Kafka Training

Duration: **5 Day**

BİLGİNÇ IT Academy
info@bilginc.com
bilginc.com

Overview

Apache Kafka Certification Training is designed to help you become a successful**Kafka Big Data Developer**. The course covers fundamental concepts such as Kafka Cluster and Kafka API, along with advanced topics including Kafka Connect, Kafka Streams, and integrations with Hadoop, Storm, and Spark.

Prerequisites

Basic knowledge of Java is required

Who Should Attend

This course is ideal for professionals working with Big Data and messaging systems:

- Developers aiming to become Kafka Big Data Developers
- Testing professionals working with messaging systems
- Big Data architects integrating Kafka into ecosystems
- Project managers handling messaging-based projects
- System administrators targeting Kafka expertise

What You Will Learn

After completing this training, you will be able to:

- Understand Kafka architecture and components
- Set up Kafka clusters with Hadoop and YARN
- Integrate Kafka with Spark and Storm
- Design high-throughput messaging systems
- Produce and consume real-time data streams
- Work with Kafka APIs and Kafka Streams
- Build real-world streaming applications

Outline

Introduction to Big Data and Apache Kafka

Goal: In this module, you will understand where Kafka fits in the Big Data space, and Kafka Architecture. In addition, you will learn about Kafka Cluster, its Components, and how to Configure a

BİLGİNÇ
IT Academy
info@bilginc.com
bilginc.com

Cluster

Skills:

- Kafka Concepts
- Kafka Installation
- Configuring Kafka Cluster

Objectives: At the end of this module, you should be able to:

- Explain what is Big Data
- Understand why Big Data Analytics is important
- Describe the need of Kafka
- Know the role of each Kafka Components
- Understand the role of ZooKeeper
- Install ZooKeeper and Kafka
- Classify different type of Kafka Clusters
- Work with Single Node-Single Broker Cluster

Topics:

- Introduction to Big Data
- Big Data Analytics
- Need for Kafka
- What is Kafka?
- Kafka Features
- Kafka Concepts
- Kafka Architecture
- Kafka Components
- ZooKeeper
- Where is Kafka Used?
- Kafka Installation
- Kafka Cluster
- Types of Kafka Clusters
- Configuring Single Node Single Broker Cluster

Hands on:

- Kafka Installation
- Implementing Single Node-Single Broker Cluster

Kafka Producer

Goal: Kafka Producers send records to topics. The records are sometimes referred to as Messages. In this Module, you will work with different Kafka Producer APIs.

Skills:

- Configure Kafka Producer
- Constructing Kafka Producer

- Kafka Producer APIs
- Handling Partitions

Objectives: At the end of this module, you should be able to:

- Construct a Kafka Producer
- Send messages to Kafka
- Send messages Synchronously & Asynchronously
- Configure Producers
- Serialize Using Apache Avro
- Create & handle Partitions

Topics:

- Configuring Single Node Multi Broker Cluster
- Constructing a Kafka Producer
- Sending a Message to Kafka
- Producing Keyed and Non-Keyed Messages
- Sending a Message Synchronously & Asynchronously
- Configuring Producers
- Serializers
- Serializing Using Apache Avro
- Partitions

Hands On:

- Working with Single Node Multi Broker Cluster
- Creating a Kafka Producer
- Configuring a Kafka Producer
- Sending a Message Synchronously & Asynchronously

Kafka Consumer

Goal: Applications that need to read data from Kafka use a Kafka Consumer to subscribe to Kafka topics and receive messages from these topics. In this module, you will learn to construct Kafka Consumer, process messages from Kafka with Consumer, run Kafka Consumer and subscribe to Topics**Skills:**

- Configure Kafka Consumer
- Kafka Consumer API
- Constructing Kafka Consumer

Objectives: At the end of this module, you should be able to:

- Perform Operations on Kafka
- Define Kafka Consumer and Consumer Groups
- Explain how Partition Rebalance occurs
- Describe how Partitions are assigned to Kafka Broker
- Configure Kafka Consumer
- Create a Kafka consumer and subscribe to Topics
- Describe & implement different Types of Commit

BİLGİNÇ IT Academy

info@bilginc.com

bilginc.com

- Deserialize the received messages

Topics:

- Consumers and Consumer Groups
- Standalone Consumer
- Consumer Groups and Partition Rebalance
- Creating a Kafka Consumer
- Subscribing to Topics
- The Poll Loop
- Configuring Consumers
- Commits and Offsets
- Rebalance Listeners
- Consuming Records with Specific Offsets
- Deserializers

Hands-On:

- Creating a Kafka Consumer
- Configuring a Kafka Consumer
- Working with Offsets

Kafka Internals

Goal: Apache Kafka provides a unified, high-throughput, low-latency platform for handling real-time data feeds. Learn more about tuning Kafka to meet your high-performance needs.

Skills:

- Kafka APIs
- Kafka Storage
- Configure Broker

Objectives: At the end of this module, you should be able to:

- Understand Kafka Internals
- Explain how Replication works in Kafka
- Differentiate between In-sync and Out-off-sync Replicas
- Understand the Partition Allocation
- Classify and Describe Requests in Kafka
- Configure Broker, Producer, and Consumer for a Reliable System
- Validate System Reliabilities
- Configure Kafka for Performance Tuning

Topics:

- Cluster Membership
- The Controller
- Replication

- Request Processing
- Physical Storage
- Reliability
- Broker Configuration
- Using Producers in a Reliable System
- Using Consumers in a Reliable System
- Validating System Reliability
- Performance Tuning in Kafka

Hands On:

- Create topic with partition & replication factor 3 and execute it on multi-broker cluster
- Show fault tolerance by shutting down 1 Broker and serving its partition from another broker

Kafka Cluster Architectures & Administering Kafka

Goal: Kafka Cluster typically consists of multiple brokers to maintain load balance. ZooKeeper is used for managing and coordinating Kafka broker. Learn about Kafka Multi-Cluster Architectures, Kafka Brokers, Topic, Partitions, Consumer Group, Mirroring, and ZooKeeper Coordination in this module.

