



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

You create an Azure Machine Learning pipeline named pipeline 1 with two steps that contain Python scripts. Data processed by the first step is passed to the second step.

You must update the content of the downstream data source of pipeline 1 and run the pipeline again.

You need to ensure the new run of pipeline 1 fully processes the updated content.

Solution: Change the value of the compute.target parameter of the PythonScriptStep object in the two steps.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

QUESTION 2

You manage an Azure Machine Learning workspace. The workspace includes an Azure Machine Learning Kubernetes compute target configured as an Azure Kubernetes Service (AKS) cluster named AKS1. AKS1 is configured to enable the

targeting of different nodes to train workloads.

You must run a command job on AKS1 by using the Azure ML Python SDK v2. The command job must select different types of compute nodes. The compute node types must be specified by using a command parameter.

You need to configure the command parameter.

Which parameter should you use?

A. environment

B. compute

C. limits

D. instance_type

Correct Answer: D

command

Create a Command object which can be used inside dsl.pipeline as a function and can also be created as a standalone command job.

Parameters include:

instance_type



Optional type of VM used as supported by the compute target.

Incorrect:

*

environment

the environment to use for this command

*

compute

the name of the compute where the command job is executed(will not be used if the command is used as a component/function)

Reference:

<https://learn.microsoft.com/en-us/python/api/azure-ai-ml/azure.ai.m>

QUESTION 3

HOTSPOT

You have an Azure Machine Learning workspace

You plan to use the Azure Machine Learning SDK for Python v1 to submit a job to run a training script.

You need to complete the script to ensure that it will execute the training script.

How should you complete the script? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point

Hot Area:



Answer Area

```
from azureml.core import Workspace, Environment, Experiment, ScriptRunConfig
```

```
ws = Workspace.from_config()
```

```
env = Environment.get(workspace=ws, name='AzureML-Minimal')
```

```
exp = Experiment(workspace=ws, name='experiment')
```

```
src = ScriptRunConfig(source_directory='./src',  
                      script='train.py',  
                      compute_target='compute-cluster',  
                      environment=env)
```

run = . (config=src)

A ▼
env
Placeholder

B ▼
submit
Placeholder

Correct Answer:



Answer Area

```
from azureml.core import Workspace, Environment, Experiment, ScriptRunConfig
```

```
ws = Workspace.from_config()
```

```
env = Environment.get(workspace=ws, name='AzureML-Minimal')
```

```
exp = Experiment(workspace=ws, name='experiment')
```

```
src = ScriptRunConfig(source_directory='./src',  
                      script='train.py',  
                      compute_target='compute-cluster',  
                      environment=env)
```

run = . (config=src)

A ▼
env
Placeholder

B ▼
submit
Placeholder

QUESTION 4

You are performing clustering by using the K-means algorithm.

You need to define the possible termination conditions.

Which three conditions can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Centroids do not change between iterations.
- B. The residual sum of squares (RSS) rises above a threshold.
- C. The residual sum of squares (RSS) falls below a threshold.
- D. A fixed number of iterations is executed.



E. The sum of distances between centroids reaches a maximum.

Correct Answer: ACD

AD: The algorithm terminates when the centroids stabilize or when a specified number of iterations are completed.

C: A measure of how well the centroids represent the members of their clusters is the residual sum of squares or RSS, the squared distance of each vector from its centroid summed over all vectors. RSS is the objective function and our goal is to minimize it.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/k-means-clustering>

<https://nlp.stanford.edu/IR-book/html/htmledition/k-means-1.html>

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 in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

An IT department creates the following Azure resource groups and resources:

Resource group	Resources
ml_resources	<ul style="list-style-type: none">• an Azure Machine Learning workspace named amlworkspace• an Azure Storage account named amlworkspace12345• an Application Insights instance named amlworkspace54321• an Azure Key Vault named amlworkspace67890• an Azure Container Registry named amlworkspace09876
general_compute	<p>A virtual machine named mlvm with the following configuration:</p> <ul style="list-style-type: none">• Operating system: Ubuntu Linux• Software installed: Python 3.6 and Jupyter Notebooks• Size: NC6 (6 vCPUs, 1 vGPU, 56 Gb RAM)

The IT department creates an Azure Kubernetes Service (AKS)-based inference compute target named aks-cluster in the Azure Machine Learning workspace.

You have a Microsoft Surface Book computer with a GPU. Python 3.6 and Visual Studio Code are installed.

You need to run a script that trains a deep neural network (DNN) model and logs the loss and accuracy metrics.

Solution: Install the Azure ML SDK on the Surface Book. Run Python code to connect to the workspace. Run the training script as an experiment on the aks-cluster compute target.

Does the solution meet the goal?

A. Yes



B. No

Correct Answer: B

Need to attach the mlvm virtual machine as a compute target in the Azure Machine Learning workspace.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/concept-compute-target>

[Latest DP-100 Dumps](#)

[DP-100 Practice Test](#)

[DP-100 Exam Questions](#)