

DP-100T01: Designing and Implementing a Data Science Solution on Azure

Course Duration: 32 Hours (4 Days)

Overview

The DP-100T01: Designing and Implementing a Data Science Solution on Azure course provides an in-depth exploration of Azure's machine learning capabilities. It covers the entire data science process from Data preparation, Model training, Model deployment, and Model management. Learners will gain practical experience with Azure Machine Learning Service and Azure Machine Learning Studio, learning how to create, train, optimize, and deploy machine learning models at scale. Throughout the course, participants will engage in hands-on labs, such as creating an Azure Machine Learning workspace, running experiments, working with Datastores and datasets, and orchestrating Machine learning workflows with Pipelines. They will also explore real-time and Batch inferencing, ensuring their models can respond promptly or handle large-scale processing. By mastering Hyperparameter tuning, Automated Machine Learning, and Model interpretation, students will be well-equipped to build responsible AI solutions. They'll also delve into the best practices for Monitoring models to maintain optimal performance over time, using tools like Application Insights and Data drift monitoring. This course is ideal for aspiring and existing data scientists looking to harness the power of Azure to streamline and enhance their Machine learning workflows.

Audience Profile

The DP-100T01 course is designed for professionals seeking to implement data science solutions on Azure's cloud platform.

- Data Scientists
- AI Engineers
- Machine Learning Engineers
- Cloud Solutions Architects
- IT Professionals with a focus on data analytics
- Software Developers interested in data science and machine learning
- Technical Leads managing data science teams
- Data Analysts aiming to advance in machine learning
- DevOps Engineers focused on ML/AI lifecycle management
- Professionals preparing for Azure Data Scientist Associate certification

Course Syllabus

Module 1: Explore and configure the Azure Machine Learning workspace

- Explore Azure Machine Learning workspace resources and assets
- Explore developer tools for workspace interaction
- Make data available in Azure Machine Learning
- Work with compute targets in Azure Machine Learning
- Work with environments in Azure Machine Learning

Module 2: Experiment with Azure Machine Learning

- Find the best classification model with Automated Machine Learning
- Track model training in Jupyter notebooks with MLflow

Module 3: Optimize model training with Azure Machine Learning

- Run a training script as a command job in Azure Machine Learning
- Track model training with MLflow in jobs
- Perform hyperparameter tuning with Azure Machine Learning
- Run pipelines in Azure Machine Learning

Module 4: Manage and review models in Azure Machine Learning

- Register an MLflow model in Azure Machine Learning
- Create and explore the Responsible AI dashboard for a model in Azure Machine Learning

Module 5: Deploy and consume models with Azure Machine Learning

- Deploy a model to a managed online endpoint
- Deploy a model to a batch endpoint

Module 5: Develop generative AI apps in Azure AI Foundry portal

- Introduction to Azure AI Foundry
- Explore and deploy models from the model catalog in Azure AI Foundry portal
- Get started with prompt flow to develop language model apps in the Azure AI Foundry
- Build a RAG-based agent with your own data using Azure AI Foundry
- Fine-tune a language model with Azure AI Foundry
- Evaluate the performance of generative AI apps with Azure AI Foundry
- Responsible generative AI

