BigScale: automatic service provisioning for Hadoop clusters

Dan Huru and Cristian Eseanu and Cătălin Leordeanu and Elena Apostol and Mariana Mocanu and Valentin Cristea
Faculty of Automatic Control and Computers, University Politehnica of Bucharest, Romania
Email:alexandru.huru2208@cti.pub.ro, eseanu.cristian@cti.pub.ro, {elena.apostol, catalin.leordeanu, mariana.mocanu, valentin.cristea}@cs.pub.ro

***Abstract*—As the number of interconnected devices grows in the IoT space, data processing systems require increased resources, robustness and flexibility. In this sense the scalability of a system becomes very important. A scalable system can process variable data volumes, requires less costs for maintenance and allows for fault tolerance and high availability. While horizontal scalability is offered by multiple Cloud providers, vertical scalability is a less addressed topic. In this article we first define the meaning and outline the benefits of doing vertical scalability. We also present a scaling solution which can automatically provision services based on the needs and resource usage of the system.**

***Keywords*—scalabilty, elasticity, resource provisioning, BigData, IoT**

I. Introduction

Today various embedded devices are capable of communicating and sharing data using the Internet. In this manner traditional web services are enriched with physical world services. In addition to the IoT vision, which gives every device an IP address and interconnects them, there is also the notion of Web of Things (WoT) which enables the devices to speak the same language.Current real-time processing is mainly done on existing web data but the extension to considerably larger amounts of data produced by multiple sensor networks requires research and design of robust and scalable processing platforms.

Scalability can be described as the capacity to handle increasing workloads [1], or the ability to improve performance when resources are added [2]. In many articles (e.g. [3] and [4]) scalability is divided into two main categories:

- Vertical scalability : adding resources to the same logical unit (e.g. to a cluster node)
- Horizontal scalability: adding multiple units of resources [4] (e.g. adding multiple nodes to a cluster)

According to [5], scalability is important if it is useful. The multiple definitions of scalability try to take into account what is useful for the domain and to prove a point about the system/algorithm/application performance, how certain workflows affect the system, cost efficiency and the ability of a system/application to scale.

While horizontal scalabilty is achieved at the infrastructure level by many Cloud providers, services can also be scaled to further optimize existing applications.

Some of the benefits of this type of scaling can be:

- Greater precision when measuring service utilization
- Enforcing SLAs or quality levels
- Cost optimization
- Personalized usage/Usage patterns
- Less interventions from cluster adminstrator

In this paper we propose a solution that does scaling for Hadoop based applications, in a Cloud environment. It employs three stategies: utilize fewer resources, maximum throughput, keep resource utilization under a threshold. The application includes automatic and manual resource allocation and a metrics monitor.

The paper has the following structure: first we review related work and state of the art in chapter II. In chapter III we propose our high-level solution while in chapter IV we showcase implementation details. Chapter V describes the experimental results and we conclude our work in chapter VI.

II. Related work

A. Node scaling

The most established Cloud Providers that achieve automatic scaling are described below.

Amazon Elastic Cloud Compute (EC2) is a web service which offers computing power to users [6]. Auto scaling in EC2 has the following components [7]: groups (a collection of EC2 instances), launch configurations (used when creating new instances) and scaling plans (they choose how to scale a group).

There are several scaling plans: maintain current number of instances running, manual scaling, scale based on a schedule and scale on demand. Maintain current number of instances running is accomplished by doing regularly health checks. If necessary, the unhealthy instance is terminated and a new one is launched instead. This is the default plan. Scale based on a schedule is done by performing time-based scaling operation Scale based on demand aka policy-based scaling. A policy is a set of rules executed by Auto Scaling in response to an alarm. An alarm is an object that monitors a metric for a specified amount of time. An Auto scaling Group can have multiple scaling policies.

Google App Engine [8] is formed by multiple services. Each service has two components: source code and configuration file. The scaling type is specified in the configuration

file and has three choices: manual scaling, basic scaling and automatic scaling. In basic scaling a instance is created when a request is received and is destroyed when the application is idle. In automatic scaling a instance is created/turned off on demand based on different application metrics.

Microsoft Azure[9] supports Azure Autoscale: dynamically add or remove instances based on schedule and/or on runtime metrics. In addition, Azure Resource Manager Rest API and/or Azure Service Management Rest API can be used for autoscaling. Azure can also use third-party services like Paraleap AzureWatch.

B. Service scaling

Periscope [10] brings QoS and autoscaling to Yarn applications. It monitors the application progress, the YARN containers and their resources and available cluster nodes to make decisions. Periscope is also capable of:

- Application re-prioritization
- Enforcing SLAs
- Enforcing guaranteed cluster capacity

The company behind Periscope has been acquired by HortonWorks[11] and the product has been integrated into CloudBreak [12]. However it's not clear whether the project has been continued and we could not find information regarding experimental results.

Qubole [13] analyzes Job Tracker information to predict future load of the cluster and make scaling decisions. It uses StarCluster [14], an open source cluster management software. Basically Qubole works in the following way:

- Nodes in the cluster send launch time to the JT
- JT monitors pending work and computes the remaining workload. If the remaining workload is higher than a preset threshold, JT will add more nodes.
- In the same way, JT will decommission nodes if the remaining workload does not justify keeping it.
- Nodes in the cluster send launch time to the JT

Qubole was designed to address scaling in a specific platform. Also no experimental results have been published so far.

III. System Overview

The system we propose can be functionally described in figure 1.

