.

Step 3: Create a new experiment

Create a new experiment by clicking +NEW at the bottom of the Machine Learning Studio window. Select EXPERIMENT > Blank Experiment.

References:

<https://docs.azuredatabricks.net/user-guide/notebooks/notebook-manage.html> <https://docs.microsoft.com/en-us/azure/machine-learning/service/quickstart-run-cloud-notebook>

---

## QUESTION 2

DRAG DROP

You need to integrate the new Bookings app and the Butler chatbot.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

### Actions

On the page where you want Butler to be used, paste the embed code of the new Bookings app.

From the Channels settings of Butler, copy the secret key and embed code for the Skype for Business channel.

Exchange the secret for a token by connecting to <https://directline.botframework.com/api/tokens>.

From the Channels settings of Butler, retrieve the secret key and embed code for the Web Chat channel.

Replace `s=YOUR_SECRET_HERE` with `t=` followed by the token.

Exchange the secret for a token by connecting to <https://webchat.botframework.com/api/tokens>.

### Answer Area



Correct Answer:

## Actions

From the Channels settings of Butler, copy the secret key and embed code for the Skype for Business channel.

Exchange the secret for a token by connecting to <https://directline.botframework.com/api/tokens>.

## Answer Area

From the Channels settings of Butler, retrieve the secret key and embed code for the Web Chat channel.

On the page where you want Butler to be used, paste the embed code of the new Bookings app

Exchange the secret for a token by connecting to <https://webchat.botframework.com/api/tokens>.

Replace s=YOUR\_SECRET\_HERE with t= followed by the token.

References: <https://docs.microsoft.com/en-us/azure/bot-service/bot-service-channel-connect-webchat?view=azure-bot-service-4.0>

### QUESTION 3

Which RBAC role should you assign to the KeyManagers group?

- A. Cognitive Services Contributor
- B. Security Manager
- C. Cognitive Services User
- D. Security Administrator

Correct Answer: A

References: <https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>

**QUESTION 4**

**HOTSPOT**

Your company is building a cinema chatbot by using the BOT Framework and Language Understanding (LUS).

You are designing the intents and the entities for LUIS. The following are utterances that customers might provide:

Which movies are playing on December 8?

What time the performance of Movie1?

You need to identify which entity types to use . The solution must minimize development effort.

Which entity type should you use for each entity? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

December 8:	<input type="checkbox"/> Composite <input type="checkbox"/> Simple <input type="checkbox"/> Prebuilt
Movie1:	<input type="checkbox"/> Composite <input type="checkbox"/> Simple <input type="checkbox"/> Prebuilt
Two adult tickets in the balcony section:	<input type="checkbox"/> Composite <input type="checkbox"/> Simple <input type="checkbox"/> Prebuilt

Correct Answer:

December 8:	<input checked="" type="checkbox"/> Composite <input type="checkbox"/> Simple <input type="checkbox"/> Prebuilt
Movie1:	<input type="checkbox"/> Composite <input type="checkbox"/> Simple <input checked="" type="checkbox"/> Prebuilt
Two adult tickets in the balcony section:	<input type="checkbox"/> Composite <input checked="" type="checkbox"/> Simple <input type="checkbox"/> Prebuilt

**QUESTION 5**

**HOTSPOT**

You are designing an Azure infrastructure to support an Azure Machine Learning solution that will have multiple phases. The solution must meet the following requirements:

1. Securely query an on-premises database once a week to update product lists.
2. Access the data without using a gateway.
3. Orchestrate the separate phases.

