

CONSTRUCTION AND APPLICATION OF AN X86 BEOWULF CLUSTER

Susan Ditmore

April 12, 2005

Computer Systems Lab

Abstract

A beowulf (linux) cluster of 15 to 20 x86 Pentium II computers with automatic load sharing will be constructed for convenient use within the Computer Systems lab. Computer classes such as the supercomputing course will be able to utilize the cluster.

Background

At the end of the summer of 2004 the systems lab was left with twenty or more Pentium II computers not powerful enough to be used as workstations. There was also a large supply of spare computer parts and spare almost working computers in the network supply closet.

The systems lab possessed a cluster of MIPS architecture computers but could not successfully port a working load distribution service to the cluster, as the most common and reliable services are tailored to x86 architecture machines. There had been previous attempts to construct an x86 cluster, but due to various reasons the attempts were unsuccessful.

Description of Methods, Algorithms, Procedures, and Process

Software

Both the system imager and Openmosix systems were implemented.

Openmosix is a kernel patch that automatically transfers and distributes processing jobs across a configured cluster of computers. It utilizes the openmosix file system, designed specifically for the file transferring involved in distributed computing.

Systemimager is a utility that allows one to copy the "image" or linux distribution and all configuration and software from one computer, called the golden client, onto an unlimited number of other computers. It also allows for "upgrades" of all the computers, downloading only what has changed.

In addition, dhcp3-server was implemented to provide dhcp service to the nodes of the cluster.

Design Criteria

The following were the specifications to which the cluster would be built:

- Must run the most recent possible version of Debian compatible with openMosix
- Must use openMosix with vanilla kernel
- Must use systemimager
- Must be able to fit image on a 4 gb harddrive, preferably 2 gb (ie, no x display)
- Must be connected to the outside TJ network and be accessible to super-computing students
- Must be secure (non-admin access restricted to head node)

Procedures

The following procedure has been followed:

1. Repair and clean old Pentium IIs.
2. Find a suitable place for the cluster (this was harder than anticipated).
3. Create vanilla kernel with appropriate OpenMosix patch (22 tries).
4. Install and test distribution on Akosmia, golden client.
5. Test and setup router.
6. Network cluster (this includes obtaining mac addresses, which was also more difficult than expected).
7. Release distribution to all nodes.
8. Test cluster with various high-performances programs.
9. Fix any potential issues and revise image.
10. Secure cluster and attach it to the network.
11. Create account forms and commence more heavy duty testing.
12. Introduce Mr. Latimer to the cluster.
13. Install MPI on the cluster.
14. Find someone to the maintain the cluster the following year.

Description of Results and Conclusions

So far each step has required a longer time than anticipated. This is due mostly to trouble with using old machines.

Many of the machines are in an extreme state of abuse and disrepair and so each one requires much attention and time to be revived. In addition the systemimager concept that was prepared was later discarded by the system administrators for another system.

In addition it took 4 tries to find a successful routing box, as the previous three all had various odd issues such as locking up after a few hours or not displaying properly.

Nevertheless, the work is proceeding.

References

<http://sourceforge.openmosix.net/>
<http://www.systemimager.org/>
<http://www.debian.org/>