



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

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.passapply.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 are developing deep learning models to analyze semi-structured, unstructured, and structured data types.

You have the following data available for model building:

1.  
Video recordings of sporting events
2.  
Transcripts of radio commentary about events
3.  
Logs from related social media feeds captured during sporting events

You need to select an environment for creating the model. Which environment should you use?

- A. Azure Cognitive Services
- B. Azure Data Lake Analytics
- C. Azure HDInsight with Spark MLlib
- D. Azure Machine Learning Studio

Correct Answer: A

Azure Cognitive Services expand on Microsoft's evolving portfolio of machine learning APIs and enable developers to easily add cognitive features such as emotion and video detection; facial, speech, and vision recognition; and speech and language understanding into their applications. The goal of Azure Cognitive Services is to help developers create applications that can see, hear, speak, understand, and even begin to reason. The catalog of services within Azure Cognitive Services can be categorized into five main pillars - Vision, Speech, Language, Search, and Knowledge.

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

---

### QUESTION 2

You manage an Azure Machine Learning workspace by using the Azure CLI ml extension v2.

You need to define a YAML schema to create a compute cluster.

Which schema should you use?

- A. <https://azuremlschemas.azureedge.net/latest/computdnstarKeichema.json>
- B. <https://azuremlschemas.azureedge.net/latest/amlCompute.schemajson>
- C. <https://azuremlschemas.azureedge.net/latest/vmCompute.schema.json>
- D. <https://azuremlschemas.azureedge.net/latest/kubernetesCompute.schema.json>



Correct Answer: B

### QUESTION 3

You are developing a hands-on workshop to introduce Docker for Windows to attendees.

You need to ensure that workshop attendees can install Docker on their devices.

Which two prerequisite components should attendees install on the devices? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Microsoft Hardware-Assisted Virtualization Detection Tool
- B. Kitematic
- C. BIOS-enabled virtualization
- D. VirtualBox
- E. Windows 10 64-bit Professional

Correct Answer: CE

C: Make sure your Windows system supports Hardware Virtualization Technology and that virtualization is enabled. Ensure that hardware virtualization support is turned on in the BIOS settings. For example:



E: To run Docker, your machine must have a 64-bit operating system running Windows 7 or higher.

References:



[https://docs.docker.com/toolbox/toolbox\\_install\\_windows/](https://docs.docker.com/toolbox/toolbox_install_windows/)

<https://blogs.technet.microsoft.com/canitpro/2015/09/08/step-by-step-enabling-hyper-v-for-use-on-windows-10/>

#### QUESTION 4

##### HOTSPOT

You create an Azure Machine Learning compute target named ComputeOne by using the STANDARD\_D1 virtual machine image.

You define a Python variable named `ws` that references the Azure Machine Learning workspace. You run the following Python code:

```
from azureml.core.compute import ComputeTarget, AmlCompute
from azureml.core.compute_target import ComputeTargetException
the_cluster_name = "ComputeOne"
try:
    the_cluster = ComputeTarget(workspace=ws, name=the_cluster_name)
    print('Step1')
except ComputeTargetException:
    config = AmlCompute.provisioning_configuration(vm_size='STANDARD_DS12_v2', max_nodes=4)
    the_cluster = ComputeTarget.create(ws, the_cluster_name, config)
    print('Step2')
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

	Yes	No
A new machine learning compute resource is created with a virtual machine size of STANDARD_DS12_v2 and a maximum of four nodes.	<input type="checkbox"/>	<input type="checkbox"/>
Any experiments configured to use <code>the_cluster</code> will run on ComputeOne.	<input type="checkbox"/>	<input type="checkbox"/>
The text <b>Step1</b> will be printed to the screen.	<input type="checkbox"/>	<input type="checkbox"/>

Correct Answer:



**Yes** **No**

A new machine learning compute resource is created with a virtual machine size of STANDARD\_DS12\_v2 and a maximum of four nodes.

Any experiments configured to use `the_cluster` will run on ComputeOne.

The text **Step1** will be printed to the screen.

Box 1: Yes

ComputeTargetException class: An exception related to failures when creating, interacting with, or configuring a compute target. This exception is commonly raised for failures attaching a compute target, missing headers, and unsupported

configuration values.

Create(workspace, name, provisioning\_configuration)

Provision a Compute object by specifying a compute type and related configuration.

This method creates a new compute target rather than attaching an existing one.

Box 2: Yes

Box 3: No

The line before `print('\\Step1\\')` will fail.

Reference:

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

## QUESTION 5

### HOTSPOT

You create an Azure Machine Learning workspace named workspace1. You assign a custom role to a user of workspace1.

The custom role has the following JSON definition:



