

BIG DATA AND HADOOP

(ELECTIVE-IV)

(Common to CSE & IT)

Course Code :13CT1136

L	T	P	C
4	0	0	3

Pre requisites: Database Management Systems, Object Oriented Programming Through Java.

Course Educational Objectives:

This course introduces the fundamental concepts of cloud and lays a strong foundation of Apache Hadoop (Big data framework).

- ❖ The HDFS file system, MapReduce frameworks are studied in detail.
- ❖ Hadoop tools like Hive, and Hbase, which provide interface to relational databases, are also covered as part of this course work.
- ❖ Analyzing data with unix tools
- ❖ Sorting. Map side and Reduce side joins.
- ❖ Implementation. Java and Mapreduce clients

Course Outcomes:

At the end of the course the student will be able to

- ❖ Understand the fundamentals of Big cloud and data architectures.
- ❖ Understand HDFS file structure and Mapreduce frameworks, and use them to solve complex problems, which require massive computation power.
- ❖ Use relational data in a Hadoop environment, using Hive and Hbase tools of the Hadoop Ecosystem..
- ❖ Understand The Hive Shell.
- ❖ Understand the Comparison with traditional databases.

UNIT-I**(12 Lectures)**

Introduction to Big Data. What is Big Data. Why Big Data is Important. Meet Hadoop. Data. Data Storage and Analysis. Comparison with other systems. Grid Computing. A brief history of Hadoop. Apache hadoop and the Hadoop EcoSystem. Linux refresher; VMWare Installation of Hadoop.

UNIT-II**(12 Lectures)**

The design of HDFS. HDFS concepts. Command line interface to HDFS. Hadoop File systems. Interfaces. Java Interface to Hadoop. Anatomy of a file read. Anatomy of a file write. Replica placement and Coherency Model. Parallel copying with distcp, Keeping an HDFS cluster balanced.

UNIT-III**(12 Lectures)**

Introduction. Analyzing data with unix tools. Analyzing data with hadoop. Java MapReduce classes (new API). Data flow, combiner functions, Running a distributed MapReduce Job.

Configuration API. Setting up the development environment. Managing configuration. Writing a unit test with MRUnit. Running a job in local job runner. Running on a cluster.Launching a job. The MapReduce WebUI.

UNIT-IV**(12 Lectures)**

Classic Mapreduce. Job submission. Job Initialization. Task Assignment. Task execution .Progress and status updates. Job Completion. Shuffle and sort on Map and reducer side. Configuration tuning. MapReduce Types. Input formats. Output formats ,Sorting. Map side and Reduce side joins.

UNIT-V**(12 Lectures)**

The Hive Shell. Hive services. Hive clients. The meta store. Comparison with traditional databases. HiveQL. Hbasics. Concepts. Implementation. Java and Mapreduce clients. Loading data, web queries.

TEXT BOOKS:

1. Tom White, Hadoop, “*The Definitive Guide*”, 3rd Edition, O’Reilly Publications, 2012.

2. Dirk deRoos, Chris Eaton, George Lapis, Paul Zikopoulos, Tom Deutsch , “*Understanding Big Data: Analytics for Enterprise Class Hadoop and Streaming Data*”, McGraw-Hill Osborne Media; 1 edition, 2011

WEB REFERENCES:

1. <http://www.cloudera.com/content/cloudera-content/cloudera-docs/HadoopTutorial/CDH4/Hadoop-Tutorial.html>
2. https://www.ibm.com/developerworks/community/blogs/Susan_Visser_Editionntry/flash_book_understanding_big_data_analytics_for_enterprise_class_hadoop_and_streaming_data?lang=en

