Microsoft Fabric Advanced Training: Real-Time Analytics, Data Science, and Visualization

Course Description

This 3-day intensive program provides a unified learning path across Microsoft Fabric's advanced capabilities. Participants will gain hands-on expertise in real-time analytics, data science workflows, and visualization techniques. The course is designed to progress from foundational concepts to advanced implementations, enabling professionals to design, deploy, and optimize enterprise-scale solutions in Fabric.

Duration : 3 Days

Pre-requisites

- Basic knowledge of SQL and data concepts
- Familiarity with cloud platforms (Azure preferred)
- Exposure to machine learning fundamentals and Python notebooks
- Optional: experience with Power BI and visualization basics

Learning Objectives

- Implement real-time analytics solutions in Microsoft Fabric
- Build, train, and deploy machine learning models using Fabric notebooks and MLflow
- Create compelling visualizations and dashboards with Fabric and Power BI
- Integrate pipelines, Data Wrangler, and orchestration tools for end-to-end workflows
- Apply monitoring, optimization, and visualization strategies in enterprise environments

Day-wise Module Coverage

Day 1 Module 1: Microsoft Fabric Overview and Unified Analytics

- Fabric architecture and ecosystem
- Lakehouse and OneLake storage concepts
- Integration of engineering, analytics, and science roles
- Real-time vs batch analytics in Fabric
- Enterprise use cases and solution patterns

Module 2: Real-Time Analytics Fundamentals

- Eventhouse and KQL databases
- Streaming ingestion pipelines
- Real-time dashboards and alerts
- Querying streaming data with KQL
- Use cases for real-time intelligence

Module 3: Foundations of Data Science in Fabric

- Data science workflows in Fabric
- Notebook-based exploration and analysis
- Preprocessing with Data Wrangler
- MLflow integration for experiment tracking
- Collaboration across teams in Fabric

Module 4: Hands-on Lab – Real-Time Data + Notebooks

- Configuring Eventhouse ingestion
- Running queries on streaming data
- Preparing datasets with Data Wrangler

- Exploring data in notebooks
- Tracking experiments with MLflow

Day 2 Module 5: Model Training and Deployment

- Training models with Fabric notebooks
- Hyperparameter tuning and experiment tracking
- Deploying models for batch and streaming predictions
- Integrating models into pipelines
- Managing model lifecycle in Fabric

Module 6: Pipelines and Workflow Orchestration

- Pipelines for ingestion and transformation
- Copy activity and parameterization
- Integration with Dataflows Gen2
- CI/CD and Git integration for pipelines
- Workflow orchestration with Airflow jobs

Module 7: Visualization Essentials

- Visualization capabilities in Fabric
- Power BI integration with Fabric datasets
- Building interactive dashboards
- Real-time visualization with Eventhouse data
- Best practices for enterprise reporting

Module 8: Hands-on Lab – Model Deployment and Visualization

- Deploying ML models in Fabric

- Creating dashboards with Power BI
- Visualizing streaming data in real-time
- Troubleshooting visualization issues
- Reviewing visualization best practices

Day 3 Module 9: Advanced Visualization Strategies

- Designing enterprise dashboards
- Advanced Power BI features in Fabric
- Custom visuals and storytelling
- Performance optimization for reports
- Case study: enterprise visualization solution

Module 10: Real-Time Analytics Deep Dive

- Advanced KQL queries and functions
- Real-time anomaly detection
- Streaming joins and aggregations
- Integrating real-time analytics with ML models
- Case study: enterprise streaming solution

Module 11: Capstone Project – End-to-End Fabric Solution

- Designing a secure, governed Fabric environment
- Building and deploying ML models
- Implementing real-time analytics dashboards
- Creating enterprise-grade visualizations
- Presenting solutions and lessons learned