CONFLUENT

TRAINING COURSE

Confluent Advanced Skills for Optimizing Apache Kafka®

Course Objectives

During this hands-on course you will learn how:

- Industry best practices to monitor, troubleshoot and tune a real-time event streaming platform built with Apache Kafka, to support your event-streaming applications

- To utilize native tools, including Confluent Control Center as well as third-party software tools

Hands-on Training

Hands-on exercises are included throughout the training to help reinforce the concepts being discussed.

Exercises include:

- Monitoring Kafka brokers, ZooKeeper and Kafka clients via JMX and Confluent Control Center, as well as Prometheus and Grafana

- Troubleshooting brokers, ZooKeeper, Producers, Consumers, Confluent ksqlDB applications, Kafka Streams applications and the Schema Registry, Tuning brokers, Producers, Consumers, Kafka Connect, ksqlDB and Kafka Streams, applications discussing production-down trouble shooting strategies

- Calculating replication throttling rates

- Evaluating partition strategies

Prerequisites

Attendees should have deep knowledge of and hands-on experience with Kafka. An understanding of additional components that contribute to a real-time event-streaming platform such as Kafka Connect, ksqlDB and Kafka Streams applications is also required.

Additionally it is suggested that students complete training beforehand to ensure familiarity with the relevant concepts. Visit www.confluent.io/training to learn the fundamentals of data streaming and Apache Kafka.

Participants are required to provide a laptop computer with unobstructed internet access to fully participate in the class.

To sign-up for one of our courses, visit us here.

Who Should Attend?

This course is designed for engineers, system administrators, and operations staff responsible for in-depth monitoring, tuning and troubleshooting an event-streaming platform powered by Kafka.

Module	Description
Major Components	This intensive hands-on course will cover the monitoring, troubleshooting, and tuning of the major components of Apache Kafka and Confluent Platform environments including: • Brokers • Producers • Consumers • Apache ZooKeeper • Kafka Streams • Kafka Connect • Confluent ksqlDB • Confluent Schema Registry
Additional Discussions	Will include: • Determining Business Needs and SLAs • Monitoring using JMX Metrics • Monitoring with Confluent Control Center and third-party tools • Recommending Troubleshooting Strategies • Using Kafka Troubleshooting tools • Analyzing Kafka Log Files • Performing a Troubleshooting Post-mortem Analysis

Visit confluent.io/training for more information.