

Course Objectives

During this self-paced, on-demand training , you will:

- Gain understanding of the Apache Kafka and Confluent Platform
- Explore use cases
- Receive an overview of Kafka's core concepts that enable it to power a highly scalable, highly available and resilient real-time event streaming platform
- Be introduced to the Confluent Platform, offering an enterprise-ready, real-time event streaming platform powered by Apache Kafka
- Begin preparation to attend Confluent Operations for Apache Kafka and/or the Confluent Skills for Building Apache Kafka training courses.

Hands-on Training

To help reinforce the concepts being taught, hands-on exercises are included throughout the training. Exercises include:

- Launching and exploring a minimal Kafka cluster
- Using Kafka command line tools to explore cluster meta data in ZooKeeper, create topics on the broker, and publish & consume messages
- Running a Java based consumer and observe consumer lag when scaling the consumer
- Configuring Kafka Connect with a MQTT Connector source to create a data pipeline
- Using Confluent Control Center to monitor your cluster and execute KSQL queries.

Course Duration

This is a one-day course available On-demand (self-paced) or Instructor-led by request.

Who Should Attend?

This course is designed for all professionals who work with a real-time event streaming platform powered by Apache Kafka. This course is a recommended prerequisite for students attending Confluent Operations for Apache Kafka and/or the Confluent Skills for Building Apache Kafka.

Course Prerequisites

Attendees are not expected to have any prior experience with Kafka. Some basic understanding of the Linux OS and experience in using a shell like Bash is beneficial.

Course Content

Motivation and Customer Use Cases

- Motivation for a paradigm change to "Event-driven"
- How Kafka is the backbone of real-time event streaming
- How other major players in the market use Kafka

- Customer Use Cases such as
 - Microservices, IoT and Edge Computing
 - Core Banking, payments engine and fraud detection
 - Cyber Data Collection and Dissemination
 - ESB Replacement
 - Data Pipelining
 - eCommerce and Customer 360
 - Mainframe offloading
 - And more...

Apache Kafka Fundamentals

- Architecture
- ZooKeeper's role
- Topics, Partitions and Segments
- The commit log and streams
- Brokers and Broker replication
- Producers Basics
- Consumers, Consumer groups and Offsets

How Kafka Works

- High-level code overview for a basic producer and a basic consumer
- High Availability through Replication
- Data Retention Policies
- Producer Design and Producer Guarantees
- Delivery Guarantees, including Exactly Once Semantic
- Partition strategies
- Consumer group rebalances
- Compacted Topics
- Troubleshooting strategies
- Security overview

Integrating Kafka into your Environment

- Get streams of data into and out of Kafka with Kafka Connect and REST Proxy
- Maintain data formats and ensure compatibility with Schema Registry and Avro
- Build real-time streaming applications with Confluent KSQL & Kafka Streams

The Confluent Platform

- The Streaming Platform as the Central Nervous System
- Deployment Models - on premise versus SaaS
- The Confluent Control Center
- Role Based Access Control (RBAC)
- The Confluent CLI
- Confluent Operator
- The Confluent Hub for Certified Connectors