

IBM Mastering Flink SQL on Confluent Cloud - Essentials

Duration: 1 Day **Course Code: MASTERFLINKG** **Delivery Method: Virtual Learning**

Overview:

Learn to design, query, and analyze real-time streaming data using Flink SQL on Confluent Cloud, enabling you to transform live events into actionable business insights.

Master real-time stream processing with Flink SQL on Confluent Cloud. This hands-on course teaches you to process live data streams, build dynamic tables, and run complex time-windowed queries- equipping you to turn raw streaming data into immediate, actionable insights within a fully managed cloud environment.

NB: This course is included in FLINKCC2G for the first 2 days.

Updated Sept. 2026

Virtual Learning

This interactive training can be taken from any location, your office or home and is delivered by a trainer. This training does not have any delegates in the class with the instructor, since all delegates are virtually connected. Virtual delegates do not travel to this course, Global Knowledge will send you all the information needed before the start of the course and you can test the logins.

Target Audience:

This course is designed for SQL practitioners who want to extend their skills to stream processing using Flink SQL on Confluent Cloud. 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:**
- Understanding the fundamentals of Apache Flink and its relevance to stream processing
- Writing and executing Flink SQL queries on Confluent Cloud
- Differentiating between streaming and batch processing
- Working with dynamic tables and understand streamtable duality
- Managing time attributes and windows for effective stream processing
- Performing complex windowed aggregations in real-time with Flink SQL
- Utilizing Flink SQL for efficient joining of streaming data
- Applying pattern-matching techniques to identify complex event sequences

Prerequisites:

Learners are recommended to have foundational SQL knowledge and a general understanding of data management and cloud computing concepts to maximize the benefits of this course.

Testing and Certification

None

Follow-on-Courses:

None recommended

Content:

Module 1: Introduction to Flink

- Origin of Stream Processing
- What is Apache Flink-
- Apache Flink's APIs
- Flink Job ; Topology
- Lab 1: Setting up the lab environment

Module 2: Getting Started with Flink SQL

- Why Flink SQL for Stream Processing-
- Flink SQL Syntax
- Streaming vs. Batch Processing
- Flink SQL on Confluent Cloud
- Lab 2: Working with Flink in Confluent Cloud

Module 3: Working with Dynamic Tables

- Traditional SQL vs. Streaming SQL
- Stream-Table Duality
- Dynamic Table Creation
- Column Types
- Stateless vs. Stateful Operators
- Lab 3: Working with Dynamic Tables

Module 4: Time ; Windows

- Event Time vs. Processing Time
- Time Attributes vs. Timestamps
- Watermarks
- Windows
- Time Functions ; Data Types
- Lab 4: Using Watermarks and Windows

Module 5: Aggregations

- Overview
- GROUP BY vs. OVER
- Aggregate Functions
- Special Aggregation Queries
- Additional Aggregation Options
- Aggregations: Important Considerations
- Lab 5: Using Aggregations in a Practical Use Case

Module 06: Joins

- Introduction to Joins
- Regular Joins
- Optimized Joins
- Window Joins
- Other Type of Joins
- Lab 6: Exploring Various Types of Joins

Appendix: Pattern Matching

- Introduction
- Understand Query (Partitioning, Order of Events, DEFINE, MEASURES, Define a Pattern, Output Mode, After Match Strategy)
- Pattern Navigation
- Time Attributes

Examples

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