

Sustainability of Big Data Servers Under Rapid Changes of Technology

Sashiraj Chandrasekaran

James Cook University, Singapore
sashiraj.chandrasekaran@my.jcu.edu.au

Insu Song

James Cook University, Singapore
insu.song@jcu.edu.au

Abstract. A big data server is a computer system designed to store and process many types of unstructured data arriving at a rapid pace. Such data captured from the Internet and Social Networks are crucial for both developed and developing countries to be able to make informed decisions in time. However, sustainability of big data infrastructures and electronic waste are big issues due to the rapid changes in technology. In this paper we evaluate the performance of big data servers on reusable computers in order to evaluate the scalability and feasibility of constructing big data servers using discarded computers that can be procured as low as \$40. In particular, we compare virtualized clusters and bare metal clusters of the low-cost recycled computing nodes for their scalability and feasibility. Virtualized environment is often considered for big data infrastructures due to more efficient management of the clusters despite of the performance overheads. Our study shows that virtualized environment is not scalable for low-cost recycled computing nodes. Our performance evaluation shows that the virtualized cluster is 66% slower than the non-virtualized cluster for read operations. For write operations, the virtualized system is 88% slower than the non-virtualized system.

Keywords: Big Data, Distributed Computing, Benchmarking, Sustainability, Electronic Waste

1 Introduction

Big data is a term used to describe the latest advancements in technology that aid the analysis of huge amounts of data [1]. Big data is different from traditional data analytics because of the nature of data being generated in the 21st century. Big data is characterized by variety, velocity and volume [2]. Data can be in various forms such as structured, unstructured or semi-structured data. Such data is churned out at a fast rate resulting in enormous volumes of data generated. A server analysing big data is expected to handle the large quantities of data generated. Since such large amounts of

data is beyond the storage and processing capacity of a single server, distributed computing systems were first conceived [3]. These systems have several machines combined together to form a cluster. The data is split into smaller pieces and distributed among the machines in the cluster. The processing power and storage capacity of each machine is used to process the chunk of data present in that machine [4].

The problem with most big data servers is the high cost of purchasing and using these servers due to rapid changes in technology. This dramatically affects the cost of big data services making such services inaccessible to small institutions and developing countries. To address this problem, this project intends to explore options to build big data servers using old discarded laptops which can be acquired at extremely low prices [4].

There are two methods that can be used to build big data servers. The first method uses a virtualized environment to manage all the machines. The second method uses a non-virtualized environment. Virtualized clusters are commonly used in data centres because of the ease of managing the computing nodes despite of some performance overhead, which is claimed to be about 25% according to a leading virtualization software company. It has also been stated in [5] that the overhead of using a hypervisor in a virtualized server environment is 5 to 7%.

In this paper we report on the performance testing results for the two methods to evaluate scalability and feasibility of using virtualized environment for building big data servers using recycled computers. Two clusters are built using the discarded laptops. The first cluster operates on a virtualized environment based on a hypervisor and the second cluster operates on a non-virtualized environment. CentOS 6.5 is used as the operating system for both clusters. The performance of both clusters is benchmarked. The results show that the virtualized environment has an overhead of 66% for read operations and 88% for write operations.

The paper is structured as follows. In Section 2, we briefly introduce Hadoop framework for big data processing and benchmarking approaches of big data servers. In Section 3, we describe how the big data servers constructed for benchmarking and how they are tested. In Section 4, the benchmarking results are presented. In Section 5, we discuss the comparison results and conclude the paper in Section 6 with remarks on the results.

2 Background

Given the advantages of big data, a large number of organisations are interested in using big data analytics. But scalability issues with existing systems and the high costs involved are big deterrents for small companies and other institutions. One of the important requirements of a big data server is the capability to adapt to changes in the database and the ability to increase the processing power and storage capacity as and when the need arises [1]. To manage and distribute data storage and processing across multiple virtual machines a software framework such as Hadoop can be used.

