THE AFFORDABLE SUPERCOMPUTER

Harrison Carranza, Aparicio Carranza and Jose Reyes Alamo
Computer Engineering Technology
New York City College of Technology of The City University of New York
186 Jay Street, Brooklyn, NY 11201

ABSTRACT - ***In order for an enterprise to resolve complex calculations and perform CPU intensive tasks, a powerful computer such as a mainframe is usually required. Mainframes are capable of performing large scale transaction processes, supporting thousands of users, and managing terabytes of data in databases. Large corporations, such as Metro North, Bank of America, The New York Times, and Rite Aid, all rely on powerful mainframes to handle their data. Although they are efficient, solutions such as mainframes come at a very high cost. Unlike large Enterprises with more resources and more money to spend, smaller Enterprises with smaller IT budget are unable to afford these powerful systems. This paper introduces MOSIX computer clustering, a cost effective alternative to mainframes.***

Keywords: MOSIX, OpenMOSIX, Supercomputer, Computer Clustering, ClusterKnoppix.

I. INTRODUCTION

Take a moment to think about how you use your computer on a daily basis. You will realize that the majority of the time that you are actually using your computer, you are performing non-intensive tasks such as browsing the web, playing solitaire, or simply editing a simple document. During these moments, your computer is barely using 5% of its overall processing power. Then there are those other times when you need an immense amount of processing power and you're already using close to 100% of your overall processing power. This situation has probably happened to you at some point on time and you probably thought to yourself, "I wish my computer had more power." Well, wouldn't it be great if you had more power at your disposal when trying to get these intensive tasks done? That is where the idea of clustering plays its role. The idea of cluster computing is to spread the amount of work that needs to be done across all nodes or computers, using the available resources of all the computers in a network.

Whether they are new or already existent, the majority of today's Enterprises will eventually grow. Therefore, these enterprises will eventually face with the need to have more processing power. When this situation comes up, often times the solution is usually for these companies to use mainframes to obtain this kind of processing power. But why waste that kind of money on a mainframe when you can use the processing power which you already have to achieve the same type of processing power?. Computer clustering allows enterprises to achieve this type of processing power through the help of clustering technology, overall providing for a cost effective alternative to mainframes.

II. WHAT EXACTLY IS COMPUTER CLUSTERING?

Computer clustering is simply a technology that allows two or more computers interconnected through a local area network to gather resources in order to perform CPU intensive tasks. These tasks can range anywhere from loading a simple spreadsheet to solving complex mathematical algorithms and handling a huge database. There are mainly two types of clusters: High Availability clusters (HA) and High Performance clusters (HP).

High availability clusters are created to provide the system with maximum redundancy. To achieve this goal, redundant nodes are maintained to be used as a backup system in case there is ever a failure within the system. Since a cluster is comprised of many nodes, even if one computer were to fail, the others would still be able to continue working.

High performance clusters, unlike that of high availability clusters, are used in order to provide maximum performance. If a system needs to perform an extremely intensive task, the cluster is able to divide the work through all the nodes in the cluster. High performance clustering is commonly used in scientific calculations and graphic rendering.

III. WHY USE COMPUTER CLUSTERING?

With the way the world's economy is today, many large businesses in the market are trying to maximize efficiency while being cost effective at the same time. The more an enterprise continues to grow, the more processing power it will require to handle transactions, huge amounts of data, users, and large databases. While a mainframe is the most common solution, it is also one of the most expensive one. Companies with a low IT budget don't have the resources to afford a powerful system solution such as a mainframe.

The alternative to mainframes is the use of clustering technology. By using clustering technology, enterprises are not forced in resorting to the use of these expensive mainframes. By interconnecting their already existing computers into a huge cluster, all the processing power which, in most cases, would not be used by users, can now be used to improve the processing capabilities of their enterprise.

There are many clustering solutions out there today but the majority of them are available only on Linux. This is due to its reliability, its security, and its free access. This will allow an enterprise to create a supercomputer without investing any part of their budget whatsoever.

