

## TRAINING COURSE

# Apache Kafka® Administration by Confluent

## 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:

- 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 Prerequisites

Attendees should have a working knowledge of the Kafka architecture, either from prior experience or the recommended prerequisite course Confluent Fundamentals for 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://cnfl.io/fundamentals-quiz>

## Course Duration

This is a three-day training course.

## Who Should Attend?

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

# Content

MODULE	DESCRIPTION
<b>Fundamentals of Apache Kafka</b>	<ul style="list-style-type: none"> <li>• Kafka as a Distributed Streaming Platform</li> <li>• The Distributed Log</li> <li>• Producer and Consumer Basics</li> </ul>
<b>Apache Kafka Architecture</b>	<ul style="list-style-type: none"> <li>• Kafka's Commit Log</li> <li>• Replication for High Availability</li> <li>• Partitions and Consumer Groups for Scalability</li> <li>• Security Overview</li> </ul>
<b>Providing Durability</b>	<ul style="list-style-type: none"> <li>• Data Replication</li> <li>• Failure Recovery</li> <li>• Log Files &amp; Offset Management</li> <li>• Exactly-Once Semantics (EOS)</li> </ul>
<b>Managing a Cluster</b>	<ul style="list-style-type: none"> <li>• Installing and Running Kafka</li> <li>• Configuration Management</li> <li>• Monitoring</li> <li>• Log Retention and Compaction</li> <li>• Commissioning and Decommissioning Brokers</li> </ul>
<b>Optimizing Apache Kafka Performance</b>	<ul style="list-style-type: none"> <li>• Monitoring, Testing, and Tuning Brokers and Kafka Clients</li> <li>• The Consumer Group Protocol</li> </ul>
<b>Apache Kafka Security</b>	<ul style="list-style-type: none"> <li>• Transport Encryption</li> <li>• Authentication</li> <li>• Securing Apache Kafka</li> <li>• Migrating to a Secure Cluster</li> </ul>
<b>Data Pipelines with Kafka Connect</b>	<ul style="list-style-type: none"> <li>• The Motivation for Kafka Connect</li> <li>• Types of Connectors</li> <li>• Kafka Connect Implementation</li> <li>• Standalone and Distributed Modes</li> <li>• Configuring the Connectors</li> </ul>
<b>Kafka in Production</b>	<ul style="list-style-type: none"> <li>• Kafka Reference Architecture for Apache Kafka and the Complete Confluent Platform</li> <li>• Capacity Planning</li> <li>• Multi Data Center Deployments</li> </ul>

Confluent offers instructor-led courses in both traditional and virtual classroom formats, as well as in an on-demand (recorded) format. Please visit <http://confluent.io/training> for more information.