

# **Cluster Construction Project**

**Fall, 2007**

**Stephen Ware, Loyola University**

**Advisor: Dr. Xuefeng Li**

## Contents

- I. Statement of Purpose
- II. Process Overview
- III. Daily Journal
- IV. Concluding Report and Suggestions for Future Projects
- V. List of Resources
- VI. Appendix A: E-Mail Correspondence Relevant to the Project
- VII. Appendix B: Posts on LinuxQuestion.org
- VIII. Appendix C: Posts on the GAMESS Usergroup
- IX. Appendix D: Posts on the Aurora Linux User Mailing List

## **I. Statement of Purpose**

The goal of this project is to convert the 15 unused “Sun Ultra 10” computes into a cluster capable of running programs designed for parallel computation at little or no cost to the Mathematics and Computer Science department at Loyola University New Orleans.

This project serves several purposes for both Loyola and me personally. Firstly, it should provide all the benefits associated with having a cluster. The Chemistry department specifically has expressed interest in having a cluster available for running GAMESS, a computation-heavy program for quantum chemistry analysis. Professors in the Mathematics department have also expressed interest in using a cluster for mathematical research. Secondly, it will put idle department resources to good use and provide a use for future retired computers. Thirdly, it will allow me to sharpen my Computer Science skills, since the unfortunate elimination of the Computer Science major has left me with no classes to take related to my field. It will also provide me with valuable research experience, which will no doubt prove useful in my future graduate career.

## II. Process Overview

This overview will provide a brief summary of the successful steps taken during this project. It is designed to assist future cluster builders or future users of my cluster by describing the process without mentioning the errors and setbacks laid out in the next section. Basically, this is a description of everything that went right.

### Phase 1: Design

Before building a cluster, one should determine its purpose and consider the costs [1]. The purpose of this project was stated in the previous section. As for costs, this cluster is being built out of computers which are not presently serving any purpose. The department was able to provide networking cables and a network hub, so no additional parts needed to be purchased. Thus, with the exception of power to run the machines and space to keep them in, the cost of this project is nominal.

When designing a cluster, one needs to consider its uses and make choices accordingly. In this case, the cluster is being designed to run GAMESS [12] and other parallel programs for academic research purposes. Thus, the GNU/Linux operating system was chosen because it is compatible with GAMESS and (due to its wide adoption at academic institutions) it is likely to be compatible with most other research programs. Linux is also free and open-source, so using it adds nothing to the cost of the project. Lastly, Linux is widely supported by online communities that can provide useful feedback when setting up the project. Specifically, I have chosen to use Aurora Linux, a distribution of Fedora Core 6 that can be installed on machines using SPARC chips. The Aurora project [9] is one of the only SPARC Linux projects still under active development, and it offers a wide variety of installation options.

### Phase 2: Preparing the Physical Resources

I began by installing Linux on all the machines. Unfortunately, due to age, most of the machines were unusable. In the end, only four machines were salvageable. I took the RAM out of the unusable machine and put it into the good ones so that each working machine now has the maximum 512 megabyte of RAM. I arranged the machines on a single table and connected them all to a network hub.

One machine was chosen to be the master node. On it, I installed a graphical desktop environment and several other important applications such as a web browser, FTP client, and firewall. The master node controls the other (slave) nodes and provides a point of contact between the cluster and the user. On the other nodes, I installed a “bare bones” Linux operating system. There is no need for the slave nodes to be individually useful, since their only purpose is to contribute to the computations of the cluster as a whole.

### Phase 3: Network Setup

For the cluster to operate, the computers need to be able to recognize and communicate with one another over a Local Area Network, or LAN [15]. Each computer is assigned a unique IP

address on the 192.168.\* space (so 192.168.0.1, 192.168.0.2, 192.168.0.3 etc) and given a unique host name. The master node's hostname is "queenbee" and the slave nodes are "workerbee01," "workerbee02," etc. These values are assigned using the following commands, which are put into the "/etc/rc.d/rc.local" shell script file:

```
ifconfig <device> <ip address> netmask 255.255.255.0
hostname <name>
```

example:

```
ifconfig eth0 192.168.0.1 netmask 255.255.255.0
hostname queenbee
```

Each node must also be told where to find the other nodes in the cluster. This is accomplished by editing the list found in "/etc/hosts" file. Each node needs an entry for each other node in the form:

```
<hostname>      <ip address>
```

example:

```
queenbee 192.168.0.1
```

Once this process has been repeated for each node in the cluster, the private LAN is up and running. Computers on the 192.168.\* address space are not routed by commercial internet routers, so no computer on this network is visible to the outside world via the internet. To give the cluster internet access, a second network card could be installed in the master node that would give it access to the internet. One must be careful when setting this up, though, because the firewall must be properly configured to block dangerous traffic from the internet without hampering traffic coming from the cluster.

Lastly, the nodes in the cluster need to be able to run commands on one another using something like RSH or SSH [16]. This step may not be required on all clusters if the programs you plan to run do not require it, but it is generally a very useful tool. RSH (Remote Shell) is a simpler protocol than SSH, and if the cluster is completely cut off from the internet, it may be safe to use. It is, however, considered outdated. As a result it has been made extremely difficult to enable. SSH (Secure Shell), is a similar protocol, but it sends encrypted messages and is thus much more secure. This added security may not always be needed (especially since it does present some overhead), but it is much easier to enable and is active by default on modern machines. I have used SSH on this cluster. Normally, using SSH to run a command on a remote computer requires you to enter the password of that computer. To avoid being asked for a password every time, one must establish a trusted connection between computers which can be used in the future. Here is the process I used. In the example, I am using "queenbee" and making a connection to "workerbee01":

1. Run the command "ssh-keygen -t -rsa"

2. When prompted for the name of the file to save the key in, type something in the form of `"/home/a/.ssh/id_rsa-queenbee"`
3. When prompted for a passphrase, do not enter one. Simply press enter.
4. If it does not already exist on workerbee01, create the directory `"~/ .ssh"`
5. Now append queenbee's new public key to the list of keys that workerbee01 knows:  
`"cat .ssh/id_rsa-queenbee.pub | ssh root@workerbee01 'cat >> .ssh/authorized_keys'"` This will require you to enter workerbee01's password.

A link to the tutorial I used for this process can be found in [16]. In some cases, one may only need to establish trusted connections between the master node and each slave node. In other cases, each node may need a trusted connection to each other node. Either way, once this process is repeated the necessary number of times, the nodes in the cluster will be able to run commands on one another without requiring each other's passwords. This is, of course, a very dangerous security risk, so make sure that your cluster is not used by anyone with malicious intent.

#### **Phase 4: Program Setup**

Now it is time to begin installing and running the various applications that the cluster will use. Unfortunately, there is no general process to describe here because each program handles parallel programming in its own way. There are a few virtual machines that attempt to accomplish general parallelization, but these are very inefficient and have a large overhead cost. For peak performance, it is best to leave the parallelization of your programs up to the programmers. In the case of GAMESS, it comes with its own parallelization software called DDI (Distributed Data Interface) that runs over TCP socket connections between nodes.

The directions for compiling GAMESS and DDI are described in detail in `"gamess/misc/readme.unix"`. They are easy to follow, but it does take a significant amount of time to complete them. To get it compiling on the SPARC architecture, I had to remove the `"-malign-double"` option from line `"comp"` shells script, since this is an Intel-specific option. Also, one should note that the default shell program on most modern versions of Linux is BASH rather than CSH, so to run a shell script from the command line, one must put a `"/"` in front of the file name. The BASH/CSH difference was also a problem when it came to making GAMESS use SSH instead of RSH. Setting the `DDI_RSH` environment variable in BASH does not work. I found the best solution is to add the line `"setenv DDI_RSH SSH"` into the `"rungms"` shell script. Per the instructions found in the GAMESS readme files, I expanded the available System V memory of each node to include the entire RAM. This can be accomplished by adding the line `"kernel.shmmax = 536870912"` to the file `"/etc/sysctl.conf"`. Lastly, one should note that even though the Sun Ultra 10 has a 64 bit processor, Aurora runs in 32 bit userland, which means that it behaves like a 32 bit machine. When choosing an operating system in the GAMESS configuration files, one should pick `"Linux32"` instead of `"Linux64."`

After this, the cluster should be ready to run!

It is important to note that some libraries are commonly used for parallelization of programs, most notably the MPI library. GAMESS does not require MPI, so I have not installed it on this

cluster, but other readers may need this resource. While MPI is commonly used, it is not always used. Thus, before you install anything on a cluster, decide which programs you plan to run on it and *then* decide what libraries to install. If, in the future, I install a program on this cluster that requires MPI, I will update this document to include a tutorial on installing and testing it.

## **III. Daily Journal**

This journal documents how I have spent my time during the project. Its goal is to provide a detailed description of the process I underwent for the benefit of future readers and to provide an account of my time for my project advisor, Dr. Li.

### **A Note About Dates:**

I am a bit of a night owl, and prefer to work on projects late at night. Many of the sessions that I document are listed as a single day, but in reality they often began on the night of the documented day and ran into the early hours of the next morning. For simplicity, I have only listed the date on which each session began.

### **August 30th, 2007**

Today I began researching various topics related to cluster construction. I read a brief description of process parallelism vs. data parallelism and browsed descriptions of various Linux parallel computing software projects including PVM (Parallel Virtual Machine) MPICH2 (Message Passing Interface) Platform LSF, Linux HA (High Availability), and Linux Virtual Server.

I also read a brief document detailing how another university had setup a cluster for its own use, and I was not surprised to find that the author considered the softwares mentioned above to be too expensive in overhead. After reading this, I began to browse the GAMESS documentation and was delighted to find that GAMESS is already designed to work on a cluster.

### **September 4th, 2007**

I continued reading about GAMESS, MPI, and the distributed virtual machine software PVM. I also read up on "Damn Small Linux," a bare bones (but sufficiently functional) Linux distribution that will probably be my operating system of choice for the cluster nodes. I also spent some time looking for other bare bones Linux distributions that might serve the same purpose as DSL.

### **September 11<sup>th</sup>, 2007**

Much to my disappointment, DSL does not offer a SPARC distribution. In fact, it seems like very few Linux distributions offer SPARC variants. Today was spent looking for a sufficiently functional, well-supported flavor of Linux that can run on a SPARC chip. I have posted on LinuxQuestions.org, a Linux user forum that has been helpful to me in the past, asking for suggestions from experienced users.

### **September 13<sup>th</sup>, 2007**

The responses to my post were not very helpful. Many flavors were suggested, but none of them publish up-to-date SPARC distributions. One user suggested cross-compiling Linux From

Scratch, but this would be a very involved process, so if I can avoid it, I will. I have discovered a website ([UltraLinux.org](http://UltraLinux.org)) that is dedicated to SPARC distributions of popular Linux flavors, but most of it seems to be out of date. Still, it has provided some helpful leads and information. If nothing else, it helped me to get more familiar with the Ultra 10 hardware.

After reading other forum posts related to mine, I have found that many users report success and satisfaction with Aurora Linux ([AuroraLinux.org](http://AuroraLinux.org)). Aurora is a SPARC port of Fedora Core 6, which I have used before. In my experience, Fedora tends to be a bloated, developer-oriented distribution, so the overhead might be too high. Still, the Fedora Project ([FedoraProject.org](http://FedoraProject.org)) is one of the most used and most widely supported distributions, so it might be worth looking into.

### **September 18<sup>th</sup>, 2007**

Since I'm not having any luck finding other updated SPARC distributions, I am going to go with Aurora. Today I downloaded the ISO images and burned them to CD in preparation for installing them in the computers in the lab.

### **September 20<sup>th</sup>, 2007**

Today I installed Aurora on one of the lab computers and began using it. I was pleasantly surprised to discover that Aurora has a "bare bones" installation option, which should be ideal for minimizing overhead in the slave nodes of the cluster. For now, I am testing out a "full installation," including normal desktop software, a web browser, and general development tools like the GCC suite. Considering the age of this machine, Linux runs very well on it. It is a little but laggy, but by and large I am pleased with its performance.

Technical Notes:

To stop a Sun Ultra 10 from using its default boot device, press STOP + A

To boot from the CD-ROM, type "boot cdrom" at the "OK" prompt

### **September 25<sup>th</sup>, 2007**

Today I ran a "bare bones" installation on another of the Sun machines. This seems to be a very good setup for the slave nodes. There is no graphical user interface to take up memory, which won't matter since the slave nodes probably will not even have monitors or keyboards attached.

I also downloaded the source code for GAMESS and began reading the documentation on building it from source.

### **September 27<sup>th</sup>, 2007**

Today I had hoped to install Linux on the rest of the Sun machines, but I ran into all kinds of trouble. Five of the machines report that they cannot detect a hard drive, one seems to have a bad graphics card, and one will not start at all. I was only able to install Linux on one of the remaining machines, bringing the final count to three. One machine has a note that reads "two

disks” attached to it, and I have not yet tested that machine, as it seems that someone was using it for something.

