

➤ **Vendor: Microsoft**

➤ **Exam Code: DP-100**

➤ **Exam Name: Designing and Implementing a Data Science Solution on Azure**

➤ **New Updated Questions from Braindump2go (Updated in January/2022)**

Visit Braindump2go and Download Full Version DP-100 Exam Dumps

QUESTION 299

You use Azure Machine Learning to train a model based on a dataset named dataset1. You define a dataset monitor and create a dataset named dataset2 that contains new data. You need to compare dataset1 and dataset2 by using the Azure Machine Learning SDK for Python. Which method of the DataDriftDetector class should you use?

- A. run
- B. get
- C. backfill
- D. update

Answer: C

Explanation:

A backfill run is used to see how data changes over time.

Reference:

<https://docs.microsoft.com/en-us/python/api/azureml-datadrift/azureml.datadrift.datadriftdetector.datadriftdetector>

QUESTION 300

You use an Azure Machine Learning workspace.

You have a trained model that must be deployed as a web service. Users must authenticate by using Azure Active Directory.

What should you do?

- A. Deploy the model to Azure Kubernetes Service (AKS). During deployment, set the `token_auth_enabled` parameter of the target configuration object to true
- B. Deploy the model to Azure Container Instances. During deployment, set the `auth_enabled` parameter of the target configuration object to true
- C. Deploy the model to Azure Container Instances. During deployment, set the `token_auth_enabled` parameter of the target configuration object to true
- D. Deploy the model to Azure Kubernetes Service (AKS). During deployment, set the `auth_enabled` parameter of the target configuration object to true

Answer: A

Explanation:

To control token authentication, use the `token_auth_enabled` parameter when you create or update a deployment. Token authentication is disabled by default when you deploy to Azure Kubernetes Service.

Note: The model deployments created by Azure Machine Learning can be configured to use one of two authentication methods:

key-based: A static key is used to authenticate to the web service.

token-based: A temporary token must be obtained from the Azure Machine Learning workspace (using Azure Active

[DP-100 Exam Dumps](#) [DP-100 Exam Questions](#) [DP-100 PDF Dumps](#) [DP-100 VCE Dumps](#)

<https://www.braindump2go.com/dp-100.html>

Directory) and used to authenticate to the web service.

Incorrect Answers:

C: Token authentication isn't supported when you deploy to Azure Container Instances.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-authenticate-web-service>

QUESTION 301

You have a Jupyter Notebook that contains Python code that is used to train a model.

You must create a Python script for the production deployment. The solution must minimize code maintenance.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Refactor the Jupyter Notebook code into functions
- B. Save each function to a separate Python file
- C. Define a main() function in the Python script
- D. Remove all comments and functions from the Python script

Answer: AC

Explanation:

C: Python main function is a starting point of any program. When the program is run, the python interpreter runs the code sequentially. Main function is executed only when it is run as a Python program.

A: Refactoring, code style and testing

The first step is to modularise the notebook into a reasonable folder structure, this effectively means to convert files from .ipynb format to .py format, ensure each script has a clear distinct purpose and organise these files in a coherent way.

Once the project is nicely structured we can tidy up or refactor the code.

Reference:

<https://www.guru99.com/learn-python-main-function-with-examples-understand-main.html>

<https://towardsdatascience.com/from-jupyter-notebook-to-deployment-a-straightforward-example-1838c203a437>

QUESTION 302

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

Answer: B

Explanation:

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 303

You use the Azure Machine Learning designer to create and run a training pipeline.

The pipeline must be run every night to inference predictions from a large volume of files. The folder where the files will be stored is defined as a dataset.

You need to publish the pipeline as a REST service that can be used for the nightly inferencing run.

What should you do?

- A. Create a batch inference pipeline

[DP-100 Exam Dumps](#) **[DP-100 Exam Questions](#)** **[DP-100 PDF Dumps](#)** **[DP-100 VCE Dumps](#)**

<https://www.braindump2go.com/dp-100.html>

- B. Set the compute target for the pipeline to an inference cluster
- C. Create a real-time inference pipeline
- D. Clone the pipeline

Answer: A

Explanation:

Azure Machine Learning Batch Inference targets large inference jobs that are not time-sensitive. Batch Inference provides cost-effective inference compute scaling, with unparalleled throughput for asynchronous applications. It is optimized for high-throughput, fire-and-forget inference over large collections of data.

You can submit a batch inference job by pipeline_run, or through REST calls with a published pipeline.

Reference:

<https://github.com/Azure/MachineLearningNotebooks/blob/master/how-to-use-azureml/machine-learning-pipelines/parallel-run/README.md>

QUESTION 304

You create a binary classification model. The model is registered in an Azure Machine Learning workspace. You use the Azure Machine Learning Fairness SDK to assess the model fairness.