IV. ADVANTAGES OF CLUSTER COMPUTING

Cost Effective – Mainframes are extremely expensive when compared to clustering technology. Most enterprises already have a full network of desktops at their disposal so there is no need to waste more of their budget on improving processing power when the processing power they need is already existent [4].

Similar Power – Depending on the number of computers in the cluster, a computer cluster can perform similarly to that of a mainframe. The price of a unit of power in a computer cluster is

less than that of a mainframe. Therefore, enterprises receive a greater rate of investment in clustering than they are able to receive from that of mainframes [4].

Modern Networking – The speed of the local area network has always been an issue. Bandwidth was once considered a luxury, but computer technology has advanced at an alarming rate. With technologies such as fast Ethernet and gigabit Ethernet at our disposal, networking lag is not an issue [4].

Scalability – Since the processing power of computer clusters are dependent on the number of nodes in the network, the more nodes you add, the more power your cluster will have. Mainframes are powerful computers, but they have a fixed processing capacity. Clusters, on the other hand, only require more nodes to be added if the enterprise were to need more power [4].

Redundancy – One of the biggest advantages to clustering is that if one node were to fail, the cluster would still be able to operate at almost maximum processing capacity with not much problem. This is due to the cluster's ability to simply transfer the work onto another node, allowing for no interruption in the system performance. As opposed to a cluster failing, if a mainframe were to fail, then the entire enterprise would have a huge problem at their hands [4].

V. CLUSTERING TECHNOLOGIES

a. MOSIX

Computer clustering is a great alternative to mainframes. However, although there are various different types of clusters, not all are ideal. The most well-known type of Linux clustering is Beowulf clusters. A Beowulf cluster simply takes two or more machines interconnected through a LAN and uses their resources to complete CPU intensive tasks, just like a regular cluster. The difference with Beowulf clusters is that in order to achieve this, it must use special clustering libraries such as PCM and MPI to use with specific cluster applications. Since programmers can write special applications specifically designed to take advantage of the computer cluster, Beowulf clusters are considered to be extremely powerful.

Although Beowulf clusters can provide huge amounts of processing power, it has its disadvantages. Beowulf clusters are completely dependent on specially designed software in order to take advantage of the cluster's resources. So, if you don't have the necessary applications to take advantage of the cluster, then the cluster is pretty much useless.

Fortunately, Beowulf isn't the only clustering implementation out there today. MOSIX is another type of clustering technology. MOSIX works in a different way than that of Beowulf in which it is not dependent on PVM, MPI, or other programming language. Instead of depending on specific applications, MOSIX is simply a kernel patch. This essentially means that any Linux distribution can take advantage of clustering capabilities. But, that is not all. Since the kernel is patched, this also means that any

task done on the node can take advantage of the clustering ability of MOSIX. The two main advantages that make MOSIX a great solution for clustering are its load balancing techniques and the way it can "transparently" migrate processes between nodes [3].

Load balancing is a technique that distributes the work efficiently throughout nodes in the cluster. The master node is responsible for managing which node completes which process. So, if one node is in charge of completing a task but then another node becomes available that is able to complete the task faster, then the process is migrated. By doing this, the use of resources in the cluster is always optimized. Now what makes MOSIX interesting with its migration of processes is that it can do it "transparently". This means that the MOSIX cluster will always be migrating processes between nodes in order to achieve maximum performance, but the nodes will never realize that it is actually occurring. As far as the nodes are concerned, they "think" that the process is still being completed on their computer. So, if a user were to use his computer to perform an intensive task, that user or the computer will never know whether the task is being performed locally or on another node. This is why essentially any application can be used with the MOSIX clustering technology.

b. OpenMOSIX

MOSIX became a proprietary software in 2001. Fortunately, however, in 2002, Moshe Bar, an original developer of the MOSIX project, started the openMOSIX project. Unlike MOSIX, openMOSIX was completely open source. Although the openMOSIX development lasted until 2008, the code is still available for anyone to continue development. The Linux PMI project is currently continuing the development of openMOSIX.

c. CLUSTERKNOPPIX

As the name implies, clusterknoppix is a Linux distribution which is closely based on knoppix. The difference between the two is that clusterknoppix is pre patched with the openMosix software. Some of its features include:

- ✔ *openMosix terminal server* – uses PXE, DHCP and TFTP to boot Linux clients through the network. This means that once the server is set up, only a network card is required in order to boot the nodes through the network [1].