Needless to say, this is a very distressing development. I am going to ask around on LinuxQuestions.org again to see if I can get some answers.

### **October 1<sup>st</sup>, 2007**

This week is the week before the Awakening retreat, so most of my time will be dedicated to preparing for the retreat. I will resume work on this project next week, after the retreat is over.

### **October 9<sup>th</sup>, 2007**

Today I took apart the cases off two of the machines: one of the working ones and one that reports having no hard drive. There is definitely a hard drive installed, and I cannot find any problems with the internal configuration of the machine. The hard drive in the working machine is of a different make and model though, so the problem could lie in not having the right drivers to support the disks in the troublesome machines.

### **October 11<sup>th</sup>, 2007**

Today I met with Dr. Kelly and collected 15 old Intel PC’s from around the department. This should add a lot of good nodes to the cluster, and will be a welcome relief from trying to deal with the poorly-supported SPARC architecture. We found what appears to be an old hub, but I don’t know if it will work. If not, I am told that Dr. Li can have one ordered through the department. This will be an essential component for networking the cluster together.

As far as the hardware goes, I have identified the make and model of one of the non-working harddrives and have made a post on LinuxQuestions.org about it.

### **October 16<sup>th</sup>, 2007**

I also built GAMESS from source today, which took quite a long time (at least an hour and a half to compile, not counting the time it took to configure it). The process is very simple and well documented in the readme file, but a couple of changes had to be made to get it to work. They are noted below. Sadly, the test files included with the source code did not yield correct results. I will worry about that later... at least I’ve made some good progress in getting GAMESS running. I think the problem may lie in the fact that the Ultra 10 is a 64 bit machine. I chose “32 Bit Linux” as the operating system when configuring GAMESS, so perhaps if I choose 64 bit, the problem will be fixed.

#### Technical Notes:

In the “comp” shell script, the g77 option “-malign-double” causes compilation to fail. I read about this option on the g77 website:

<http://gcc.gnu.org/onlinedocs/gcc-3.4.6/g77/Optimize-Options.html#Optimize-Options>

It seems to be something that is Intel specific, so I removed it (comp, line 1129). This fixed the problem and compilation was successful. Removing this option may also be the reason that I am getting incorrect results when running GAMESS.

In Aurora, to run a shell script in the current working directory from the command line, one must include “./” in front of it. So the command “comp” will not run the comp shell script. One needs to use “./comp”. The authors of the GAMESS shell scripts seem to have been using a version of Linux that does not require the dot-slash prefix, so this must be added in several places for the shell scripts to work. Most notably, line 23 in “runall” needs to be fixed.

### **October 18<sup>th</sup>, 2007**

Sadly, I got not useful responses to my post about the unusable harddrives. I have not been able to find these drivers anywhere online. If I cannot get them working, it will be a major setback.

Today I tried rebuilding GAMESS using 64 Bit Linux as the target. Unfortunately, the g77 Fortran compiler (which comes with Aurora) cannot be used to parts of the compilation. The configuration files tell me that I will need gFortran, which is not installed. I spent some time trying to download and build gFortran, but no luck. I am worried that it will be very difficult to get gFortran working, since a compiler needs to build machine-specific code, and the SPARC chip is so poorly supported. I’ve made a post about this on [LinuxQuestions.org](http://LinuxQuestions.org) also.

I also began looking at the Intel machines today, and that was also very disappointing. I fear that they may be too old to be useful. One machine that I opened has only 16 megabytes of RAM, which is not enough to run a graphical installer for Linux. Even if I were to get some of these machines working and attached to the cluster, it is doubtful whether or not they would be useful. The SPARC machines, on the other hand, are fairly powerful. They all have 256 megabytes of RAM, and since some of the machines do not work, I can probably gut them and put their RAM chips into the working machines. With some luck, I should be able to get a handful of machines each with working hardware and 512 megabytes of RAM. Now if only I could get GAMESS to work on the SPARC machines...

I also spoke with Dr. Li over e-mail about ordering a switch for the network.

### **October 22<sup>nd</sup>, 2007**

Today I discovered a resource which should prove more helpful than [LinuxQuestions.org](http://LinuxQuestions.org), which has failed me every time so far. The Aurora user mailing list is, to my surprise, active. I sent out a question about compiling GAMESS in 64 bit mode, and got a response almost instantly.

As it turns out, even though the machines are 64 bit machines, Aurora runs in what is called “32 bit userland,” meaning that all applications run as if on a 32 bit machines. So my original choice of “32 Bit Linux” as the GAMESS target was right. I suppose now I’ll have to try re-compiling it in 32 bit mode and hope that I can figure out why none of the test results came out correctly.

I plan to ask the Aurora user mailing list about the hard drive problems I have been having.

**October 25<sup>th</sup>, 2007**

Today was a major breakthrough, and a badly-needed morale boost. I have been running into problems around every corner of this project for weeks, so I was delighted when I finally got GAMESS to compile and give correct results. I started again from scratch, reconfigured, and re-compiled it. I also examined the “comp” shell script more closely. I think I may have missed some of the configuration settings the first time. Unfortunately, the settings in “comp” are distributed throughout the file instead of being concentrated at the top of the file like all the rest.

There are still a few problems though. Two of the 42 tests did not yield correct results, and I am not sure why. I will need to look into this and figure what went wrong in these cases. Also, I have not gotten any useful feedback from the Aurora user mailing list about the harddrive issues, so I am still stuck there.

**October 30<sup>th</sup>, 2007**

Today I asked a friend of mine who uses GAMESS in her research to examine the two failed tests. I read the files also, and I think it might be a problem with this version of Fortran (which does not support certain number types higher than 99). I will need to look into this further before I can make any more conclusive decisions.

More good news: the old hub that Dr. Kelly and I found works. I have successfully networked two of the working SPARC computers together, and will try to get a basic cluster running soon.

I also spent some time reading about the setup of a Beowulf cluster. A Beowulf cluster is one that is made up of dedicated, low-cost, commercially available machines running Linux and designed for computation. That describes my project perfectly, so this online book should be very helpful:

[http://www.phy.duke.edu/~rgb/Beowulf/beowulf\\_book/beowulf\\_book/index.html](http://www.phy.duke.edu/~rgb/Beowulf/beowulf_book/beowulf_book/index.html)

**November 1<sup>st</sup>, 2007**

I am taking the Computer Science Subject GRE this weekend, so I have had very little time to work on anything else for the last few days. I have gotten a few replies from the Aurora mailing list that look helpful, so I will try those out as soon as possible.

**November 6<sup>th</sup>, 2007**

Some of the replies that I got from the mailing list have helped. I still can't get the machines that claim to have no harddrives to recognize the devices installed in them, but at least now I have a way to determine which machines can and cannot recognize their harddrives. This will allow me to focus my efforts on the working machines. I make take apart the machines which cannot recognize their harddrives and distribute their RAM chips and other useful hardware to the working machines.

I have had some progress with the machines which recognize their harddrives but generated errors during the Linux installation process. The problem seems to lie in the harddrive formats. By reformatting their harddrives (using the Linux programs fdisk and parted), I have gotten much farther along in the installation process, and with some luck, I hope to finally get Linux installed on them. This should increase the number of usable machines by two or three.

I have also started reading about the process of setting up a Local Area Network, which I will need to do with these machines once they are properly configured. My plan for the rest of this week is to finish physically configuring the working machines, finish installing Linux on them, and begin networking them together into a LAN. Once this is done, I can begin testing GAMESS on a distributed network.

#### Technical Notes:

The Sun OBP command "probe-ide" lists all the IDE devices that the system recognizes. This is useful for determine if the system can see the harddrive that is installed. The machines that I have not been able to use due to not having recognizable hard drives will report no IDE devices available, except perhaps the CD-ROM drive.

#### **November 9<sup>th</sup>, 2007**

Today I tried each of the Sun machines again, just to make sure that the unusable ones were, in fact, unusable. I took the RAM out of the ones in the worst condition and put it into the working machines. Now, all of the working machines have 512 megabytes of RAM. In the end, I have four working machines, and Aurora Linux has been installed on all of them. I am still holding out hope for a fifth machine, as well as a couple of other machines that do not seem to recognize their harddrives.

I also physically arranged the working machines today. Now all 4 working machines (and the fifth machine that I am still holding out hope for) are all on the same table and all networked together with the hub. Next week, I can begin configuring the LAN.

#### **November 12<sup>th</sup>, 2007**

The LAN is up and running. Each machine has been given an IP address and a hostname in the address space 192.168.0.x. I've configured each machine to use sufficiently large System V shared memory, via the instructions found in games/ddi/readme.ddi. I've also setup ssh keys between all the machines so that they can launch processes on their peers without being prompted for a password.

I've started fiddling with executing GAMESS in parallel, but no luck yet. I have gotten a variety of different errors, but I think they stem from the fact that GAMESS assumes you are using csh, whereas Fedora uses bash. I am going to post on the GAMESS usergroup to see if this problem can be solved.

#### **November 15<sup>th</sup>, 2007**

I have now got the bash/csh problems worked out, and I think I am getting closer to successfully executing GAMESS in parallel on the cluster. I have copied the GAMESS binaries to each node in the cluster, along with the necessary c libraries. DDI (Distributed Data Interface), which is responsible for executing GAMESS in parallel via SSH and TCP sockets, boots up and successfully contacts the other nodes in the cluster... but I am getting some troublesome TCP errors. In the process of debugging all this, I discovered that along with setting the IP address of each node, I also need to set their host names with the hostname command. That fixed several of the name resolution errors, but not all. I've posted again on LinuxQuestions.org asking for some feedback on the LAN setup. Hopefully someone there will be able to help me figure out what is going wrong.

## **Thanksgiving Break**

**November 29<sup>th</sup>, 2007**

Today was a big day! I finally got GAMESS to run over the cluster. There were two problems: the master node's firewall and some kind of compatibility issue with the shell scripts (bash / csh conflicts I think). To fix the first problem, I simply disabled the firewall on the head node. When the head node is on the 192.168.\* address space, it is no longer visible on the rest of the school network or on the internet at large, so I do not foresee any security issues.

The shell script problem was very troublesome, but I finally found a way around it. Rather than trying to set the DDI\_RSH=ssh shell variable in bash (via the NAME=VALUE method), I added a setenv statement into the rungms shell script that comes with GAMESS. This worked, and GAMESS finally started using SSH instead of RSH. These two fixes allowed the cluster to work as a whole on a GAMESS job. For the very small text examples, the cluster actually took longer to finish the job than a single node, but that is probably due to the large overhead. Next I will test the cluster on a larger file to see how much faster it can process the data.

**December 4<sup>th</sup>, 2007**

Today I setup the cluster to run a larger GAMESS job: first on a single node and then on the cluster as a whole. The script will record the time it takes to complete the job in both modes so that I can compare the two and discover how much the process has been sped up thanks to parallel processing. I will check the cluster tomorrow to see if it has finished the job.

**December 5<sup>th</sup>, 2007**

If my initial reports are correct, the job actually took significantly longer to run on the cluster (half an hour on a single node, and over 24 hours on the cluster). I am not entirely sure that my setup was correct, so I have run the job again. I will check the results tomorrow.

If the cluster is actually slowing things down, there are two things I can think to try. Firstly (I went ahead and did this today), I unplugged the cable connecting the network hub to the school network. Since it is a hub, not a switch, the extra noise from the school network might have been slowing communication between nodes. The second option, which I will look into if needed, is

to set the cluster up to run over RSH instead of SSH. SSH is the more secure (and thus, more widely supported) child of RSH, and RSH has fallen out of favor. However, in a closed network that is not accessible from the internet at large, the encryption of SSH is not needed, and may just be adding to the overhead. Root access with RSH is very difficult to enable, but if all else fails, I will try to get it working.

**December 6<sup>th</sup>, 2007**

My initial reports were off. The reason that the first job ran faster on a single node was because it encountered a run time error and exited before finished the job. Unfortunately, GAMESS error messages are not very descriptive and I do not understand the quantum chemistry well enough to check the input file myself. I will ask one of my colleagues in the Chemistry department who uses GAMESS regularly if she can spot the error in the file. It is possible that the error arose from not having enough memory, which might explain why the same error did not come up when running the job over the cluster.

I ran a different sample job, and this one completed normally. On one node, the job took roughly 2100 seconds. On the cluster, the job took roughly 2030 seconds, indicating a speedup of about 60 seconds. This is not a very impressive speedup. Also, from analyzing the output files, the computation efficiency was much lower when running over the cluster. This is to be expected, but the values seem too low for comfort. Perhaps given very large input files (one that would take several weeks to complete), the cluster would speed up the process even more, but I do not have time to test jobs that requires more than 24 hours to run.

