

# Apache Kafka® Administration by Confluent

Version: 5.3.0-v1.1.0

Date: November 2019

## Course Objectives

In this three-day, hands-on course you will learn how to build, manage, and monitor Kafka 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:

- Using Kafka's command-line tools
- Automating configuration
- Using Kafka's administrative tools
- Tuning Producer and Consumer performance
- Securing the cluster
- Building data pipelines with Kafka Connect

## Course Duration

This is a three day, instructor led class but the value of the course lives well beyond that time. Your course materials include guidance made possible from the collective wisdom of the experts at Confluent. Your lab exercise book can be used to replay exercises and experiment on your own in a local docker-compose or VM environment long after the class is over.

## Who Should Attend?

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

## Course Prerequisites

Attendees should have a working knowledge of the Kafka architecture, either from prior experience or the recommended prerequisite course [Confluent Fundamentals of Apache Kafka®](#). It is also important to have strong knowledge of Linux/Unix and understand basic TCP/IP networking concepts. Familiarity with the Java Virtual Machine (JVM) is helpful.

To evaluate your Kafka knowledge for this course, please complete the free, anonymous self-assessment here: <https://confluent.io/training>

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

Confluent offers this course in Instructor-led (in-person or online) and On Demand formats. Please visit <https://confluent.io/training> for more information.

For inquiries about onsite training, please email [training-admin@confluent.io](mailto:training-admin@confluent.io)

## Course Content

### Fundamentals of Apache Kafka

- Kafka as a Distributed Streaming Platform
- The Distributed Log
- Producer and Consumer Basics

### Kafka's Architecture

- Kafka's Commit Log
- Replication for High Availability
- Partitions and Consumer Groups for Scalability
- Security Overview

### Providing Durability

- Data Replication
- Failure Recovery
- Log Files and Offset Management
- Exactly-Once Semantics (EOS)

### Managing a Cluster

- Installing the Confluent Platform
- Configuration Management
- Monitoring
- Log Retention and Compaction
- Commissioning and Decommissioning Brokers

## Optimizing Kafka's Performance

- Monitoring, Testing, and Tuning Brokers and Kafka Clients
- The Consumer Group Protocol

## Kafka Security

- Transport Encryption
- Authentication
- Securing Apache Kafka and the Complete Confluent Platform
- Migrating to a Secure Cluster

## Data Pipelines with Kafka Connect

- The Motivation for Kafka Connect
- Types of Connectors
- Kafka Connect Implementation
- Standalone and Distributed Modes
- Configuring the Connectors

## Kafka in Production

- Kafka Reference Architecture for Apache Kafka and the Complete Confluent Platform
- Capacity Planning
- Multi Data Center Deployments