Designing and Implementing a Data Science Solution on Azure (DP-100)

Description

In this course, **Designing and Implementing a Data Science Solution on Azure (DP-100)**, learn to master machine learning solutions on the cloud with Azure Machine Learning. Our training guides you through **data ingestion, model training, deployment, and monitoring of solutions within the Microsoft Azure ecosystem**.

Perfectly designed for those with existing knowledge of Python and machine learning, this training will enable you to optimize your projects using the best available tools and techniques. Discover how to transform your data into decisions with a clear strategy and advanced skills in **Azure Machine Learning**.

Course Content
Module 1: Design a data ingestion strategy for machine learning projects

- Identify your data source and format
- Choose how to serve data to machine learning workflows
- Design a data ingestion solution

Module 2: Design a machine learning model training solution

- Identify machine learning tasks
- Choose a service to train a machine learning model
- Decide between compute options

Module 3: Design a model deployment solution

- Understand how model will be consumed
- Decide on real-time or batch deployment

Module 4: Explore Azure Machine Learning workspace resources and assets

- Create an Azure Machine Learning workspace
- Identify Azure Machine Learning resources
- Identify Azure Machine Learning assets

- Train models in the workspace

Module 5: Explore developer tools for workspace interaction

- Explore the studio
- Explore the Python SDK
- Explore the CLI

Module 6: Make data available in Azure Machine Learning

- Understand URIs
- Create a datastore
- Create a data asset

Module 7: Work with compute targets in Azure Machine Learning

- Choose the appropriate compute target
- Create and use a compute instance
- Create and use a compute cluster

Module 8: Work with environments in Azure Machine Learning

- Understand environments
- Explore and use curated environments
- Create and use custom environments

Module 9: Find the best classification model with Automated Machine Learning

- Preprocess data and configure featurization
- Run an Automated Machine Learning experiment
- Evaluate and compare models

Module 10: Track model training in Jupyter notebooks with MLflow

- Configure MLflow for model tracking in notebooks
- Train and track models in notebooks

Module 11: Run a training script as a command job in Azure Machine Learning

- Convert a notebook to a script
- Run a script as a command job
- Use parameters in a command job

Module 12: Track model training with MLflow in jobs

- Track metrics with MLflow
- View metrics and evaluate models

Module 13: Run pipelines in Azure Machine Learning

- Create components
- Create a pipeline
- Run a pipeline job

Module 14: Perform hyperparameter tuning with Azure Machine Learning

- Define a search space
- Configure a sampling method
- Configure early termination
- Use a sweep job for hyperparameter tuning

Module 15: Deploy a model to a managed online endpoint

- Explore managed online endpoints
- Deploy your MLflow model to a managed online endpoint
- Deploy a model to a managed online endpoint

Module 16: Deploy a model to a batch endpoint

- Understand and create batch endpoints
- Deploy your MLflow model to a batch endpoint
- Deploy a custom model to a batch endpoint
- Invoke and troubleshoot batch endpoints

Lab / Exercises

- Official Microsoft Labs

Documentation

- Access to Microsoft Learn (online learning content)

Exam

- This course prepares you to the exam DP-100: Designing and Implementing a Data Science Solution on Azure
- If you wish to take this exam, please select it when you add the course to your basket

Participant profiles

- Data Specialist

Prerequisites

- Have followed the course: Microsoft Azure Fundamentals (AZ-900)
- Understanding of data science including how to prepare data, train models, and evaluate competing models to select the best one
- How to program in the Python programming language and use the Python libraries: pandas, scikit-learn, matplotlib, and seaborn

Objectives

- Use Azure services to develop, train, and deploy machine learning solutions
- Design a data ingestion strategy
- Train machine learning models
- Deploy models in real-time or batch
- Explore and utilize Azure Machine Learning resources
- Run machine learning pipelines
- Optimize models with hyperparameter tuning

Description

Designing and Implementing a Data Science Solution on Azure (DP-100)

Niveau

Intermédiaire

Classroom Registration Price (CHF)

3200

Virtual Classroom Registration Price (CHF)

3000

Duration (in Days)

4

Reference

DP-100T01