I will continue to run some test jobs to measure the speedup of the cluster, but overall it is rather disappointing. Still, this has been a valuable learning experience, and it should pave the way for more useful research if better hardware becomes available to the department.

## IV. Concluding Report and Suggestions for Future Projects

The goals of the project were only partially successful. I succeeded in building a working cluster to run GAMESS (and other programs as needed), but the real-world usefulness of the cluster is doubtful. The educational goals of the project were, however, completely successful, and I was able to put idle department resources to good use.

Of the 15 Sun machines, only 4 were usable, and this has significantly reduced the potential power of the cluster. The most common problems were harddrive failures, so if the harddrives in the malfunctioning machines could be fixed or replaced, the cluster could probably be doubled or tripled in size. It has been suggested to me that updating the OBP (Sun equivalent of BIOS) on the malfunctioning machines may cause the harddrives to be recognized, but I find this doubtful. Presumably, all of these computers worked when they were initially purchased, and to my knowledge their hardware has never been updated. Thus, they should have all the necessary drivers to recognize their harddrives. I suspect that the problems are physical and will require extensive cleaning or replacement of the malfunctioning drives.

Dr. Kelly also gave me over a dozen old Intel computers to add to the cluster, but none of them were usable. Some computers had as little as 64 megabytes of RAM, so even if they could be added to the cluster, they would not be able to contribute significantly to its work. The overhead required to distribute parts of the jobs to these machines would probably be more than the computational benefit gained. Even the Sun machines, which all have 512 megabytes of RAM, are not providing very high performance.

There are a couple ways one could improve the cluster. Firstly, the hub connecting them all into a LAN could be replaced with a switch. Reducing the noise on the network would save some computational overhead. Also, assuming the cluster is not exposed to the internet and remains on the 192.168.\* address space, RSH could be used instead of SSH. RSH is not encrypted, so it would allow for slightly more efficient communication between nodes, and the security risks associated with RSH would not be a problem.

If the cluster were to be put to regular use, it should have an accessible public interface. To achieve this, one would need to purchase and install a second Ethernet card on the master node so that it could be made available to the school's network without removing itself from the 192.168.\* cluster LAN.

I do not, however, suggest that the above steps be taken. Unless the power of the cluster is significantly expanded (by adding more and/or better machines) it will not be very useful in research. Most users would be better off running research work on a single modern machine rather than a cluster of four old ones. This research has been very valuable to me because I now have the practical experience necessary to build a cluster, but sadly the hardware I was working with was just too outdated or dysfunctional to render a useful research tool. I have written this detailed guide so that if, in the future, better hardware becomes available, I or someone else reading this guide will have a much easier time building a useful cluster. For example, Loyola replaces its lab computers on a regular basis. When the time comes for a lab to be replaced, rather than sell the old machines, they could be organized into a cluster. There are roughly two

dozen computers in most of the labs in Monroe Hall, and two dozen working, semi-modern computers would make a very formidable cluster. Hopefully I have provided my readers with the information required to build such a cluster, if the opportunity ever arises.

So while the practical goals of the project were only partially successful, the educational experience has been excellent, and I hope that my work has paved the way to more productive work in the future.

I would like to include a special word of thanks to Dr. Xuefeng Li, my advisor, who made this project possible and who offered advice to me throughout its execution. I would also like to thank Dr. Michael Kelly for allowing me to use the Sun lab and its computers in this project. This has been a valuable opportunity to continue my Computer Science training despite a lack of major-related classes.

## V. List of Resources

### Articles Related to Cluster Construction

- [1] Engineering a Beowulf Style Cluster  
[http://www.phy.duke.edu/~rgb/Beowulf/beowulf\\_book/beowulf\\_book/index.html](http://www.phy.duke.edu/~rgb/Beowulf/beowulf_book/beowulf_book/index.html)
- [2] Linux Cluster HOWTO  
[http://www.ram.org/computing/linux/linux\\_cluster.html](http://www.ram.org/computing/linux/linux_cluster.html)
- [3] Computer Cluster  
[http://en.wikipedia.org/wiki/Computer\\_cluster](http://en.wikipedia.org/wiki/Computer_cluster)

### Cluster Software

- [4] MPICH2 – Message Passing Interface Standard 2  
<http://www.mcs.anl.gov/research/projects/mpich2/>
- [5] The Linux Virtual Server Project  
<http://www.linuxvirtualserver.org/>
- [6] PVM – Parallel Virtual Machine  
[http://www.csm.ornl.gov/pvm/pvm\\_home.html](http://www.csm.ornl.gov/pvm/pvm_home.html)
- [7] Platform LSF  
<http://www.platform.com/Products/Platform.LSF.Family/Platform.LSF/Home.htm>
- [8] Linux HA (High Availability)  
<http://www.linux-ha.org/>

### Aurora SPARC Linux Resources

- [9] Project Homepage  
<http://auroralinux.org/>
- [10] The User Mailing List  
<http://lists.auroralinux.org/mailman/listinfo/aurora-sparc-user>
- [11] The User Mailing List Archives  
<http://lists.auroralinux.org/pipermail/aurora-sparc-user/>

### GAMESS Resources

- [12] Project Homepage  
<http://www.msg.chem.iastate.edu/gamess/>
- [13] Google User Group  
<http://groups.google.com/group/gamess>

### General Linux Resources

- [14] Linux Questions  
<http://www.linuxquestions.org/>
- [15] Linux Networking Tutorial  
<http://www.aboutdebian.com/network.htm>
- [16] How to setup SSH Login Without a Password  
[http://www.linuxproblem.org/art\\_9.html](http://www.linuxproblem.org/art_9.html)

## VI. Appendix A: E-Mail Correspondence Relevant to the Project

**August 20<sup>th</sup>, 2007**

From: Stephen Ware  
To: Dr. Kelly and Dr. Li

Dr. Kelly and Dr. Li,

Hello from Birmingham, Alabama! This is Stephen Ware, one of those die-hard CS students. I hope your summers have been good... mine's been great! I got a fantastic internship working for an innovative software firm, and I've learned a lot. I've already had two potential offers for jobs once I've got my degree, and I've also taken the GRE, so my post-undergraduate future looks bright! I wanted to send you both a note before I come back in the hopes that we could get the independent study that I discussed at the end of last year worked out.

Before I left for the summer, Dr. Li kindly took interest in the project and I am hoping that he will agree to mentor it.

My proposal is to convert the old Sun machines sitting in the unused Sun lab into a cluster of computers that can be put at the disposal of the Chemistry and Math departments. I intend to install a SPARC version of UNIX on all the computers and then network them into a cluster that will focus on high performance and run computationally-intensive jobs (as opposed to high availability, which is the goal of some clusters). The power of the finished cluster will not be on par with many professional clusters in the computing world, but it will most certainly be of greater power and use to us than the Sun machines are in their present form. Because the cluster will be UNIX based, and because UNIX has been compiled for virtually every architecture imaginable, the cluster should be highly-scalable and will thus provide a worth-while place to retire old machines once they are no longer adequate as personal computers in the labs.

This project will server several purposes. Firstly, it will provide a valuable resource to Loyola University, specifically to the College of Humanities and Natural Sciences. The Chemistry department already makes use of a highly computationally-intensive program called GAMESS for chemical analysis (which it currently outsources to a cluster at another university), and several professors in the Mathematics department have expressed interest in using a cluster for mathematical research. Secondly, this project will allow me as a Computer Science student to continue my studies despite the regrettable loss of our University's CS department. I have had all the required Computer Science course work needed to graduate and have had experience with both UNIX and GAMESS, so I feel capable of carrying out the proposed project within a semester's time. Also, a colleague and friend of mine, Mary Spulak, has spent her summer in Zurich, Switzerland studying Computational Chemistry, so her assistance could make this project even more fruitful. Lastly, this project would put several mostly-unused machines to good used, and would provide a useful place to retire future machines, thus extending the usefulness of our University's already existing resources.

This project as I envision it would have 5 main stages:

1. Design the cluster and decide which clustering software is best for the job at hand
2. Install the UNIX operating system and clustering software on all machines and network them together into a cluster
3. Test the cluster's performance and tweak it based on the design specifications
4. Install GAMESS and any other software to be used on the cluster, along with a user-friendly interface that can be made available over the school network
5. Run actual test cases on the software to measure performance and ensure satisfactory completion of the requirements.

Dr. Kelly, I hope that this proposal meets your requirements for an independent study class. I was very much hoping to have this on my schedule for this semester. If it does not, please let me know so that I can either refine my outline or find another class to take.

Dr. Li, I hope that this project is still to your liking and that you will still agree to mentor it.

Thank you both very much for considering this. I look forward to your replies and to seeing you both in person next week.

Stephen Ware  
Computer Science Senior

## **August 21<sup>st</sup>, 2007**

From: Dr. Kelly  
To: Stephen Ware

Stephen, sounds like a great proposal. First thing is to get coordinated with Dr. Li, as he is on sabbatical this semester.

Dr. Kelly

---

From: Dr. Li  
To: Stephen Ware

> Dr. Kelly and Dr. Li,  
>  
> Hello from Birmingham, Alabama! This is Stephen Ware, one of those  
> die-hard CS students. I hope your summers have been good... mine's  
> been great! I got a fantastic internship working for an innovative  
> software firm, and I've learned a lot. I've already had two potential  
> offers for jobs once I've got my degree, and I've also taken the GRE,  
> so my post-undergraduate future looks bright! I wanted to send you  
> both a note before I come back in the hopes that we could get the  
> independent study that I discussed at the end of last year worked out.  
> Before I left for the summer, Dr. Li kindly took interest in the  
> project and I am hoping that he will agree to mentor it.  
>

Sounds great. I'll be happy to help as much as I can. Because I am on sabbatical this fall, we'll be communicating via email

most of the times. I'll make some remarks below regarding your proposal.

> My proposal is to convert the old Sun machines sitting in the unused  
> architecture imaginable, the cluster should be highly-scalable and  
>.....  
> will thus provide a worth-while place to retire old machines once they  
> are no longer adequate as personal computers in the labs.  
>

A cluster that utilizes old Sun SPARC workstations as well as other PCs in the future will be great. Seems like you are going to install Solaris on the Sun's, aren't you? And for PCs that will be added into the cluster in the future, what OS are you planning to use?

> 1. Design the cluster and decide which clustering software is best for  
> the job at hand  
>

That's when you decide the OS for the cluster.

> 2. Install the UNIX operating system and clustering software on all  
> machines and network them together into a cluster  
>

What clustering software have you had in mind? Will the cluster need a separate head node? What are the specs for a head node?

Feel free to contact me should new hardware be needed for the project.

Regards,  
--Xuefeng Li, (504)865-3340(phone)

---

From: Stephen Ware  
To: Dr. Kelly and Dr. Li

Hello again from Birmingham!

Dr. Li - The operating system I intend to use would be some flavor of UNIX... possibly Solaris, though I was leaning toward Linux (maybe a Beowulf cluster? see [beowulf.org](http://beowulf.org)). I'm leaning toward some sort of Linux because it exists on all architectures (Intel, PowerPC, SPARC), which would allow us to add any kind of machine to the cluster, even if it wasn't made by Sun. I'll research the various alternatives and decide on one as the first step of the project. As for hardware, the only thing I will need is a hub and some network cables, and I have a feeling we can find those laying around the labs. If not, they are pretty cheap. Thank you very much for agreeing to help with this project!

Dr. Kelly - If everything is in order, can we go ahead and add this to my schedule? I'm not exactly sure how that is done since I have never done an independent study before. If it's possible, I'd like to get it on there before classes start.

Thanks again to you both,  
-Stephen

**August 22<sup>nd</sup>, 2007**

From: Dr. Kelly  
To: Stephen Ware

Stephen, to get the independent study processed, there is a form that has to be completed. As far as I know you will have to be here, as will Dr. Li, to complete and sign the form.

---

From: Dr. Li  
To: Stephen Ware

> Hello again from Birmingham!  
>  
> Dr. Li - The operating system I intend to use would be some flavor of  
> UNIX... possibly Solaris, though I was leaning toward Linux (maybe a  
> Beowulf cluster? see [beowulf.org](http://beowulf.org)). I'm leaning toward some sort of  
> Linux because it exists on all architectures (Intel, PowerPC, SPARC),  
> which would allow us to add any kind of machine to the cluster, even  
> if it wasn't made by Sun. I'll research the various alternatives and  
> decide on one as the first step of the project. As for hardware, the  
> only thing I will need is a hub and some network cables, and I have a  
> feeling we can find those laying around the labs. If not, they are  
> pretty cheap. Thank you very much for agreeing to help with this  
> project!

