

Hadoop and its Role in Facebook: An Overview

Divya Ramadoss

Research Scholar, Department of Computer Science
Prist University
Thanjavur, India
divyarmdss@gmail.com

Chandra P

Department of Research – Ph.D. Computer Science
Tirupur Kumaran College for Women
Tirupur, India
chandragunasekar@gmail.com

Abstract— One of the challenges in storing and processing the data and using the latest internet technologies has resulted in large volumes of data. The technique to manage this massive amount of data and to pull out the value, out of this volume is collectively called Big data. Over the recent years, there has been a rising interest in big data for social media analysis. Online social media have become the important platform across the world to share information. Facebook, one of the largest social media site receives posts in millions every day. One of the efficient technologies that deal with the Big Data is Hadoop. Hadoop, for processing large data volume jobs uses MapReduce programming model. This paper provides a survey on Hadoop and its role in facebook and a brief introduction to HIVE.

Keywords-Bigdata, Hadoop, HDFS, MapReduce, Hive

I. INTRODUCTION

In recent years, the growth of social networking is vast and a lot of data is generated. The generated data has a wide variety of use mainly for business intelligence. This huge volume of data has created a new field in data processing, called Big data and it is marked among top ten strategic technologies [1]. Big data means really a huge data, it is a collection of large datasets that cannot be processed using conventional computing techniques. Big data is not just a data, relatively it has become a complete subject, which involves various tools, techniques, and frameworks.[2]. The processing of this massive data can be completed using distributed computing and parallel processing mechanism. Hadoop [3] is an open-source framework that allows to store and process big data in a distributed environment across a group of computers by means of simple programming models. It is designed to range up from single servers to thousands of machines, each offering local calculation and storage.

In section II we discuss in more detail about Hadoop's components HDFS and MapReduce. In section III we discuss Hadoop Ecosystem. Section IV discusses Hadoop's role in Facebook. Further section V gives a brief overview of Hive. Finally, section VI concludes the paper followed by the references.

II. HADOOP AT A GLANCE

Apache Hadoop is a framework planned to work with the huge amount of datasets which are much larger in amount than the normal systems can handle. Hadoop distributes this information across a set of machines. The actual function of Hadoop hails from the truth its ability should versatile with hundreds or thousands of computers each containing several processor cores. Most of the giant enterprises believe that within a few years more than half of the world's data will be

stored in Hadoop [4]. Furthermore, Hadoop joint with Virtual Machine gives more realistic outcomes. It consists of i) Hadoop Distributed File System (HDFS): a scattered file system to achieve storage and fault tolerance and ii) Hadoop MapReduce a powerful parallel programming model which progress the vast quantity of data by means of distributed computing across the clusters.

HDFS- Hadoop Distributed file system

Google File System (GFS) is a versatile distributed file system (DFS) developed by Google Inc. and to hold Google's escalating data processing requirements. It offers fault tolerance, reliability, scalability, availability, and performance to large networks and connected nodes. It is optimized to hold Google's different information use and storage needs, such as its search engine, which generates enormous amounts of data that must be stored. It is also called as GoogleFS. Files are divided into chunks of 64 MB and are usually attach to or read and only extremely rarely overwritten or shrunk. In contrast with traditional file systems, GFS is designed and enhanced to run on data centers to provide enormously high data throughputs, low latency and to handle individual server failures. HDFS[5] holds large files across multiple machines.It provides easier access to files. These files are stored in a superfluous fashion to release the system from data losses in case of breakdown.It also makes applications accessible to parallel processing.

