Available Online at http://www.recentscientific.com

CODEN: IJRSFP (USA)

International Journal of Recent Scientific Research
Vol. 10, Issue, 05(C), pp. 32358-32361, May, 2019

ISSN: 0976-3031

International Journal of Recent Scientific Research

DOI: 10.24327/IJRSR

Research Article

WORKLOAD CHARACTERIZATION OF HADOOP CLUSTER

C.M.Arun, D Sathya and V.Monica

Erode Sengunthar Engineering College, Erode

DOI: http://dx.doi.org/10.24327/ijrsr.2019.1005.3453

ARTICLE INFO

Article History:
Received 10th February, 2019
Received in revised form 2nd March, 2019
Accepted 26th April, 2019
Published online 28th May, 2019

Key Words:

HDFS is a tool to implement Hadoop. The default scheduler used is a YARN scheduler. It is concerned with five priority levels

ABSTRACT

As the organizations start to use data intensive cluster computing systems like Hadoop for more applications, there is a growing need to share clusters between users. To address the conflict between locality and fairness various algorithm are proposed. Map Reduce is becoming the high-tech computing paradigm for processing large-scale datasets on a large cluster with large number of nodes. It has been most useful in various applications such as e-commerce, Web search, social networks, and scientific computation. Tounderst and the characteristics of Map Reduce workloads is the key to achieve better configuration decisions and improving the system throughput. To achieve better performance, a map reduce scheduler must avoid unnecessary data transmission by enhancing the data locality. A map reduce which is found to be the offline computing
Engine solves the issues of too large data to fit into a single machine. This mapreducefunction comprises of Job Tracker and Task Tracker, where the Job Tracker is concerned
with the division of the given input dataset into chunks and sends to the individual nodes. Map Reduce is a programming model which supports distributed and parallel computing on the data intensive applications. HDFS is a tool to implement Hadoop. The default scheduler used is a YARN scheduler. It is concerned with five priority levels. Improving the performance of the map reduce function will be in the contest of improvement in system latency, memory settings, input output bandwidth, job parallelization. The factors that affect the performance of the Hadoop are Hardware, Mapreduce, HDFS, Shuffle tweaks.

INTRODUCTION

Big Data is a term also named as “veracity” is applied to data sets whose size is beyond the capability of commonly used software tools to capture, manage, and process. The sheer size of the data, combined with complexity of analysis and commercial to create value from it, has led to a new class of technologies and tools to adopt it. The term Big Data tends to be useful in multiple ways, often referring to both the type of data being managed as well as the technology used to store and process it. These technologies originated from companies such as Google, Amazon, Face book and Linked-In, where they were developed for each company’s own use in order to analyze the massive amounts of social media data they were dealing with. Due to the nature of these companies, the emphasis was on low cost scale-out commodity hardware and open source software

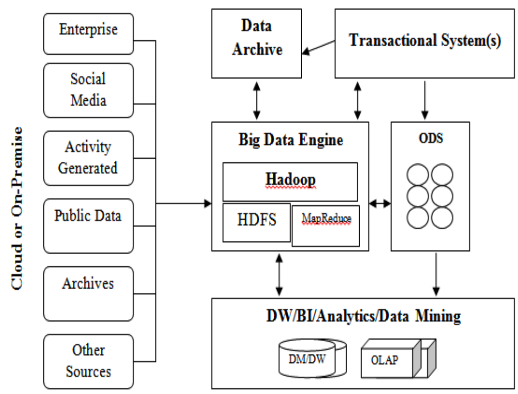

The world of Big Data is increasingly being defined by the 3 Vs. But now it has given defined by 4Vs, these ‘Vs’ become a reasonable test as to whether a Big Data approach is the right one to adopt for a new area of analysis.
The Vs are:

Corresponding author:* **C.M.Arun
Erode Sengunthar Engineering College, Erode

- ✓ Volume
- ✓ Velocity
- ✓ Variety

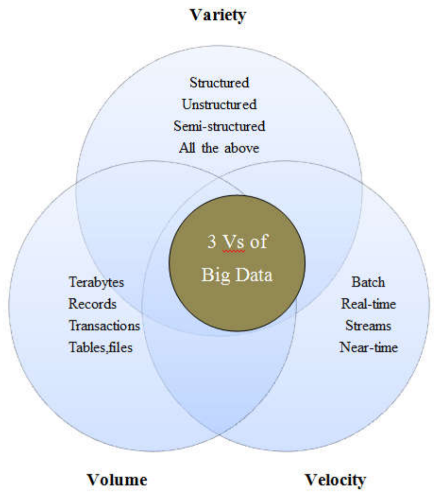

Volume

The size of the data. As there is continuous advancement in technology, numbers get quickly outdated so it's better to deal with volume in a relative sense. The volume of data is an order of magnitude or larger than anything previously encountered in industry, then probably dealing with Big Data.