>  
Most existing clusters indeed use Linux. Consequently, most clustering software are written with Linux as the primary target. That means installation of clustering software on Linux would be easier than to install it on other flavors of UNIX.

Please keep me posted.

> Dr. Kelly - If everything is in order, can we go ahead and add this to  
> my schedule? I'm not exactly sure how that is done since I have never  
> done an independent study before. If it's possible, I'd like to get  
> it on there before classes start.

>  
As indicated by the Dr. Kelly's reply email, we need to get together some day to complete and sign the independent study form. I am open every day. Please let me know a day that you will be coming to complete the form, and I'll show up to sign it.

---

From: Stephen Ware  
To: Dr. Li

Is this Friday (the 24th) ok? I will be driving back to New Orleans tomorrow, so I can come by the office to sign whatever forms need to

be signed on Friday. Let me know if that will work. I can come any time that day, and it will give me a chance to get it on my schedule before classes start.

- Stephen

---

From: Dr. Li  
To: Stephen Ware

Sounds good. Should we meet at 10AM on Friday?

### **August 23<sup>rd</sup>, 2007**

From: Stephen Ware  
To: Dr. Li

Yes, 10 AM on Friday sounds perfect. I'll see you there!  
-Stephen

### **August 31<sup>st</sup>, 2007**

From: Dr. Li  
To: Stephen Ware

Hello, Stephen.

I came across the following today. Interesting to read.

<http://www.calvin.edu/~adams/research/microwulf/>  
<http://www.calvin.edu/news/releases/2007-08/microwulf.htm>

How are things going so far? Please keep me posted should you come up with anything extra ordinary: new ideas, good ideas, failed ideas, great articles, great websites, etc. It could be slow at the beginning because preparation is normally the time-consuming part.

It would be nice if you can also keep a diary on the web regarding this experience.

---

From: Stephen Ware  
To: Dr. Li

Thanks very much for the link... that is very interesting. It might be possible to imitate what they have done using our existing hardware.

Right now the only progress I have made is research. I've been reading some articles on different clustering software and also the GAMESS documentation. As it turns out, GAMESS has some built-in cluster support. So now I am faced with a question: Do I design the cluster to run GAMESS optimally using its built-in cluster support, or do I abstract the cluster layer out and have GAMESS treat the cluster

as a single, powerful computer? GAMESS would presumably run faster if I used its built-in load-balancing functionality, but then the cluster would need some other configuration to run other programs. What do you think?

## September 4<sup>th</sup>, 2007

From: Dr. Li  
To: Stephen Ware

> Right now the only progress I have made is research. I've been  
> reading some articles on different clustering software and also the  
> GAMESS documentation. As it turns out, GAMESS has some built-in  
> cluster support. So now I am faced with a question: Do I design the  
> cluster to run GAMESS optimally using its built-in cluster support, or  
> do I abstract the cluster layer out and have GAMESS treat the cluster  
> as a single, powerful computer? GAMESS would presumably run faster if  
> I used its built-in load-balancing functionality, but then the cluster  
> would need some other configuration to run other programs. What do  
> you think?  
>

There are several factors involved in the above questions. But please keep in mind your objective in this project. You would like to choose a design so that you'll have a working cluster at the end. Please keep it simple.

On the other hand, because you've chosen high performance over high availability, I suggest that you continue along this line. A cluster with a highest FLOPS possible would look very nice on your resume.

Have you heard of MPI (<http://www-unix.mcs.anl.gov/mpi/>)? It is widely supported and used by parallel computational professionals (<http://www-unix.mcs.anl.gov/mpi/whodidmpitab.html>). Please research on MPI and let me know of your opinions about it.

---

From: Dr. Li  
To: Stephen Ware

<http://www.computerworlduk.com/technology/mobile-wireless/mobile-wireless/news-analysis/index.cfm?articleid=743>

Regards,  
--Xuefeng Li, (504) 865-3340 (phone)

## September 5<sup>th</sup>, 2007

From: Stephen Ware  
To: Dr. Li

I did a bit more reading about GAMESS and MPI, and I think it's time to start testing the hardware. On Thursday, I plan to install "Damn Small Linux" (DSL) on one of the Sun machines in the lab. DSL is a

Linux distribution designed to run on very low memory computers... it is very "bare bones," but still provides an easy interface to add new programs as needed. I think this will be a good operating system to start with. After that, I'll look into networking the machines together.

I'm working on some way to keep an online journal of my work. I'll let you know as soon as it's available.

## **September 25th, 2007**

From: Stephen Ware  
To: Dr. Li

Hi Dr. Li,

Just wanted to let you know that the cluster project is coming along nicely. I ran into a lot more trouble than I expected when trying to load Linux onto the Sun machines, but in the end I got it worked out. So far I have one machine running a full suite of desktop software, which I use for web browsing and reading while I am in the lab. I have another machine that is a bare bones Linux kernel on which I plan to run some initial tests.

Right now I'm in the process of compiling and linking GAMESS for these Sun machines. Once that is done, I will test out the program on the bare bones machine and clock how long it takes to run a test job. Once that is done, I'll start networking the computers together and running some tests to see how much performance is improved.

The flavor of Linux I am using is the Aurora project, which is a SPARC adaptation of Fedora (Red Hat) Linux. It's very easy to set up and comes with a lot of excellent software.

Hope your research is going well!  
-Stephen

## **September 26<sup>th</sup>, 2007**

From: Dr. Li  
To: Stephen Ware

> Hi Dr. Li,

>

> Just wanted to let you know that the cluster project is coming along  
> nicely. I ran into a lot more trouble than I expected when trying to  
> load Linux onto the Sun machines, but in the end I got it worked out.  
> So far I have one machine running a full suite of desktop software,  
> which I use for web browsing and reading while I am in the lab. I  
> have another machine that is a bare bones Linux kernel on which I plan  
> to run some initial tests.

>

Glad to hear from you. You've accomplished so much so fast.  
OS installation is the first thing one needs to do, and it is

the most critical step in the process. Fine-tuning of other software will then be possible from now on.

> Right now I'm in the process of compiling and linking GAMESS for these  
> Sun machines. Once that is done, I will test out the program on the  
> bare bones machine and clock how long it takes to run a test job.  
> Once that is done, I'll start networking the computers together and  
> running some tests to see how much performance is improved.

>  
> The flavor of Linux I am using is the Aurora project, which is a SPARC  
> adaptation of Fedora (Red Hat) Linux. It's very easy to set up and  
> comes with a lot of excellent software.

>  
> Hope your research is going well!  
> -Stephen  
>

That sounds good. Have you been able to keep a journal all along? To make this project meaningful in the long run, it would be nice if instructions could be written so that future students may follow your foot-steps to continue and improve the functionality of the cluster.

Feel free to ask if you have any questions regarding software and/or hardware.

Regards,  
--Xuefeng Li, (504) 865-3340 (phone)

---

From: Stephen Ware  
To: Dr. Kelly

Hi Dr. Kelly,

Do we have any old hubs (or switches) laying around that I can use for the cluster I'm building out of the old Sun machines? I'm almost ready to start networking the computers together, so I want to make sure we have the necessary hardware. If we don't have any old ones in the storage closets, one can be purchased for very cheap.

Thanks!  
-Stephen Ware

---

From: Dr. Kelly  
To: Stephen Ware

Stephen, stop by my office and we can check the various locations where we keep old hardware. Also, have you asked Dr. Li?

**October 9<sup>th</sup>, 2007**

From: Stephen Ware  
To: Dr. Li

Dr. Li, I am having a lot of trouble with the Sun machines in the lab.  
4 of the machines report that they have no hard drives, one will not start, one seems to have a bad graphics card (among other problems), and one has the note "two disks" wirtten on it's piece of paper.

As it stands, only 3 computers are currently usable. I have not yet tested the one labled "two disks." Any ideas what's wrong?

-Stephen

**October 10<sup>th</sup>, 2007**

From: Dr. Li  
To: Stephen Ware

> Dr. Li, I am having a lot of trouble with the Sun machines in the lab.  
> 4 of the machines report that they have no hard drives, one will not  
> start, one seems to have a bad graphics card (among other problems),  
> and one has the note "two disks" wirtten on it's piece of paper.  
>

First, a Sun Sparc workstation is almost like an ordinary PC.  
So feel free to open it up and inspect. But be careful,  
make sure to unplug the power cord before performing any  
work inside the computer boxes.

As far as I know, these workstations should be original,  
i.e., no one has changed any hardware device inside them.  
However, I was paying attention to them in the last  
couple of years. Dr. Khalaf might have done something to them.

> As it stands, only 3 computers are currently usable. I have not yet  
> tested the one labled "two disks." Any ideas what's wrong?  
>

In this case, you may open up the working computer. Use it  
as a model to compare to others. Make sure all hardware  
is still inside the box.

If there is any hardware missing, just give up on that one.  
If no hardware is missing, it indicates the hard drive  
is no longer working. Feel free to swap hard drives among  
the workstations.

If more computers are needed for constructing a cluster,  
there are more retired Pentium PCs (between 233-500Mhz)  
available for this purpose. Please ask Dr. Kelly about  
them.

Hope some of the computers are still salvageable.

Regards,  
--Xuefeng Li, (504) 865-3340 (phone)

---

From: Dr. Kelly  
To: Stephen Ware

Stephen, Dr. Li has informed me that you are having trouble getting a number of usable machines from the lab. Stop by and see me soon, as we have a number of other SUN's and old PC's in other locations.

Dr. Kelly

---

From: Stephen Ware  
To: Dr. Kelly

I'll probably come by tomorrow... I work on the project on Tuesdays and Thursdays, and it's finally time to start networking the computers together.

### **October 18<sup>th</sup>, 2007**

From: Stephen Ware  
To: Dr. Li

Dr. Li, I got together with Dr. Kelly and he gave me a bunch of old Intel machines that were laying around. I'm currently in the process of installing Linux on them.

I'm having trouble with the Sun machines. For one thing, I cannot find the proper hard drive drivers anywhere. On the machines with working hard drives, I can't get GAMESS to compile correctly. I did manage to get it to compile (finally), but all the test runs came out failed. I think the problem has something to do with the Sun machines being 64 bit. Anyway, I am going to try setting up a small, 3 machine cluster with the Intel machines first, and then I will try to add the Sun machines.

Also, Dr. Kelly asked me to ask you if I needed any hardware. I am going to need either a hub or a switch to connect all the computers. Dr. Kelly said that you would know what to do for ordering one. My (optimistic) guess is that we'd need one that can handle 20 computers. Please let me know if you can order one through the Math department for me.

Hope you are doing well!  
-Stephen

---

From: Dr. Li  
To: Stephen Ware

> Dr. Li, I got together with Dr. Kelly and he gave me a  
> bunch of old Intel machines that were laying around. I'm  
> currently in the process of installing Linux on them.  
>  
> I'm having trouble with the Sun machines. For one thing,  
> I cannot find the proper hard drive drivers anywhere. On  
> the machines with working hard drives, I can't get GAMESS  
> to compile correctly. I did manage to get it to compile  
> (finally), but all the test runs came out failed. I think  
> the problem has something to do with the Sun machines

> being 64 bit. Anyway, I am going to try setting up a  
> small, 3 machine cluster with the Intel machines first,  
> and then I will try to add the Sun machines.  
>

It used to be that Sun was the machine that one could easily get drivers for. Now Intel PCs took over.

Is the OS in 64-bit mode? Will GAMESS compile in 64-bit mode? It would be nice if we can actually take full advantage of the 64-bit architecture available from the hardware.

>  
> Also, Dr. Kelly asked me to ask you if I needed any  
> hardware. I am going to need either a hub or a switch to  
> connect all the computers. Dr. Kelly said that you would  
> know what to do for ordering one. My (optimistic) guess  
> is that we'd need one that can handle 20 computers.  
> Please let me know if you can order one through the Math  
> department for me.

>  
>  
I've attached URLs for several gigabit switches below. Please let me know which one you preferred.

<http://www.cdwg.com/shop/products/default.aspx?EDC=1012601>

<http://www.cdwg.com/shop/products/default.aspx?EDC=512294>

<http://www.cdwg.com/shop/products/default.aspx?EDC=1030729>

<http://www.cdwg.com/shop/products/default.aspx?EDC=652859>

<http://www.cdwg.com/shop/products/default.aspx?EDC=674865>

---

From: Stephen Ware  
To: Dr. Li

> It used to be that Sun was the machine that one could  
> easily get drivers for. Now Intel PCs took over.  
>  
> Is the OS in 64-bit mode? Will GAMESS compile in 64-bit  
> mode? It would be nice if we can actually take full  
> advantage of the 64-bit architecture available from  
> the hardware.

