

Federated Cloud-based Big Data Platform in Telecommunications

Chao Deng

China Mobile Research Institute
Beijing 100053, China

dengchao@chinamobile.com

Yujian Du

China Mobile Research Institute
Beijing 100053, China

duyujian@chinamobile.com

Ling Qian

China Mobile Research Institute
Beijing 100053, China

qianling@chinamobile.com

Zhiguo Luo

China Mobile Research Institute
Beijing 100053, China

luozhiguo@chinamobile.com

Meng Xu

China Mobile Research Institute
Beijing 100053, China

qianling@chinamobile.com

Shaoling Sun

China Mobile Research Institute
Beijing 100053, China

sunshaoling@chinamobile.com

ABSTRACT

China Mobile is the biggest telecommunication operator in the world, with more than 600 million customers and an ever increasing information technology (IT). To provide better service to 600 million customers and reduce the cost of IT systems, China Mobile adopted a centralized IT strategy based on cloud computing. The big data issue becomes the most significant challenge to the cloud computing based China Mobile IT structure. This paper presents the China Mobile's big data platform based on the cloud. This platform integrates the big data storage, the development and deployment of big data ETL (Extract, Transfer, Load) and DM (Data Mining) into a unified framework. This big data analysis platform can effectively support the analytical tasks in telecommunications, but it can also help China Mobile provide public cloud computing service. In this paper, we introduce the detailed architecture of China Mobile's platform and discuss its performance.

Categories and Subject Descriptors

C.2.4 [Distributed Systems]: Distributed applications.

General Terms

Algorithms, Performance, Design, Experimentation

Keywords

Telecommunication; big data; Cloud Computing; Hadoop; SaaS

1. INTRODUCTION

China Mobile Communication Corporation (CMCC) takes a hierarchical structure based on branches. In fact, the China Mobile has 31 branches, where a province is an individual operator branch, and the 31 branches constructed IT infrastructure and applications individually. The 31 branches upload their business

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Workshop on Cloud Services, Federation, and the 8th Open Cirrus Summit, , September 21 2012, San Jose, CA, USA.
Copyright 2012 ACM 978-1-4503-1267-7/12/09 \$10.00.

data to the headquarters by various periods.

For data-intensive business companies, the increase of user scale and the rising of business applications have brought huge challenges to the IT infrastructure and information processing system. For example, the storage volume of Service Support system at CMCC has reached 8000TB at the beginning of 2012, and the collected data for business analysis reached additional 7000TB. In CMCC, a middle level branch has more than 20 million users, and its CDR (Calling Detail Record) can be 12~16TB in 12 months, and its signal data can be 1TB per day. To provide efficient decision-making proposal, reporting, the OLAP data, and the marketing strategy, the traditional data analysis systems confront both the storage and the computing bottlenecks. The traditional data analysis systems commonly are constructed in a centralized infrastructure based on a few expensive UNIX servers, and then many analysis applications are developed using database or data warehouse products installed on UNIX servers. However, the centralized infrastructure results in low scalability and high cost, due to the massive data to be stored and analyzed. For example, a popular commercial business intelligence system uses clustering algorithms that can only support 1 million user's data for knowledge discovery processing, which lags significantly behind real demand.

China Mobile has total more than 600 million customers. To provide better service to the increasing number of customers within the 31 branches results in great challenges to updating and expanding the IT infrastructure in the branches. Hence, China Mobile adopted a centralized IT strategy based on federated cloud infrastructure to reduce the cost of IT systems. China Mobile has selected 3 provinces with beneficial climate to build the big data centers, and selected centers in 2 provinces as their operations centers. The big data issue becomes the most significant challenge after the centralized strategy is performed by federated cloud infrastructure. The federated cloud consists of large numbers of X86 PC server-based clusters. To reduce the cost of developing various applications on big data, it is required for China Mobile to develop and deploy a unified data analysis platform in the operations centers utilizing federated cloud strategies.

Cloud computing is an emerging technology to help data-intensive companies deal with the large-scale data analysis requirement with low cost and high performance. Cloud computing provides IT supporting infrastructure with flexible scalability, large-scale storage and high performance computing. A typical and successful cloud computing use is demonstrated by Google.

