IBM Mastering Flink SQL on Confluent Platform

Duration: 5 Days Course Code: FLINKCPG

Overview:

Learn how to use Flink SQL on Confluent Platform to process, enrich, and correlate streaming data, enabling scalable real-time analytics and event-driven applications.

This two-day course provides comprehensive training on real-time stream processing using Flink SQL on Confluent Platform. Hands-on labs deliver practical experience with dynamic tables, watermarks, and event-time windowing. Participants learn to analyze live data streams using complex real-time aggregations and joins.

Throughout the course, hands-on exercises reinforce the topics being discussed. Exercises include: Setting up a Flink SQL lab environment on Confluent Platform (Kubernetes + CFK + CMF)Working with Flink in Confluent Platform using the Flink SQL CLI and Control CenterCreating dynamic tables and comparing change log modesImplementing watermarks and time windowsPerforming aggregations in a practical use caseExploring joins (regular, interval, temporal, windowed)

NB: If you work with Confluent Cloud please consider MASTERFLINKG FLINKCC2G

Updated Sept2026

Target Audience:

This course is designed for SQL practitioners who want to extend their skills to stream processing using Flink SQL on Confluent Platform. It is ideal for data engineers, analysts, and developers who are familiar with SQL and need to apply it to real-time data streams.

Objectives:

- **After this course participants should be able to:**
- Understand the fundamentals of Apache Flink and its role in stream processing
- Write and execute Flink SQL queries on Confluent Platform
- Differentiate between streaming and batch processing
- Work with dynamic tables and understand stream-table duality
- Manage time attributes, watermarks, and windows for event-time processing
- Perform complex windowed aggregations in real-time with Flink SQL
- Join, enrich, and correlate data across multiple streaming sources

Prerequisites:

This course is recommended for participants who have familiarity with SQL syntax and operations and have a basic understanding of Kafka (topics & partitions) and stream processing. Participants are required to provide a laptop computer with unobstructed internet access to fully participate in the class.

Testing and Certification

None

Follow-on-Courses:

None recommended

Content:

Module 01: Introduction to Flink

- Lab 01: Setting up the Lab Environment
- Origin of Stream Processing
- What is Apache Flink-
- Apache Flink's APIs
- Flink Job ; Topology

Module 02: Getting Started with Flink SQL

- Lab 02: Working with Flink in Confluent Platform
- Why Flink SQL for Stream Processing-
- Flink SQL Syntax
- Stream vs. Batch Processing
- Stateless ; Stateful Operators
- Flink SQL on Confluent Platform

Module 03: Working with Dynamic Tables

- Lab 03: Working with Dynamic Tables
- Traditional SQL vs. Streaming SQL
- Stream-Table Duality
- Dynamic Table Creation
- Table Configuration ; Patterns
- Column Types
- Processing Schema-less Events
- Scan Modes

Module 04: Time ; Windows

- Lab 04: Using Watermarks and Windows
- Event Time vs. Processing Time
- Time Attributes vs. Timestamps
- Watermarks
- Windows
- Time Functions ; Data Types

Module 05: Aggregations

- Lab 05: Using Aggregations in a Practical Use Case
- Overview
- GROUP BY vs. OVER
- Aggregate Functions
- Special Aggregation Queries
- Additional Aggregation Options
- Aggregations: Important Considerations

Module 06: Joins

- Lab 06: Exploring various types of Joins
- Introduction to Joins
- Regular Joins
- Optimized Joins
- Window Joins
- Other types of Joins

Additional Information:

Official course book provided to participants

Further Information:

For More information, or to book your course, please call us on Head Office 01189 123456 / Northern Office 0113 242 5931

info@globalknowledge.co.uk

www.globalknowledge.com/en-gb/

Global Knowledge, Mulberry Business Park, Fishponds Road, Wokingham Berkshire RG41 2GY UK