2.1 Hadoop Framework For Big Data Processing

Hadoop was developed by Apache to provide reliable, scalable software frameworks for distributed computing. Hadoop integrates data storage, processing and management into a single software solution [6]. Hadoop is designed to detect and handle failures at the application layer. This ensures that the cluster of computers managed by Hadoop is always available. Each of the machines in a Hadoop cluster is available for storage as well as computation. Hadoop framework consists of the common utilities to support the Hadoop modules. These utilities include Hadoop Distributed File System (HDFS) which handles the storage and retrieval of data files as well as YARN and MapReduce which are used for parallel processing of data. The key component that makes Hadoop scalable is the Hadoop Distributed File System [7]. Data in a Hadoop cluster is broken down into smaller blocks and distributed among the machines in the cluster. Data is stored in the hard drives of the machines instead of a Network Attached Storage or Storage Area Network. The data stored in each of the machines is processed by MapReduce. MapReduce consists of Map and Reduce tasks performed by Hadoop. The Map tasks break down data into smaller key/value pairs. The Reduce tasks take the key/value pairs and combine them into smaller tuples. The MapReduce jobs are managed by two processes. The Task Tracker manages the map and reduce operations on the compute nodes. The Job Tracker monitors the jobs performed and manages the allocation of tasks across different machines [8]. In addition to MapReduce and HDFS, Hadoop offers several other services for more specific big data operations, some of which are discussed. Hive is used to manage and query large datasets stored in distributed systems. HBase is a column oriented non-relational database which is built to handle Java-based applications. HBase performs functions on HDFS similar to that performed by Google's BigTable. Pig is a programming language used to analyse large data sets. Pig also provides a runtime environment. Hadoop services also include monitoring tools such as Ganglia and Nagios. Ganglia is a distributed monitoring service, developed by University of California and aimed at high performance computing systems. Nagios is used to monitor clusters as well as network devices such as switches and routers [9].

With the number of Hadoop services in place, it is a challenge to install and manage all of these services. Software frameworks like Ambari aid in installing, configuring, managing and monitoring Hadoop services. To facilitate better management of scalable systems, many organizations use virtual clusters.

2.2 Virtualization

Virtual cluster is a set of virtual machines which have been configured to work as a single system. A virtual machine emulates the lower levels of computer abstraction for the higher levels. A software process called Virtual Machine Monitor (VMM) or hypervisor manages the hardware resources of the real machine. Virtual machines have the ability to create a new guest environment using the resources on the host [10]. The guest environment can be configured according to the user requirements.

Multiple guest systems can be deployed on the host system using virtual clusters. The VMM helps allocate resources among the different virtual machines found on the host machine. Resource utilization on a virtual system is flexible [11]. The state of a virtual machine can be saved and moved to a different host system for better load balancing on the clusters [12].

There are several advantages of using Hadoop on virtual clusters. Virtualization eliminates the need for specialised hardware for master services. The same machines used for slave services can be used for master services also. Deployment of virtual machines is often very fast and can be automated using specialised virtualization software like Project Serengeti. Virtualization also allows for multiple versions of Hadoop to run on the same cluster [13].

One of the approaches using virtualized clusters was developed by [14]. The paper aimed to study the effect of different configurations of the data node, name node, job tracker and task tracker in a virtualized big data server. The data node is where the HDFS file system is stored on the local file system. The name node maintains the file system meta data of the data stored in the data node [15]. They used a single physical server to host 4 virtualized Hadoop clusters and compared the performance between different configurations of the clusters. A Dell PowerEdge T420 server with two Intel Xeon E5-2420 CPUs with 6 cores, operating at 1.9 GHz was used. The setup had 32GB of RAM and 4TB of storage capacity in the form of four 1TB hard disks (Western Digital-SATA).