Before compiling GAMESS, there are several settings files that you must configure. When set to "linux32," GAMESS will compile... but all the test runs that I ran came out false. When set to "linux64", you do not have the option of using the g77 fortran compiler, which comes with Aurora Linux. None of the allowed compiles for 64-bit are installed, and I haven't been able to install get any of them (the compilers) installed yet. I will keep trying this later, but in the interest of time, I am going to set something up using the widely-supported intel machines. Once I know more about the process, it should make fixing the Sun machines easier.

> I've attached URLs for several gigabit switches below.  
> Please let me know which one you preferred.  
Hmm... those are much more expensive than I thought they would be.  
Dr. Kelly and I did find something that looks sort of like those switches, so before we buy something, I will try it out to see if it really is a switch. My intuition is that the thing Dr. Kelly and I found either is not a switch or is not large enough to support all the machines, but just in case it is, let's delay buying this for just a while. I will make testing this my first priority once I have Linux working on enough of the intel machines.

**October 29<sup>th</sup>, 2007**

From: Stephen Ware  
To: Dr. Li

Hi Dr. Li!

I finally have some good news! After a lot of work, forum posting, and mailing list reading, I've gotten GAMESS to compile! I built it from source successfully on the sun computers and ran the test cases. 41 out of the 42 test cases gave correct results. I'm not really sure what was wrong with the one that didn't work... I will ask one of my friends who uses GAMESS if she can figure out what is wrong with the results.

Just thought that I would let you know that I can finally move forward with this project after this long delay! This delay took a lot longer than I expected, so I will be spending a lot of time in the lab making up for lost time over the next few weeks. When are you getting back to Loyola? I'd like to show it to you.

The good news is that the DDI (Distributed Data Interface) part of GAMESS, which allows it to run over a cluster using UNIX sockets compiled successfully too. That means that once I get these computers hooked up to one another, it should run in parallel (fingers crossed).

I've narrowed the problem with the other sun machines down to a lack of hard drive driver support. Here is a post that I made on the Aurora mailing list... no one's given me any useful feedback yet, but I thought I would send it to you just in case:

<http://lists.auroralinux.org/pipermail/aurora-sparc-user/2007-October/004816.html>

I'll keep you updated as things move forward. Hope you're going well!  
-Stephen

**October 30<sup>th</sup>, 2007**

From: Dr. Li  
To: Stephen Ware

> Hi Dr. Li!  
>  
> I finally have some good news! After a lot of work, forum  
> posting, and mailing list reading, I've gotten GAMESS to  
> compile! I built it from source successfully on the sun  
> computers and ran the test cases. 41 out of the 42 test  
> cases gave correct results. I'm not really sure what was  
> wrong with the one that didn't work... I will ask one of  
> my friends who uses GAMESS if she can figure out what is  
> wrong with the results.  
>

That is great! I am happy for you too. The experience of  
building something from scratch is more important in  
this case.

>  
> I've narrowed the problem with the other sun machines down  
> to a lack of hard drive driver support. Here is a post  
> that I made on the Aurora mailing list... no one's given  
> me any useful feedback yet, but I thought I would send it  
> to you just in case:  
>  
>

<http://lists.auroralinux.org/pipermail/aurora-sparc-user/2007-October/004816.html>

>  
I went through the posting, and there is one issue that occurred  
to me. It mentioned the disk drive in the Ultra 10 is ATA, but  
it is indeed EIDE, which is a format before ATA came into being.  
You may want to find out if there is any difference between  
ATA and EIDE. In other words, driver support for EIDE and for  
ATA may or may not be the same.

> I'll keep you updated as things move forward. Hope you're  
> going well!  
>

I will be back to New Orleans for the Thanksgiving break for a  
whole week. If you are in town during that week, we may get  
together to look at your work. Or else we can meet in  
January.

Remember to keep a journal for yourself as well as for any  
future students. Your postings and answers from others are  
also good for that purpose.

Regards,  
--Xuefeng Li, (504)865-3340(phone)

---

From: Stephen Ware  
To: Dr. Li

I will be going home for the Thanksgiving break, but if there is some  
time when we are both here, than I'll be happy to show you what I've  
got. Meeting in January is fine also. Just out of curiosity, how  
will the grading for this work?

---

From: Dr. Li  
To: Stephen Ware

> I will be going home for the Thanksgiving break, but if  
> there is some time when we are both here, than I'll be  
> happy to show you what I've got. Meeting in January is  
> fine also. Just out of curiosity, how will the grading  
> for this work?  
>  
>

Well, I'll arrive in New Orleans on Nov. 17, and leave  
New Orleans on Nov. 25. If it's too much trouble to arrange  
a meeting during that week, we sure can meet in January.

As for the grading, please write a report documenting all  
that you've done, including successes and failures. Don't  
forget to attach the postings and answers from the Aurora  
group and other groups.

**October 31<sup>st</sup>, 2007**

From: Stephen Ware  
To: Dr. Li

Dr. Li,

Here is my journal so far for the cluster project. Right now, it is  
just the journal, but when the project is done, I will also add in all  
the other things mentioned in the table of contents. Let me know if  
this is what you had in mind. I've managed to maintain a 4.0 GPA so  
far in my Loyola career, and now that I'm looking at grad schools, I'm  
doing everything in my power to make sure I don't blow it on my last  
year. Let me know if this journal needs anything, and how you think  
the project is going so far.

-Stephen

---

From: Dr. Li  
To: Stephen Ware

The format of your journal is great. You've done a great job  
so far.

I would like to suggest that you add a "References" section  
before the appendices. Do list all references in the References  
section. That includes all books, manuals, technical reports,  
papers, and any online web pages. In other word, make writing  
this journal a professional experience. You will need this  
experience in the very near future.

Once you've compiled the reference list, you can cite a reference  
whenever needed and/or required. For example, when something  
important is mentioned for the first time, there should be

a citation to it. This is out of professionalism as well as ease of reading for your readers.

Here is a particular example. When PVM is mentioned for the first time in your journal, instead of writing "PVM (Parallel Virtual Machine)", you should change it to "PVM (Parallel Virtual Machine) [15]" where [15] is the reference #15 that points to the source concerning PVM. This source could be a book, a paper, a website, or just anything that OFFICIALLY identifies PVM.

Never hesitate to add detail into your journal. It is always easier to delete than to add.

And please do not worry about your grade for this course. As I've said before, you've done a great job so far. I am confident you'll be doing fine for this project.

Regards,  
--Xuefeng Li, (504) 865-3340 (phone)

### **November 15<sup>th</sup>, 2007**

From: Stephen Ware  
To: Dr. Li

Dr. Li,

The final count of working Sun machines is 4, though I'm still holding out hope that I can rescue a 5th one. I've harvested RAM from the unusable machines and put it into the working machines, bringing each machine to 512 megabytes of RAM. I've got all the working machines networked together on a private LAN, and have got SSH key pairs generated so that they can communicate and run processes across the cluster. I'm in the process of getting GAMESS to work... it's almost there, I think... just a few more bugs to work out. Just thought I would let you know how things are coming.

Happy Thanksgiving a little early,  
-Stephen

---

From: Dr. Li  
To: Stephen Ware

Sounds great. And happy Thanksgiving to you, too.

Regards,  
--Xuefeng Li, (504) 865-3340 (phone)

### **December 4<sup>th</sup>, 2007**

From: Stephen Ware  
To: Dr. Li

Dr. Li, I got the cluster working! I finally got it to run GAMESS over all four nodes simultaneously. Right now it is running a larger job so that I can test the speed of the cluster.

There's about two weeks left in the semester, just enough time to do a few more things on this project. Is there anything specifically that you would like to see done?

Cheers,  
Stephen

## **December 5<sup>th</sup>, 2007**

From: Dr. Li  
To: Stephen Ware

> Dr. Li, I got the cluster working! I finally got it to run GAMESS  
> over all four nodes simultaneously. Right now it is running a larger  
> job so that I can test the speed of the cluster.  
>

That's fantastic! I can feel the joy in you.

> There's about two weeks left in the semester, just enough time to do a  
> few more things on this project. Is there anything specifically that  
> you would like to see done?  
>

Please email me your latest version of report/journal.  
I'll then be able to tell if we can do anything further.

Regards,  
--Xuefeng Li, (504)865-3340(phone)

## **December 6<sup>th</sup>, 2007**

From: Stephen Ware  
To: Dr. Li

Here's the up-to-date journal.

---

From: Dr. Li  
To: Stephen Ware

> Here's the up-to-date journal.  
>  
Got it. It is indeed up-to-date.

But please DO add all relevant citations/references  
at the end to complete the journal.

## VII. Appendix B: Posts on LinuxQuestions.org

Please note that in all of the following forum threads, my username was “sgware.”

### **Thread: Best distro for a cluster of old Sun Ultra 10's?** **(Posted September 12<sup>th</sup>, 2007)**

**sgware:**

I'm converting a bunch of old Sun Ultra 10 computers (SPARC chips, I believe) into a Linux cluster. Since they are old machines, I want a no-nonsense distro with a small memory footprint. I've only ever worked with Intel chips, so I really have no idea where I should look for this one. Right now I'm trying out Gentoo... but I wanted to ask here since you folks are always so knowledgeable and helpful.

Thanks!  
-Stephen

**jacook:**

Vector Linux  
<http://www.vectorlinux.com/>  
DSL  
<http://www.damnsmalllinux.org/>  
Puppy Linux  
[http://www.puppylinux.org/user/viewpage.php?page\\_id=3](http://www.puppylinux.org/user/viewpage.php?page_id=3)  
Slackware  
<http://www.slackware.com/>  
BeaFanatIX  
<http://bea.cabarel.com/>  
Elive:  
<http://www.elivecd.org/>  
Luit Linux  
<http://luitlinux.sarovar.org/>  
Also Try Frozen Tech-Live CD List:  
<http://www.livecdlist.com/?pick=All&...ing&sort=&sm=1>

Jake

**sgware:**

I've looked at several of those, but to my knowledge, many of them don't exist for SPARC chips... namely DSL and Puppy, both of which I have used before. Slackware has a SPARC flavor, but is it's memory footprint as small?

If I can use those distros you mentioned on SPARC, please let me know... maybe I am just missing something.

**weibullguy:**

Cross Linux from Scratch

## **Thread: Seagate hd not recognized by Aurora (Fedora 6) on Sun Ultra 10 (Posted October 11<sup>th</sup>, 2007)**

**sgware:**

This is a very specialized case, so thanks in advance to anyone who can help.

I'm trying to put Linux on an old Sun Ultra 10, which means it's a SPARC chip. I'm using Aurora Linux, which is a SPARC port of Fedora. When I run the installer, I get an error message complaining that the machine has no recognizable hard drive. The hard drive in the machine is a Seagate Metalist 9140... ATA I believe.

Does anyone know how I can get Linux to recognize this hard drive? I've tried googling around for the drivers, but have had no luck. From what I understand, this is a fairly common hd model... so I was hoping Fedora would support it out of the box.

Any suggestions (or links to SPARC compatible drivers)?

**Simon Bridge:**

SATA?

IIRC: at the boot prompt do this -  
linux all-generic-ide pci=nommconf

**sgware:**

You mean when I am asked to boot into either graphical or text-based install mode? I will try this and let you know what it says soon (I would do it right now, but the computers as in a lab at my University so I only have access to them at certain times).

**Simon Bridge:**

That would be at the "boot:" prompt. Usually you would just hit enter for a graphical install.

**sgware:**

Ok, I tried this, but I'm still getting the same error message about needing to manually select drivers.

**Simon Bridge:**

Quote the exact message.  
You may end up needing a different distro.

Skimming the Aurora website, I get the impression that this distro has a way to go before it can be considered stable. To add to this, FC6 had problems with some HDD controllers... particularly SATA.

**sgware:**

Are there any other good SPARC distro's?

**Simon Bridge:**

You mean, other than those in the link? They include debian and SUSE.

Admittedly, the "news" on the UltraLinux site is not recent, but the ports have their own pages. Just google for a mainstream distro and "spark" and see if you cannot find something recent... slackware? As far as I can tell, ubuntu 7.10 only has a server port for sparc. This platform is increasingly neglected.

I guess, there's always Solaris.

**sgware:**

That's how I came by Aurora. I looked at or tried several of the others (Debian, Slackware) before settling on Aurora, the only one that actually worked. I will look around some more and see if I can find anything else.

In the mean time, if anyone knows a good source for Seagate drivers, please let me know.

**Simon Bridge:**

It's not the drive, it's the controller (if I've understood you right). Find a pata (regular ide) drive and do an initial install to that. The full installed distro should have what you want, so you can move your /home to the seagate. Or wait until aurora produce a release based on f7 or 8.