Nowadays, the commercial PC cluster along with the distributed parallel computing environment is widely accepted as a popular solution to implementing the low-cost and high performance target. The combination of Google File System (GFS) [1] and MapReduce programming framework represents the necessary synergy between data distribution and parallel computation. Hadoop [2] is an open source implementation of MapReduce and GFS. Hadoop has been widely selected as an infrastructure to provide cloud computing service, such as Yahoo [3], Cloudera [4] and Taobao [5]. Many large-scale data analysis systems also turn to Hadoop and MapReduce, since Google successfully implemented and demonstrated data analysis applications based on MapReduce [6][7]. The recent Gartner report [8] indicates that the Hadoop-based large-scale data analysis solution is disrupting the traditional data analysis landscape. The open source project Mahout [9] provides many MapReduce style data mining algorithms based on Hadoop. Some data warehouse and BI products support data collection, processing and storage with HDFS (Hadoop distributed file system), such as Teradata [10] and Pentaho [11].

To reach the goal of low-cost and high effectiveness operation, China Mobile Research Institute (CMRI) has started the cloud computing project, named *Big Cloud*. Big Cloud includes a series of systems, supporting IaaS, PaaS and SaaS. The infrastructure of Big Cloud is based on commercial PC server cluster and Hadoop platform. BC-PDM (Big Cloud – Parallel Data Mining) is one of the core systems in SaaS service layer of Big Cloud. There are many successful data analysis applications of CMCC developed

2. BC-PDM

The detailed architecture, functionality, features and performance of BC-PDM are introduced in this section.

2.1 Architecture of BC-PDM

The architecture of BC-PDM system is depicted in Figure 1. We distinguish three layers: the cloud computing cluster layer, the BC-PDM service layer, and large-scale data analysis applications layer.

The role of each layer is explained as follows:

- 1) The Cloud computing platform layer is built on the X86 PC servers cluster. Hadoop is deployed among these PCs to provide the HDFS, MapReduce, and SQL JDBC support. HDFS can provide an efficient and reliable distributed data storage as well as file control service needed by applications involving large databases. MapReduce is an easy-to-use parallel programming paradigm suitable for large scale, data intensive parallel computing applications, which at the same time offers load balancing and fault tolerance. Hive, HBase and HugeTable offers online ad hoc querying through the JDBC API.
- 2) The BC-PDM layer provides the parallel data analysis capabilities. Data load and export module supports variety of data sources, such as URL, database, and file systems. Data Exchange Transfer and Management module supports remote file transfer, access with data in databases, and data