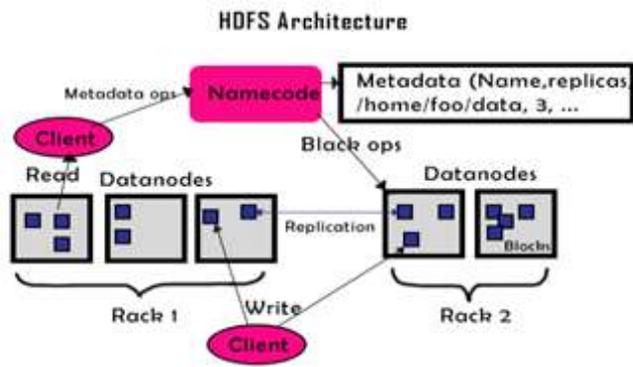


Figure 1.HDFS Architecture

Hadoop Distributed File System is composed of master-slave architecture. It consists of Namenode and datanode. Namenode handles the file system namespace and also executes system operations like as renaming, closing, and opening files and directories. Datanode performs read-write operations as per user request. It also performs operations such as block creation, deletion, and replication according to the instructions of the name node. The block size is 64MB and is replicated into 3 copies. The second copy is stored on the local rack itself while the third on the remote rack. A rack is nothing but just a collection of data nodes.

Hadoop MapReduce Overview

MapReduce[6] initially developed by Google. It is a programming model and it is a major element for processing and generating big data sets with a parallel, distributed algorithm on a cluster. The two important functions in MapReduce are Map and Reduce. The raw data is given as input to the Map phase which process and gives accurate results. The input is stored in HDFS in the form of files. This file is passed to the mapper function line by line. It processes the data and creates some small portion of data. Reduce stage is the combination of the Shuffle stage and the Reduce stage. The main aim of this stage is to process the data that comes from the mapper and it produces a new set of output, which will be stored in the HDFS.

The Hadoop MapReduce is a framework that process vast amount of data in parallel on a large cluster. It comprises of Master node (JobTracker) and Worker nodes(TaskTrackers). The JobTracker is loaded with jobs as inputs which in turn converts them into a number of Map and Reduce tasks. These tasks are allocated to the TaskTrackers and examine the execution of these tasks and at the end when all tasks are finished. The user is intimated about job completion. HDFS offers fault tolerance and reliability by storing and replicating the inputs and outputs.

III. HADOOP ECOSYSTEM

The Hadoop Framework also contains different projects which include the following

- Apache HBase: A column-oriented, non-relational distributed key/value data store which is built to run on top of the HDFS platform. It is designed to store structured data in tables that could have billions of rows and millions of columns in distributed compute clusters.
- Apache Hive: A data warehouse infrastructure built on top of Hadoop that addresses how data is summarized, query and analyze large datasets stored in Hadoop files. It is not designed to offer real-time queries, but it can support text files, and sequence files.
- Apache Pig: It is a procedural language provides a high-level parallel mechanism for the programming of MapReduce jobs to be executed on Hadoop clusters. It includes Pig Latin, which is a scripting language. It is geared towards processing of data in a parallel manner.
- Apache Zookeeper: It is an Application Program Interface (API) that provides high-performance coordination service for distributed applications.
- Apache Sqoop: It is a command-line interface that facilitates moving bulk data from Hadoop into relational databases and other structured data stores.
- Apache Flume: A scattered service for gathering, cumulating, and moving a large amount of data.

IV. FACEBOOK: DISCARDED MYSQL

With more than 2.1 billion monthly active users [7], Facebook is one of the largest social media sites receives posts in millions in every day. It ends up gathering large amounts of data. The storage and processing these data was the biggest challenge, in the aspects to improve the user experience. It is achieved by using an open-source project called Hadoop. Facebook has largest Hadoop cluster[8] with 2300 nodes. The single HDFS cluster[9] comprising of 21PB of storage, 2000 machines with 12TB capacity, 32GB RAM and 15 MapReduce tasks per machine. The features of the Facebooks Datawarehousing Hadoop cluster are 12TB of compressed data, 800TB of compressed data is scanned and 25000 MapReduce jobs per day. It also comprised of 65 millions of files in HDFS.

Facebook developed its first application is Facebook Messages, which is based on Hadoop. The major reasons for opting for Hadoop architecture instead of RDBMS environment is that High write throughput, massive datasets, unpredictable growth in data.