What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

To connect to the on-premises data:

	▼
A point-to-site VPN connection	
A site-to-site VPN connection	
Azure App Service Hybrid Connections	

To orchestrate the phases:

	▼
A Machine Learning experiment	
Azure Machine Learning Studio	
Machine Learning pipelines	

To control the orchestration:

	▼
Azure Automation	
Azure Databricks	
Azure Notebooks	

Correct Answer:



# Answer Area

To connect to the on-premises data:

- A point-to-site VPN connection
- A site-to-site VPN connection
- Azure App Service Hybrid Connections

To orchestrate the phases:

- A Machine Learning experiment
- Azure Machine Learning Studio
- Machine Learning pipelines

To control the orchestration:

- Azure Automation
- Azure Databricks
- Azure Notebooks

Box 1: Azure App Service Hybrid Connections With Hybrid Connections, Azure websites and mobile services can access on-premises resources as if they were located on the same private network. Application admins thus have the flexibility to simply lift-and-shift specific most front-end tiers to Azure with minimal configuration changes, extending their enterprise apps for hybrid scenarios.

Incorrect Answer: The VPN connection solution both use gateways.

Box 2: Machine Learning pipelines Typically when running machine learning algorithms, it involves a sequence of tasks including pre-processing, feature extraction, model fitting, and validation stages. For example, when classifying text documents might involve text segmentation and cleaning, extracting features, and training a classification model with cross-validation. Though there are many libraries we can use for each stage, connecting the dots is not as easy as it may look, especially with large-scale

datasets. Most ML libraries are not designed for distributed computation or they do not provide native support for pipeline creation and tuning.

Box 3: Azure Databricks

References:

<https://azure.microsoft.com/is-is/blog/hybrid-connections-preview/>

<https://databricks.com/glossary/what-are-ml-pipelines>

---

## QUESTION 6

You plan to build an application that will perform predictive analytics. Users will be able to consume the application data by using Microsoft Power BI or a custom website.



You need to ensure that you can audit application usage.

Which auditing solution should you use?

- A. Azure Storage Analytics
- B. Azure Application Insights
- C. Azure diagnostics logs
- D. Azure Active Directory (Azure AD) reporting

Correct Answer: D

References: <https://docs.microsoft.com/en-us/azure/active-directory/reports-monitoring/concept-audit-logs>

---

## 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 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 migrate all the entries to the LargePersonGroup object for Ben Smith.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

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>

---

**QUESTION 8**

**DRAG DROP**

You develop a custom application that uses a token to connect to Azure Cognitive Services resources.

A new security policy requires that all access keys are changed every 30 days.

You need to recommend a solution to implement the security policy.

Which three actions should you recommend be performed every 30 days? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

**Actions**

- Retrieve a token from the Cognitive Services endpoint
- Generate new keys in the Cognitive Services resources
- Generate new keys in Azure Key Vault
- Update the custom application to use the new authorization
- Retrieve a token from the Azure Key Vault endpoint

**Answer Area**



Correct Answer:

## Actions

Generate new keys in Azure Key Vault

Retrieve a token from the Azure Key Vault endpoint

## Answer Area

Generate new keys in the Cognitive Services resources

Retrieve a token from the Cognitive Services endpoint

Update the custom application to use the new authorization

Step 1: Generate new keys in the Cognitive Service resources

Home > SampleTextAnalyticsWestUS - Keys

### SampleTextAnalyticsWestUS - Keys

Cognitive Services

Search (Ctrl+/)

Regenerate Key1 Regenerate Key2

Notice: It may take up to 10 minutes for the newly

NAME
SampleTextAnalyticsWestUS

KEY 1

KEY 2

RESOURCE MANAGEMENT

Keys

Quick start

Step 2: Retrieve a token from the Cognitive Services endpoint

Step 3: Update the custom application to use the new authorization

Each request to an Azure Cognitive Service must include an authentication header. This header passes along a subscription key or access token, which is used to validate your subscription for a service or group of services.

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/authentication>

---

### QUESTION 9

Your company plans to deploy an AI solution that processes IoT data in real-time.

You need to recommend a solution for the planned deployment that meets the following requirements:

1.

Sustain up to 50 Mbps of events without throttling.

2.

Retain data for 60 days. What should you recommend?

