



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.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

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 2

DRAG DROP

You have an Azure Machine Learning workspace that contains a training cluster and an inference cluster.



You plan to create a classification model by using the Azure Machine Learning designer.

You need to ensure that client applications can submit data as HTTP requests and receive predictions as responses.

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 a real-time inference pipeline and run the pipeline on the compute cluster.
- Create a batch inference pipeline and run the pipeline on the compute cluster.
- Deploy a service to the compute cluster.
- Create a pipeline that trains a classification model and run the pipeline on the compute cluster.
- Deploy a service to the inference cluster.

Answer area

Correct Answer:

Actions

- Create a real-time inference pipeline and run the pipeline on the compute cluster.
- Create a batch inference pipeline and run the pipeline on the compute cluster.
- Deploy a service to the compute cluster.
- Create a pipeline that trains a classification model and run the pipeline on the compute cluster.
- Deploy a service to the inference cluster.

Answer area

- Create a pipeline that trains a classification model and run the pipeline on the compute cluster.
- Create a batch inference pipeline and run the pipeline on the compute cluster.
- Deploy a service to the inference cluster.

QUESTION 3

You register a model that you plan to use in a batch inference pipeline.

The batch inference pipeline must use a ParallelRunStep step to process files in a file dataset. The script has the ParallelRunStep step runs must process six input files each time the inferencing function is called.

You need to configure the pipeline.

Which configuration setting should you specify in the ParallelRunConfig object for the ParallelRunStep step?

- A. process_count_per_node= "6"
- B. node_count= "6"
- C. mini_batch_size= "6"
- D. error_threshold= "6"

Correct Answer: B



node_count is the number of nodes in the compute target used for running the ParallelRunStep.

Incorrect Answers:

A: process_count_per_node

Number of processes executed on each node. (optional, default value is number of cores on node.)

C: mini_batch_size

For FileDataset input, this field is the number of files user script can process in one run() call. For TabularDataset input, this field is the approximate size of data the user script can process in one run() call.

Example values are 1024, 1024KB, 10MB, and 1GB.

D: error_threshold

The number of record failures for TabularDataset and file failures for FileDataset that should be ignored during processing. If the error count goes above this value, then the job will be aborted.

Reference:

<https://docs.microsoft.com/en-us/python/api/azureml-contrib-pipeline-steps/azureml.contrib.pipeline.steps.parallelrunconfig?view=azure-ml-py>

QUESTION 4

You create an Azure Machine Learning workspace. You use Azure Machine Learning designer to create a pipeline within the workspace.

You need to submit a pipeline run from the designer.

What should you do first?

- A. Create an experiment.
- B. Create an attached compute resource.
- C. Create a compute cluster.
- D. Select a model.

Correct Answer: B

Create a new workspace (already done) Create the pipeline

1.

Sign in to ml.azure.com, and select the workspace you want to work with.

2.

Select Designer -> Classic prebuilt

3.



Select Create a new pipeline using classic prebuilt components.

4.

Click the pencil icon beside the automatically generated pipeline draft name, rename it to Automobile price prediction. The name doesn't need to be unique.

Set the default compute target

A pipeline jobs on a compute target, which is a compute resource that's attached to your workspace. After you create a compute target, you can reuse it for future jobs.

Note: Create a new workspace

You need an Azure Machine Learning workspace to use the designer. The workspace is the top-level resource for Azure Machine Learning, it provides a centralized place to work with all the artifacts you create in Azure Machine Learning.

Reference: <https://learn.microsoft.com/en-us/azure/machine-learning/tutorial-designer-automobile-price-train-score>

QUESTION 5

DRAG DROP You previously deployed a model that was trained using a tabular dataset named training-dataset, which is based on a folder of CSV files. Over time, you have collected the features and predicted labels generated by the model in a folder containing a CSV file for each month. You have created two tabular datasets based on the folder containing the inference data: one named

predictions-dataset with a schema that matches the training data exactly, including the predicted label; and another named features-dataset with a schema containing all of the feature columns and a timestamp column based on the filename, which includes the day, month, and year.

You need to create a data drift monitor to identify any changing trends in the feature data since the model was trained. To accomplish this, you must define the required datasets for the data drift monitor.

Which datasets should you use to configure the data drift monitor? To answer, drag the appropriate datasets to the correct data drift monitor options. Each source may be used once, more than once, or not at all. You may need to drag the

split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point. Select and Place:



Target datasets

training-dataset

predictions-dataset

features-dataset

Answer Area

Baseline dataset

Target dataset

Target dataset

Target dataset

Correct Answer:

Target datasets

features-dataset

Answer Area

Baseline dataset

training-dataset

Target dataset

predictions-dataset

Box 1: training-dataset

Baseline dataset - usually the training dataset for a model.

Box 2: predictions-dataset

Target dataset - usually model input data - is compared over time to your baseline dataset. This comparison means that your target dataset must have a timestamp column specified. The monitor will compare the baseline and target datasets.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/how-to-monitor-datasets>



VCE & PDF

PassApply.com

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

2024 Latest passapply DP-100 PDF and VCE dumps Download

[DP-100 Practice Test](#)

[DP-100 Study Guide](#)

[DP-100 Exam Questions](#)