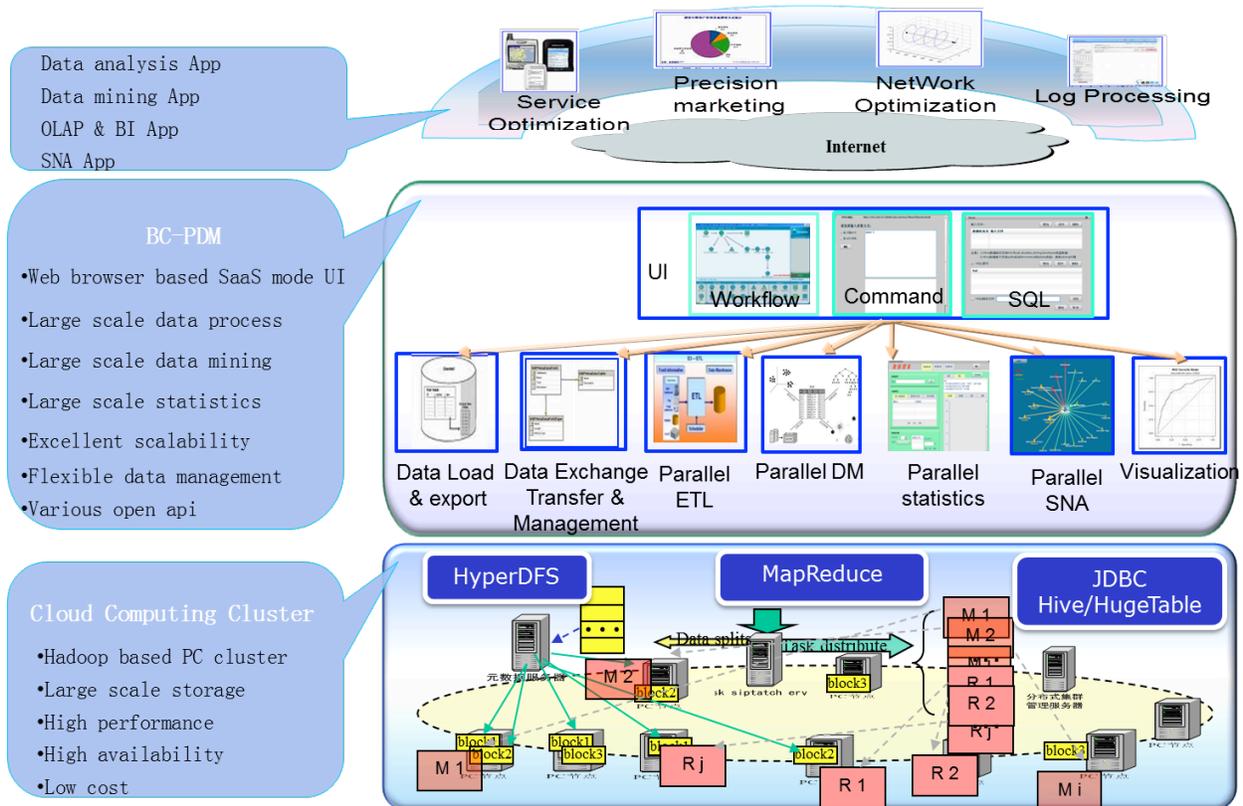


Figure 1. BC-PDM architecture.

on the BC-PDM system.

management remote access control. Parallel ETL and parallel DM modules consist of various large-scale data refining and mining algorithms with MapReduce parallelization. Parallel

statistics module provides large-scale data exploration operation. Parallel SNA module provides large-scale social network analysis algorithms with MapReduce parallelization. Visualization module presenting analysis results to users. A UI module offers workflow, command line, SQL GUI interface to users.

- 3) The Data analysis application layer consists of various ETL, OLAP, BI, Data mining and SNA applications. All these applications could be developed on the BC-PDM.

2.2 Functionality of BC-PDM

To support various large-scale data analysis requirements, the current version of BC-PDM has provided the parallel ETL, data mining, statistics, SNA, and SQL operations. These capabilities covered almost all popular operations used in our large-scale data analysis.

The Parallel ETL provides 7 classes and 45 operations for data extraction and transformation on large-scale data. The 7 classes are cleaning, transfer, count, aggregation, sampling, sets and update class. The cleaning class includes type-checking, primary key and foreign key checking, missing value checking, duplication and max-min value delete. The transfer class includes case-when, internalization, type-change, normalization and anti-normalization, ID-add, field exchange, and PCA. The count class includes column-generation, group-by and basic statistics. The aggregation class includes delete, join, dimension-table join, sort, where and select. The sampling class includes hierarchical sampling, random sampling and dataset split. The sets class includes sets-difference and sets-union. The update class includes normal-update and insert-update.

The parallel data mining algorithms provides 3 classes and 12 algorithms popularly used in data analysis. The 3 classes are classification, clustering and association. The classification algorithms include C45 decision tree, K-NN, NaiveBayes classifier, BP Neural Network, LR logistic regression. The clustering algorithms include K-means, DBSCAN and Clara. The association algorithms include Apriori, FPGrowth, Awfits and sequence association.

The Parallel SNA provides social network detection of individual node and groups. Furthermore, it can also provide group evolvement tracking.

In addition, the Parallel SQL receives user's SQL shell input, and then submits these SQL command to Hive, HBase or HugeTable to execute through JDBC. Thus the original application developed by SQL can easily be adapted to the BC-PDM.