- ✔ No CD-ROM drive or hard drive is needed for clients to use clusterknoppix. OpenMosix's auto discovery allows new nodes to automatically join the cluster without any configuration. The only requirement is that it must be connected to the same network as the master node [1].

- ✔ *Cluster Management tools* – openMosix userland and openMosixview help with Cluster management. Every node in the network has root access to every other node through the use of ssh, RSA keys and MFS support [1].

The biggest advantage of clusterknoppix is that every node on the system can run X server, providing a full GUI KDE interface for users [1]. The main purpose of clusterknoppix is to provide the user with the means to create a fully working cluster through the use of the knoppix Linux distribution, as well as providing a complete GUI with enough software for most users on a desktop. All of this is achieved without even having to install the operating system onto the desktop computer, although this is still available as an option. This means that users can use clusterknoppix without even having to alter their current operating system. Clusterknoppix seems to be one of the most popular methods of achieving a cluster due to its easy access and being open source. Another advantage of clusterknoppix is that many of CHAOS' security enhancements are already implemented into clusterknoppix.

VI. CHAOS

Chaos is another Linux distribution patched with the openMosix kernel software. The purpose of Chaos is to provide an extremely lightweight clustering solution that can be highly distributed. Unlike clusterknoppix, Chaos' intended purpose is not to provide the user with a full blown desktop, but to create a slave node with maximum performance. No bloat ware is included with Chaos. The entire image only consists of the Linux kernel and it is only a few hundred kb in size [3]. However, it is not ideal to run a cluster only using the Chaos distribution. Chaos was only meant to be used as a slave node. Therefore, it excels at aiding a master node such as one created through clusterknoppix. When using Chaos, all the applications necessary to complete the work are only required to be on the master node.

