

Week 8: Cloud Applications: MapReduce, Spark and Apache Kafka

1. MapReduce: Paradigm, Programming Model, Applications, Scheduling, Fault-Tolerance, Implementation Overview, Examples
2. Introduction to Spark: Resilient Distributed Datasets (RDDs), RDD Operations, Spark applications: Page Rank Algorithm, GraphX, GraphX API, GraphX working
3. Introduction to Kafka: What is Kafka, Use cases for Kafka, Data model, Architecture, Types of messaging systems, Importance of brokers