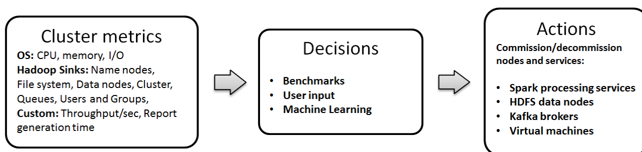

Fig. 1: Functional diagram

The application regularly polls the cluster for OS and service level metrics. Then it generates comissioning/decomissioning decisions which are sent to the cluster manager. This translates into storage/processing services being added or dismissed dinamically. In this paper specifically, the software scales YARN Node managers and Data nodes.

In our implementation the cluster manager role is perfomed by Apache Ambari. Ambari is an open source framework for provisioning, managing, and monitoring Apache Hadoop clusters [15], [16]. It provides a REST API and a web interface for monitoring and management. The REST API exposes calls to comission/decomission Hadoop services.

A. Architecture

A more detailed view of the workflow shows six main components: the Metrics Monitor, Hadoop Sinks, Metrics Collector, Ambari Server, Ambari Rest API, BigScale and Hadoop Cluster. The picture below shows the relationship between them.

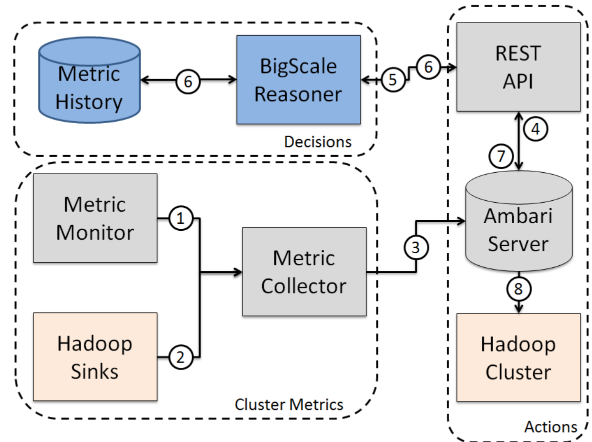

Fig. 2: Main workflow

The steps performed by the applications are the following:

- 1 - Metric monitors send system-level metrics to Metric Collector
- 2 - Hadoop Sinks send Hadoop-level metrics to Metric Collector
- 3 - Metric Collector stores and aggregates metrics
- 4, 5 - Metrics are retrieved by BigScale through a REST API exposed by Ambari
- 6 - Based on retrieved metrics and historical data determine what service components must be scaled
- 7 - Send requests to Ambari server to scale services
- 8 - Ambari performs the corresponding action on the Hadoop Cluster

The resource provisioning application dynamically scales slave components (Data nodes and Name nodes) to ensure the required infrastructure for a YARN application and to

conform with system administrator requirements. In order to offer different options to users, there are three automatic scaling strategies.

1) Use fewer resources - This is done by having equal number of running Node Managers and Data Nodes
2) Balanced with minimum resources - The goal is to recommission/decommission Node Managers to maintain a certain level of resource utilization. The condition for recommissioning a Node Manager is that the memory utilization is above threshold t1 or CPU utilization is above threshold t2; and for decommissioning a Nodemanager is that the memory utilization is below threshold t3 and CPU utilization is below threshold t4. The CPU and Memory threshold (t1, t2, t3, t4) may differ.
3) Highest performance - Allocate all processing power, Node Managers must be running - to achieve less running time for an application.

In strategy 1 the number of data nodes is equal to the number of node managers and new node managers are added if the containers exceed their allocated processing power. In strategy 2 scaling is done by setting thresholds for node managers. A new node manager is recommissioned if they exceed a certain threshold for a longer period of time. They are decommissioned if the resource utilization is below the specified threshold. In strategy 3 the running time of the applications is reduced by using all the available resources.

Data nodes behave the same regardless of the strategy: they scale up and down if the space and threshold requirements are not satisfied. Also, if there are no applications running, the recommission operation for data node/namenode is prohibited. For the decommissioning command, there is a different free space available condition for Data Nodes.

The number of Node Managers scales down to be equal to the number of Data Nodes. This operation is necessary for the system to be more cost-effective.

The application collects metrics offered by Ambari Metrics from the master services (YARNs Resource Manager and HDFSs Name node) and depending on the scaling strategy it sends commands to the Ambari Server. Commands are targeted for slave components of the master components mentioned above (master Resource Manager to slave Node Manager and master Name node to slave Data Node).

Decommissioning and recommissioning data nodes is done by sending a request through Ambari REST API to the Name node to include/exclude the data nodes hosts. Decommissioning and recommissioning Node Manager is done also by sending a request through Ambari REST API to Resource Manager to include/exclude node managers hosts. In addition, Resource Manager stops the decommissioned host. So when a Node Manager is recommissioned it is restored to its previous state.

The diagram below 3 shows how the automatic scheduling application works. For the scaling operations to function properly, a configuration file must be read in order to know the Ambari server address, cluster name, scaling strategy and other parameters. Once the configuration file is read, we enter a loop in which we make Node Manager and Data Node scaling operations.

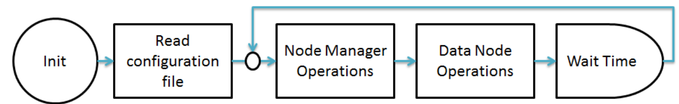

Fig. 3: Application flow

The wait time is used to collect enough relevant metrics. We collect 1 metric per second and if there are enough metrics and the scaling conditions are satisfied, the scaling operations are executed.

B. *Decision making*

We describe in detail how the main function of the application takes decision. The following abbreviations are used: NM (Node Manager), DN(Data Node).

In the Node Manager Operations module the scaling operations are based on metrics taken from the YARN Resource Manager. If there is no application running, the number of Node Managers in-service and started must be equal to the number of Data Nodes in-service and started. By doing this we keep resource utilization to a minimum. If there are applications running, then we take into consideration metrics and scaling strategies 4.