### **Thread: Easy way to get gfortran? (Posted: October 21<sup>st</sup>, 2007)**

**sgware:**

I have a fair bit of experience with Linux actually, but this is a newbie question. On Fedora 6, is there a convenient command line method for getting and installing the gfortran compiler? I'm using a SPARC port of FC6, which doesn't come with gfortran by default, and the package manager seems to have problems. IIRC, there is some really simple command line tool (yum perhaps?) that will download, compile, and install packages for you. I'd like to do with that gfortran.

**homey:**

I just use the search feature of yum to see if a package is available...

```
$ yum search gcc-gfortran
gcc-gfortran.i386 : Fortran 95 support
```

The package name maybe abit different for sparc.

Then, you can tell yum to install it...

```
yum install gcc-gfortran
```

**sgware:**

Hmm... I tried that, but I don't think it worked.

```
> yum search gcc-gfortran
> error: Macro % has illegal name (%define)
> Searching Packages:
> Setting up Repo:  aurora-extras
> http://atlantis.ausil.us/extras/2/sparc/repodata/repomd.xml: [Errno 4]
IOError: HTTP Error 404: Not Found
> Trying other mirror.
> Cannot open/read repomd.xml file for repository: aurora-extras
failure: repodata/repomd.xml from aurora-extras: [Errno 256] No more mirrors
to try.
```

**Thread: Anaconda dies after selecting language and keyboard layout on Sun Ultra 10 using FC 6**  
**(Posted: October 21<sup>st</sup>, 2007)**

**sgware:**

I'm installing Aurora [Linux](#) (a port of Fedora 6 to SPARC) on a Sun Ultra 10, and the graphical installer launches fine. But after I select my language and keyboard type, I get an unhandled exception and the installer exits. It gives me a really long error message.

The top line is:

```
> File "/var/tmp/anaconda-10.1.0.2//usr/lib/anaconda/gui.py", line 789, in
nextClicked self.dispatch.gotoNext()
```

The last line is:

```
> SystemError: (2, 'No such file or directory')
```

Does anyone have any idea what's causing this and how to fix it?

Thanks!

-Stephen

**Thread: Auto execute shell script on startup, right before use**  
**(Posted: November 12<sup>th</sup>, 2007)**

**sgware:**

How can I create (or add to an already existing) shell script that will be executed after the rest of the system is done initializing?

Basically, I have some configuration options that I need to set for the computer in question every time it starts up... and some of them rely on the networking services already being configured.

The OS is Aurora SPARC Linux, which is just Fedora 6 for SPARC chips, using bash as the default. If it matters, there is no GUI [system](#) installed on this particular machine... it's a bare-bones command line only box.

I know that several such scripts already exist. I need the one that executes

last, so that all the rest of the system configuration will have already taken place.

**jschiwal:**

Sounds like:  
/etc/rc.d/rc.local

**sgware:**

Prefect, thank you!

### **Thread: Is it enough to set the IP Address and Hostname in a LAN? (Posted: November 15<sup>th</sup>, 2007)**

**sgware:**

Hi LQ,

I'm in the process of building a Beowulf cluster. All of the nodes in the cluster are running Linux and are networked together, and I am using the 192.168.0.x address range for my LAN.

I have added the following two lines to each computer's /etc/rc.d/rc.local file, which is a shell script that gets executed after all the other setup happens.

Code:

```
ifconfig eth0 192.168.0.<number> netmask 255.255.255.0  
hostname node<number>
```

where <number> is the number of the computer.

I've added the necessary aliases to the /etc/hosts files, so that each computer can identify the others by hostname. The nodes use SSH to communicate with each other, and I have set up proper SSH authentication so that no password is required when nodes communicate with each other or run remote commands.

The application I am trying to run on the cluster is returning various TCP errors. Are steps mentioned above enough to effectively change the identities of the computers on the network, or is there more I need to do?

The errors I am getting (if they are any help) are:

Code:

```
rmcd: getaddrinfo: Temporary failure in name resolution
```

```
TCP connect error: Unknown error message.
```

```
TCP connect error: return value errno=43
```

```
TCP: Connect failed. node01 -> masternode:32809
```

This is on Aurora Linux, which is Fedora 6 for SPARC chips, running bash.

Thanks very much for your help,  
Stephen

**cyberfishee:**

the gateway ip?

**jschiwal:**

Look in your /etc/resolve.conf and /etc/host.conf and /etc/nsswitch.conf files.

```
cat /etc/host.conf
#
# /etc/host.conf - resolver configuration file
#
# Please read the manual page host.conf(5) for more information.
#
#
# The following option is only used by binaries linked against
# libc4 or libc5. This line should be in sync with the "hosts"
# option in /etc/nsswitch.conf.
#
order hosts, bind
Make sure that hosts is listed and it is first.
```

In /etc/nsswitch.conf, make sure that files is listed first.

```
hosts:          files mdns4_minimal [NOTFOUND=return] dns
networks:       files dns
Check that you have the libraries "libnss_files" and "libnss_compat" installed.
Example:
/lib64/libnss_files.so.2
```

Make sure that you can ping every host by name.

I'm not as familiar with working with clusters. Could the program be trying to set up unix pipe between the hosts instead of going through ssh? Something about the overhead of going through ssh tunnels in a cluster seems counter productive to me, but you have far more experience with clusters than I do ( which for me is none ). Do you have ssh tunnels set up for certain ports that this program uses?

### **Thread: Use chs instead of bash? (Posted: November 12<sup>th</sup>, 2007)**

**sgware:**

I have a shell script that I need to run from csh, rather than bash. Is that possible?

I have tried entering a csh shell from bash and running it from there, and I have also tried putting "#! /bin/csh" at the top of the file... but I don't think it is working, because I am getting errors in the form "bash: blah no such file or directory."

How can I use straight up csh?

This is on Aurora SPARC Linux, which is just Fedora Core 6 for SPARC chips.

**jschiwal:**

Do you have csh installed? Is it located at /usr/bin/ or /bin?

Try changing your shell to csh (using chsh) and run the script from this shell. I've never used csh and don't know whether an error would start with csh: or bash: because the script was launched from the bash shell. Also, an obvious point is to check the syntax of the script.

**wang gincai:**

```
just do
```

```
$ su -c yum install csh
```

```
$su -c chsh  
reply root password  
$/bin/csh
```

## VIII. Appendix C: Posts on the GAMESS Usergroup

### Thread: GAMESS Cluster Node Setup (Posted: September 27<sup>th</sup>, 2007)

**Stephen Ware:**

Hello GAMESS Group :)

My name is Stephen and I'm in the process of converting a bunch of old Sun Ultra10 workstations into a Linux cluster to run GAMESS. I'm using the Aurora Project's Fedora distribution, so even though I am on a SPARC chip, these computers perform almost exactly the same as other Red Hat-ish distributions I have worked with.

I've just finished building GAMESS from source (with ddikick) on one of these machine, so now I should be ready to start running GAMESS jobs... but I have a questions about the cluster setup. I appologize is this is a stupid question, but how does one go about setting up slave nodes for the cluster? Does each individual box in the cluster need its own copy of GAMESS, or is it possible to install GAMESS on a single machine? Will the master node dole out processes to the slave nodes... or how does that work? Remote Procedure Calls? SSH?

Also, are there any tools or tricks to make sure that a slave node is working? Is there any way to make sure that it is running and contributing work to the overall process? I imagine the Unix 'ps' command should show a GAMESS process if one exists, but I figured while I'm asking dumb questions, I might as well throw this one out there too.

Thanks very much for reading,  
Stephen Ware

Loyola University, New Orleans

**Stephen Ware:**

I also appologize for my spelling in the above post... still getting used to the keyboards on these Sun machines.

**Victor Rosas Garcia:**

Hi Stephen,  
You can have GAMESS stored in the master node, and export its directory using NFS (I'm assuming a single binary will run on all the nodes). Make sure the directory for temporary files is on a local hard disk for each node, because these files can grow to huge sizes and sending all that info through a network would slow things down too much. I don't have experience with actually running GAMESS in parallel, so, cannot help you there.

Hope this helps.

Victor M. Rosas García, PhD  
Coordinador del Posgrado en Ciencias  
Facultad de Ciencias Químicas, UANL  
e-mail: [rosas.vic...@gmail.com](mailto:rosas.vic...@gmail.com)  
Tel: (81) 8329-4010 ext. 6253  
Fax: (81) 8376-5375

**Thread: Compile in 64 bit mode with g77?**  
**(Posted: October 20<sup>th</sup>, 2007)**

**Stephen Ware:**

I noticed in "comp" that the only fortran compilers available for "linux64" are gfortran, pgf77, and pathf90... not g77. I have g77 installed on this machine, so is it possible to use it?

Just for some context, this is on a Sun Ultra 10, which has a 64 bit sparc chip, running Fedora Core 6 Linux.

**Jörg Saßmannshausen:**

Dear Stephen,

I would not recommend using g77 for 64bit code. FC6 should have gfortran installed which gives reasonable binaries.

On Intel (and AMD) chips I would recommend the Intel compiler (works for AMD as well). I am not sure whether it will work on sparc chips, you might one to test it.

Try and get the gfortran on your machine which should not be too much of a problem I would guess.

I hope that helps  
Best wishes  
Jörg

**Stephen Ware:**

Thank you for your speedy reply. As it turns out, even though the chip is a 64 bit chip, this version of Linux runs in "32 bit userland," which apparently means that everything except the kernel operate in 32 bit mode. Sorry for the the trouble, but I think I need to keep trying to compile the source in 32 bit mode.

Thanks again,  
Stephen

**Thread: f771: error: invalid option `align-double'**  
**(Posted: October 23<sup>rd</sup>, 2007)**

**Stephen Ware:**

When trying to compile the GAMESS source code, I keep getting this error in "compall.log."

```
g77 -c -O2 -malign-double -fautomatic -Wno-globals -fno-globals blas.f
f771: error: invalid option `align-double'
```

I looked through the g77 documentation, and found this:  
<http://gcc.gnu.org/onlinedocs/gcc-3.4.6/g77/Optimize-Options.html#Opt...>

Which mentions that the align-double option is only for Intel chips. The computer I am trying to compile GAMESS on is a Sun Ultra 10, which has a SPARC chip, not an Intel chip. How can I safely fix this error without messing up the GAMESS source?

Thanks very much for your help,  
Stephen

### **Thread: Choose Sun32/64 or Linux32? (Posted: October 23<sup>rd</sup>, 2007)**

**Stephen Ware:**

Hi Gamess Group,

I am in the process of installing GAMESS on some old Sun Ultra 10 workstations. I have installed Aurora Linux on these machines, which is a version of Fedora Core 6 (Redhat project) adapted for the SPARC chip. The Ultra 10 has a 64 bit chip, but Aurora runs in "32 bit userland," meaning that for all intents and purposes, the machine is a 32 bit machine.

My question is: when choosing the target type for compilation, should I choose "linux32" or one of the Sun options, "sun32" or "sun64?" The machine is a Sun machine, but it's running Linux... so I don't really know which one to pick.

Thanks very much for your help,  
Stephen

### **Thread: Does the executing node need to be included in <nodelist> when running ddikick? (Posted: November 12<sup>th</sup>, 2007)**

**Stephen Ware:**

When running ddikick.x and specifying the hostnames of the nodes to use for the process, do I need to include the hostname (localhost) for the computer that is executing GAMESS if I want it to be used in the calculation?

**Peter:**

Yes. But you should give it's hostname instead of just 'localhost', otherwise other nodes won't be able to communicate with it.

### **Thread: Using SSH in bash? (Posted: November 13<sup>th</sup>, 2007)**

**Stephen Ware:**

Hi GAMESS group,

I've got GAMESS compiled and linked, along with ddikick, and have run the test jobs. I've finally got my cluster set up, and am ready to execute GAMESS in parallel on the cluster... but I can't seem to get it to work in bash (the default shell in Fedora).

I think the problem might have something to do with setting the environment variable DDI\_RSH. Readme.ddi states that to use ssh, one needs to set the environment variable DDI\_RSH to 'ssh.' The command provided in the readme file "setenv DDI\_RSH ssh" does not work... perhaps because this is a c-shell command? I am not sure... but I do know that in bash, one can see an environment variable with "DDI\_RSH=ssh" followed by "export DDI\_RSH", so this is what I did. But I don't think it is being recognized.

