

Kafka by Confluent - Confluent Stream Processing using Apache Kafka Streams

Download Whitepaper: Accelerate Your Modernization Efforts with a Cloud-Native Strategy
Get Your Free Copy Now

Course Number: AK-STREAM

Duration: 2 days

Overview

Course Description

In this two-day, hands-on training course, you will learn how to use Confluent KSQL to transform, enrich, filter and aggregate streams of real-time data using a SQL-like language. You will learn to use the Streams library to build applications, and then learn how to test, monitor, secure, scale and integrate them.

Skills Gained

During this instructor-led, hands-on course, you will learn to:

- Identify common patterns and use cases for real-time stream processing
- Understand the high level architecture of Apache Kafka® Streams
- Write real-time applications with the Kafka Streams API to filter, transform, enrich, aggregate, and join data streams

- Describe how KSQL combines the elastic, fault-tolerant, high-performance stream processing capabilities of Kafka Streams with the simplicity of a SQL-like syntax
- Author KSQL queries that showcase its balance of power and simplicity
- Test, secure, deploy, and monitor Kafka Streams applications and KSQL queries

Who Can Benefit

This course is designed for application developers, architects, DevOps engineers, and data scientists who need to interact with Kafka clusters to create real-time applications using Kafka Streams for filtering, transforming, enriching, aggregating, and joining data streams to discover anomalies, analyze behavior, or monitor complex systems.

Prerequisites

Attendees should be familiar with developing professional apps in Java (preferred), .NET, C#, Python, or another major programming language.

Attendees should also have a working knowledge of the Kafka client development, either through:

- Prior experience, or by taking the recommended course prerequisites:
- Confluent Fundamentals for Apache Kafka and Confluent Developer Skills for Building Apache Kafka.
- Participants are required to provide a laptop computer with unobstructed internet access to fully participate in the class.

Hands-on Training

Throughout the course, you will interact with hands-on lab exercises to reinforce stream processing concepts. Some exercises include:

- Scaling a Kafka Streams Application
- Anatomy of a Kafka Streams Application
- Working With JSON
- Windowing & Aggregation
- Joining Two Streams
- Using the Processor API
- Integration Tests Using Embedded Kafka
- Using JConsole to Monitor a Streams App
- Securing a Kafka Streams Application

Audience

Course Details

Course Outline

Introduction to Kafka Streams

- Gain a better understanding of the fundamentals of Apache Kafka
- Delve into how Apache Kafka uses the group management protocol to balance resources
- Give a description of some Stream Processing concepts

Working with Kafka Streams

- Describe the anatomy of a Kafka Streams application
- Write a streams application employing components of the Kafka Streams DSL:
- Stateless transformations
- Stateful transformations
- Optimizations

Complex Stateful Processing Operations

- Review the concept of time in Apache Kafka Streams
- Use the Kafka Stream stateful operations:
- Windowing
- Aggregations
- Joins

Advanced Concepts

- Handle late-arriving events
- Explore foreign key joins in Kafka Streams
- Use the Kafka Streams Processor API

Testing, Monitoring, and Troubleshooting

- Perform testing with Apache Kafka Streams
- Evaluate monitoring options with:
- Confluent Control Center
- JMX metrics
- Discuss common errors and troubleshooting approaches

Deployment

- Discuss deployment strategies:
- Parallelism
- Capacity planning
- Elasticity
- Fault tolerance

Security

- Explore how to secure your Kafka Streams applications:
- Security overview
- Access Control Lists (ACLs) examples
- Kafka Streams security configurations