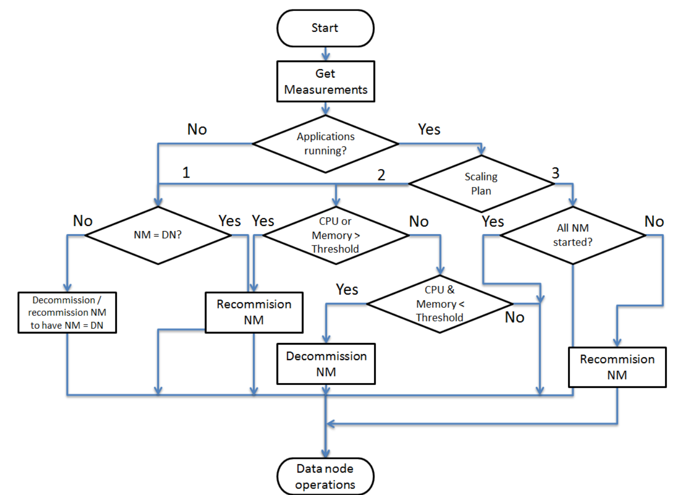

Fig. 4: Node Manager Operations

In the case of Data Node autoscaling, there is a single scaling strategy because storage scaling is needed to keep up with increasing amounts of data. According to the figure above, if there is no YARN application there is a restriction for recommissioning a Data Node even if the conditions are satisfied. In this case a Data Node can be decommissioned if the space is higher than fs1.

If there are applications running, the next conditions are applied. If the used percentage is below threshold t1 and the free space is less than fs3 then a Data Node is recommissioned.

This condition is required if the free space is very low and the threshold conditions are not met, because if this condition does not exist, the time remaining to fill all the HDFS space may be lower than the time to recommission a Data Node and the Name node to take notice of that change. If used space percentage is higher than threshold t1 and free space is lower than fs4, a Data Node is recommissioned. If there is much free space and used space percentage is above threshold t1, recommissioning is a mistake, because we have more than enough space for the current workload to run for a longer period of time. Note that fs4 is much higher than fs1. If the free space is higher than fs2 then a Data Node is decommissioned. It is recommended that fs2 should be higher than fs1.

The figure below 5 shows how Data node Operations are decided. The values fs1, fs2, fs3,fs4 and t1 are taken from the configuration file.

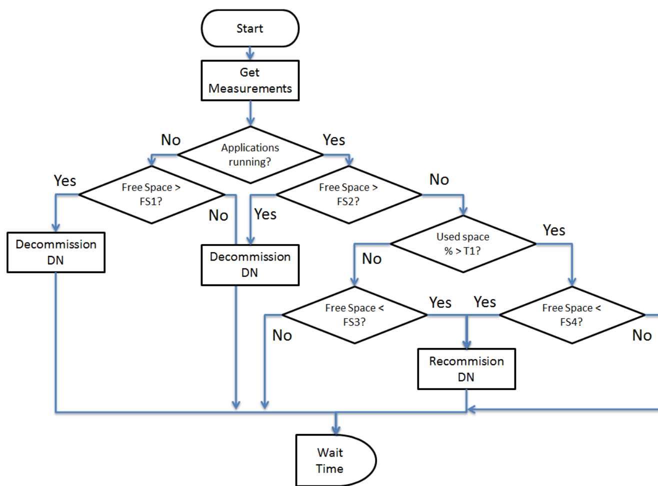

Fig. 5: Data Node Operations

In conclusion, in order for these conditions to work as intended for Data node scaling operations, the following relationships between free spaces specified in the configuration file must be true: fs2 should be higher than fs1, fs3, fs4 and fs3 must be less than fs4.

The tool leaves room for administrator intervention. The manual allocation operation flow is: it reads from the configuration file the Ambari server address, cluster name, etc. and afterwards it enters a loop. In that loop, it waits for the user input command. When a command is entered, it executes that command and returns the possible output.

The main parameters that need to be defined in the master configuration file are:

- Ambari server authentication
- Ambari server address and cluster name
- Used space percentage and free space threshold for DN scaling operations
- CPU and Memory utilization threshold for NM scaling operations
- Overhead time between two successive scaling operations
- Number of measurements taken
- Scaling strategy

A scaling option can be done if:

- the constraints stated in the scaling strategies are accomplished
- the measurements are equal to the specified number of measurements
- enough time has passed since the previous scaling operation

The time is measured after the scaling operation has finished. After each scaling operation the measurements are reset in order to avoid previous values influencing future scaling decisions. The time between two measurements is one second and measurements are stored in a local list.

One advantage of gathering more measurements is that the spikes in the measurements gathered have little effect on the scaling decision. In addition, the overhead time between two successive scaling operations is needed to make sure the master components notice the system change and can update their info accordingly. To make sure the scaling operation is finished (decommission/recommission), it is necessary to do polling, because the Ambari REST API sends back as a response an HTTP link with the progress of the operation. In other words, Ambari Server sends through REST API an asynchronous response.

In order to know the status (STARTED/ DECOMMISSIONED) of the slave components and to not overwhelm the Ambari server with requests, at the beginning of the program a request is sent to Ambari to store the status of slave components in a local list; after each scaling operation we update the local list.

If there are no running YARN application then the wait time between two successive Node Manager components is 0, because it doesnt depend on any metrics given by the Resource Manager.

To ensure that the final result of executing multiple tasks for a certain slave component is correct, all tasks run in sequential order for a slave type. In other words, Data node tasks and Node Manager tasks can run in parallel, but each task belonging to a certain slave type is executed sequentially.

IV. Implementation Details

A. Workload types

In this subsection we analyze the projected results when different types of YARN applications run in an Apache Ambari environment. Regardless of the workload, an increase in in the number of Node Manager components will result in a decrease in the running time of the YARN application.

