![QuickStart - Over 35 Years Of Technology Training]

Document Generated: 09/19/2024

Learning Style: Virtual Classroom

Provider: Microsoft

Difficulty: Intermediate

Course Duration: 4 Days

Azure Data Scientist Certification: Designing and Implementing a Data Science solution on Azure (DP-100)

If you enroll in this course at the listed price, you receive a Free Official Exam Voucher for the DP-100 Exam. This course does not include Exam Voucher if enrolled within the Master Subscription, however, you can request to

purchase the Official Exam Voucher separately.

About this Course:

In today's world where every industry relies heavily on a strong tech department and where efficient management of data is a must for all organizations, having an Azure data engineering certification is a sure way to ensure a well paying job at any firm; whether big or small. An average Azure Data Scientist earns around **$128,627** per annum which is a great wage in view of the current economy.

This course is designed to train participants how to use Azure services. Azure services can be used to create, train and deploy machine learning solutions. This course includes an overview of data science services available on Azure. It also includes a detailed insight into Azure Machine learning service which is the major data science service by Azure . The students will be taught how to use Azure Machine learning service to automate data.

This course deals specifically with Azure and does not guarantee a training in the basis of data science in general. A prior assumption is that students signing up for this course know this beforehand.

Course Objective:

Upon completion of this course, the participant should have an advanced skill set and a sound working knowledge of the following principals while also be able to;

- Learn how to practice data science on Azure
- Learn how to practice data science on Azure using Azure Machine learning service
- Gain the skills needed to automate Machine learning with Azure Machine learning service

Audience:

This particular course is aimed at the following audience;

- Data scientists
- Those tasked with the role of training and deploying machine learning models

Prerequisites:

The following prerequisites are absolutely necessary to be eligible to take this course;

- Basic understanding of Azure fundamentals
- Basic knowledge of data sciences including the skills needed to prepare data, train and evaluate models ultimately selecting the most efficient one
- Know programming using Python programming language
- Basic knowledge of using python libraries namely; Pandas, Sickit-learn,

matplotlib and seaborn

Course Outline:

Module 1: Getting Started with Azure Machine Learning

In this module, you will learn how to provision an Azure Machine Learning workspace and use it to manage machine learning assets such as data, compute, model training code, logged metrics, and trained models. You will learn how to use the web-based Azure Machine Learning studio interface as well as the Azure Machine Learning SDK and developer tools like Visual Studio Code and Jupyter Notebooks to work with the assets in your workspace.

Lessons

- Introduction to Azure Machine Learning
- Working with Azure Machine Learning

Lab : Create an Azure Machine Learning Workspace

After completing this module, you will be able to

- Provision an Azure Machine Learning workspace
- Use tools and code to work with Azure Machine Learning

Module 2: Visual Tools for Machine Learning

This module introduces the Automated Machine Learning and Designer visual tools, which you can use to train, evaluate, and deploy machine learning models without writing any code.

Lessons

- Automated Machine Learning
- Azure Machine Learning Designer

Lab : Use Automated Machine Learning

Lab : Use Azure Machine Learning Designer

After completing this module, you will be able to

- Use automated machine learning to train a machine learning model
- Use Azure Machine Learning designer to train a model

Module 3: Running Experiments and Training Models

In this module, you will get started with experiments that encapsulate data processing and model training code, and use them to train machine learning models.

Lessons

- Introduction to Experiments
- Training and Registering Models

Lab : Train Models

Lab : Run Experiments

After completing this module, you will be able to

- Run code-based experiments in an Azure Machine Learning workspace
- Train and register machine learning models

Module 4: Working with Data

Data is a fundamental element in any machine learning workload, so in this module, you will learn how to create and manage datastores and datasets in an Azure Machine Learning workspace, and how to use them in model training experiments.

Lessons

- Working with Datastores
- Working with Datasets

Lab : Work with Data

After completing this module, you will be able to

- Create and use datastores
- Create and use datasets

Module 5: Working with Compute

One of the key benefits of the cloud is the ability to leverage compute resources on demand, and use them to scale machine learning processes to an extent that would be infeasible on your own hardware. In this module, you'll learn how to manage experiment environments that ensure consistent runtime consistency for experiments, and how to create and use compute targets for experiment runs.

Lessons

- Working with Environments
- Working with Compute Targets

Lab : Work with Compute

After completing this module, you will be able to

- Create and use environments

- Create and use compute targets

Module 6: Orchestrating Operations with Pipelines

Now that you understand the basics of running workloads as experiments that leverage data assets and compute resources, it's time to learn how to orchestrate these workloads as pipelines of connected steps. Pipelines are key to implementing an effective Machine Learning Operationalization (ML Ops) solution in Azure, so you'll explore how to define and run them in this module.

Lessons

- Introduction to Pipelines
- Publishing and Running Pipelines

Lab : Create a Pipeline

After completing this module, you will be able to

- Create pipelines to automate machine learning workflows
- Publish and run pipeline services

Module 7: Deploying and Consuming Models

Models are designed to help decision making through predictions, so they're only useful when deployed and available for an application to consume. In this module learn how to deploy models for real-time inferencing, and for batch inferencing.

Lessons

- Real-time Inferencing
- Batch Inferencing
- Continuous Integration and Delivery

Lab : Create a Real-time Inferencing Service

Lab : Create a Batch Inferencing Service

After completing this module, you will be able to

- Publish a model as a real-time inference service
- Publish a model as a batch inference service
- Describe techniques to implement continuous integration and delivery

Module 8: Training Optimal Models

By this stage of the course, you've learned the end-to-end process for training, deploying, and consuming machine learning models; but how do you ensure your model produces the best predictive outputs for your data? In this module, you'll explore how you can use hyperparameter tuning and automated machine learning

to take advantage of cloud-scale compute and find the best model for your data.

Lessons

- Hyperparameter Tuning
- Automated Machine Learning

Lab : Use Automated Machine Learning from the SDK

Lab : Tune Hyperparameters

After completing this module, you will be able to

- Optimize hyperparameters for model training
- Use automated machine learning to find the optimal model for your data

Module 9: Responsible Machine Learning

Data scientists have a duty to ensure they analyze data and train machine learning models responsibly; respecting individual privacy, mitigating bias, and ensuring transparency. This module explores some considerations and techniques for applying responsible machine learning principles.

Lessons

- Differential Privacy
- Model Interpretability
- Fairness

Lab : Explore Differential provacy

Lab : Interpret Models

Lab : Detect and Mitigate Unfairness

After completing this module, you will be able to

- Apply differential provacy to data analysis
- Use explainers to interpret machine learning models
- Evaluate models for fairness

Module 10: Monitoring Models

After a model has been deployed, it's important to understand how the model is being used in production, and to detect any degradation in its effectiveness due to data drift. This module describes techniques for monitoring models and their data.

Lessons

- Monitoring Models with Application Insights

- Monitoring Data Drift

Lab : Monitor Data Drift

Lab : Monitor a Model with Application Insights

After completing this module, you will be able to

- Use Application Insights to monitor a published model
- Monitor data drift

Credly Badge:

 Display your Completion Badge And Get The Recognition You Deserve.

Add a completion and readiness badge to your Linkedin profile, Facebook page, or Twitter account to validate your professional and technical expertise. With badges issued and validated by Credly, you can:

- Let anyone verify your completion and achievement by clicking on the badge
- Display your hard work and validate your expertise
- Display each badge's details about specific skills you developed.

Badges are issued by QuickStart and verified through Credly.

Find Out More or See List Of Badges