Project Serengeti, which is an integrated big data extension found in vSphere, was used for deploying and managing the Hadoop clusters. The virtual machines were split into 5 types- name node vm, data node vm, compute master node (job tracker), compute worker node (task tracker) and worker node, which is the combined data node and compute worker node on one single virtual machine. The clusters were benchmarked for different configurations of the virtual machines. One advantage of this setup is that the number of data and compute nodes in the clusters can be scaled up dynamically according to the requirement. While this method has a lot of useful information, there are some limitations which the experiment does not address. One of the limitations in this method is that with each additional vm in the physical machine the overhead caused by the vm goes up. Another disadvantage is that the system uses only one physical machine. Since there is no physical network communication between different physical machines, the speed of the virtual machines over a network cannot be studied. The variations in speed caused by the network constraints are also not studied. To benchmark the performance, two MapReduce applications from the Hi-Bench benchmarking suite developed by Intel were used. Word count and TestDFSIO applications were used for benchmarking. It was observed that adding more data nodes to the data-compute type node cluster increased reading speed by 39%. However, the performance of virtualized clusters with respect to non-virtualized clusters is not tested by this experiment.

2.3 Bench Marking Big Data Servers

There are many different tools that can be used to evaluate the performance of computer systems. Evaluation of performance with comparison to a standard is called benchmarking. Benchmark tools thus help in evaluating and comparing performance parameters of large computer systems. Hi-Bench benchmarking suite was used by [14] for their study. Another study by [16] focuses on the BigBench benchmarking tool. BigBench deals with structured, unstructured as well as semi-structured data components often associated with big data. The semi-structured data used for testing BigBench is user clicks on a website. Product reviews found online are used for unstructured data and a data generator is designed to produce scalable volumes of raw data for structured data. The Teradata Aster database is used to test the accuracy of the benchmarking tool. A 200GB test data set is first generated and loaded and BigBench is then used to test the data set. The benchmarking tool thus aimed to address the variety, velocity and volume components of big data.

3 Methodology

3.1 Theory and Hypothesis

In this paper, we explore the possibility of building big data servers using recycled computers to solve the issues of sustainability of big data infrastructure and reducing electronic wastes. There are two approaches to building big-data infrastructure: virtualized environment and non-virtualized environment. Virtualized environment offers more flexibility in term of management whereas non-virtualized environment provides better performance. According to one study [5], virtualized infrastructure has 5-7% performance overhead. Due to flexibility and reduced management cost, virtualized environment is preferred in industry [5].

Therefore, we compare virtualized big-data servers with non-virtualized big-data servers built using recycled computers for the first time. The main hypothesis is that the virtualized big-data server has a similar 5-7% performance overhead over non-virtualized server.