1) CPU and I/O: The Wordcount program, like the Pi calculation program, is interesting because is has a CPU and a I/O component. In this case, it is expected a reduction in time by recommissioning Node Managers but the monitoring is trickier in some situations because CPU usage level is dependent of I/O part.

2) I/O intensive: For an I/O intensive program, like Teragen or TESTDFSIO, adding Node Managers will reduce time and lower the resource utilization levels. The workload's objective is to write large quantities of data so adding Data Nodes will suffice. In most cases, adding a Data Node is followed by adding a new Node Manager and therefore the resource utilization will probably be lower. Because of the Hadoop write-once policy we expect applications like Teragen not to produce spikes in CPU and Memory utilization levels.

B. Experimental setup

Our experimental setup is based on OpenStack [17], an open-source cloud solution used for management and deploying IaaS infrastructure. It can scale up to 1 million physical machines, up to 60 million virtual machines and billions of stored objects. The cluster we used for prototyping has the following configuration:

VCPUs	RAM (MB)	Disk (GB)	State
4	4096	24	Active
1	1536	16	Active
1	1536	16	Active
1	1536	16	Active
2	4096	10	Active
4	4096	24	Active
1	1536	16	Active
1	1536	16	Active

Although we achieve promising results with 8 machines, part of future work is to extend the experiments to larger clusters like GRID 5000 [26].

The metrics are collected as follows:

- Node Managers - CPU and Memory Utilization
- Data Nodes - Used space percentage and free space

On top of the cluster we installed an Apache Ambari Server[15] with the following services: HDFS[18], YARN[19], MapReduce[20], Ambari Metrics[21] and Zookeeper Server[22]. Zookeeper is used for coordinating distributed applications.

V. Experimental results

A. Test applications

The experiments used in this chapter express different workloads. The dynamic scheduling applications responsibility is to handle these workload properly. The following test applications are used: Teragen, Wordcount and Pi.

1) Teragen: Teragen application is an I/O write application that generates specified quantities of data.([34]) thus the nodemanager and data node scaling operations are tested by generating large quantities of data.

2) WordCount: The wordcount application sums up the number of appearances of each word in the input text([35]). In this experiment, we want to see how a workload of a program that utilises both CPU and I/O is handled.

3) Pi: The pi benchmark approximates the value of pi using quasi-Monte Carlo method.([36]). The experiment utilizes this program only for demonstrating the decrease of application running time when Node Managers are recommissioned.

B. Assumptions and results

In this subsection the truthfulness of the assumptions made and how the application behaves in different benchmarks are tested.

1) More computational resources decrease running time: if we add more Node Manager components the application running time will decrease.

For this test we used the PI benchmark, because in other benchmarks the I/O part may interfere and, as a consequence, the results may be inconsistent. We obtained the following results:

Nr Node Managers	Runtime(seconds)
2	204.341
4	108.144
8	66.881

The figures below 6 show a graphic representation of resource utilization when an application is run three times, each time with a different number of node managers. First time it is executed with 2 NM, second time with 4 and third time with 8.

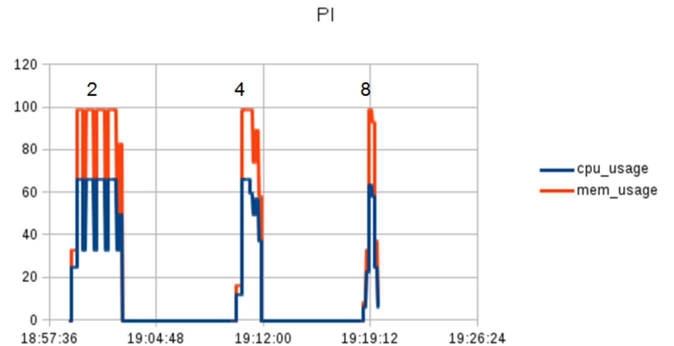

Fig. 6: Pi CPU and memory usage

2) Scale in when no application is running: if there are no applications running, the slave components will scale in or out so the number of node manager components match the number of data node components and data node components will scale in if the free space condition specified in the configuration file is fulfilled

In the current setup, free space threshold is set to 20 GB, 8 active Node Manager components, 5 active Data Node components and the dfs replication is set to 2. There is a difference as shown in the figure below, between the time a Data Node was decommissioned and the time the remaining

free space decreased. This happens because the name node needs some time to keep up with the system changes.

The node manager components will scale in order to match the number of data nodes 7.

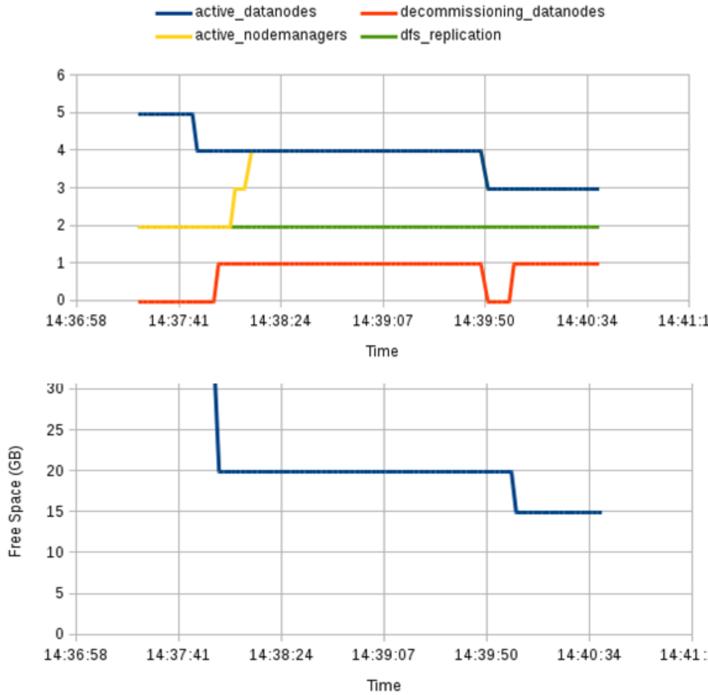

Fig. 7: Components scale in

Even though the free space condition is fulfilled, we cannot decommission a Data Node, because the number of decommissioning data nodes is equal to dfs replication plus 1. This condition is necessary for having a backup of the data in other active Data Node components.