Velocity

The rate at which data is being received and has to be processed is becoming much more real-time. While it is unlikely that any real analysis will need to be completed in the same time period, delays in execution will inevitably limit the effectiveness, limit interventions or lead to sub-optimal processes. For example, some kind of discount offer to a customer based on their location is less likely to be successful if they have already walked some distance past the store.

Variety

There are two aspects of variety to consider: syntax and semantics. In the past these have determined the extent to which data could be reliably structured into a relational database and content exposed for analysis.

Value

'Value' offers a particular challenge to IT in the current harsh economic climate. It is difficult to attract funds without certainly of the ROI and payback period. The tractability of the problem is closely related to the issue as problems that are inherently more difficult to solve will carry greater risk, making project funding more uncertain.

One aspect that most clearly distinguishes Big Data from the relational approach is the point at which data is organized into a schema. In the relational approach, data is placed into a schema when it is initially written to the database, where as in a Big Data approach data is only organized into a schema immediately prior to analysis as it is read from disk. Thus Big Data can be said to be 'schema on read' where as relational technology is 'schema on write'.

Big Data is often seen as being more agile approach because in 'schema on read' data is only structured immediately prior to analysis, but the approach also has hidden costs and risks which must be managed. For example, with 'schema on read', data quality is very dependent on the developer responsible for de-serializing / tokenizing the data from disk and this cost is potentially repeated for each program. It may also be difficult to find developers who are sufficiently knowledgeable about data streams written many years ago.

Data Scientists will typically use a broad range of technologies such as Data Mining, Statistical and graphical analysis depending on the problem being tackled. It seems inevitable that in the future analysis tools will most likely work seamlessly across technologies, obfuscating the underlying storage technology. The dominant Big Data technologies in use today commercially are Apache's Hadoop and No-SQL databases. No-SQL databases are typically part of the real-time event detection process deployed to inbound channels but can also be seen as an enabling technology behind analytical capabilities such as contextual search applications.

Literature Survey

Matchmaking: a New Map Reduces Scheduling Technique.

Speculative execution prevents a job from being delayed by the worst performing node.GOOGLE has announced that this mechanism can improve a job's response time by 44%.To make speculative execution effective in heterogeneous environment, LATE scheduler and SAMR algorithm. Map Reduce Cluster's data locality can be improved by prefetching. Estimate job's execution times and tries to let jobs satisfy their deadlines by scheduling resources according to the estimated finishing times. The difference is if a job cannot finish before the hard deadline, the scheduler will not execute the job and will instead inform the user to adjust the job deadline.

Failure Data Analysis of A Large-Scale Heterogeneous Server Environment

The very rapid initial declines in the distribution functions, where the majority of the functions, where the majority of the probability mass occurs for very small values. The delay rates of the tails flatter out for larger values, although the tails are obviously bounded at maximum values. Use of AUTO SLEX, as it does not identify changes in the correlation structure, because the partition is based on the spectral value of zero frequency. The failure rates are still significant and highly variable. In particular, the system error and failure patterns are clearly nonstationary and they consist of relatively long time intervals that are stationary, many spanning more than a day.

Pfair Scheduling of Periodic Tasks with Allocation Constraints on Multiple Processor.

A special type of allocation constraints in which each fixed task can be assigned to one of disjoint and dedicated processor sets.HPA may lead to a task missing, its deadline, but the task does not miss its deadline.The property of HPA guarantees some fairness not too is bad for allocating resources.HPA can be used in not only real-time task scheduling, but also packet scheduling in performance-guaranteed communications.

Llf Schedulability Analysis on Multiprocessor Platforms.

The structure of LLF schedulability test consists of a set of necessary conditions for a deadline miss. Reduce the number of pre-emptions. Some tie-breaking rules into LLF.

Practical Power Proportionality for data Center Storage

The low gear experiment shows that provisioning and gear selection methodology and gear selection methodology is reasonable. The performance penalty is slightly more pronounced for the up shift / reclaims experiment. Reclaim process is very faster. To redesign the data layout to allow for the possibility of servers to be off. To introduce new service to the data center named DVL. Failure scenario that introduces a tradeoff between power savings and availability.

