



DP-100^{Q&As}

Designing and Implementing a Data Science Solution on Azure

Pass Microsoft DP-100 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/dp-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

DRAG DROP

You have an Azure Machine Learning workspace that contains a CPU-based compute cluster and an Azure Kubernetes Services (AKS) inference cluster. You create a tabular dataset containing data that you plan to use to create a classification model.

You need to use the Azure Machine Learning designer to create a web service through which client applications can consume the classification model by submitting new data and getting an immediate prediction as a response.

Which three 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

Create and run a batch inference pipeline on the compute cluster.

Deploy a real-time endpoint on the inference cluster.

Create and run a real-time inference pipeline on the compute cluster.

Create and run a training pipeline that prepares the data and trains a classification model on the compute cluster.

Use the automated ML user interface to train a classification model on the compute cluster.

Create and start a Compute Instance.

Answer Area



Correct Answer:

**Actions**

Create and run a batch inference pipeline on the compute cluster.

Deploy a real-time endpoint on the inference cluster.

Use the automated ML user interface to train a classification model on the compute cluster.

Answer Area

Create and start a Compute Instance.

Create and run a training pipeline that prepares the data and trains a classification model on the compute cluster.

Create and run a real-time inference pipeline on the compute cluster.

Step 1: Create and start a Compute Instance

To train and deploy models using Azure Machine Learning designer, you need compute on which to run the training process, test the model, and host the model in a deployed service.

There are four kinds of compute resource you can create:

Compute Instances: Development workstations that data scientists can use to work with data and models.

Compute Clusters: Scalable clusters of virtual machines for on-demand processing of experiment code.

Inference Clusters: Deployment targets for predictive services that use your trained models.

Attached Compute: Links to existing Azure compute resources, such as Virtual Machines or Azure Databricks clusters.

Step 2: Create and run a training pipeline..

After you've used data transformations to prepare the data, you can use it to train a machine learning model. Create and run a training pipeline

Step 3: Create and run a real-time inference pipeline

After creating and running a pipeline to train the model, you need a second pipeline that performs the same data transformations for new data, and then uses the trained model to inference (in other words, predict) label values based on its

features. This pipeline will form the basis for a predictive service that you can publish for applications to use.

Reference:

<https://docs.microsoft.com/en-us/learn/modules/create-classification-model-azure-machine-learning-designer/>

QUESTION 2**HOTSPOT**



Complete the sentence by selecting the correct option in the answer area.

Hot Area:

Answer Area

	▼
SSD	
FPGA	
GPU	
Power BI	

is required for a Deep Learning Virtual Machine (DLVM) to support Compute Unified Device Architecture (CUDA) computations.

Correct Answer:

Answer Area

	▼
SSD	
FPGA	
GPU	
Power BI	

is required for a Deep Learning Virtual Machine (DLVM) to support Compute Unified Device Architecture (CUDA) computations.

A Deep Learning Virtual Machine is a pre-configured environment for deep learning using GPU instances.

QUESTION 3

You create a workspace by using Azure Machine Learning Studio.

You must run a Python SDK v2 notebook in the workspace by using Azure Machine Learning Studio.

You need to reset the state of the notebook.

Which three actions should you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.



- A. Stop the current kernel.
- B. Change the compute.
- C. Reset the compute.
- D. Navigate to another section of the workspace.
- E. Change the current kernel.

Correct Answer: BCE

These actions will reset the notebook state and will reset all variables in the notebook. Action, Result

*

(B) Switch compute

Notebook automatically uses the new compute.

*

(C) Reset compute

Starts again when you try to run a cell

*

(E) Change the kernel Notebook uses new kernel

*

Stop compute No cells will run

*

Open notebook in Jupyter or JupyterLab Notebook opened in a new tab.

Incorrect:

These actions won't change the notebook state or the values of any variables in the notebook:

Action, Result

*

(not A) Stop the kernel

Stops any running cell. Running a cell will automatically restart the kernel.

*

(not D) Navigate to another workspace section

Running cells are stopped.

Reference:



<https://learn.microsoft.com/en-us/azure/machine-learning/how-to-run-jupyter-notebooks>

QUESTION 4

You are planning to register a trained model in an Azure Machine Learning workspace.

You must store additional metadata about the model in a key-value format. You must be able to add new metadata and modify or delete metadata after creation.

You need to register the model.

Which parameter should you use?

- A. description
- B. model_framework
- C. tags
- D. properties

Correct Answer: D

`azureml.core.Model.properties`:

Dictionary of key value properties for the Model. These properties cannot be changed after registration, however new key value pairs can be added.

Reference:

<https://docs.microsoft.com/en-us/python/api/azureml-core/azureml.core.model.model>

QUESTION 5

You train and register an Azure Machine Learning model

You plan to deploy the model to an online endpoint

You need to ensure that applications will be able to use the authentication method with a non-expiring artifact to access the model.

Solution: Create a managed online endpoint and set the value of its `auth.mode` parameter to `aml.token`.

Deploy the model to the online endpoint.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B



VCE & PDF

GeekCert.com

<https://www.geekcert.com/dp-100.html>

2024 Latest geekcert DP-100 PDF and VCE dumps Download

[DP-100 Practice Test](#)

[DP-100 Study Guide](#)

[DP-100 Exam Questions](#)