

Confluent Operations Training for Apache Kafka

Course Objectives

In this three-day hands-on course you will learn how to build, manage, and monitor clusters using industry best-practices developed by the world's foremost Apache Kafka® experts. You will learn how Kafka and the Confluent Platform work, how their main subsystems 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:

- » Cluster installation
- » Basic cluster operations
- » Viewing and interpreting cluster metrics
- » Recovering from a Broker failure
- » Performance-tuning the cluster
- » Securing the cluster

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 three-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.

Introduction

The Motivation for Apache Kafka

- Systems Complexity
- Real-Time Processing is Becoming Prevalent
- Kafka: A Stream Data Platform

Kafka Fundamentals

- An Overview of Kafka
- Kafka Producers
- Kafka Brokers
- Kafka Consumers
- Kafka's Use of ZooKeeper
- Comparisons with Traditional Message Queues

Providing Durability

- Basic Replication Concepts
- Durability Through Intra-Cluster

Replication

- Writing Data to Kafka Reliably
- Broker Shutdown and Failures
- Exactly Once Semantics (EOS)
- Controllers in the Cluster
- The Kafka Log Files

Managing a Kafka Cluster

- Installing and Running Kafka
- Monitoring Kafka
- Basic Cluster Management
- Log Retention and Compaction
- An Elastic Cluster

Optimizing Kafka Performance

- Batching for Performance
- Producer Performance
- Broker Performance
- Broker Failures and Recovery Time
- Load Balancing Consumption
- Consumption Performance
- Performance Testing

Kafka Security

- SSL for Encryption and Authentication
- SASL for Authentication
- Securing ZooKeeper and the REST Proxy
- Migration to a Secure Cluster

Integrating Systems with Kafka

- Offset Management

Designing for High

Availability Connect

- Motivation for Kafka Connect
- Kafka Reference Architecture
- Types of Connectors
- Brokers
- Kafka Connect Implementation
- ZooKeeper
- Standalone and Distributed Modes
- Connect
- Configuring the Connectors
- Schema Registry
- REST Proxy
- Comparison with Other Systems
- Multiple Data Centers

Confluent offers public training courses and private, onsite events. Please visit <http://confluent.io/training> for the public course schedule.

For inquiries about onsite events, email training-admin@confluent.io