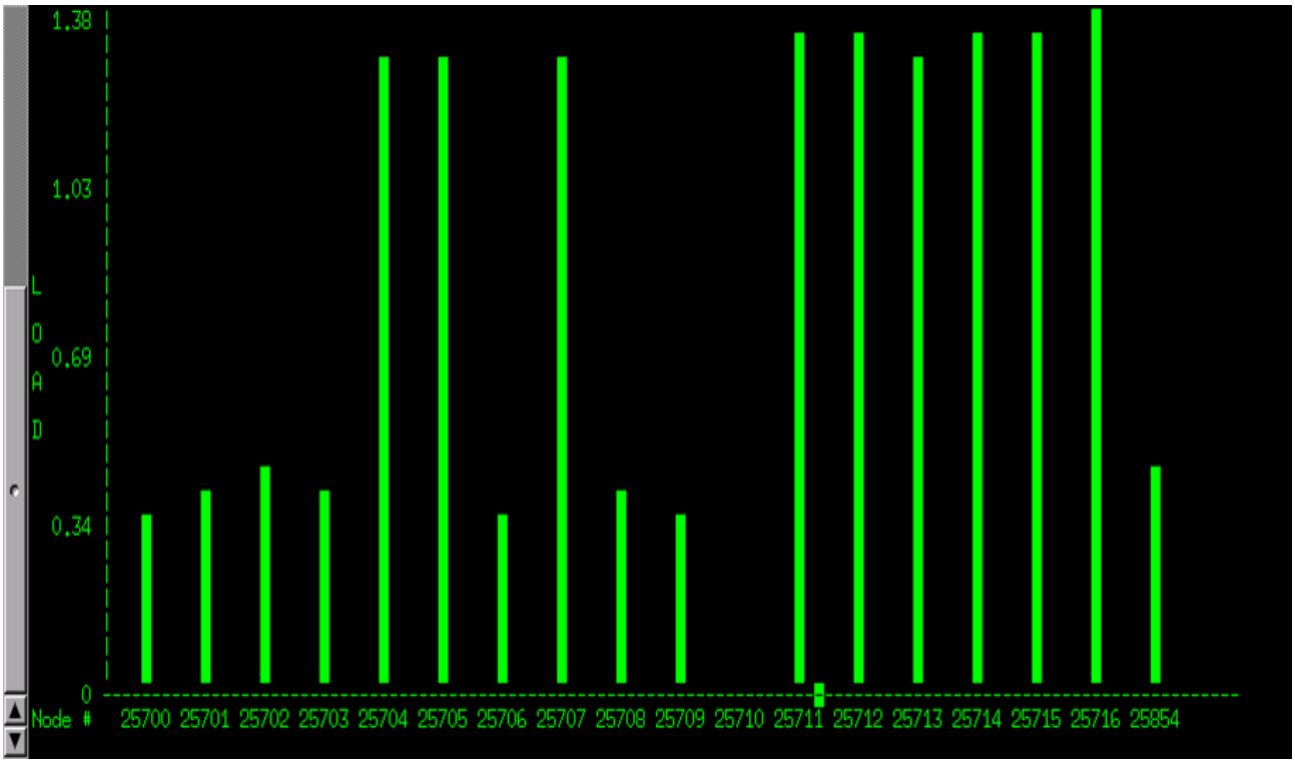

Figure 1: Since CHAOS does not use a GUI, applications such as MOSMON must be used which allows the user to monitor the cluster through the terminal.

VII. A SIMPLE LOAD TEST

Although using openMOSIX to implement a cluster in a work environment is a great alternative to mainframes, not all enterprises will be willing to take the risk of changing their entire IT infrastructure. This is where clusterknoppix comes in. It is extremely affordable to build a cluster with any number of nodes by using clusterknoppix. This allows an IT department to test out the openMOSIX clustering software to decide whether or not it is worth implementing into their environment.

Since clusterknoppix runs as a live CD, there is no risk of affecting any existing partitions that the current desktops of an enterprise currently contain. IBM currently offers a simple tutorial on creating and testing out a computer cluster using the MOSIX software. A demonstration is provided to show a simple cluster setup.

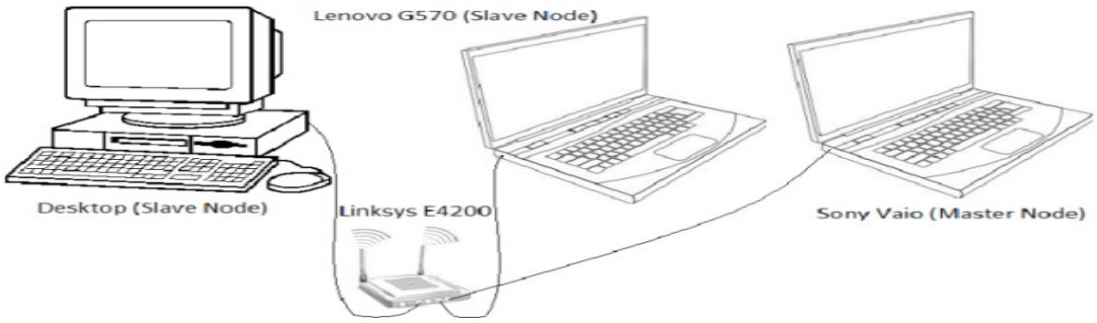

Fig. 2: Setuprequires 3 PCs and a router/switch.

Using clusterknoppix and three systems connected through a router, an example of the openMOSIX clustering capabilities is demonstrated. A simple script provided by IBM was used to load the cluster:

```
awk ' BEGIN{for(i=0;i<10000;i++)
      for(j=0;j<10000;j++);}'
```