HBase is a part of Hadoop and it uses Hadoop File System and Facebook chose it in part because it has experience scaling and

debugging HDFS. The new Facebook messaging uses HBase to store the data and metadata for the messages as well as the indices needed to search among these messages. Even before the new messaging system was moved out, Facebook was shuffling with about 14TB messages a month and 11TB chat messages. HBase doesn't store everything, instead, Haystack handles attachments and overly large messages. Facebook also picked HBase because it deals with automatic failover. It simply means when the data is broken down into pieces and distributed across a cluster and this completely separate data on each machine. When one server fails, the work can be taken from other machines. The newest Facebook message combines existing ones with chats, SMS, and e-mail. It also has the threading model stores the messages for each user. There are more than 350 million users sending messages person to a person per month. With this rate, the volume of ingested data would be very large from day one and it grows. This can be easily dealt with HBase. One of the major requirement is that messages would not be deleted unless explicitly done by the user. Messages are not read always, as such messages must be made available at all times with minimal latency. This is not possible in a system like MySQL but in HBase.

V HIVE – AN NEW APPROACH TO HADOOP

Apache Hive is data warehousing tool work on the top of Hadoop. It provides SQL like language called HQL (Hive Query Language) to perform analysis on data which may be structured or unstructured. The actual query processing is achieved by the Hive compiler by converting the HQL queries to Map Reduce jobs. The developer need not write difficult Map Reduce programs. Hive supports Data Definition Language, Data Manipulation Language, and User-Defined Functions.

Apache Hive takes the benefit of both SQL systems and Hadoop framework. It is mostly used in the fields of Data warehousing and Ad- hoc analysis.

Features

- Indexing for accelerated processing
- Different storage types support such as plain text, RCFile, HBase, ORC, and others.
- Metadata storage in RDBMS will reduce the time taken to perform semantic checks during execution of queries.
- SQL like queries get implicitly converted into MapReduce, Tez or Spark jobs.
- Built in user defined functions to manipulate strings, dates, and other data mining tools.

Hive Architecture

Fig 2 depicts the Hive architecture and the process flow of a query submitted to Hive

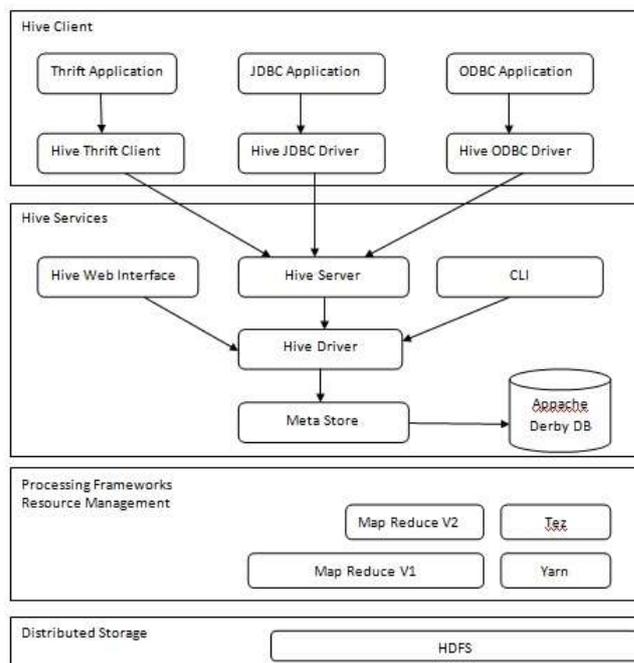


Figure 2. Hive Architecture

Fig 2. depicts the components of Hive are Hive Clients, Hive Services, Processing frameworks resource management and distributed storage. Hive supports the languages like Java, C++, Python etc. By using the drivers like JDBC, ODBC, and Thrift. So one can write hive client applications written in a language of their own. Hive provides many services like Hive Command Line Interface, Hive Web Interfaces, Hive Server, Hive Driver and Metastore.