Workload Characterization in a High energy data grid and Impact on Resource Management.

To quantify reduction in data transfer and increase in byte hit rate when using filecules for prefetching LFU-GRV algorithm. Filecules identified using jobs executed during one month for prefetching data into the cache during the next month. Modelling workloads for data intensive scientific collaborations. For designing data management techniques adapted to multiple file processing.

Workload Analysis of A Cluster in A Grid Environment

Globus is used to co-allocate the nodes on different clusters. Bag of tasks applications which interact together in a pipeline way by files stored on SES. There is a strong correlation between the successive jobs running time but it seems unlikely that a general model for duration highly depends on algorithms and data used by users. The main merit is that it can be used numerically with non-constant parameters at the expense of difficult sense. The job maximum run times provided by users are essentially inaccurate is dependent on scheduling. Software manager jobs may be regarded as more urgent than other jobs type. Since with this content, sending jobs with an estimated runtime could be replaced by sending jobs with an urgency parameter. A major conflict occurs when the sum of all computation flow rate is greater than the site capacity. The open problem is finding the best suited scheduling policy.

User Group based Workload Analysis and Modeling

Increases write throughput and reduces latency. Buffering gives a significant throughput benefit. New workloads can be easily created, including generalized workloads to examine system fundamentals. Injecting faults are not straight forward different systems have different components and different unique failure models.

Job Scheduling for Multi-User Mapreduce Clusters

Task is free, scans through jobs in order of priority and submits. Gain in throughput and response time. Intermediate results produced by map will not be deleted until job ends. Poor response time to shorter jobs in the presence of larger jobs. Performance of FAIR scheduler is found to be less. It breaks the locality of data. Starvation occurs.

Towards Characterizing Cloud Backend Workloads: Insights from Google Compute Clusters.

To minimize the number of workloads. Inter task communication is found to be reliable. To address characterization of the task arrival process and extend task classification to consider job constraints.

Design Insights for Mapreduce from Diverse Production Workloads

Analyze more workloads over longer time periods and additional statistics analysis is done. Automate analysis and monitoring of tools and creation of a map reduce workload taxonomy. Map Reduce has evolved to the point where performance claims should be qualified with the underlying workload assumptions. System engineers should regularly re-access designs priorities subject to changing use cases. Pre-requisites to these efforts are workload replay tools and a public workload repository, so that engineers can share insights across deployments.

Objective of the Project

To setup the Hadoop infrastructure and to run the sample mapreduce program. To study the performance of the normal Hadoop cluster and to analyze the performance of the sample Map Reduce program.To propose a new technique to support the performance enhancement of Hadoop cluster. To compare the performance improvement of prior and the newly proposed technique in the case of system latency and the memory settings.

Existing System

The system environment configured for Hadoop infrastructure is found to have YARN scheduler in default. The Map Reduce chunks are splitted as 64 blocks of default size which may affect the memory settings. There are various inferences identified from the literature survey which will pave the way to improve the performance of the Map Reduce enabled between the various systems involving Hadoop cluster. This system tends to have the features of YARN scheduler in it. Various scheduler have been studied and the inferences have been made

Proposed System

To incorporate the Hadoop infrastructure that has been set up, and to work up on the workload an efficient strategy will be adopted in future, in order to improve the performance of the Hadoop cluster. The proposed system will have the sufficient accomplishments for the factors such as the system latency and the memory settings. As a result the workload will be characterized and evaluated with the performance of the existing system.

Modules

There are three modules used in this system, which are described below

- ✓ Setting up of the HADOOP infrastructure.
- ✓ Running the Map Reduce program.
- ✓ Analyzing the performance of the HADOOP cluster.
- ✓ Implementation of the new scheduling algorithm.

Setting up of the HADOOP infrastructure

Setting up the HADOOP infrastructure includes various specific requirements such as,

- ✓ Configure a system with Ubuntu Operating system
- ✓ Installation of JDK in Ubuntu
- ✓ Installation of HADOOP.

Configure a system with Ubuntu Operating system.

The system should be configured with Ubuntu Operating system inorder to setup the HADOOP infrastructure.
Installation of JDK in Ubuntu

JDK should be installaed in the Ubuntu, since the MapReduce programs that we are going run are JAVA programs.
Installation of HADOOP

HADOOP could be installed after cross checking the JDK installation. HADOOP installation takes several steps and various configurations need to be carried out for successful HADOOP installation. Finally check for HADOOP, whether installed.

Analyzing the Performance of Map reduce Cluster

Running the sample Map Reduce program in the HADOOP and verifying the performance of it.

