



Kafka : Confluent Administration

Duration 3 day(s) (CONFLUENT-OP-02)

Build, Manage and Monitor Kafka clusters

Official Training



Description

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.

Goals

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.

Public

This course is designed for developers and operational teams who need to deploy, administer, and configure a Kafka cluster.

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.

Structure

50% Theory, 50% Practice

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