3.2 Experiment Setup

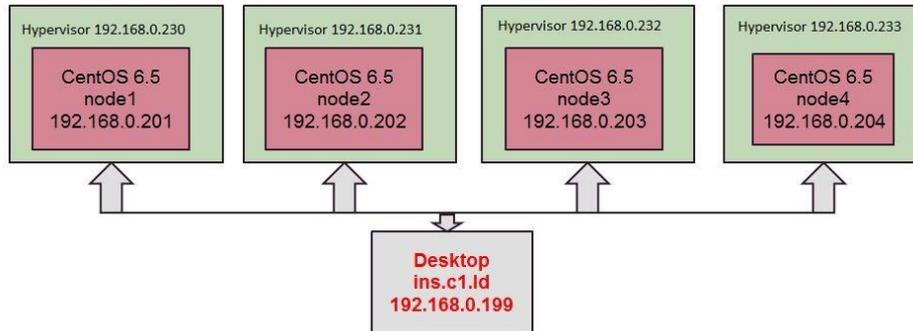


Fig. 1. Virtualized cluster

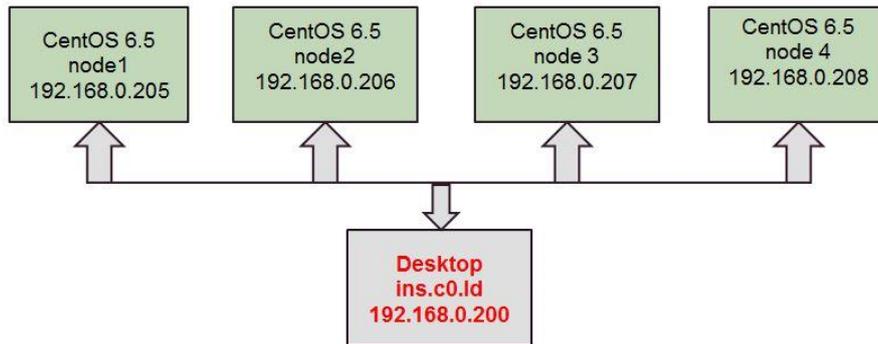


Fig. 2. Non-virtualized cluster

In order to compare the performance of virtualized big-data clusters and non-virtualized big-data clusters, the following experiments are set up.

There are 2 clusters of 4 computer nodes each as shown in Fig. 1 and Fig 2. On the first cluster (Fig. 1), virtualized big-data clusters of 4 nodes are connected. On the second cluster (Fig. 2), non-virtualized big-data clusters of 4 nodes are connected. For the virtualized cluster, a hypervisor software is first installed on the 4 computers of the virtualized cluster before installing a guest operating system (CentOS 6.5). All nodes had the following hardware specifications: Intel Core i3 CPU M 350 processor, 2.27 GHz, 3.7GB RAM and 200GB hard disk storage. The nodes are connected using two Gigabit LAN switches. The nodes are connected using CAT-6 24AWG Ethernet LAN cables. Each node is equipped with a 2.27 GHz Dual Core Intel CPU with Hyper Threading (i.e., 4 virtual cores per node), RAM memory of 3.7GB and 200GB of hard disk storage.

Once CentOS6.5 is installed on all nodes, the the network is configured to ensure that all computers could connect to each other within their clusters. A separate machine is used to serve as the install node of Hadoop. A GUI version of CentOS6.5, with Firefox, is configured on the install node. The Ambari server is then set up on the install node. The Ambari server is then used to install the following Hadoop services: HDFS, Yarn, Tez, Nagios, Ganglia, HBase, and Pig services. Hadoop configuration on Hadoop 2.1 stack was used in this experiment. The first node is used as the master node and the other 3 nodes are used as the slave nodes. The master node has the Name Node, Secondary Name Node, History Server and App Timeline server running on it. The name node contained the metadata of the filesystem. The secondary name node serve as back up for the name node in the event of a system crash. The slave node has the data node, node manager, region server and the clients for the Hadoop services. Once the Hadoop services are installed, the cluster is ready for testing.

For the test data, we use synthetic data to perform the benchmark tests. Synthetic data is generated by using the write function of the Hadoop DFSIO benchmark test.

To compare the performance of the two clusters, the TestDFSIO benchmarking tool is used. TestDFSIO is a distributed I/O benchmarking tool which can be used to test the speed at which data can be read and written to the disks. TestDFSIO is selected over other bench-marking tools as this is inbuilt in the Hadoop system and input/output data is given by the benchmark. The number of files and the size of each file are specified and the TestDFSIO write function is used first to generate the data set. The time taken to generate the data set, the throughput speed (in Mb/s) and the average input/output rate (in Mb/s) is stored in the log file. The TestDFSIO read function is then used to perform the read test on the data generated during the write test. The results of the read test are appended to the log file. TestDFSIO erase function is used to erase the data set. A new data set is created with the TestDFSIO write function and the process is repeated 4 times to ensure consistency of data.

