

Apache Kafka



ZETLAN TECHNOLOGIES

Online Course

www.zetlantech.com

Help Desk: +91 8680961847



COURSE MODULES

Module 1: Introduction to Apache Kafka

- What is Apache Kafka?
- Kafka Architecture Overview
- Use Cases of Kafka
- Components of Kafka (Producers, Consumers, Brokers, Topics, Partitions, Zookeeper)
- Kafka vs. Traditional Messaging Systems

Module 2: Kafka Installation & Setup

- Installing Kafka on Windows/Linux/Mac
- Setting up Zookeeper and Kafka Brokers
- Configuring Kafka Properties
- Running Kafka using CLI Commands
- Basic Kafka Commands (Creating Topics, Producing & Consuming Messages)

Module 3: Kafka Producers and Consumers

- Understanding Producers & Consumers
- Producer API (Key, Partitioner, Acknowledgment)
- Consumer API (Polling, Commit Offsets, Consumer Groups)
- Kafka Serialization & Deserialization (String, JSON, Avro)
- Consumer Offsets & Auto Commit

Module 4: Kafka Topics & Partitions

- Topic Creation & Management
- Partitioning Strategies
- Replication in Kafka
- Log Retention & Cleanup Policies
- High Availability & Fault Tolerance

Module 5: Kafka Streams & Real-time Processing

- Introduction to Kafka Streams API
- Stream Processing vs. Batch Processing
- Stateless & Stateful Transformations
- Windowing & Aggregations
- Interactive Queries



Module 6: Kafka Connect & Integrations

- What is Kafka Connect?
- Source & Sink Connectors
- Connecting Kafka with Databases (MySQL, PostgreSQL)
- Connecting Kafka with Big Data Systems (Hadoop, Spark, Elasticsearch)

Module 7: Kafka Security & Authentication

- SSL/TLS Encryption
- SASL Authentication
- Role-Based Access Control (RBAC)
- Securing Kafka with ACLs

Module 8: Kafka Monitoring & Performance Tuning

- Monitoring Kafka Clusters (JMX, Prometheus, Grafana)
- Optimizing Kafka Performance (Batch Size, Compression, Acks)
- Troubleshooting Kafka Issues

Module 9: Advanced Kafka Topics

- Exactly Once Semantics
- Rebalancing Strategies
- Schema Registry & Avro
- Kafka in Cloud (AWS MSK, Azure Event Hubs, Google Pub/Sub)
- Multi-Cluster Kafka & Cross-Data Center Replication

Module 10: Hands-on Projects

- Implementing a Real-Time Data Pipeline
- Kafka Integration with Spark for Stream Processing
- Building a Kafka-based Microservices Architecture
- Data Ingestion using Kafka & Kafka Connect