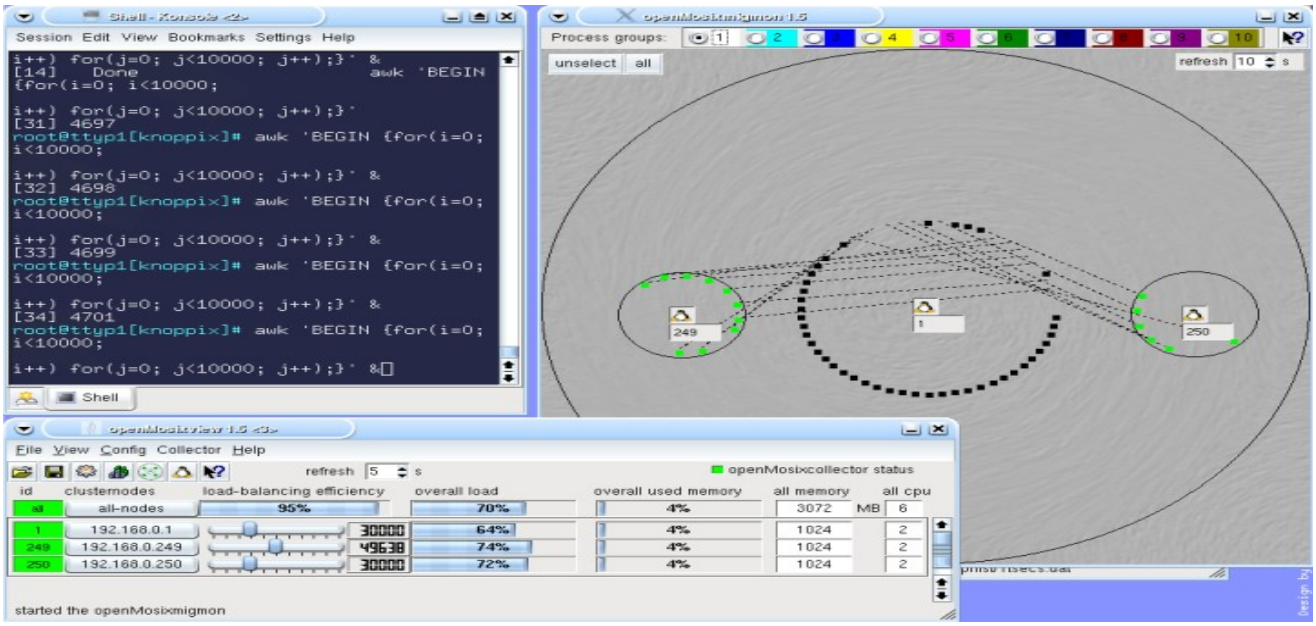

Fig.3: Clusterknoppix comes pre-packaged with a number of applications to monitor and manage the cluster.

This simple script was run 17x on each system as well as the cluster. It basically tells the computer to count from 1 to 10,000 a total of 10,000 times. Although not a complex application, it loads the system to use its maximum processing power, forcing its CPU usage to 100%.