2.3 Features of BC-PDM

As a successful practice of cloud computing based data analysis system, the BC-PDM has several attractive features:

- 1) Web2.0 based SaaS (Software as a Service) mode: allow user to connect to internet and develop data analysis applications on Web-browser. Thus, the user doesn't need to buy analysis software or install.
- 2) MapReduce based parallel algorithms: All the ETL data processing and mining algorithms are developed via MapReduce mode. Thus, the large-scale data computing performance is offered.

- 3) The PC server cluster based infrastructure: all the storage and computing are based on the commercial PC server. Thus the low-cost and high scalability target are reached.
- 4) Open architecture: the BC-PDM software defines open architecture that allows the user to develop new algorithm component and add it to BC-PDM. Thus, the user can easily act as a new capability contributor.
- 5) Various API: each parallel algorithm in the BC-PDM provides Java api, Webservice API, and Command line API. Thus the third-party providers can develop their own applications via various modes. In addition, the various api also provide a good selection of implementing PaaS platform.
- 6) Flexible data management: The data exchange, transfer and management in BC-PDM can help user transfer data from remote to HDFS, load data from DFS into user directory of BC-PDM, and remote control the access of private data and meta information.
- 7) Workflow and scheduler: Allow user design analysis application by drag and drop operation. The workflows between different users could be shared and reused with authority. The scheduler helps user assign existing applications to run under predefined time and conditions.
- 8) CWM and PMML output: the ETL operations and data mining models in BC-PDM are exported as CWM standards and PMML description. Thus the process and model from BC-PDM can be reused in other systems that support the CWM and PMML.

2.4 Performance of BC-PDM

To evaluate the scalability of BC-PDM for large-scale data analysis, the performance of ETL and mining algorithms on PC server cluster is tested. The number of Hadoop based PC nodes is increased from 32 to 64 and 128. And the experimental results are shown in Figure 2 and Figure 3. The experiment results indicate that the ETL processing has excellent scalability; meanwhile, the parallel data mining algorithm has acceptable scalability.

As shown in Figure 2, there are 11 ETL operations with good scalability, whose speedup ratio increases nearly linearly with the number of nodes.

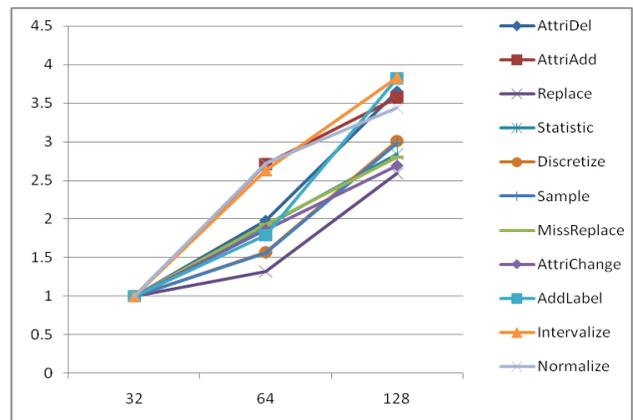


Figure 2. Scalability of ETL operations in BC-PDM.

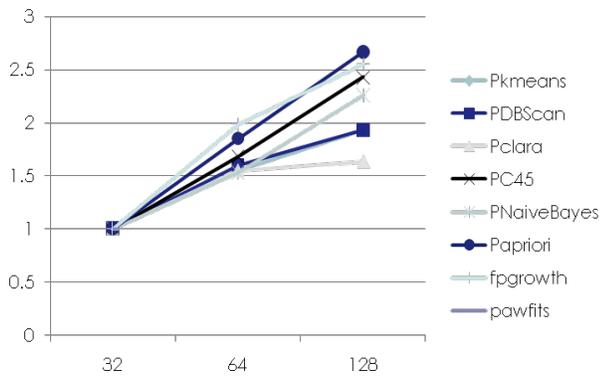


Figure 3. Scalability of mining algorithms in BC-PDM.

Figure 3 indicates that some data mining algorithms (i.e. FP-growth, PC45 and PNaiveBayes) achieved desired scalability; other algorithms also have suitable scalability.