CONCLUSION

Thus, inorder to setup the Hadoop infrastructure there are three steps followed. They are installation of Ubuntu, followed by installation of JDK. Since the MapReduce programs in the Hadoop are Java programs JDK is needed. After the successful installation of JDK we can install Hadoop. This includes various steps such as creation of user, generation of key for the purpose of the authentication and the authorization.

Then followed by the configuration set up. Formatting the Namenode will enable the startup of our work. Then running the sample mapreduces program and analyzing the performance of it. From the literature survey various inferences have been made and some of the factors have been identified such as memory settings, shuffle tweaks, system latency.

Future Enhancement

Future scope of it is to determine reducing the cost and power consumption which can be highly useful in the future. Our current analysis are based on a simple three cluster model, further systematic studies of more generalized multi-cluster networks are needed. Thus far we have concentrated on the homogeneous sensor networks with a single powerful processing center (sink). In our future work, we would rather focus on the heterogeneous wireless sensor networkswith multiple resource-rich actors for carrying out energy consuming tasks.

References

1. J. Dean And S. Ghemawat, “Mapreduce: Simplified Data Processing On Large Clusters,” In Osdi, 2004, Pp. 137–150.
2. Y. Chen, S. Alspaugh, D. Borthakur, And R. H. Katz, “Energy Efficiency For Large-Scale Mapreduce Workloads With Significant Interactive Analysis,” In Eurosys. Acm, 2012, Pp. 43–56.
3. J. D. C. Little, “A Prof For The Queuing Formula L=Hw,” Operations Research, Vol. 9, 1961.
4. X. Liu, J. Han, Y. Zhong, C. Han, and X. He, “Implementing webGISon hadoop: A case study of improving small file I/O performance onHDFS,” in CLUSTER, 2009, pp. 1–8.
5. M. Zaharia, D. Borthakur, J. S. Sarma, K. Elmeleegy, S. Shenker, andStoica, “Delay scheduling: a simple technique for achieving localityand fairness in cluster scheduling,” in EuroSys, 2010, pp. 265–278.
6. C. He, Y. Lu, and D. Swanson, “Matchmaking: A new mapreducescheduling technique,” in CloudCom, 2011, pp. 40–47.
7. R. K. Sahoo, A. Sivasubramaniam, M. S. Squillante, and Y. Zhang, “Failuredata analysis of a large-scale heterogeneous server environment,” inDSN, 2004, p. 772.
8. D. Liu and Y.-H. Lee, “Pfair scheduling of periodic tasks with allocationconstraints on multiple processors,” in IPDPS, 2004.
9. J. Lee, A. Easwaran, and I. Shin, “LLF schedulability analysis onmultiprocessor platforms,” in IEEE Real-Time Systems Symposium,2010, pp. 25–36.
10. E. Thereska, A. Donnelly, and D. Narayanan, “Sierra: practical powerproportionality for data center storage,” in EuroSys, 2011, pp. 169–182.
11. A.Iamnitchi, S. Doraimani, and G. Garzoglio, “Workload characterizationin a high-energy data grid and impact on resource management,”Cluster Computing, vol. 12, no. 2, pp. 153–173, 2009.
12. E. Medernach, “Workload analysis of a cluster in a grid environment,”in Job Scheduling Strategies for Parallel Processing, 2005, pp. 36–61.
13. B. Song, C. Ernemann, and R. Yahyapour, “User group-based workloadanalysis and modelling,” in CCGRID, 2005, pp. 953–961.
14. M. Zaharia, D. Borthakur, J. S. Sarma, S. Shenker, and I. Stoica, “Jobscheduling for multi-user mapreduce clusters,” Univ. of Calif., Berkeley,CA, Technical Report No. UCB/EECS-2009-55, Apr. 2009.
15. A.K. Mishra, J. L. Hellerstein, W. Cirne, and C. R. Das, “Towardscharacterizing cloud backend workloads: insights from google computeclusters,” SIGMETRICS Performance Evaluation Review, vol. 37, no. 4,pp. 34–41, 2010.
16. B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears,“Benchmarking cloud serving systems with YCSB,” in SoCC, J. M.Hellerstein, S. Chaudhuri, and M. Rosenblum, Eds., 2010, pp. 143–154.
17. Y. Chen, S. Alspaugh, and R. H. Katz, “Design insights for mapreducefrom diverse production workloads,” University of California, Berkeley,Tech. Rep. UCB/EECS-2012-17, 2012.