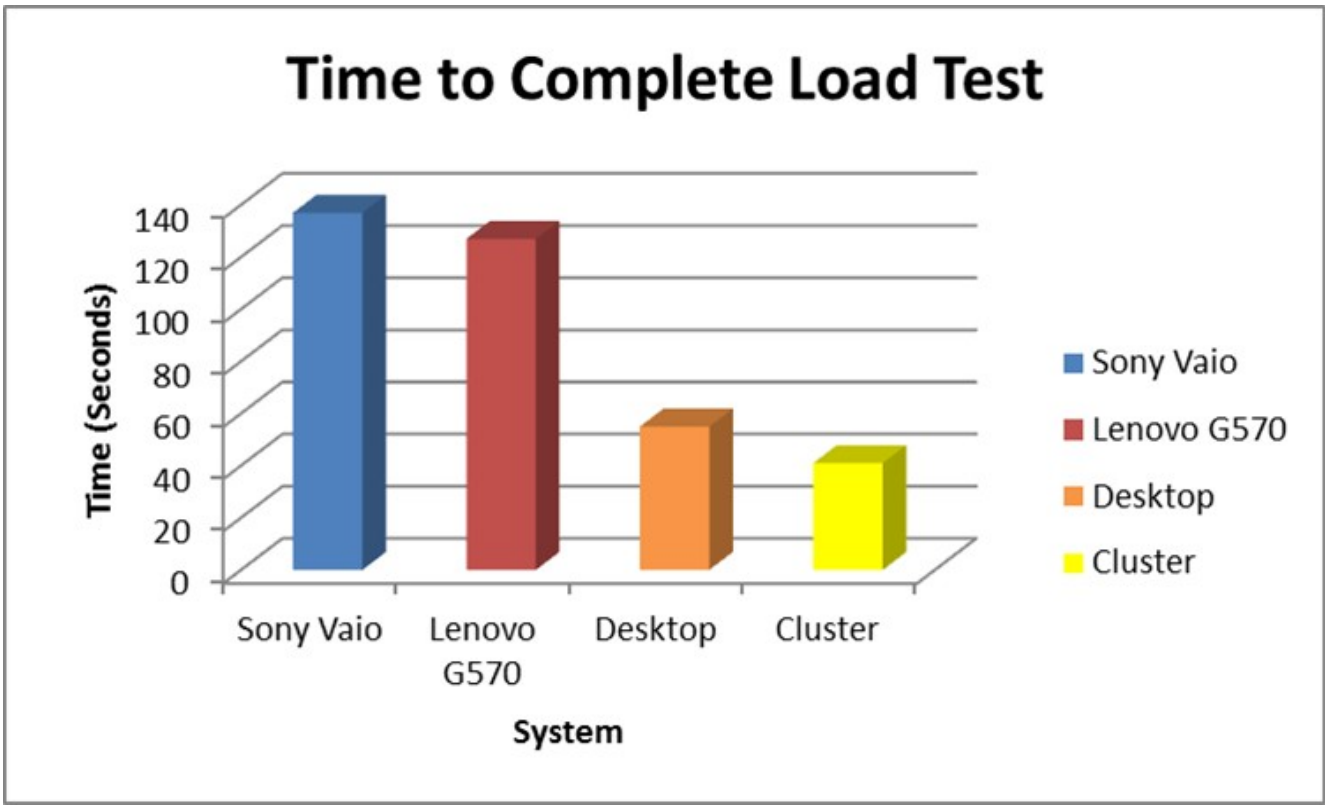

Fig. 4: The load test results show the performance increase of a simple task that the cluster achieves.

The computers used were as follows:

- *Lenovo G570*
  - ✔ Intel B940 @ 2.00GHz
  - ✔ 3GB RAM
- *Sony Vaio*
  - ✔ Intel Pentium P6100 @ 2.00GHz
  - ✔ 4GB RAM
- *Desktop PC*
  - ✔ Intel i5-2500k @ 4.5GHz
  - ✔ 16GB RAM

Sony Vaio	137 seconds
Lenovo G570	127 seconds
Desktop PC	55 seconds
cluster	41 seconds

Fig. 5: The Sony Vaio received 70% performance increase when processing through the cluster while the Desktop PC only received 25%. The more computers that are added, the more powerful the cluster.

VIII. CONCLUSION

For enterprises with a limited IT budget, openMOSIX clustering proves to be a great alternative to mainframes. MOSIX is also a great proprietary solution. Through the use of this technology, CPU intensive tasks that require a supercomputer can be successfully completed without having to use expensive mainframes. Clustering is extremely cost effective but not only that, it is moderately not too complicated to implement into the current setup of many enterprises with an already existing network.

A simple test has been shown to demonstrate how feasible it is to implement a cluster using MOSIX. Many enterprises today already rely on various Linux distributions such as Ubuntu, openSUSE, Fedora, and CentOS. If a company were to convert the entire network of an enterprise into a cluster, the entire user base wouldn't even notice the difference.

REFERENCES

[1] Bilbao, Javier, et al. "Easy clustering with openMOSIX." World Scientific and Engineering Academy and Society (WSEAS) (2005): n. pag. WSEAS E-Library. www.wseas.us/elibrary/conferences/2005athens/cscc/papers497-787.pdf

[2] Robbins, Daniel. "Linux clustering with MOSIX." developerWorks. IBM, 20 Dec. 2001. www.ibm.com/developerworks/linux/tutorials/l-mosix/index.html

[3] Sharma, Mayank. "Craft a load- balancing cluster with ClusterKnoppix." developerWorks. IBM, 22 Dec. 2004. www.ibm.com/developerworks/linux/library/l-clustknop/index.html

[4] "Cluster Computing / Computer Clusters." Best Price Computers. ***www.bestpricecomputers.co.uk/glossary/cluster-computing.htm***