3) Scaling while application are running: Teragen Mapreduce application is designed to write a large file for the terasort benchmark.

Scaling Strategy 1. The following setup is employed: 2 Data Nodes, 2 Node Manager, free space 9.9 GB, free space threshold 10 GB if used space is above 80% and 5 GB if less, node manager and data node add cooldown 30 s

In 8 it can be observed that a node manager is recommissioned after a data node is recommissioned.

It also shows that when a data node is recommissioned, the free space increases. A data node is recommissioned, in this scenario, if the free space reaches under 15 GB free and used space is above 80%. The second condition, free space under 5 GB and used percent under 80% is never fulfilled. The free space continues to go down even if the data nodes are recommissioned because it takes some time for Name node to take notice of the system change. This will happen in all strategy plans if the data node scaling conditions are satisfied.

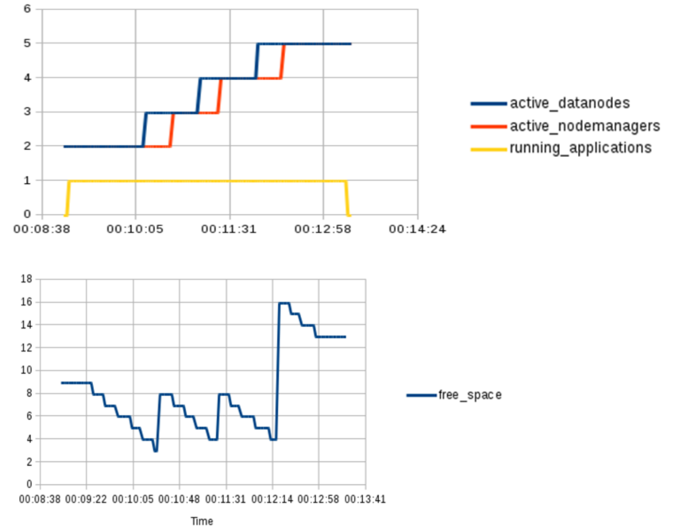

Fig. 8: Node manager recommissioned after data node

Figure 9 shows that the increase of Node Managers will decrease for this program the CPU and memory usage.

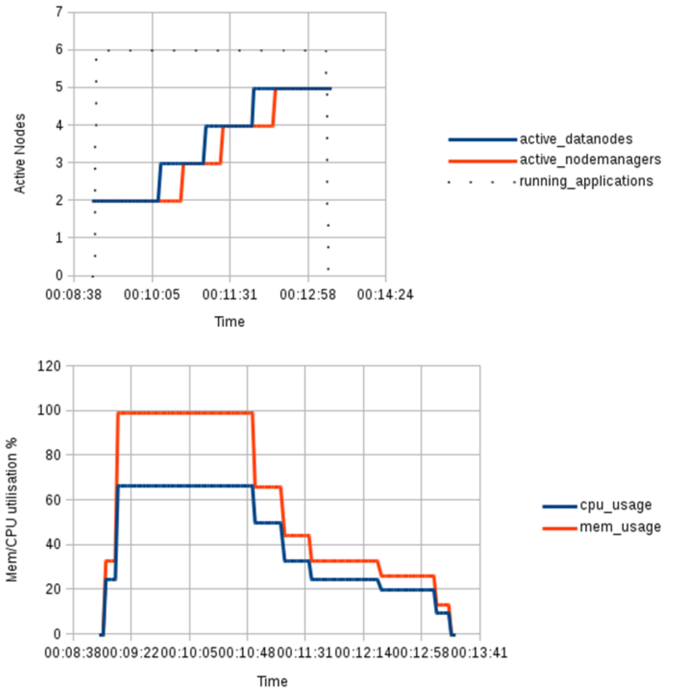

Fig. 9: Scaling out

Scaling Strategy 2. The following setup is employed: 2 Data Nodes, 2 Node Managers, free space 9.9 GB, free space threshold 11 GB if used space percent is above 80% and 5 GB if less, node manager and data node add cooldown 30 s, recommission node manager CPU threshold 50% and memory 60%.

Figure 10 shows that the Node Managers scale before Data Nodes in the first 2 cases, because their scaling condition is fulfilled before they are forced to scale to match the number of active data node.

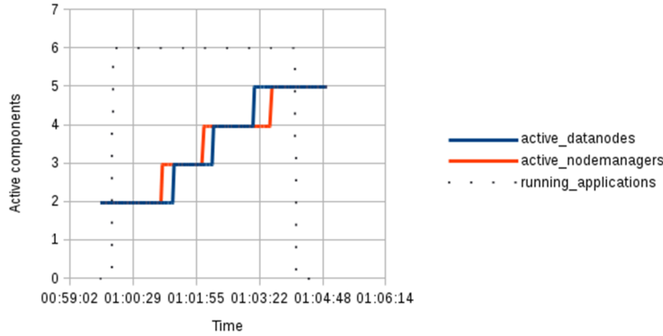

Fig. 10: Active Nodes

Figure 11 shows how the CPU and memory usage lowers as we recommission Node Managers.