4 Results

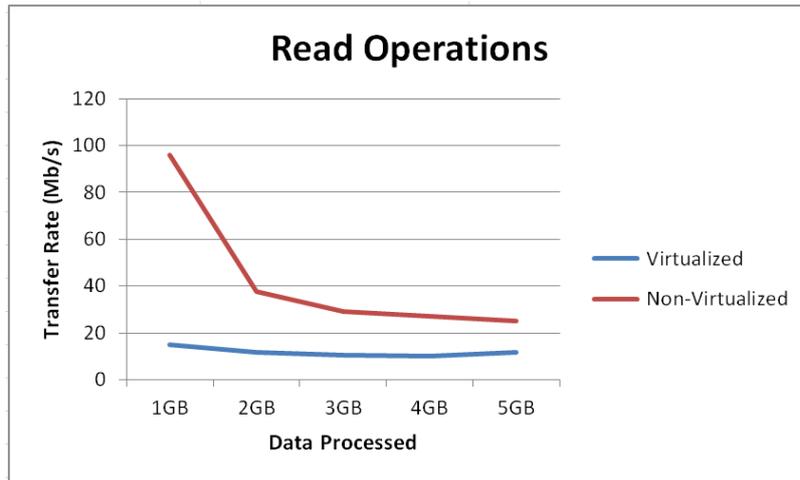


Fig. 3. Throughput Read Operations: Graph showing speed for read operations

Fig. 3 shows the performance results of the clusters for read data transfer operations. It can be observed from the figure that the rate for read data transfer throughput is fairly constant for the virtualized-cluster. However, the transfer rate changes drastically with file size for the non-virtualized cluster. On average, the read transfer rate of the virtualized cluster is 66% less than that of the non-virtualized cluster.

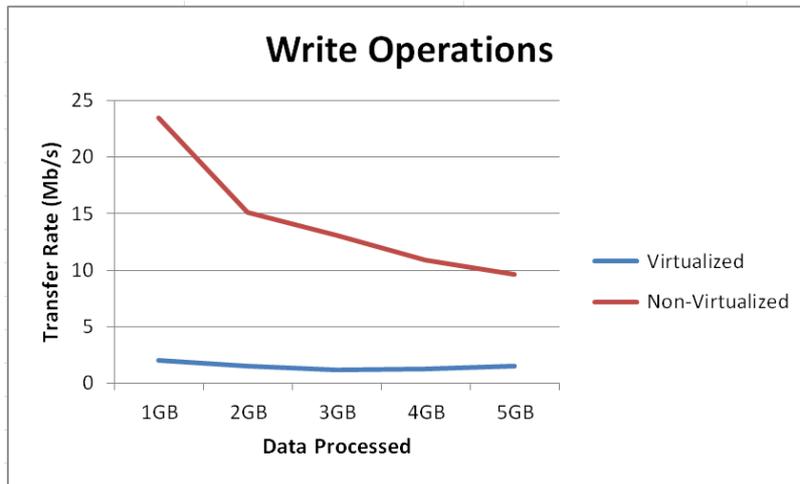


Fig. 4. Throughput Speed for write Operations: Graph showing throughput speed for write operations.

Fig. 4 shows the write data transfer throughputs of the clusters. The throughput rate for the write operation also remains fairly constant for the virtualized cluster, but changes dramatically over the data size for the non-virtualized cluster. On average, the write transfer rate of the virtualized cluster is 88% less than that of the non-virtualized cluster.

5 Discussion

The results obtained by this experiment are quite unexpected because of the high overhead found for this experimental setup. The virtualized cluster is performing considerably slower than expected when tested against a non-virtualized cluster. However, it is noted that the hardware used for this setup is considerably different from the hardware traditionally used for virtualization and benchmarking. Conventional benchmarking tests are run on servers which are made to handle a high amount of traffic whereas this experimental setup used laptops which are a couple of years old. Generally, laptops are not subject to the same performance standards as servers. The CPU used in a server is significantly different from the CPU used in conventional laptop computers. CPUs used in servers have more cache and more CPU cores while compared to the desktop CPUs. Server CPUs are also designed to run at 100% capacity for extended amounts of time whereas desktop CPUs are not. It must also be noted that servers do not need to allocate resources to manage devices such as monitor, keyboard, mouse and other peripherals. Furthermore, the RAM used by the laptops is also different from the ECC RAMs used by server systems.

With respect to virtualized clusters, some of the methods used by [14] used big data extensions such as Serengeti. Serengeti is developed by VMWare and it aids in rapid deployment of virtualized hardware. The functions performed by Serengeti are similar to the functions performed by Ambari in our setup. But it is possible that while Ambari, developed by apache, is fast and efficient on non-virtualized environments, Serengeti is better suited for virtualized clusters deployed using hypervisors.

