



Confluent Operations Training for Apache Kafka

Course Objectives

In this two-day hands-on course you will learn how to build and manage clusters using industry best-practices developed by the world's foremost Apache Kafka experts. You will learn how Kafka and the Confluent Platform work, their main subsystems, how they interact, and how to set up, manage, monitor, and tune your cluster.

Hands-On Training

Throughout the course, hands-on exercises reinforce the topics being discussed. Exercises include:

- ▶ Basic cluster operations
- ▶ Viewing and interpreting cluster metrics
- ▶ Recovering from a Broker failure
- ▶ Performance-tuning the cluster
- ▶ Changing the number of partitions in a topic

Who Should Attend?

This course is designed for engineers, system administrators, and operations staff responsible for building, managing, monitoring, and tuning Kafka clusters.

Course Duration

This is a two-day training course.

Course Prerequisites

Attendees should have a strong knowledge of Linux/Unix, and understand basic TCP/IP networking concepts. Familiarity with the Java Virtual Machine (JVM) is helpful. Prior knowledge of Kafka is helpful, but is not required.



Course Contents

Introduction

Basic Kafka Concepts

- Core Kafka Features
- Comparisons with Traditional Message Queues

Intra-Cluster Replication

- Basic Replication Concepts
- Replica Recovery and Failure Detection

An Inside Look at Kafka's Components

- Inside the Kafka Producer
- Inside the Kafka Consumer
- Inside the Kafka Broker

Log Retention and Compaction

- Log Compaction

Hardware and Runtime Configurations

- Hardware and Capacity Planning
- Runtime Configuration Settings
- Dealing with Oversized Messages

Monitoring and Alerting

- ZooKeeper and OS-Level Monitoring
- Key Kafka Metrics

Cluster Administration

- Key Cluster Administration Tasks
- Replicating Clusters with MirrorMaker

Kafka Security

- SSL for Encryption and Authentication
- SASL for Authentication
- Authorization
- Migration to a Secure Cluster

Conclusion

Appendix: Integrating Systems with Kafka Connect

- The Motivation for Kafka Connect
- Using Kafka Connect in Distributed Mode
- Tracking Offsets
- Converters
- Connector Configuration
- Comparing Kafka Connect with Other Options