

Course Content: Apache Kafka, Kafka Streams and ksqlDB

Introduction to Big data and Apache Kafka

- Introduction to Big Data
- Big Data Customer Scenarios
- What is Kafka?
- Need for Kafka
- Core Concepts of Kafka
- Kafka Architecture
- Where is Kafka Used

About Kafka Cluster

- Understanding the components of Kafka Cluster
- Installation of Kafka Cluster
- Configuring Kafka Cluster
- Producer of Kafka
- Consumer of Kafka
- Producer and Consumer in Action

Kafka Java Programming

- Understanding the components of Kafka Cluster
- Kafka SDK List
- Creating Kafka Project
- Java Producer
- Java Producer Callbacks
- Java Producer with Keys
- Java Consumer
- Java Consumer – Graceful Shutdown
- Java Consumer inside Consumer Group
- Java Consumer Incremental Cooperative Rebalance and Static Group Membership
- Java Consumer Incremental Cooperative Rebalance – Practice
- Java Consumer Auto Offset Commit Behavior

Kafka Operations and Performance Tuning

- Offset
- Design
- Hardware
- Kafka Monitoring and Issues
- Kafka Performance Tuning
- Reading data from Kafka
- Demo - Java Producer
- Mixed Paradigm-Functional Programming

Kafka Wikimedia Producer and Advanced Producer Configurations

- Wikimedia Producer Project Setup
- Wikimedia Producer Implementation
- Wikimedia Producer Run
- Wikimedia Producer – Producer Config Intros
- Producer Acknowledgements Deep Dive
- Producer Retries
- Idempotent Producer
- Safe Kafka Producer Settings
- Wikimedia Producer Safe Producer Implementation
- Kafka Message Compression

OpenSearch Consumer and Advanced Consumer Configurations

- OpenSearch Consumer – Project Overview
- OpenSearch Consumer – Project Setup
- Setting Up OpenSearch on Docker
- OpenSearch Consumer Implementation
- Consumer Delivery Semantics
- OpenSearch Consumer Implementation - Idempotence
- Consumer Offsets Commit Strategies
- OpenSearch Consumer Implementation – Delivery Semantics
- OpenSearch Consumer Implementations – Batching Data
- Consumer Offset Reset Behavior
- OpenSearch Consumer Implementation – Replaying Data

The Kafka Streams

- Difference vs other libraries (Spark Streaming, Flink, NiFi...)
- Underlying concepts (topologies, processors...)
- Basic and Advanced transformations API – Stateless and Stateful
- Grouping, Aggregations, Joins
- Exactly-Once Semantics and Capabilities
- Production Settings and Configurations.

ksqlDB

- ksqlDB and KSQL Setup
- ksqlDB and KSQL Basics
- ksqlDB and KSQL intermediate
- ksqlDB and KSQL Extensions - UDF and UDAF
- ksqlDB and KSQL in Production