CONFLUENT

Datasheet

2-DAY COURSE

Mastering Flink SQL on Confluent Cloud

Unlock the Potential of Real-Time Data Processing

Flink SQL on Confluent Cloud enables robust stream processing by combining the power of Apache Flink with the scalability and reliability of Confluent Cloud. This comprehensive training is focused on using Flink SQL within Confluent Cloud for stream processing.

The 2-day course is focused on:

- 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 stream-table 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
- Appling pattern-matching techniques to identify complex event sequences

Participants will learn to utilize Flink SQL to process real-time data streams, create dynamic tables, manage time windows, and perform complex queries and pattern matching. The course demonstrates how Flink SQL can transform raw data into actionable insights within a fully managed cloud environment.

By the end of this course, participants will have a solid understanding of how to leverage Flink SQL for stream processing within Confluent Cloud, enabling them to transform real-time data streams into meaningful insights.

Who Should Attend?

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.

This 2-day course aims to enable participants to efficiently utilize Flink SQL for stream processing tasks, leveraging the capabilities of Confluent Cloud to manage and process streaming data effectively.

Prerequisites

This 200-level course is recommended for participants who have:

- Familiarity with SQL syntax and operations.
- Basic understanding of Kafka (topics & partitions) and stream processing

Hands-on Training

The hands-on lab exercises in this course include:

- Setting up Flink SQL environments on Confluent Cloud
- Writing and executing Flink SQL queries
- Creating and managing dynamic tables
- Implementing time windows and time-based functions
- Performing aggregations and joins on streaming data
- Developing pattern-matching queries to detect event patterns in data streams

Topics in Detail

Prerequisite - Fundamental course

Module 01 - Introduction to Flink

- 01a. Origin of Stream Processing
- 01b. What is Apache Flink?
- 01c. Apache Flink's APIs
- 01d. Flink Job & Topology

Module 02 - Getting Started with Flink SQL

- 02a. Why Flink SQL for Stream Processing?
- 02b. Flink SQL Syntax
- 02c. Streaming vs. Batch Processing
- 02d. Flink SQL on Confluent Cloud

Module 03 - Working with Dynamic Tables

- 03a. Traditional SQL vs. Streaming SQL
- 03b. Stream-Table Duality
- 03c. Dynamic Table Creation
- 03d. Column Types
- 03e. Stateless vs. Stateful Operators

Module 04 - Time & Windows

- 04a. Event Time vs. Processing Time
- 04b. Time Attributes vs. Timestamps
- 04c. Watermarks
- 04d. Windows
- 04e. Time Functions & Data Types

Module 05 - Aggregations

- 05a. Overview
- 05b. GROUP BY vs. OVER
- 05c. Aggregate Functions
- 05d. Special Aggregation Queries
- 05e. Additional Aggregation Options
- 05f. Aggregations: Important Considerations

Module 06 - Joins

- 06a. Introduction to Joins
- 06b. Regular Joins
- 06c. Optimized Joins
- 06d. Window Joins
- 06e. Other Type of Joins

Module 07 - Pattern Matching

- 07a. Introduction
- 07b. Understanding the Query:
  - Partitioning
  - Order of Events
  - DEFINE
  - MEASURES
  - Define a Pattern
  - Output Mode
  - After Match Strategy
- 07c. Pattern Navigation
- 07d. Time Attributes
- 07e. Examples