- A. Apache Kafka
- B. Microsoft Azure IoT Hub
- C. Microsoft Azure Data Factory
- D. Microsoft Azure Machine Learning

Correct Answer: A

Apache Kafka is an open-source distributed streaming platform that can be used to build real-time streaming data pipelines and applications.

References: <https://docs.microsoft.com/en-us/azure/hdinsight/kafka/apache-kafka-introduction>

---

### QUESTION 10

You have an app that records meetings by using speech-to-text capabilities from the Speech Services API.

You discover that when action items are listed at the end of each meeting, the app transcribes the text inaccurately.

You need to improve the accuracy of the meeting records.

What should you do?

- A. Add a phrase list
- B. Create a custom wake word
- C. Parse the text by using the Language Understanding (LUIS) API
- D. Train a custom model by using Custom Translator

Correct Answer: D

Speech Services API with subscription to the Microsoft Text Translation API enables you to use Custom Translator to use your own data for more accurate translations.

References: <https://www.microsoft.com/en-us/translator/business/customization/>

---

## QUESTION 11

### HOTSPOT

Your company is building a cinema chatbot by using the Bot Framework and Language Understanding (LUIS).

You are designing of the intents and the entities for LUIS.

The following are utterances that customers might provide:

1.

Which movies are playing on December 8?

2.

What time is the performance of Movie1?

3.

I would like to purchase two adult tickets in the balcony section for Movie2.

You need to identify which entity types to use. The solution must minimize development effort.

Which entry type should you use for each entity? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

December 8:

	▼
Composite	
Simple	
Prebuilt	

Movie1:

	▼
Composite	
Simple	
Prebuilt	

Two adult tickets in the balcony section:

	▼
Composite	
Simple	
Prebuilt	

Correct Answer:

## Answer Area

December 8:

	▼
Composite	
Simple	
Prebuilt	

Movie1:

	▼
Composite	
Simple	
Prebuilt	

Two adult tickets in the balcony section:

	▼
Composite	
Simple	
Prebuilt	

Box 1: Prebuilt

Datetime is prebuilt.

Language Understanding (LUIS) provides prebuilt entities. When a prebuilt entity is included in your application, LUIS includes the corresponding entity prediction in the endpoint response.

Box 2: Simple

Box 3: Composite

A composite entity is made up of other entities, such as prebuilt entities, simple, regular expression, and list entities. The separate entities form a whole entity.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-reference-prebuilt-entities>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/reference-entity-composite>

---

### QUESTION 12

You plan to design a solution for an AI implementation that uses data from IoT devices.



You need to recommend a data storage solution for the IoT devices that meets the following requirements:

1.

Allow data to be queried in real-time as it streams into the solution.

2.

Provide the lowest amount of latency for loading data into the solution. What should you include in the recommendation?

- A. a Microsoft Azure Table Storage solution
- B. a Microsoft Azure HDInsight R Server cluster
- C. a Microsoft Azure HDInsight Hadoop cluster
- D. a Microsoft Azure SQL database that has In-Memory OLTP enabled

Correct Answer: C

You can use HDInsight to process streaming data that's received in real time from a variety of devices.

Internet of Things (IoT)

You can use HDInsight to build applications that extract critical insights from data. You can also use Azure Machine Learning on top of that to predict future trends for your business.

By combining enterprise-scale R analytics software with the power of Apache Hadoop and Apache Spark, Microsoft R Server for HDInsight gives you the scale and performance you need. Multi-threaded math libraries and transparent parallelization in R Server handle up to 1000x more data and up to 50x faster speeds than open-source R, which helps you to train more accurate models for better predictions.

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/apache-hadoop-introduction>

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

## Try our product !

**100%** Guaranteed Success

**100%** Money Back Guarantee

**365** Days Free Update

**Instant Download** After Purchase

**24x7** Customer Support

Average **99.9%** Success Rate

More than **800,000** Satisfied Customers Worldwide

Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

## Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p><b>One Year Free Update</b> Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p><b>Money Back Guarantee</b> To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p><b>Security &amp; Privacy</b> We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information &amp; peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.