6 Conclusion and Future work

This project was aimed to test the viability of using recycled computers for big data servers. We benchmarked a virtualized cluster against a non-virtualized cluster. The bench mark results shows that the overhead of virtualized environment is much greater than expected: 66% and 88% overheads in read and write operations respectively.

There are various paths for exploration in the future. A different virtualization solution can be used to compare the performance with a non-virtualized environment. The number of data nodes used in the system can also be increased to see if there is a performance difference for bigger systems for the same file sizes. Big data extensions such as Serengeti can also be used to see if the use of extensions can improve the performance of virtualized clusters.

References

1. Cuzzocrea A, Song I-Y, Davis KC Analytics over large-scale multidimensional data: the big data revolution! In: Proceedings of the ACM 14th international workshop on Data Warehousing and OLAP, 2011. ACM, pp 101-104
2. Zaslavsky A, Perera C, Georgakopoulos D (2013) Sensing as a service and big data. arXiv preprint arXiv:13010159
3. DeWitt D, Gray J (1992) Parallel database systems: the future of high performance database systems. *Communications of the ACM* 35 (6):85-98
4. Chen M, Mao S, Liu Y (2014) Big data: A survey. *Mobile Networks and Applications* 19 (2):171-209
5. Buell J (2013) Virtualized Hadoop Performance with VMware vSphere 5.1. <http://www.vmware.com/files/pdf/vmware-virtualizing-apache-hadoop.pdf>. Accessed 26-09 2015
6. Hu H, Wen Y, Chua T-S, Li X (2014) Toward scalable systems for big data analytics: A technology tutorial. *Access, IEEE* 2:652-687
7. Dede E, Govindaraju M, Gunter D, Canon RS, Ramakrishnan L Performance evaluation of a mongodb and hadoop platform for scientific data analysis. In: Proceedings of the 4th ACM workshop on Scientific cloud computing, 2013. ACM, pp 13-20
8. Venner J (2009) *Pro Hadoop*. Apress,
9. Lakhe B (2014) Monitoring in Hadoop. In: *Practical Hadoop Security*. Springer, pp 119-141
10. Sailer R, Jaeger T, Valdez E, Caceres R, Perez R, Berger S, Griffin JL, Van Doorn L Building a MAC-based security architecture for the Xen open-source hypervisor. In: *Computer security applications conference, 21st Annual, 2005*. IEEE, pp 10 pp.-285
11. Zhang X, Keahey K, Foster I, Freeman T (2005) Virtual cluster workspaces for grid applications. sl: TR-ANL/MCS-P1246-0405
12. Foster I, Freeman T, Keahy K, Scheftner D, Sotomayer B, Zhang X Virtual clusters for grid communities. In: *Cluster Computing and the Grid, 2006. CCGRID 06. Sixth IEEE International Symposium on, 2006*. IEEE, pp 513-520
13. McDougall R (2012) Project Serengeti: There's a Virtual Elephant in my Datacenter. <https://blogs.vmware.com/cto/project-serengeti-theres-a-virtual-elephant-in-my-datacenter/>. Accessed 20.10.2015 2015
14. Ivanov T, Zicari RV, Izberovic S, Tolle K (2014) Performance Evaluation of Virtualized Hadoop Clusters. arXiv preprint arXiv:14113811
15. Shvachko K, Kuang H, Radia S, Chansler R The hadoop distributed file system. In: *Mass Storage Systems and Technologies (MSST), 2010 IEEE 26th Symposium on, 2010*. IEEE, pp 1-10
16. Ghazal A, Rabl T, Hu M, Raab F, Poess M, Crolotte A, Jacobsen H-A BigBench: towards an industry standard benchmark for big data analytics. In: *Proceedings of the 2013 ACM SIGMOD international conference on Management of data, 2013*. ACM, pp 1197-1208