

Training: Cloudera

Cloudera Streaming Analytics: Using Apache Flink and SQL Stream Builder on CDP

CLOUDERA
Training Partner

TRAINING GOALS:

During this two-day instructor-led training course, participants will learn development and operations for Cloudera Streaming Analytics, a framework for low-latency processing and analytics powered by Apache Flink and Cloudera's innovative SQL Stream Builder.

Through extensive hands-on exercises, students will gain experience deploying and managing a Flink cluster, developing and running Flink applications, and using SQL Stream Builder's continuous SQL to perform analytics on streaming data.

During this course, you learn how to:

- Deploy a Flink cluster using Cloudera Manager
- Develop Flink batch and streaming applications
- Run and view Flink jobs
- Transform data streams
- Use watermarks and windows to analyze streaming data
- Analyze data with Cloudera SQL Stream Builder
- Monitor Flink application metrics

CONSPECT:

- **Overview**
 - Introduction to Apache Flink and Stream Processing
 - Typical Use Cases
 - Related Products
- **Service Deployment**
 - Planning Requirements
 - Installation
 - Flink Dashboard
 - Exercise: Running a Flink Program
- **Flink Basics**
 - Execution Environment

- Flink Application Structure
- Create a Flink Project
- Build a Flink Program
- Exercise: Building a Simple Flink Program
- **Architecture**
 - Logical
 - Physical
 - Parallelism
 - Fault Tolerance
 - Data Storage
- **DataStream API**
 - DataStream API Overview
 - Data Types and Serialization
 - Sources and Sinks
 - Transformations
 - Exercise: Batch Processing Using Flink
- **DataStream API (continued)**
 - Exercise: Creating a Flink Streaming Application
 - Using Kafka as a Source and Sink
 - Exercise: Creating a Streaming Application Using a Kafka Source
- **Flink SQL and Table API**
 - Streaming Concepts
 - Programming Options
 - Integrations
 - Exercise: Using Flink SQL and Kafka
- **Stateful Stream Processing**
 - Stateful Streaming Applications
 - Checkpoints
 - Event Time Processing
 - Watermarks
 - Windows
 - Exercise: Tumbling Windows with Event Time
- **Cloudera SQL Stream Builder**
 - Overview
 - Streaming SQL Console
 - Data Providers

- SQL Stream Jobs
- Exercise: Using SQL Stream Builder
- **Monitoring**
 - Metrics
 - Logging
 - Backpressure
 - Resource Configuration
 - Exercise: Monitoring

REQUIREMENTS:

This course is designed for those who have experience with administration and application development on the Cloudera platform. Students must have at least basic familiarity with Java and Linux. Our [Cloudera Training for Apache Kafka](#) course, or equivalent experience with Apache Kafka, is a recommend prerequisite.

Difficulty level



TRAINER:

Certified Cloudera Instructor.