```
{  
  "Name": "MyRole",  
  "IsCustom": true,  
  "Description": "New custom role description.",  
  "Actions": ["*"],  
  "NotActions": [  
    "Microsoft.MachineLearningServices/workspaces/write",  
    "Microsoft.MachineLearningServices/workspaces/computes/*/write",  
    "Microsoft.MachineLearningServices/workspaces/computes/*/delete",  
    "Microsoft.Authorization/*/write"  
  ],  
  "AssignableScopes": [  
    "/subscriptions/<subscription_id>/resourceGroups/resourcegroup1/providers/  
    Microsoft.MachineLearningServices/workspaces/workspace1"  
  ]  
}
```

Instructions: For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
The user can perform all actions in the workspace	<input type="radio"/>	<input type="radio"/>
The user can delete a compute resource in the workspace	<input type="radio"/>	<input type="radio"/>
The user can write metrics to the workspace	<input type="radio"/>	<input type="radio"/>

Correct Answer:





## Answer Area

Statements	Yes	No
The user can perform all actions in the workspace	<input type="radio"/>	<input checked="" type="radio"/>
The user can delete a compute resource in the workspace	<input type="radio"/>	<input checked="" type="radio"/>
The user can write metrics to the workspace	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: No

The actions listed in NotActions are prohibited.

If the roles include Actions that have a wildcard (\*), the effective permissions are computed by subtracting the NotActions from the allowed Actions.

Box 2: No

Deleting compute resources in the workspace is in the NotActions list.

Box 3: Yes

Writing metrics is not listed in NotActions.

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/overview#how-azure-rbac-determines-if-a-user-has-access-to-a-resource>

### QUESTION 6

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.

You are creating a model to predict the price of a student's artwork depending on the following variables:

the student's length of education, degree type, and art form.



You start by creating a linear regression model.

You need to evaluate the linear regression model.

Solution: Use the following metrics: Relative Squared Error, Coefficient of Determination, Accuracy, Precision, Recall, F1 score, and AUC.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

Relative Squared Error, Coefficient of Determination are good metrics to evaluate the linear regression model, but the others are metrics for classification models.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/evaluate-model>

### QUESTION 7

You have the following Azure subscriptions and Azure Machine Learning service workspaces:

Subscription	Workspace	Comment
385bdf5e-4cef-4ad4-b977-3f86d92727c9	ml-default	This is the default subscription.
5a5891d1-557a-4234-9b83-2e90412b1068	ml-project	The information required to uniquely identify this workspace is stored in the file config.json in the same folder as the Python script.

You need to obtain a reference to the ml-project workspace. Solution: Run the following Python code:

```
from azureml.core import Workspace
ws = Workspace(workspace_name="ml-project")
```

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

### QUESTION 8

#### 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	B
env	submit
Placeholder	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)

<b>A</b> ▼	<b>B</b> ▼
env	submit
Placeholder	Placeholder

### QUESTION 9

You need to select a pre built development environment for a series of data science experiments. You must use the R language for the experiments.

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

NOTE: Each correct selection is worth one point.

- A. MI.NET Library on a local environment
- B. Azure Machine Learning Studio
- C. Data Science Virtual Machine (OSVM)
- D. Azure Data bricks
- E. Azure Cognitive Services



Correct Answer: ABD

**QUESTION 10**

**HOTSPOT**

You are developing a deep learning model by using TensorFlow. You plan to run the model training workload on an Azure Machine Learning Compute Instance.

You must use CUDA-based model training.

You need to provision the Compute Instance.

Which two virtual machines sizes can you use? To answer, select the appropriate virtual machine sizes in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Virtual machine size

Name ↑	vCPUs	GPUs	RAM	Resource disk
BASIC_A0	1		0.75 GB	20 GB
STANDARD_D3_V2	4		14 GB	200 GB
STANDARD_E64_V3	64		432 GB	1,600 GB
STANDARD_M64LS	64		512 GB	2,000 GB
STANDARD_NC12	12	2	112 GB	680 GB
STANDARD_NC24	24	4	224 GB	1,440 GB

Correct Answer:



## Answer Area

### Virtual machine size

Name ↑	vCPUs	GPUs	RAM	Resource disk
BASIC_A0	1		0.75 GB	20 GB
STANDARD_D3_V2	4		14 GB	200 GB
STANDARD_E64_V3	64		432 GB	1,600 GB
STANDARD_M64LS	64		512 GB	2,000 GB
STANDARD_NC12	12	2	112 GB	680 GB
STANDARD_NC24	24	4	224 GB	1,440 GB

CUDA is a parallel computing platform and programming model developed by Nvidia for general computing on its own GPUs (graphics processing units). CUDA enables developers to speed up compute-intensive applications by harnessing the power of GPUs for the parallelizable part of the computation.

Reference: <https://www.infoworld.com/article/3299703/what-is-cuda-parallel-programming-for-gpus.html>

## QUESTION 11

You train and register a machine learning model. You create a batch inference pipeline that uses the model to generate predictions from multiple data files.

You must publish the batch inference pipeline as a service that can be scheduled to run every night.

You need to select an appropriate compute target for the inference service.

Which compute target should you use?

- A. Azure Machine Learning compute instance
- B. Azure Machine Learning compute cluster
- C. Azure Kubernetes Service (AKS)-based inference cluster
- D. Azure Container Instance (ACI) compute target

Correct Answer: B

Azure Machine Learning compute clusters is used for Batch inference. Run batch scoring on serverless compute. Supports normal and low-priority VMs. No support for real-time inference.

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

## QUESTION 12

You use the Azure Machine Learning SDK for Python v1 and notebooks to train a model. You create a compute target,



an environment, and a training script by using Python code.

You need to prepare information to submit a training run.

Which class should you use?

- A. ScriptRun
- B. ScriptRunConfig
- C. RunConfiguration
- D. Run

Correct Answer: B

A ScriptRunConfig is used to configure the information necessary for submitting a training job as part of an experiment.

Reference: <https://learn.microsoft.com/en-us/azure/machine-learning/v1/how-to-set-up-training-targets>

## QUESTION 13

### HOTSPOT

You create an Azure Machine learning workspace and load a Python training script named tram.py in the src subfolder. The dataset used to train your model is available locally. You run the following script to tram the model:

```
ws = Workspace.from_config()
experiment = Experiment(workspace=ws, name='nlp-experiment-train')

cpu_cluster_name = "cpu-cluster"
try:
    cpu_cluster = ComputeTarget(workspace=ws, name=cpu_cluster_name)
except ComputeTargetException:
    compute_config = AmlCompute.provisioning_configuration(vm_size='STANDARD_D2_V2', max_nodes=4)
    cpu_cluster = ComputeTarget.create(ws, cpu_cluster_name, compute_config)

cpu_cluster.wait_for_completion(show_output=True)

config = ScriptRunConfig(source_directory='./src',
                        script='train.py',
                        compute_target='cpu-cluster')

env = Environment.from_conda_specification(
    name='pytorch-env',
    file_path='./azureml/pytorch-env.yml'
)
config.run_config.environment = env

run = experiment.submit(config)
```



instructions: For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point

Hot Area:

Statements	Yes	No
The script will use local compute resources and a new Azure Machine Learning compute will be created upon failure.	<input type="radio"/>	<input type="radio"/>
The dataset used during the training phase is automatically loaded in a new datastore.	<input type="radio"/>	<input type="radio"/>
A new environment object is created from a local Conda environment.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Statements	Yes	No
The script will use local compute resources and a new Azure Machine Learning compute will be created upon failure.	<input type="radio"/>	<input checked="" type="radio"/>
The dataset used during the training phase is automatically loaded in a new datastore.	<input type="radio"/>	<input checked="" type="radio"/>
A new environment object is created from a local Conda environment.	<input type="radio"/>	<input checked="" type="radio"/>

#### QUESTION 14

You plan to use automated machine learning by using Azure Machine Learning Python SDK v2 to train a regression model. You have data that has features with missing values, and categorical features with few distinct values.

You need to control whether automated machine learning automatically imputes missing values and encode categorical features as part of the training task.

Which enum of the automl package should you use?

A. ForecastHorizonMode





- B. RegressionModels
- C. FeaturizationMode
- D. RegressionPrimaryMetrics

Correct Answer: C

AutoMLConfig Class

Represents configuration for submitting an automated ML experiment in Azure Machine Learning.

This configuration object contains and persists the parameters for configuring the experiment run, as well as the training data to be used at run time.

Constructor

Featurization str or FeaturizationConfig

\\'auto\\' / \\'off\\' / FeaturizationConfig Indicator for whether featurization step should be done automatically or not, or whether customized featurization should be used.

Column type is automatically detected. Based on the detected column type preprocessing/featurization is done as follows:

Categorical: Target encoding, one hot encoding, drop high cardinality categories, impute missing values.

\*-> Numeric: Impute missing values, cluster distance, weight of evidence.

DateTime: Several features such as day, seconds, minutes, hours etc.

Text: Bag of words, pre-trained Word embedding, text target encoding.

Reference:

<https://learn.microsoft.com/en-us/python/api/azureml-train-automl-client/azureml.train.automl.automlconfig.automlconfig>  
<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-use-automated-ml-for-ml-models>

---

## QUESTION 15

DRAG DROP

You create an Azure Machine Learning workspace.

You must implement dedicated compute for model training in the workspace by using Azure Synapse compute resources. The solution must attach the dedicated compute and start an Azure Synapse session.

You need to implement the compute resources.

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 compute clusters by using Azure Machine Learning studio.

Create a linked service by using Azure Synapse studio.

Create a linked service by using Azure Machine Learning studio.

Create an Azure Synapse workspace by using the Azure portal.

Create an Apache Spark pool by using the Azure portal.

## Answer area

Correct Answer:



### Actions

Create compute clusters by using Azure Machine Learning studio.

Create a linked service by using Azure Synapse studio.

### Answer area

Create a linked service by using Azure Machine Learning studio.

Create an Azure Synapse workspace by using the Azure portal.

Create an Apache Spark pool by using the Azure portal.