

Complete Apache Kafka Administration (Kafka Operations + Security + Monitoring)

Prerequisites: Basic knowledge of java

Duration: 6 Days (8 Hrs./Day)

Course Objective: This course provides a comprehensive introduction to Big Data concepts and Apache Kafka, including its architecture, key components, and real-world use cases. Participants will gain hands-on experience with Kafka clusters, Java-based Kafka applications, performance tuning, and security measures. By the end of the course, students will be proficient in setting up and managing Kafka systems, monitoring performance with Prometheus and Grafana, and implementing secure Kafka solutions.

Apache Kafka version: Latest.

Lab Requirement: Koenig DC/Linux (CentOS 9) (customizable).

Module 1 - 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?

Module 2 - About Kafka Cluster

Understanding the Components of Kafka Cluster

Producer of Kafka

Consumer of Kafka

Lab: Installation of Kafka Cluster

Lab: Configuring Kafka Cluster

Lab: Producer and Consumer in Action

Module 3 - Kafka with Java

Understanding the Components of Kafka Cluster

Kafka SDK List

Lab: Creating Kafka Project with Java

Lab: Java Producer



Lab: Java Producer Callbacks

Lab: Java Producer with Keys

Lab: Java Consumer

Lab: Java Consumer with Graceful Shutdown

Lab: Java Consumer in Consumer Group

Lab: Java Consumer with Incremental Cooperative Rebalance and Static Group Membership

Lab: Java Consumer Auto Offset Commit Behavior

Module 4 - Kafka Operations and Performance Tuning

Offset

Design

Hardware

Kafka Monitoring and Issues

Kafka Performance Tuning

Module 5 - Kafka Wikimedia Producer and Advanced Producer Configurations

Introduction to Kafka Producer

Lab: Wikimedia Producer Project Setup

Lab: Wikimedia Producer Implementation with Java

Lab: Wikimedia Producer Run with Java

Lab: Producer Config Intros

Lab: Producer Acknowledgements

Lab: Producer Retries

Lab: Idempotent Producer

Lab: Safe Kafka Producer Settings

Lab: Wikimedia Producer Safe Producer Implementation

Lab: Kafka Message Compression

Module 6 - OpenSearch Consumer and Advanced Consumer Configurations

Introduction to OpenSearch

Lab: OpenSearch Consumer with Java

Lab: Setting Up OpenSearch on Docker

Lab: OpenSearch Consumer Implementation with Java

Lab: Consumer Delivery Semantics

Lab: OpenSearch Consumer Implementation Idempotence

Lab: Consumer Offsets Commit Strategies

Lab: OpenSearch Consumer Implementation Delivery Semantics



Lab: OpenSearch Consumer Implementations Batching Data

Lab: Consumer Offset Reset Behavior

Lab: OpenSearch Consumer Implementation Replaying Data

Module 7 - Kafka Monitoring Setup with Prometheus

Overview of Monitoring

Lab: Setting up Prometheus + Kafka Broker 1

Lab: Broker 2 & 3 Prometheus Setup

Lab: KRaft & Prometheus Setup

Module 8 - Kafka Monitoring Setup with Grafana

Overview of Monitoring

Lab: Setup Grafana on Administration Server

Lab: Setup Kafka Dashboard on Grafana

Lab: Adding Graphs to Grafana

Lab: Adding more graphs to Grafana

Lab: Observing Grafana when we kill a broker

Module 9 - Introduction to Kafka Security

Kafka Authentication Basics

Kafka Authentication with SSL and SASL_SSL

Authorization

Encryption

Securing KRaft

Audit Logs

Security Recommendations

Module 10 – Kafka security hands on

Lab: Creating a Certificate Authority (CA)

Lab: SSL Setup in Kafka

Lab: SSL Setup for Clients

Lab: SSL Authentication

Lab: Hands-On Kerberos - Part 1: Setup Azure Virtual Machines

Lab: Hands-On Kerberos - Part 2: Principals & Keytabs

Lab: Hands-On Kerberos - Part 3: Kafka Configuration

Lab: Hands-On Kerberos - Part 4: Client Configuration

Lab: JAAS file / config