Skills:

- Administer Kafka

Objectives: At the end of this module, you should be able to

- Understand Use Cases of Cross-Cluster Mirroring
- Learn Multi-cluster Architectures
- Explain Apache Kafka's MirrorMaker
- Perform Topic Operations
- Understand Consumer Groups
- Describe Dynamic Configuration Changes
- Learn Partition Management
- Understand Consuming and Producing
- Explain Unsafe Operations

Topics:

- Use Cases - Cross-Cluster Mirroring
- Multi-Cluster Architectures
- Apache Kafka's MirrorMaker
- Other Cross-Cluster Mirroring Solutions
- Topic Operations
- Consumer Groups
- Dynamic Configuration Changes
- Partition Management
- Consuming and Producing
- Unsafe Operations

Hands on:

- Topic Operations
- Consumer Group Operations
- Partition Operations
- Consumer and Producer Operations

Kafka Monitoring and Kafka Connect

Goal: Learn about the Kafka Connect API and Kafka Monitoring. Kafka Connect is a scalable tool for reliably streaming data between Apache Kafka and other systems.

Skills:

- Kafka Connect
- Metrics Concepts
- Monitoring Kafka

Objectives: At the end of this module, you should be able to:

- Explain the Metrics of Kafka Monitoring
- Understand Kafka Connect
- Build Data pipelines using Kafka Connect
- Understand when to use Kafka Connect vs Producer/Consumer API
- Perform File source and sink using Kafka Connect
- **Topics:**
- Considerations When Building Data Pipelines
- Metric Basics
- Kafka Broker Metrics
- Client Monitoring
- Lag Monitoring
- End-to-End Monitoring
- Kafka Connect
- When to Use Kafka Connect?
- Kafka Connect Properties

Hands on:

- Kafka Connect

Kafka Stream Processing

Goal: Learn about the Kafka Streams API in this module. Kafka Streams is a client library for building mission-critical real-time applications and microservices, where the input and/or output data is stored in Kafka Clusters.

Kafka Training

Duration: **5 Day**

BİLGİNÇ
IT Academy
info@bilginc.com
bilginc.com

Skills:

- Stream Processing using Kafka

Objectives:

- At the end of this module, you should be able to,
- Describe What is Stream Processing
- Learn Different types of Programming Paradigm
- Describe Stream Processing Design Patterns
- Explain Kafka Streams & Kafka Streams API

Topics:

- Stream Processing
- Stream-Processing Concepts
- Stream-Processing Design Patterns
- Kafka Streams by Example
- Kafka Streams: Architecture Overview

Hands on:

- Kafka Streams
- Word Count Stream Processing

Integration of Kafka With Hadoop, Storm and Spark

Goal: In this module, you will learn about Apache Hadoop, Hadoop Architecture, Apache Storm, Storm Configuration, and Spark Ecosystem. In addition, you will configure Spark Cluster, Integrate Kafka with Hadoop, Storm, and Spark.

Skills:

- Kafka Integration with Hadoop
- Kafka Integration with Storm
- Kafka Integration with Spark

Objectives: At the end of this module, you will be able to:

- Understand What is Hadoop
- Explain Hadoop 2.x Core Components
- Integrate Kafka with Hadoop
- Understand What is Apache Storm
- Explain Storm Components
- Integrate Kafka with Storm
- Understand What is Spark
- Describe RDDs
- Explain Spark Components
- Integrate Kafka with Spark

Topics:

Kafka Training

Duration: **5 Day**

BİLGİNÇ
IT Academy
info@bilginc.com
bilginc.com

- Apache Hadoop Basics
- Hadoop Configuration
- Kafka Integration with Hadoop
- Apache Storm Basics
- Configuration of Storm
- Integration of Kafka with Storm
- Apache Spark Basics
- Spark Configuration
- Kafka Integration with Spark

Hands On:

- Kafka integration with Hadoop
- Kafka integration with Storm
- Kafka integration with Spark

Integration of Kafka With Talend and Cassandra

Goal: Learn how to integrate Kafka with Flume, Cassandra and Talend.

Skills:

- Kafka Integration with Flume
- Kafka Integration with Cassandra
- Kafka Integration with Talend

Objectives:At the end of this module, you should be able to,

- Understand Flume
- Explain Flume Architecture and its Components
- Setup a Flume Agent
- Integrate Kafka with Flume
- Understand Cassandra
- Learn Cassandra Database Elements
- Create a Keyspace in Cassandra
- Integrate Kafka with Cassandra
- Understand Talend
- Create Talend Jobs
- Integrate Kafka with Talend

Topics:

- Flume Basics
- Integration of Kafka with Flume
- Cassandra Basics such as and KeySpace and Table Creation
- Integration of Kafka with Cassandra
- Talend Basics
- Integration of Kafka with Talend

Hands On:

Kafka Training

Duration: **5 Day**

BİLGİNÇ
IT Academy
info@bilginc.com
bilginc.com

- Kafka demo with Flume
- Kafka demo with Cassandra
- Kafka demo with Talend