Hive Data Model

Data in Hive can be divided into three types: On a granular level, the types are Table, Partition, and Bucket. Tables are same like in RDBMS[10] and it is likely to perform filter, project, join and union. There are two types of tables available, Managed and external tables. The data in the managed table are managed by Hive itself. For external tables, Hive is not responsible for the table data. Hive separates tables into partitions for grouping similar type of data based on a column or partition key. The partitions will be identified by using the partition key and this allows having a faster query on the partitions.

Each partitioned or unpartitioned table may be separated into Buckets based on the hash function of a column in a table. So each bucket is just a file in the partition directory or the table directory. Performing bucketing to a partition is for two reasons. One is a map side join requires the data belonging to a

unique join key to be present in the same partition. The next reason is Bucketing makes the sampling process more efficient and allowing decreased query processing time.

Advantages of Hive:

- Hive reduces writing complex map reduce programs. So, people who are not from programming background can use Hive to write simple commands.
- It is very much Extensible and Scalable to cope up with the growing volume and variety of data, without affecting the performance of the system.
- It is an efficient Extract, Transform and Load Tool.
- Hive supports different types of languages to write Hive client applications. So, the programmer can choose by his / her own choice of language to write the client application.
- The Meta data information is stored in an RDBMS. So, it significantly reduces the time to perform semantic checks during query execution.

Hive in FaceBook

The data processing infrastructure of face book before 2008, was built around a data warehouse based on some commercial Relational Database Management System. At that time it was sufficient to process the data needs of face book. But, when the data needs started growing very fast, the challenge of managing and processing the data increases. Actually, the data scaled from 15 TB data set to 2 PB data from 2007 to 2009. Many Facebook products like Audience Insights, Facebook Lexicon, Facebook Ads, etc are involving analysis of data. To overcome this problem Face book needs a solution and decided with Hadoop framework.

The ability of SQL to suffice for most of the analytic requirements and the scalability of Hadoop gave birth to **Apache Hive** that allows performing SQL like queries on the data present in HDFS. Later, the Hive project was open sourced in August' 2008 by Facebook and is freely available as Apache Hive today.

VI CONCLUSION

Big data is in greater demand among the giant companies nowadays. It gives a lot of benefits to the business. In this paper, we have discussed the technology that handles the Big Data is Hadoop. It provides a framework for large scale parallel processing using a distributed file system and the map-reduce programming paradigm. The paper talks about the use of Hadoop and HBase at Facebook. Due to the restrictions in MapReduce, the concept of Hive came into existence. Hive is an open source technology. Hadoop and Hive is used for data processing in Facebook.

REFERENCES

- [1] Orlando, "Gartner identifies the top 10 strategic technologies for 2012", <http://www.gartner.com/it/page.jsp?id=1826214>.
- [2] Bigdata, https://www.tutorialspoint.com/hadoop/hadoop_big_data_overview.htm
- [3] Hadoop, <https://www.tutorialspoint.com/hadoop/index.htm>.
- [4] Kala Karun. A, Chitharanjan.K, —"A Review on Hadoop – HDFS Infrastructure Extensions", Proceedings of 2013 IEEE Conference on Information and Communication Technologies (ICT 2013).
- [5] Hdfs, https://www.tutorialspoint.com/hadoop/hadoop_hdfs_overview.htm
- [6] MapReduce, <https://en.wikipedia.org/wiki/MapReduce>
- [7] Zephoria, "The Top 20 valuable Facebook statistics—updated" <https://zephoria.com/top-15-valuable-facebook-statistics>
- [8] <https://www.quora.com/What-are-some-of-the-largest-Hadoop-clusters-to-date>
- [9] <http://hadoopblog.blogspot.in/2010/05/facebook-has-worlds-largest-hadoop.html>
- [10] Ashish Bakshi, "https://www.edureka.co/blog/hivetutorial/#story_of_hive", eureka.



Figure 3. Evolution of Hive in FaceBook