IBM Mastering Flink SQL on Confluent Cloud - Extended

Cursusduur: 2 Dagen **Cursuscode: FLINKCC2G**

Beschrijving:

Learn how to develop and optimize real-time data streaming solutions using Flink SQL on Confluent Cloud, including aggregations, joins, event pattern detection, and advanced state management.

This intensive, three-day course provides in-depth training on building end-to-end, real-time streaming pipelines using Flink SQL on Confluent Cloud. Participants transition from basic stream processing concepts to advanced streaming architecture, gaining practical experience in creating high-throughput, low-latency analytics pipelines.

Throughout the course, hands-on exercises reinforce the topics being discussed. Exercises include: Setting up Flink SQL on Confluent CloudWriting and managing Flink SQL queriesCreating dynamic tables and comparing change log modesImplementing time windows and watermarksPerforming aggregations in a practical use caseExploring joins (regular, interval, temporal, windowed)Pattern matching to detect complex event sequencesImplementing and deploying Java UDFsWorking with advanced statement patterns

NB: This course includes the content of MASTERFLINKG

Updated Sept 2026

Doelgroep:

This course is designed for SQL practitioners whowant to extend their skills to stream processing usingFlink SQL on Confluent Cloud.It is ideal for data engineers, analysts, and developerswho are familiar with SQL and need to apply it toreal-time data streams.

Doelstelling:

- **After this course participants should be able to:**
- Build real-time streaming pipelines using Flink SQL
- Work with dynamic tables and understand stream-table duality
- Master time windows, watermarks, and event-time processing
- Perform windowed aggregations and analytics on live data streams
- Join, enrich, and correlate data across multiple streaming sources
- Detect complex event patterns in real-time
- Extend Flink SQL capabilities with custom UDFs
- Use advanced patterns such as Statement Sets, Snapshots, EOS, and State Management
- Connect Flink SQL to external databases and AI/ML services

Vereiste kennis en vaardigheden:

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.

Examens en certificering

None

Vervolgcursussen:

None

Cursusinhoud:

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 Cloud
- Why Flink SQL for Stream Processing-
- Flink SQL Syntax
- Streaming vs. Batch Processing
- Stateless vs. Stateful Operators
- Flink SQL on Confluent Cloud

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

Module 07: Pattern Recognition

- Lab 07: Pattern Matching with MATCH_RECOGNIZE
- Introduction to Pattern Recognition
- Defining Patterns
- Output and Match Navigation
- Navigating Along the Pattern Examples

Module 08: UDFs

- Lab 08: Implementing a User-Defined Function (UDF)
- UDFs in Confluent Cloud
- Writing UDFs in Java

Module 09: Advanced Statement Patterns

- Lab 09: Advanced Statement Patterns
- Unioning Multiple Topics
- EXECUTE STATEMENT SET for Fan-In/Fan-Out
- Exactly-Once Semantics for Kafka Sinks
- Statement Evolution with Carry-over Offsets
- Snapshot Queries
- State Management
- Monitoring and Managing Statements with DLQ

Module 10: Connections

- Connections in Confluent Cloud
- External Tables
- AI/ML with Connections

Extra informatie:

Official course book provided to participants

Nadere informatie:

Neem voor nadere informatie of boekingen contact op met onze Customer Service Desk 030 - 60 89 444

info@globalknowledge.nl

www.globalknowledge.com/nl-nl/

Iepenhoeve 5, 3438 MR Nieuwegein