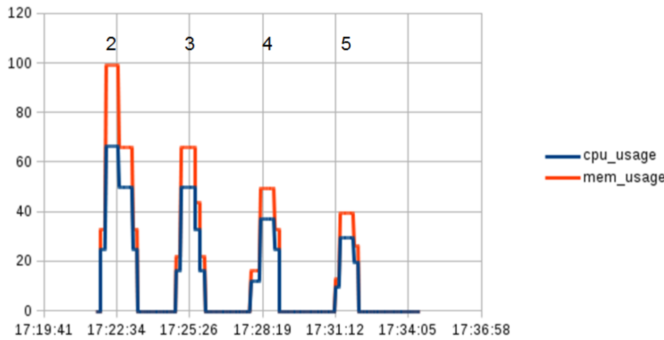

Fig. 11: Scaling out node managers

Scaling Strategy 3. The following setup is employed: 2 Data Nodes, 2 Node Managers, free space 9.9 GB, free space threshold 11 GB if used space is above 80% and 5 GB if less, maximum Node Managers 8.

In this strategy plan, the node managers will increase at the maximum capacity, independent of the CPU and memory metrics. In 12 the number of active Node Managers is 7, not 8, because the program finished running before the final node manager is recommissioned.

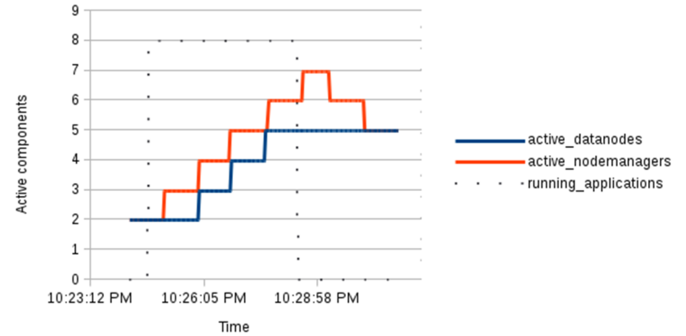

Fig. 12: Active Nodes

Figure 13 shows the effect of successive recommissioning node managers: the time decreases and the CPU and memory usage lowers.

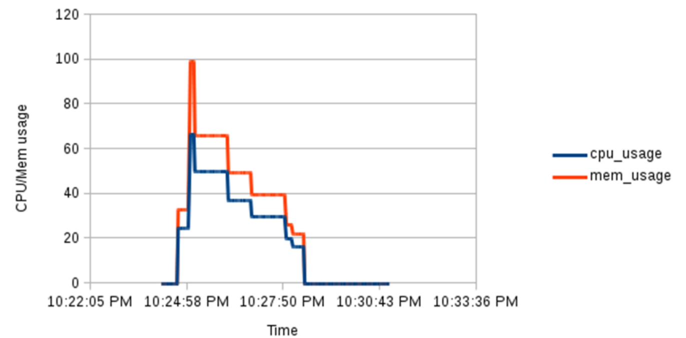

Fig. 13: CPU and memory usage

For a better view, in 14 and 15 all strategies are compared. To easier observe the difference, the time line was changed so that all applications executions start at the same time.

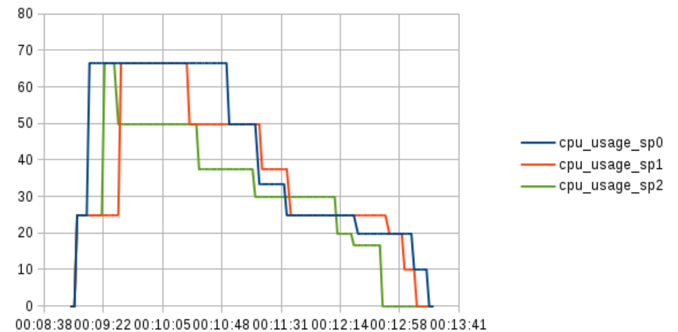

Fig. 14: CPU usage for all three strategies

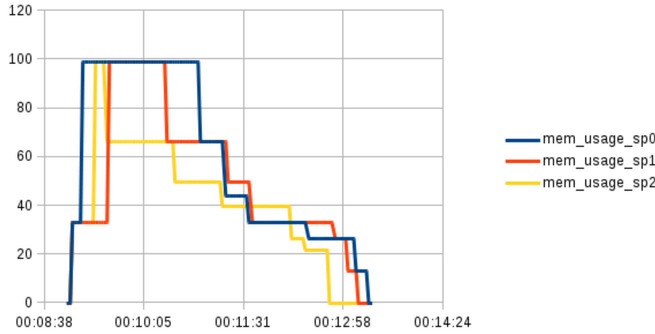

Fig. 15: Memory usage comparison

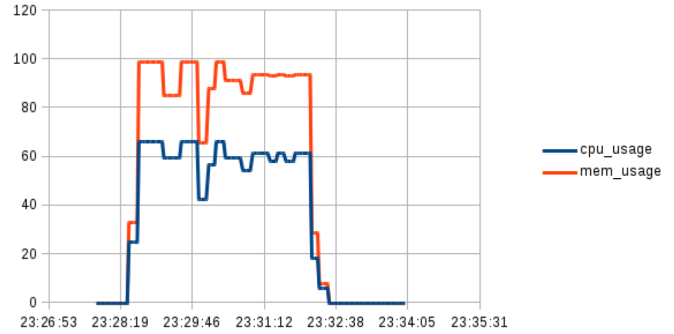

Fig. 17: CPU and Memory utilization in scaling strategy 2

As we can observe in all the Teragen figures, there are no spikes and the line is smooth. An explanation for this event is that the policy of the Hadoop one-writer-many-readers limits the number of writers that are concurring for the writing access to just one. The time is reduced, depending on how many Node Managers are added and how soon they are added.

The strategy that utilizes the least amount of resources is scaling strategy 2. However, it can recommission more Node Managers than necessary. For this reason it is not recommended for this type of program.