Here is the input I am using:

```
DDI_RSH=ssh
export DDI_RSH
/gamess/ddikick.x /gamess/gamess.01.x /gamess/tests/exam01.inp -ddi 4
4 queenbee workerbee01 workerbee02 workerbee03
```

And the output:

```
Distributed Data Interface kickoff program.
Initiating 4 compute processes on 4 nodes to run the following
command:
/gamess/gamess.01.x tests/exam01.inp
connect to address 192.168.0.3: Connection refused
Trying krb4 rsh...
connect to address 192.168.0.3: Connection refused
trying normal rsh (/usr/bin/rsh)
connect to address 192.168.0.2: Connection refused
Trying krb4 rsh...
connect to address 192.168.0.2: Connection refused
trying normal rsh (/usr/bin/rsh)
workerbee02: Connection refused
workerbee01: Connection refused
ddikick.x: Timed out while waiting for DDI processes to check in.
ddikick.x: Fatal error detected.
The error is most likely to be in the application, so check for
input errors, disk space, memory needs, application bugs, etc.
ddikick.x will now clean up all processes, and exit...
ddikick.x: Sending kill signal to DDI processes.
ddikick.x: Execution terminated due to error(s).
```

**Jörg Saßmannshausen:**

Dear Stephen,

ah, the old ssh problem again. :-)

You need to make sure that you can ssh into any node `_without_` using a password. There are a few ways of doing that:

- Either ask your system administrator to do that globally (for example, using `shosts.equivalent` and host-based keys)
- or, if the above is not possible, you need to generate yourself a keypair and copy the PUBLIC KEY (!) on all the nodes. If, what I assume, the nodes share all the same /home directory, then you would do:
  - `ssh-keygen -t dsa` (leave the line for the password empty!!)
  - `sh .ssh`
  - `cat id_dsa.pub >> authorized_keys`
  - if you just generated the `authorized_keys` file do: `chmod 600 authorized_keys`

Whichever way you so, you should be able to do:  
`ssh nodel`

without getting asked a password.

As for the other question, bash is using `'export VARIABLE=VALUE'` and csh uses `'setenv VARIABLE VALUE'`.

You always can check it with: `echo $VARIABLE`

Also, I would symlink `rsh` to `ssh`, makes things easier ;-)

I hope that helps  
Jörg

**Stephen Ware:**

Thanks for the speedy reply!

> You need to make sure that you can ssh into any node `_without_` using a  
> password.  
I've got this working.

> Also, I would symlink `rsh` to `ssh`, makes things easier ;-)  
That sounds like it would solve the problem... but is the syntax for using RSH the same as for using SSH? I just want to make sure that I won't confuse my programs if they are expecting RSH but get SSH instead.

**Thread: Unknown TCP Errors occur when trying to run on cluster  
(Posted: November 20<sup>th</sup>, 2007)**

**Stephen Ware:**

I have a Linux cluster set up to run GAMESS, and everything seems to be in order... but when I try to run GAMESS, I get strange TCP errors. All the computers on the network can find each other (I can

ping any computer on the network from any other computer), and I have SSH set up so that the head node can ssh into each slave node without a password.

Does anyone know where these strange errors are coming from?

Here is exam01.log, which reports the errors:

```
----- begin exam01.log -----
----- GAMESS execution script -----
This job is running on host queenbee
under operating system Linux at Tue Nov 20 18:22:44 CST 2007
Available scratch disk space (Kbyte units) at beginning of the job is
Filesystem          1K-blocks      Used Available Use% Mounted on
/dev/mapper/VolGroup00-LogVol100
                        8095672    3649756   4034680   48% /
cp -i exam01.inp /gamess/scr/exam01.F05
unset echo
setenv ERICFMT /gamess/ericfmt.dat
<snip>
setenv FCIINT /gamess/scr/exam01.F91
unset echo
/gamess/ddikick.x /gamess/gamess.02.x /gamess/input/exam01 -ddi 4 4
queenbee workerbee01 workerbee02 workerbee03 -scr /gamess/scr
Distributed Data Interface kickoff program.
Initiating 4 compute processes on 4 nodes to run the following
command:
/gamess/gamess.02.x /gamess/input/exam01
TCP connect error: Unknown error message.
TCP connect error: return value errno=43
TCP: Connect failed. workerbee01 -> queenbee:32775.
A fatal error occurred on DDI Process 1.
TCP connect error: Unknown error message.
TCP connect error: return value errno=43
TCP: Connect failed. workerbee01 -> queenbee:32775.
A fatal error occurred on DDI Process 5.
TCP connect error: Unknown error message.
TCP connect error: return value errno=43
TCP: Connect failed. workerbee02 -> queenbee:32775.
A fatal error occurred on DDI Process 2.
TCP connect error: Unknown error message.
TCP connect error: return value errno=43
TCP: Connect failed. workerbee02 -> queenbee:32775.
A fatal error occurred on DDI Process 6.
workerbee03: Connection refused
ddikick.x: Timed out while waiting for DDI processes to check in.
ddikick.x: Fatal error detected.
The error is most likely to be in the application, so check for
input errors, disk space, memory needs, application bugs, etc.
ddikick.x will now clean up all processes, and exit...
DDI Process 0: terminated upon request.
DDI Process 4: terminated upon request.
ddikick.x: Sending kill signal to DDI processes.
ddikick.x: Execution terminated due to error(s).
unset echo
----- accounting info -----
Tue Nov 20 18:23:45 CST 2007netstat -anep | grep 32775
```

```

tcp      0      0 0.0.0.0:32775
0.0.0.0:*          LIST EN      0          10872      3092/
ddikick.x
tcp      0      0 192.168.0.1:32775
192.168.0.1:54582  ESTABLISHED 0          10889      3092/
ddikick.x
tcp      0      0 192.168.0.1:54582
192.168.0.1:32775  ESTABLISHED 0          10886      3097/
gamess.02.x
tcp      0      0 192.168.0.1:32775
192.168.0.1:54581  ESTABLISHED 0          10884      3092/
ddikick.x
tcp      0      0 192.168.0.1:54581
192.168.0.1:32775  ESTABLISHED 0          10881      3096/
gamess.02.x

```

Files used on the master node queenbee were:

```
-rw----- 1 root 4001 1136 Nov 20 18:22 /gamess/scr/exam01.F05
```

Files from workerbee01 are:

```
ls: /gamess/scr/exam01.*: No such file or directory
```

Files from workerbee02 are:

```
ls: /gamess/scr/exam01.*: No such file or directory
```

Files from workerbee03 are:

```
ls: /gamess/scr/exam01.*: No such file or directory
```

```
1.068u 0.344s 1:03.78 2.1%      0+0k 0+0io 0pf+0w
```

```
----- end exam01.log -----
```

And here is a report of what's happening on port

```

----- begin terminal output -----
netstat -anep | grep 32775
tcp      0      0 0.0.0.0:32775
0.0.0.0:*          LIST EN      0          10872      3092/
ddikick.x
tcp      0      0 192.168.0.1:32775
192.168.0.1:54582  ESTABLISHED 0          10889      3092/
ddikick.x
tcp      0      0 192.168.0.1:54582
192.168.0.1:32775  ESTABLISHED 0          10886      3097/
gamess.02.x
tcp      0      0 192.168.0.1:32775
192.168.0.1:54581  ESTABLISHED 0          10884      3092/
ddikick.x
tcp      0      0 192.168.0.1:54581
192.168.0.1:32775  ESTABLISHED 0          10881      3096/
gamess.02.x
----- end terminal output -----

```

**Stephen Ware:**

Additional information:

This is a cluster of SUN Ultra 10's that are running Aurora Linux, which is a port of Fedora Core 6 to the SPARC architecture.

**Jörg Saßmannshausen:**

Dear Stephen,

two things are crossing my mind:

- you are not running a firewall here between the nodes?
- make sure that you actually got the same scratch directory on every node, as:

```
ls: /gamess/scr/exam01.*: No such file or directory
```

looks very conspicuous to me. Usually the directory where gamess rests in and the directory where the input files are have to be the same on all nodes. So if that is your home directory you should be fine if you are exporting that to all nodes, but if you installed gamess somewhere else (say: /usr/local/gamess) you have to make sure that the binaries are the same on all nodes. Same goes for scratch disc, unless you are using /tmp make sure that they are identical on all nodes (like: /gamess/scr). Somehow I got the feeling it is this what goes wrong and it is not a TCP problem.

But thanks for the detailed analysis of the TCP network, it makes helping much easier (and for me more enjoyable) if one gets plenty of information provided :-)

I hope it helps a bit  
Best wishes  
Jörg

**Stephen Ware:**

> - you are not running a firewall here between the nodes?  
Hmm... I will check this. I know that the slave nodes do not have firewalls, but the head node does, and that might be what is causing the problem, since the errors seem to happen when the slaves nodes try to connect to the head node.

> - make sure that you actually got the same scratch directory on every node, The directories and binaries are exactly the same on all nodes... but do I need to copy the input file to each node? Right now, I have the input file only on the head node... perhaps I need to copy it to all nodes before running the program?  
Thank you very much for your help. I will have access to the lab again in a few days and will let you know if either of these fixes the problem.

**Stephen Ware:**

Hmm... well I disabled the firewall on the head node, and that seems to have helped a little, but now I am getting a different, equally ambiguous error.

All computers are set up so that they can ping and ssh into each other without passwords, so I don't know why the connection is being confused. I do not have RSH working... instead, I just got rid of RSH and aliased it to SSH. Could that be the problem? Maybe I need to try using RSH instead?

```
----- Start exam01.log -----
```

```
----- GAMESS execution script -----
```

```
This job is running on host queenbee
```

```
under operating system Linux at Thu Nov 29 21:37:07 CST 2007
```

```
Available scratch disk space (Kbyte units) at beginning of the job is
```



```
cp -i exam01.inp /gamess/scr/exam01.F05
unset echo
setenv ERICFMT /gamess/ericfmt.dat
<snip/>
setenv FCIINT /gamess/scr/exam01.F91
unset echo
/gamess/ddikick.x /gamess/gamess.02.x /gamess/input/exam01 -ddi 4 4
queenbee workerbee01 workerbee02 workerbee03 -scr /gamess/scr
  Distributed Data Interface kickoff program.
  Initiating 4 compute processes on 4 nodes to run the following
  command:
    /gamess/gamess.02.x /gamess/input/exam01
ddikick.x: execvp failed in Kickoff_Remote.
Possible remedies include:
1. rsh may not be on your path, insert
    setenv DDI_RSH /usr/bin/rsh
2. check .rhost authentication file
3. if rsh is not allowed on your system, insert
    setenv DDI_RSH /full/path/to/ssh
4. check path leading to ddikick.x,
    and remote access to ddikick.x on all nodes.
Error: execvp(/etc/shh,args) failed (ENOENT).
Error: execvp(/etc/shh,args) failed (EFAULT).
Error: execvp(/etc/shh,args) failed (ENOTDIR).
Error: execvp(/etc/shh,args) failed (ENOEXEC).
Error: execvp(/etc/shh,args) failed (ETXTBSY).
Error: execvp(/etc/shh,args) failed (ENAMETOOLONG).
Error: execvp(/etc/shh,args) failed (errno=unknown).
ddikick.x: execvp failed in Kickoff_Remote.
Possible remedies include:
1. rsh may not be on your path, insert
    setenv DDI_RSH /usr/bin/rsh
2. check .rhost authentication file
3. if rsh is not allowed on your system, insert
    setenv DDI_RSH /full/path/to/ssh
4. check path leading to ddikick.x,
    and remote access to ddikick.x on all nodes.
Error: execvp(/etc/shh,args) failed (ENOENT).
Error: execvp(/etc/shh,args) failed (EFAULT).
Error: execvp(/etc/shh,args) failed (ENOTDIR).
Error: execvp(/etc/shh,args) failed (ENOEXEC).
Error: execvp(/etc/shh,args) failed (ETXTBSY).
Error: execvp(/etc/shh,args) failed (ENAMETOOLONG).
Error: execvp(/etc/shh,args) failed (errno=unknown).
ddikick.x: Timed out while waiting for DDI processes to check in.
ddikick.x: Fatal error detected.
The error is most likely to be in the application, so check for
input errors, disk space, memory needs, application bugs, etc.
ddikick.x will now clean up all processes, and exit...
DDI Process 0: terminated upon request.
DDI Process 4: terminated upon request.
ddikick.x: Sending kill signal to DDI processes.
ddikick.x: Execution terminated due to error(s).
unset echo
----- accounting info -----
Thu Nov 29 22:54:06 CST 2007
Files used on the master node queenbee were:
```

```
-rw----- 1 root 4001 1136 Nov 29 22:53 /gamess/scr/exam01.F05
Files from workerbee01 are:
ls: /gamess/scr/exam01.*: No such file or directory
Files from workerbee02 are:
ls: /gamess/scr/exam01.*: No such file or directory
Files from workerbee03 are:
ls: /gamess/scr/exam01.*: No such file or directory
1.128u 0.288s 1:04.38 2.1%      0+0k 0+0io 2pf+0w
----- End exam01.log -----
```

**Stephen Ware:**

That doesn't seem to have helped much, but if I add  
setenv DDI\_RSH /usr/bin/ssh  
to the rungms script, I am back to getting "ddikick.x: Timed out while  
waiting for DDI process to check in."  
Any ideas? I've tried everything I can think of and I'm pulling my  
hair out over this.

