

## Course duration

- 2 days

## Course Benefits

- Understand the use of Kafka for high performance messaging.
- Identify the usages for Kafka in Microservices.
- Explain the benefits of Kafka patterns.
- Differentiate between messaging and message brokers.
- Describe Kafka messaging environments.
- Develop producers and consumers for Kafka.
- Recognize how Kafka enables Cloud-native applications.
- Summarize characteristics and architecture for Kafka.
- Demonstrate how to process messages with Kafka.
- Design distributed high throughput systems based on Kafka.
- Describe the built-in partitioning, replication and inherent fault-tolerance of Kafka.

## Course Outline

1. Introduction to KAFKA
  1. Messaging Architectures – What is Messaging?
  2. Messaging Architectures – Steps to Messaging
  3. Messaging Architectures – Messaging Models
  4. What is Kafka?
  5. What is Kafka? (Contd.)
  6. Kafka Overview
  7. Need for Kafka
  8. When to Use Kafka?
  9. Kafka Architecture
10. Core concepts in Kafka
11. Kafka Topic
12. Kafka Partitions
13. Kafka Producer
14. Kafka Consumer
15. Kafka Broker
16. Kafka Cluster
17. Why Kafka Cluster?
18. Sample Multi-Broker Cluster
19. Overview of ZooKeeper
20. Kafka Cluster & ZooKeeper
21. Who Uses Kafka?

- 22. Summary
- 2. Using Apache Kafka
  - 1. Installing Apache Kafka
  - 2. Configuration Files
  - 3. Starting Kafka
  - 4. Using Kafka Command Line Client Tools
  - 5. Setting up a Multi-Broker Cluster
  - 6. Using Multi-Broker Cluster
  - 7. Kafka Connect
  - 8. Kafka Connect – Configuration Files
  - 9. Using Kafka Connect to Import/Export Data
  - 10. Creating a Spring Boot Producer
  - 11. Adding Kafka dependency to pom.xml
  - 12. Defining a Spring Boot Service to Send Message(s)
  - 13. Defining a Spring Boot Controller
  - 14. Testing the Spring Boot Producer
  - 15. Creating a Nodejs Consumer
  - 16. Summary
- 3. Building Data Pipelines
  - 1. Building Data Pipelines
  - 2. Considerations When Building Data Pipelines
  - 3. Timeliness
  - 4. Reliability
  - 5. High and Varying Throughput
  - 6. High and Varying Throughput (Contd.)
  - 7. Data Formats
  - 8. Data Formats (Contd.)
  - 9. Transformations
  - 10. Transformations (Contd.)
  - 11. Security
  - 12. Failure Handling
  - 13. Coupling and Agility
  - 14. Ad-hoc Pipelines
  - 15. Loss of Metadata
  - 16. Extreme Processing
  - 17. Kafka Connect Versus Producer and Consumer
  - 18. Kafka Connect Versus Producer and Consumer (Contd.)
  - 19. Summary
- 4. Integrating Kafka with Other Systems
  - 1. Introduction to Kafka Integration
  - 2. Kafka Connect
  - 3. Kafka Connect (Contd.)
  - 4. Running Kafka Connect
  - 5. Key Configurations for Connect workers:
  - 6. Kafka Connect API
  - 7. Kafka Connect Example – File Source
  - 8. Kafka Connect Example – File Sink

9. Kafka Connector Example – MySQL to Elasticsearch
10. Kafka Connector Example – MySQL to Elasticsearch (Contd.)
11. Write the data to Elasticsearch
12. Building Custom Connectors
13. Kafka Connect – Connectors
14. Kafka Connect - Tasks
15. Kafka Connect - Workers
16. Kafka Connect – Workers (Contd.)
17. Kafka Connect - Converters and Connect's data model
18. Kafka Connect - Offset management
19. Alternatives to Kafka Connect
20. Alternatives to Kafka Connect (Contd.)
21. Introduction to Storm
22. Other Components of Spark
23. Integrating Storm with Kafka
24. Integrating Storm with Kafka – Sample Code
25. Integrating Storm with Kafka
26. Introduction to Hadoop
27. Hadoop Components
28. Integrating Hadoop with Kafka
29. Hadoop Consumers
30. Hadoop Consumers (Contd.)
31. Hadoop Consumers (Contd.)
32. Hadoop Consumers – Produce Topic
33. Hadoop Consumers – Fetch Generated Topic
34. Summary
5. Kafka Security
  1. Kafka Security
  2. Encryption and Authentication using SSL
  3. Encryption and Authentication using SSL (Contd.)
  4. Configuring Kafka Brokers
  5. Configuring Kafka Brokers – Optional Settings
  6. Authenticating Using SASL
  7. Authenticating Using SASL – Configuring Kafka Brokers
  8. Authenticating Using SASL – Configuring Kafka Brokers (Contd.)
  9. Authorization and ACLs
  10. Authorization and ACLs (Contd.)
  11. Securing a Running Cluster
  12. Securing a Running Cluster (Contd.)
  13. ZooKeeper Authentication
  14. ZooKeeper Authentication (Contd.)
  15. Summary
6. Monitoring Kafka
  1. Introduction
  2. Metrics Basics
  3. JVM Monitoring
  4. Garbage collection

5. Garbage Collection (Contd.)
6. Java OS monitoring
7. OS Monitoring
8. OS Monitoring (Contd.)
9. Kafka Broker Metrics
10. Under-Replicated Partitions
11. Active controller count
12. Request handler idle ratio
13. Intelligent Thread Usage
14. All topics bytes in
15. All topics bytes out
16. All topics messages in
17. Partition count
18. Leader count
19. Offline partitions
20. Request metrics
21. Request Metrics (Contd.)
22. Logging
23. Logging (Contd.)
24. Client Monitoring
25. Producer Metrics
26. Overall producer metrics
27. Overall producer metrics (Contd.)
28. Per-broker and per-topic metrics
29. Consumer Metrics
30. Fetch Manager Metrics
31. Per-broker and per-topic metrics
32. Consumer coordinator metrics
33. Quotas
34. Quotas (Contd.)
35. Lag Monitoring
36. Lag Monitoring (Contd.)
37. End-to-End Monitoring
38. Summary
7. Kafka Operational Aspects
  1. Kafka Administration
  2. Kafka Cluster Management
  3. Kafka Replica Distribution
  4. Partition Reassignment
  5. Partition Reassignment (Contd.)
  6. Kafka Topic Management
  7. Kafka Topic Management (Contd.)
  8. Kafka Cluster Mirroring
  9. Kafka Cluster Mirroring (Contd.)
  10. Integration with Other Tools
  11. Summary

## Class Materials

Each student will receive a comprehensive set of materials, including course notes and all the class examples.

### Class Prerequisites

Experience in the following *would be useful* for this Microservices class:

- Basic understanding of messaging, cloud, development, architecture, and virtualization would be beneficial.