4) Wordcount: Setup: 2 Data Nodes, 2 Node Managers, free space 9.9 GB, free space threshold 11 GB if used space percent is above 80% and 5 GB if less, 8 maximum Node Managers

In the current setup, data node scaling operations are not needed, because the scaling conditions are not satisfied (the free space available does not lower down enough). As a consequence, applying scaling strategy 1 has no effect on the system.

There is little difference between scaling strategy 2 and 3, as it can be seen in 16 and 17.

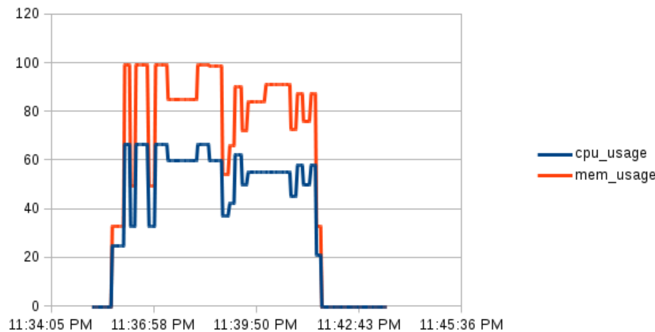

Fig. 16: CPU and Memory usage in scaling strategy 1

The difference is given by the time needed to gather the measurements, satisfying Node Manager thresholds and the overhead between two successive scaling operations in scaling strategy 2.

VI. Lessons and Limitations

A data node is not recommended to be decommissioned if there are $dfs.replication - 1$ nodes in the decommissioning state, because there is a possibility that a certain file would not exist on any available nodes if the decommissioning operation is done

If there is no application running there must be an add prohibition for slave components. This option would not be so relevant if the application has a prediction component. In the current state, there is no way to know when a new application will start so the best action is to keep resource utilization to a minimum.

The reaction time of master components to system changes must be taken into consideration when the metrics are obtained, because some metrics may be compromised during this time interval and should not be taken into consideration.

There must be a time interval between two scaling operations applied on the same component type so not to overwhelm the system with requests and to leave time for Resource Manager to allocate those new resources, master components to notice the system changes and to collect enough metrics for calculated decision.

After a scaling operation the previously obtained metrics must be reset so not to influence future metrics. The metrics used to measure computation resources utilization have a certain error, because there is a level of abstraction in place and there is no way to access what happens inside a container. The metrics we used are formed by making a ratio between the number of resources allocated and the number of resources available.

The number of machines that can be added by a privileged user is limited by the OpenStack infrastructure.

The decommissioning data nodes operation can last an extremely long time, because all the files that are stored in the decommissioned data node host must be moved to other data nodes to maintain the replication factor.

If a slave component is deleted, then the master component must be restarted. This is happening because Resource Manager and Name node does not run in high availability mode. A consequence is that there is an overhead time caused by restarting the master and the slave components, in order to determine them to resume their previous state. Luckily, if a component is added, there is no such restriction.

Restarting the master component can take a very long time. For example, if the Name node is restarted, it must index all the files stored in the HDFS and this operation will be time-consuming.

In Hadoop only one writer is allowed at a certain moment, but there can be many readers [23] This will limit the throughput for write operations, but will increase it for read operations.

Ambari Metrics Collector runs in embedded mode (default option), because the cluster has a small size (eight nodes). By running in distributed mode, metrics are stored in HDFS and therefore there is an additional network overhead that will limit the applications throughput. Additionally, the Name node restart will take longer because it has to index all the files in the HDFS. The metrics are obtained through Hadoop services master components: YARNs Resource Manager and HDFSs Name node. The same information can be obtained through the slave services: Data Node and Node Managers. The upside is that there is no overhead time to take notice of system changes (e.g. decommissioning a data node, recommissioning a Node Manager), but the downside is that the time for getting the metrics through increases with the growth of the slaves.

The time for getting a metric through Ambari REST is fairly long. In order to reduce the time of a regular Ambari REST get command, a partial request is used to select and retrieve only the metrics from the master components that are important to decide what scaling operations should be done.

The default deletion time of a file in Hadoop is 360 seconds. We changed the value to 0 to reduce the time between tests.

The small files are a huge problem for Hadoop, because it spends a lot of time managing their metadata information and as pointed [23] the memory usage is also high (63.53%). That happens because the metadata information is stored in the Name node memory [24]. Another issue for small files is that Name node restricts the number of files stored in the HDFS and according to [24], accessing a large number of these files results in a bottleneck in Name node. Furthermore, high latency is expected when reading small files [25] and the throughput falls below expectations.

VII. Conclusions and Future Work

In this article we have argued for the importance of scalability in distributed systems, specifically vertical scalability and its implications. We have outlined the benefits of this approach and proposed a solution capable of automatic scaling in and out of a Hadoop cluster. We have described 3 scaling strategies and ran experiments on multiple types of workloads. The experiments demonstrate that such an approach is feasible and can integrate easily with other cluster components.

Although we achieved promising results, we intend to extend our experiments to larger clusters such as GRID5000 [26]. The experiments will also involve multiple concurrent applications competing for the same resources. This will imply the development of a scheduling algorithm and the encryption of the transmitted data. The scheduling algorithm will need to employ a check-pointing mechanism in order to resume the applications once they are able to run.

Future work will also involve scaling services such as processing, messaging and storage. The application will also be transformed to allow for more abstract scaling expressions. In this sense the administrator will input QoS/SLA parameters (e.g data processing throughput or response time) and BigScale will adjust the cluster to those purposes. The underlying infrastructure will also receive horizontal scaling recommendations.