**Stephen Ware:**

Nevermind! I had forgotten to reset my IP address and hostname from  
the last restart.

So, the final solution to this problem is rather than trying to set  
the DDI\_RSH environment variable before running GAMESS, just put it  
into the rungms script. That seems to have fixed the problem for  
me :).

## **IX. Appendix D: Post on the Aurora User Mailing List**

**Thread: My machine is 64 bit, but does it think it is 32 bit?**

**(Posted: October 22<sup>nd</sup>, 2007)**

**Stephen Ware:**

Hi Aurora group!

I've just recently downloaded Aurora and installed it on some old Sun Ultra 10's at my school in the hopes of making them into a cluster. The program I want to compile and run is called GAMESS, which is a quantum chemistry program that can run in parallel. Because there is no pre-compiled version for SPARC, I am building it from source. In the configuration process, I have to choose several options based on the operating system of the computer. My first instinct was to choose "64 bit linux" since the Ultra 10 is a 64 bit machine, but as I was reading around online about other SPARC linux ports, I realized that perhaps some applications are running in 32 mode. GAMESS is mostly fortran with a tiny bit of C, and I am compiling it from the command line. Should I choose the 64 bit options, or the 32 bit option?

Thanks very much,  
Stephen

**Dominik 'Rathann' Mierzejewski:**

32bit code, unlike on x86, runs much faster than 64bit on SPARC. So if you don't need to access large heaps of memory or if you don't need 64bit precision, then go for the 32bit option. It'll be faster.

Regards,  
R.

**Jos van der Ende:**

Linux on UltraSPARC systems normally runs with a 64 bit kernel and 32 bit userland, since there is rarely a need for an application to use 64 bit unless it does some heavy computations or address a lot (>4G) of memory. The later should not be an issue since an U10 has only 1G max, unless you invested in some exotic memory I am not aware of.

If you are going to do heavy computations, I'd suggest using something better than ancient hardware like an Ultra 10. Even when clustered, the performance will probably not be too great. (Though distcc is doing nicely on my Ultra 5s and Blade 100s.)

So I'd advice using 32 bit. There's a good chance 64 bit won't run or even compile, unless there is a 64 bit version of glibc installed.

**Dennis Gilmore:**

The entire userland for aurora is 32 bit. you will find the only 64 bit thing installed is the kernel.

so you will need to compile 32 bit

Dennis

**Thread: Anaconda dies after language and keyboard selection  
(Posted: October 22<sup>nd</sup>, 2007)**

**Stephen Ware:**

Hi again Aurora group, and thank very much for you speedy and helpful replies to my last question. It's so good to know that there is still support for SPARC Linux somewhere!

As I mentioned earlier, I'm in the process of installing Aurora on some old Sun Ultra 10's that are just taking up space in my school's computer lab. But I'm having problems with some of them. On one two of them, Anaconda doesn't seem to work properly. The graphical installer launches, but after selecting my language and keyboard layout, I get an unhandled exception. The error message is quite long, but this is the top line:

```
File "/var/tmp/anaconda-10.1.0.2//usr/lib/anaconda/gui.py", line 789,  
in nextClicked self.dispatch.gotoNext()
```

And the bottom line:

```
SystemError: (2, 'No such file or directory')
```

I have tried running the installer with "linux ide=nodma," but that didn't help. Does anyone know what is causing this problem, or how I can avoid it?

Thanks very much,  
Stephen

**Tom "spot" Callaway:**

What kind of install is this? Also, what version of Aurora?

~spot

**Stephen Ware:**

Apologies if this reply appears out of thread... I just changed my settings from "digest" to "no digest," so I'm not sure if I'm going to get the original reply (though I have read them on the archives).

spot - I'm not sure what you mean by "what kind of install." Anaconda freezes before asking me to select packages or configuration options. To be honest I don't exactly remember which version of Aurora I downloaded... it's Fedora Core 6, if that helps any. If you need to know the exact version, let me know and I will figure it out based on the machines that I have successfully installed on.

Also, I noticed that the content of the various mirrors is different. Can you please send me a link to the most recent, stable release of Aurora? In case I got an old version, I'd like to re-download and re-burn the disks.

Thanks very much,  
Stephen

**Tom "spot" Callaway:**

Did you install from CD? DVD? TFTP?

Aurora 2.0 is the latest stable release at the moment. Make sure that you check the md5sums of the ISOs if you did a CD/DVD install.

~spot

**Stephen Ware:**

Ok, I'm still having trouble with Anaconda... but it's different trouble, so hopefully that means that I've made some progress. I'm using Aurora 2.0 built on FC3, I have verified the media, and I have used these same disks to install Aurora on other machines, so I do not think it is a problem with the disks. The machine in question is a Sun Ultra 10.

Before, the installer would crash right after it started "Searching for Linux Installations," which happens right after keyboard and language selection. I have been able to get past that point by booting into rescue mode and formatting the harddrive using fdisk and parted. Now, I get all the way through the installation process... package selection and everything (I chose "minimal" install)... but when it starts the actual installation, I get an error with a long traceback. Because the system shuts down after this error, I can't copy and paste it here... so here's the first couple of line copied by hand:

```
File "/var/tmp/anaconda-10.1.0.2//usr/lib/anaconda/gui.py", line 1074,  
in handleRenderCallback self.currentWindow.renderCallback()  
File "/var/tmp/anaconda-10.1.0.2//usr/lib/anaconda/iw/progress_gui.py",  
line 242, in renderCallback self.intf.icw.nextClicked()  
...
```

I have tried installing normally and while using "ide=nodma." I get the same error either way. Can anyone tell me the cause of this error and/or how to fix it?

Thanks very much,  
Stephen

**Thread: Support for a Seagate Metalist 9140 harddrive?  
(Posted: October 22<sup>nd</sup>, 2007)**

**Stephen Ware:**

Greetings again Aurora group. I hope it's not bad form to ask so many question in a row, but I was so encouraged by your speedy response to my last question that I can't help it!

Context: I'm trying to build a Linux cluster out of some old Sun Ultra 10's using Aurora. Unfortunately, some of the hard drives in the machines do not seem to be supported. On about five of the machines, before the graphical installer launches, I am told that the computer does not have a recognizable harddrive and that I will need to select one manually. The computers in question have Seagate Metalist 9140 drives... ATA I believe. I have looked online for drivers and asked about them LinuxQuestions.org, but have had no luck. I assume these disks were fairly standard hardware, since these computers have never been upgraded to my knowledge. Does anyone know where I can get drivers or otherwise fix this problem?

Thank you very much for your support,  
Stephen

**Tom "spot" Callaway:**

Hmm. Those drives should work fine, I wonder if OpenBoot doesn't like them for some reason.

Keep in mind that on an Ultra 10, you need to boot the installer with ide=nodma.

Are there any errors on the other VTs? Maybe, if we're lucky, the kernel is trying to tell us something about those drives.

~spot

**Stephen Ware:**

I have tried booting with ide=nodma, but I got the same error. Also, I'm not sure what a VT is... sorry... I'm new to this!

**T. Ribbrock:**

I have had cases recently with some older drives in an U10 where the drives could not be found by the OBP (and subsequently OS), no matter what I did. Appartently, the on-board IDE controller does have some oddities...

Cheerio,

Thomas

**Thread: f771: error: invalid option 'align-double'  
(Posted: October 25<sup>th</sup>, 2007)**

**Stephen Ware:**

I've installed Aurora Linux on an old Sun Ultra 10, and am attempting to build from source a quantum chemistry analysis program called GAMESS (<http://www.msg.ameslab.gov/gamess/>). When I try to build the source, I keep getting the error:

```
f771: error: invalid option `align-double'
```

I looked this up on the g77 webpage:

<http://gcc.gnu.org/onlinedocs/gcc-3.4.6/g77/Optimize-Options.html#Optimize-Options>

which says that this option applies to Intel x86 architectures only. I tried simply removing this option from the shell scripts that build GAMESS, but the resulting binaries did not function properly. I am not sure if removing this option is what caused the problem or not, but if anyone is familiar with this error, please let me know what I can do to fix it without affecting the GAMESS code.

Also, on a sidenote, I'm still looking for drivers for the Seagate Metalist 9140. If anyone knows where I can find them, please send me a link. Aurora does not seem to recognize my Metalist harddrives.

Thanks very much,  
Stephen

**Owen Mann:**

I would suggest exchanging a working disk with a non-working one. If the working disk boots the previously non-working machine, try the bad one in a PC. There are no special drivers for disks - an ATA drive is an ATA drive.

Another idea is to see if the firmware (aka OBP aka OpenBoot) sees the drive; I don't have a lot of experience with that so I'll have to leave the details to someone else.

**Jos van der Ende:**

*> I would suggest exchanging a working disk with a non-working one. If the  
> working disk boots the previously non-working machine, try the bad one  
> in a PC. There are no special drivers for disks - an ATA drive is an ATA  
> drive.*

Indeed. The IDE controller needs a driver, but the individual disks don't.

The driver needed for an U5/10 is cmd64x, which has been in the kernel tree for years and should have all the bugs ironed out by now. :-)

*> Another idea is to see if the firmware (aka OBP aka OpenBoot) sees the  
> drive; I don't have a lot of experience with that so I'll have to leave  
> the details to someone else.*

The OBP command to do that is probe-ide. This should give a list of attached IDE disks.

I do have a lot of experience with the OBP. ^^

I have noticed in the past that sometimes Aurora has problems with disks that were previously formatted by Solaris, perhaps due to differences in the geometry used.

One thing you can try is erase the partition table and let fdisk create a brand new one:

```
dd if=/dev/zero of=/dev/hda bs=1k count=512
fdisk /dev/hda
```

This will of course destroy all data that happens to be on the disk.

**Stephen Ware:**

> *The OBP command to do that is probe-ide. This should give a list of attached IDE disks.*

> *I do have a lot of experience with the OBP. ^^*

Hmm... this could be my problem. When I ran that command, it didn't recognize anything attached to it except the cd rom drive. Does this indicate a hardware problem?

> *One thing you can try is erase the partition table and let fdisk create a brand new one:*

>

> *dd if=/dev/zero of=/dev/hda bs=1k count=512*

> *fdisk /dev/hda*

Where do I enter this command? I have tried it both on the OBP and at the boot prompt for installing SPARC, but neither worked.

**Stephen Ware:**

> *Where do I enter this command? I have tried it both on the OBP and at the boot prompt for installing SPARC, but neither worked.*

Nevermind on this question... the commands work after booting into rescue mode. Reformatting the harddrives on the computers that can recognize that they have drives allows me to get a little further in the installation process. I'm still getting errors though. If I can't get them fixed, I'll start a new thread.

**Jos van der Ende:**

> > *The OBP command to do that is probe-ide. This should give a list of attached IDE disks.*

> > *I do have a lot of experience with the OBP. ^^*

>

> *Hmm... this could be my problem. When I ran that command, it didn't recognize anything attached to it except the cd rom drive. Does this indicate a hardware problem?*

Mmm, that might indeed be. You'd best check the cabling and especially the master-slave settings of the drives. I've seen this behaviour before and it was caused by faulty jumper settings. You should not use cable select, just one master and one slave per IDE bus.

Another possibility is that the OBP does not recognize the drive. Ultra 10 systems officially support only 4.2-, 9- and 20 GB drives. Try running .version (including the leading dot) on the OBP prompt. If it's not already at 3.31.0, upgrading the firmware might help. If you need help with that, let me know.

Yet another possibility is that the drive is dud. Try replacing it.

**Stephen Ware:**

```
.version yields:  
Release 3.19 Version 4  
OBP 3.19.4  
POST 3.0.7
```

There's certainly not 3.31! I'm not really sure that it's a firmware problem, since I believe the drive currently in the computer was the one that came with it... but it would be worth a try anyway. I'm willing to try anything. What do I need to do to upgrade the firmware?

**Jos van der Ende:**

You need to download the firmware patch from Sun. You can find it on [sunsolve.sun.com](http://sunsolve.sun.com). Select desktop, then Ultra 10. There will be a link to the latest firmware patch, which is 106121-18. The patch comes with installation instructions. Basically you need to boot from a special file.

There are three ways to do that:

- Put the file in an UFS filesystem
- Burn the file to a CD-ROM
- Put it on a jumpstart server

The last one is my preferred method.

**Thread: Where to get an Ethernet card for an old Ultra 10?  
(Posted: November 15<sup>th</sup>, 2007)**

**Stephen Ware:**

I need to add a second Ethernet port to an old Ultra 10. I assume it works more or less way as an Intel machine... buy the card and hook it up to the motherboard? Does anyone know a good source from which to get such a part?

**venn:**

3com! Any

**Jos van der Ende:**

The Ultra 10 has a PCI bus just like (older) PCs, and just a normal PCI ethernet card from your favorite PC parts outlet should work.