You develop a training script for the model on a local machine.

You need to load the model fairness metrics into Azure Machine Learning studio.

What should you do?

- A. Implement the `download_dashboard_by_upload_id` function
- B. Implement the `create_group_metric_set` function
- C. Implement the `upload_dashboard_dictionary` function
- D. Upload the training script

Answer: C

Explanation:

import azureml.contrib.fairness package to perform the upload:

from azureml.contrib.fairness import upload_dashboard_dictionary, download_dashboard_by_upload_id

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-machine-learning-fairness-aml>

QUESTION 305

You have a dataset that includes confidential data. You use the dataset to train a model.

You must use a differential privacy parameter to keep the data of individuals safe and private.

You need to reduce the effect of user data on aggregated results.

What should you do?

- A. Decrease the value of the epsilon parameter to reduce the amount of noise added to the data
- B. Increase the value of the epsilon parameter to decrease privacy and increase accuracy
- C. Decrease the value of the epsilon parameter to increase privacy and reduce accuracy
- D. Set the value of the epsilon parameter to 1 to ensure maximum privacy

Answer: C

Explanation:

Differential privacy tries to protect against the possibility that a user can produce an indefinite number of reports to eventually reveal sensitive data. A value known as epsilon measures how noisy, or private, a report is. Epsilon has an inverse relationship to noise or privacy. The lower the epsilon, the more noisy (and private) the data is.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/concept-differential-privacy>

QUESTION 306

Hotspot Question

You are using an Azure Machine Learning workspace. You set up an environment for model testing and an environment for production.

The compute target for testing must minimize cost and deployment efforts. The compute target for production must provide fast response time, autoscaling of the deployed service, and support real-time inferencing.

You need to configure compute targets for model testing and production.

Which compute targets should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Environment

Compute target

Testing

	▼
Local web service	
Azure Kubernetes Services (AKS)	
Azure Container Instances	
Azure Machine Learning compute clusters	

Production

	▼
Local web service	
Azure Kubernetes Services (AKS)	
Azure Container Instances	
Azure Machine Learning compute clusters	

Answer:

Answer Area

Environment

Compute target

Testing

	▼
Local web service	
Azure Kubernetes Services (AKS)	
Azure Container Instances	
Azure Machine Learning compute clusters	

Production

	▼
Local web service	
Azure Kubernetes Services (AKS)	
Azure Container Instances	
Azure Machine Learning compute clusters	

Explanation:

Box 1: Local web service

The Local web service compute target is used for testing/debugging. Use it for limited testing and troubleshooting. Hardware acceleration depends on use of libraries in the local system.

Box 2: Azure Kubernetes Service (AKS)

Azure Kubernetes Service (AKS) is used for Real-time inference.

Recommended for production workloads.

Use it for high-scale production deployments. Provides fast response time and autoscaling of the deployed service

Reference:

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

QUESTION 307

Drag and Drop Question

You are using a Git repository to track work in an Azure Machine Learning workspace.

You need to authenticate a Git account by using SSH.

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.

Actions

Generate a public/private key pair
Add the private key to the Git account
Clone the Git repository by using an SSH repository URL
Add the public key to the Git account
Create a new Azure Key Vault resource

Answer Area



Answer:

[DP-100 Exam Dumps](#) [DP-100 Exam Questions](#) [DP-100 PDF Dumps](#) [DP-100 VCE Dumps](#)

<https://www.braindump2go.com/dp-100.html>

Actions

Add the private key to the Git account

Create a new Azure Key Vault resource

Answer Area

Generate a public/private key pair

Add the public key to the Git account

Clone the Git repository by using an SSH repository URL

**Explanation:**

Authenticate your Git Account with SSH:

Step 1: Generating a public/private key pair

Generate a new SSH key

1. Open the terminal window in the Azure Machine Learning Notebook Tab.

2. Paste the text below, substituting in your email address.

```
ssh-keygen -t rsa -b 4096 -C "your_email@example.com"
```

This creates a new ssh key, using the provided email as a label.

> Generating public/private rsa key pair.

Step 2: Add the public key to the Git Account

In your terminal window, copy the contents of your public key file.

Step 3: Clone the Git repository by using an SSH repository URL

1. Copy the SSH Git clone URL from the Git repo.

2. Paste the url into the git clone command below, to use your SSH Git repo URL. This will look something like:

```
git clone git@example.com:GitUser/azureml-example.git
```

Cloning into 'azureml-example'.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/concept-train-model-git-integration>