2.5 Performance of China Mobile ETL Application

China Mobile's traditional Business Analysis Support System (BASS) based on the Data Warehouse, and the ETL applications occupy more IT resources than other tasks. Therefore, the ETL applications become a significant factor for increasing the IT cost of BASS. In order to reduce the Data Warehouse burden from ETL applications, we have used BC-PDM to complete the ETL tasks. The GPRS records aggregation is a common ETL application in the China Mobile BASS Data Warehouse. With the fast development of mobile internet and smart phone in China, the GPRS records are significantly increasing. In one China Mobile branch, the task to complete the aggregation of GPRS records of one month includes that the ETL application should group 2TB record data by user id, plan type, access type and brand type dimensions, and then summarize. To output and report the summation value by the end of each month, the traditional application based on our Data Warehouse needs to run 1~2 hours

per day to get one day's value, thus the total time is 30 hours for aggregating one month GPRS records. When we use Groupby and Join components to form the GPRS aggregation application on BC-PDM, which is deployed on 10 PC servers, the total time for one month's TB data is reduced to 2 hours 5 minutes. The difference between traditional 30 hours and BC-PDM's 2 hours shows BC-PDM could significantly improve the performance of practical ETL application in China Mobile BASS. As well as the cost of 10 PC servers used by BC-PDM is only 1/10 of the cost of traditional Data Warehouse and Unix servers.

3. USAGE OF BC-PDM

China Mobile is planning to use BC-PDM provide public cloud computing service.

3.1 Basic Usage of BC-PDM

A snapshot of the web browser based workflow GUI using the mode BC-PDM is shown in Figure 4. Our workflow GUI allows the designing of data analysis application via a graphical layout by allowing drag-and-drop of various operation components. The GUI consists of three functional layouts.

The top area is component bar, and all the parallel ETL, parallel data mining, parallel statistics, visualization, and data management capabilities of BC-PDM are exhibited as components for user selection.

The middle layout is application design area, and the user can drag the desired operation from the component bar, configure parameters, concatenate with other components, and form the workflow consistent with the business process. Additionally, an SQL shell and command shell are allowed to be included in order to input and combine into the workflow as parallel execution threads. Rthrough this mechanism efficiency can be improved and the user can integrate schedulers for the execution of the workflows.

The bottom layout is the state monitor window. The Hadoop cluster state, such as nodes number, HDFS data volume, and nodes logs, are shown. The task progress and MapReduce progress of each job submitted to Hadoop are also captured and