Currently, the application performs regular polling operations of the cluster metrics. As part of a modular solution, we will employ an event-based architecture to reduce the resulting overhead. Commissioning and decommissioning will be done in a parallel fashion, in order to speed up the short-term capability of the scaling process.

Machine learning techniques will be employed to profile the workloads and make predictions on the resource needs of the cluster. In this sense we intend to use fewer thresholds and input parameters so that the administrator will delegate the scaling responsibility entirely to the application.

References

[1] Garcia, Daniel F., G. Rodrigo, Joaqun Entrialgo, Javier Garcia, and Manuel Garcia. "Experimental evaluation of horizontal and vertical scalability of cluster-based application servers for transactional workloads." In 8th International Conference on Applied Informatics and Communications (AIC'08), pp. 29-34. 2008.

[2] Agrawal, Divyakant, Amr El Abbadi, Sudipto Das, and Aaron J. Elmore. "Database scalability, elasticity, and autonomy in the cloud." In International Conference on Database Systems for Advanced Applications, pp. 2-15. Springer Berlin Heidelberg, 2011.

[3] Mei, Lijun, Wing Kwong Chan, and T. H. Tse. "A tale of clouds: paradigm comparisons and some thoughts on research issues." In Asia-Pacific Services Computing Conference, 2008. APSCC'08. IEEE, pp. 464-469. Ieee, 2008.

[4] Anandhi, R., and K. Chitra. "A challenge in improving the consistency of transactions in cloud databases-scalability." International Journal of Computer Applications 52, no. 2 (2012).

[5] Hill, Mark D. "What is scalability?." ACM SIGARCH Computer Architecture News 18, no. 4 (1990): 18-21.

[6] Amazon EC2, [https://aws.amazon.com/documentation/ec2/

[7] Vaquero, Luis M., Luis Rodero-Merino, and Rajkumar Buyya. "Dynamically scaling applications in the cloud." ACM SIGCOMM Computer Communication Review 41, no. 1 (2011): 45-52.

[8] Google App Engine https://cloud.google.com/appengine/docs/python/an-overview-of-app-engine

[9] Microsoft Azure, https://azure.microsoft.com/en-us/documentation/articles/best-practices-auto-scaling/

[10] Benedict, Shajulin, Ventsislav Petkov, and Michael Gerndt. "Periscope: An online-based distributed performance analysis tool." In Tools for High Performance Computing 2009, pp. 1-16. Springer Berlin Heidelberg, 2010.

[11] HortonWorks http://hortonworks.com/

[12] Whelan, Christopher W., Jeffrey Tyner, Alberto L'Abbate, Clelia Tiziana Storlazzi, Lucia Carbone, and Kemal Snmez. "Cloudbreak: accurate and scalable genomic structural variation detection in the cloud with MapReduce." arXiv preprint arXiv:1307.2331 (2013).

[13] Qubole, https://www.qubole.com/blog/product/industrys-first-auto-scaling-hadoop-clusters/

[14] StarCluster http://star.mit.edu/cluster/

[15] Ambari https://cwiki.apache.org/confluence/display/AMBARI/Ambari

[16] Dagli, Mikin K., and Brijesh B. Mehta. "Big Data and Hadoop: A Review." International Journal of Applied Research in Engineering and Science 2, no. 2 (2014): 192.

[17] Sefraoui, Omar, Mohammed Aissaoui, and Mohsine Eleuldj. "OpenStack: toward an open-source solution for cloud computing."

[18] Borthakur, Dhruba. "HDFS architecture guide." HADOOP APACHE PROJECT http://hadoop. apache. org/common/docs/current/hdfs design. pdf (2008): 39.

[19] Vavilapalli, Vinod Kumar, Arun C. Murthy, Chris Douglas, Sharad Agarwal, Mahadev Konar, Robert Evans, Thomas Graves et al. "Apache hadoop yarn: Yet another resource negotiator." In Proceedings of the 4th annual Symposium on Cloud Computing, p. 5. ACM, 2013.

[20] Dean, Jeffrey, and Sanjay Ghemawat. "MapReduce: simplified data processing on large clusters." Communications of the ACM 51, no. 1 (2008): 107-113.

[21] Wadkar, Sameer, and Madhu Siddalingaiah. "Apache ambari." In Pro Apache Hadoop, pp. 399-401. Apress, 2014.

[22] Hunt, Patrick, Mahadev Konar, Flavio Paiva Junqueira, and Benjamin Reed. "ZooKeeper: Wait-free Coordination for Internet-scale Systems." In USENIX Annual Technical Conference, vol. 8, p. 9. 2010.

[23] Liu, Xuhui, Jizhong Han, Yunqin Zhong, Chengde Han, and Xubin He. "Implementing WebGIS on Hadoop: A case study of improving small file I/O performance on HDFS." In 2009 IEEE International Conference on Cluster Computing and Workshops, pp. 1-8. IEEE, 2009.

[24] Chandrasekar, S., R. Dakshinamurthy, P. G. Seshakumar, B. Prabavathy, and Chitra Babu. "A novel indexing scheme for efficient handling of small files in hadoop distributed file system." In Computer Communication and Informatics (ICCCI), 2013 International Conference on, pp. 1-8. IEEE, 2013.

[25] Dong, Bo, Qinghua Zheng, Feng Tian, Kuo-Ming Chao, Rui Ma, and Rachid Anane. "An optimized approach for storing and accessing small files on cloud storage." Journal of Network and Computer Applications 35, no. 6 (2012): 1847-1862.

[26] Bolze, Raphal, Franck Cappello, Eddy Caron, Michel Dayd, Frdric Desprez, Emmanuel Jeannot, Yvon Jgou et al. "Grid'5000: a large scale and highly reconfigurable experimental grid testbed." International Journal of High Performance Computing Applications 20, no. 4 (2006): 481-494.