IBM Confluent Stream Processing using Apache Kafka Streams

Duration: 3 Days Course Code: STREAMSG

Overview:

Harness the power of continuous data processing by building high-performance, event-driven applications with the Apache Kafka® Streams API and ksqlDB.

Learn to design scalable, stateful stream architectures that execute low-latency joins, windowed aggregations, and real-time transformations. Transform reactive batch workflows into continuous data pipelines that deliver immediate operational insights and maximize business value.

Throughout the course, hands-on exercises reinforce the topics being discussed. Exercises include: Scaling a Kafka Streams Application Anatomy of a Kafka Streams ApplicationWorking With JSON Windowing & Aggregation Joining Two Streams Using the Processor API Integration Tests Using Embedded KafkaUsing JConsole to Monitor a Streams App Securing a Kafka Streams Application

Updated 16/09/2026

Target Audience:

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.

Objectives:

- The lessons and activities in this course enable participants to build the skills to:
- Identify common patterns and use cases for real-time stream processing
- Describe 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 Kafka Streams provide elastic, fault-tolerant and high-performance stream processing capabilities
- Test, secure, deploy, and monitor Kafka Streams applications

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

Testing and Certification

- *None*

Follow-on-Courses:

- *None recommended*

Content:

Module1: 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

Module 2: Working with Kafka Streams

- Describe the anatomy of a Kafka Streams application
- Write streams application employing components of the Kafka Streams DSL (Stateless/stateful transformations Optimizations)

Module 3: Complex Stateful Processing Operations

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

Module 4: Advanced Concepts

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

Module 5: 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

Module 6: Deployment

- Discuss deployment strategies (Parallelism, Capacity planning, Elasticity, Fault tolerance)

Module 7: Security

- Explore how to secure Kafka Streams applications (Security overview, Access Control Lists examples)

Additional Information:

Official course book provided to participants

Further Information:

For More information, or to book your course, please call us on 0800/84.009

info@globalknowledge.be

www.globalknowledge.com/en-be/