

Apache Kafka

Duration: 4 days (8hrs/day)

Prerequisites: Basic knowledge of Linux and Java.

Course Objective: This Apache Kafka course introduces the fundamentals of Apache Kafka and its core concepts such as topics, partitions, replication, and message queues. It also covers Kafka's architecture, configuration, integration, and security. Additionally, the course explores the Kafka Producer API, Consumer API, Streams API, and Connect API.

Apache Kafka Version: Latest

Lab Requirement: Koenig DC (CentOS 7)

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 Java Programming

Understanding the components of Kafka Cluster

Kafka SDK List

Lab: Creating Kafka Project

Lab: Java Producer

Lab: Java Producer Callbacks

Lab: Java Producer with Keys

Lab: Java Consumer

Lab: Java Consumer Lab: Graceful Shutdown

Lab: Java Consumer inside Consumer Group

Lab: Java Consumer Incremental Cooperative Rebalance and Static Group Membership

Lab: Java Consumer Incremental Cooperative Rebalance

Lab: Java Consumer Auto Offset Commit Behaviour

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

Lab: Wikimedia Producer Run

Lab: Wikimedia Producer

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

Lab: OpenSearch Consumer

Lab: Setting Up OpenSearch on Docker

Lab: OpenSearch Consumer Implementation

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 Behaviour

Lab: OpenSearch Consumer Implementation Replaying Data