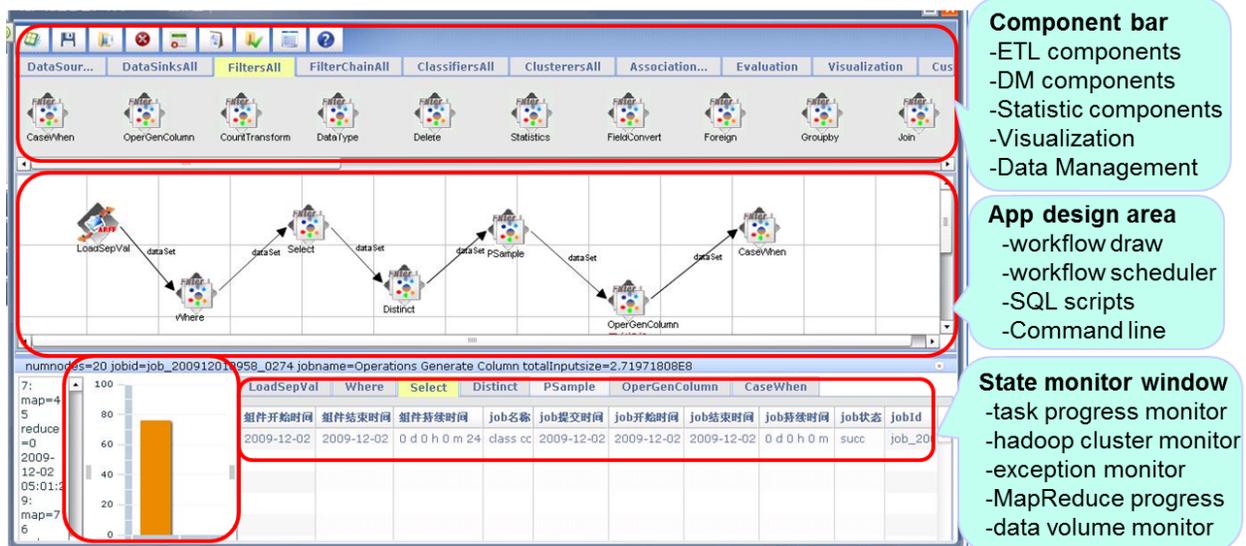


Figure 4. BC-PDM workflow GUI.

shown in the window. Furthermore, exceptions are caught and shown, too.

3.2 Steps in Using Big Cloud by a Public User

A public user should first register as a user of Big Cloud, and then the BC-PDM system could be accessed and authorized. The basic steps of a public registration using BC-PDM system include:

- 1) Prepare the data at the user-end
- 2) Download and install the ftp-client provided by BC-PDM at the user-end
- 3) Upload the prepared data to remote HDFS by using the ftp client, and the data will be stored in the user directory of BC-PDM system.
- 4) Use web browser to design data analysis application at BC-PDM for data processing and mining
- 5) Check the result at BC-PDM, and download the result to user-end via the ftp client

4. SUMMARY AND FUTURE WORK

The BC-PDM platform has successfully been used for large-scale data analysis applications with MapReduce and HDFS. Analysis applications in telecommunication industries are a preliminary proof that BC-PDM is a promising schema of helping data-intensive problems to be solved with low-cost, high performance and high scalability IT target. To broaden the usage of data analysis systems based on the cloud computing infrastructure, the future work in BC-PDM focuses not only on the more functionality features and open issues, but also on a user incentives plan.

The new functionality features of BC-PDM will include complex statistics operations, web mining needed operations, OLAP needed operations, reporting service and BI platform service. The open issues that need to be urgently resolved include the hard disk error and I/O exceptions of HDFS PC server cluster that we observed, and the MapReduce limitation on improving the parallel performance of iterative data mining algorithms.

The user incentives plan includes: 1) collect large-scale public datasets, such as Google N-Gram datasets and human gene datasets, and offer free access and free data processing for

registered public user; 2) originate data analysis competition with prize at the SDN (software developer network) community.

5. ACKNOWLEDGMENTS

We thank Bill Huang for his guidance and inspiration, and Dejan Milojicic, from HP Labs, for his help in improving the quality of the presentation of the paper.

6. REFERENCES

- [1] S. Ghemawat, H. Gobioff, and S.-T. Leung, "The Google file system" in In Proceedings of 19th ACM Symposium on Operating Systems Principles, October 2003. LOCATION MISSING, DATE MISSING
- [2] Hadoop, an open source implementing of MapReduce and GFS, <http://hadoop.apache.org>.
- [3] Hadoop at Yahoo! <http://developer.yahoo.com/hadoop/>
- [4] Cloudera services on Hadoop. <http://www.cloudera.com/hadoop/>
- [5] Hadoop at tldata.org. <http://www.tldata.org/archives/category/cloud-computing/hadoop>
- [6] J. Dean and S. Ghemawat, "MapReduce: Simplified data processing on large clusters," in In Proceedings of OSDI'04: Sixth Symposium on Operating System Design and Implementation, December 2004.
- [7] Sridhar Ramaswamy, "Extreming Data Mining", Google Keynote speech in SIGMOD 2008. LOCATION MISSING. DATE MISSING
- [8] Gartner report. Hadoop and MapReduce: Big Data Analytics. <http://www.gartner.com/technology/media-products/reprints/cloudera/vol1/article1/article1.html>
- [9] Mahout, open source project on data mining algorithms based MapReduce, <http://lucene.apache.org/mahout/>.
- [10] Teradata exchanges with Hadoop. <http://developer.teradata.com/tag/hadoop>.
- [11] Pentaho for Hadoop. <http://www.pentaho.com/products/hadoop/>.