NMSU Astronomy Linux cluster: Users guide

Jon Holtzman

August 12, 2015

Entire document in a single Postscript file.

1 Introduction

We have established a cluster of machines in the Astronomy department that are running the Linux operating system. While there are other non-Linux machines (a few Sun workstations, some Windows PCs, and a few Macs), the Linux cluster provides the main computing environment for the academic members of the department.

Linux is a open-source Unix operating system that has been developed by a large number of people. Several different groups distribute the Linux package. We are running RedHat/CentOS Enterprise Linux on our computers; this products has a long support lifetime, and we generally upgrade machines on a several year timescale as convenient.

Every Linux workstation has its own Linux installation, and all packages that come with the installation are located locally on each machine. For software that does not come with the RedHat/CentOS distribution (e.g., astronomy-related software), the software is centrally installed on one of the main Linux servers, astrodisk. The disk with this software is exported and mounted by all of the other machines, so it is accessible to everyone without requiring multiple installations.

The Linux machines are set up as an NIS domain for user accounts, mail aliases, and a few other things. What this means is that account information is kept on a few NIS servers, and all of the other machines query a server for account information. As a result, once you have a Linux account, it works on **all** of the Linux machines; your password will be the same, regardless of what machine you happen to log into.

2 Basic astronomy cluster usage

2.1 Establishing an account and setting the password

Before you can log in, you will need to get an account and a password. Note that an *astronomy* account is distinct from an *NMSU* account; the NMSU account is what you will use for registering for classes and university business. For best handling of email (see below), it is strongly preferred that your astronomy account name match the name of your NMSU account. The best way to insure this is to get your NMSU account set up first. This is done by going to my.nmsu.edu, and using the Activate link to activate your account; it will ask for your desired account name, and will let you know if that name is already taken by someone else. Once you have your account name, then we can set up as astronomy account with the same name: ask Jon Holtzman. When we set up the account, we'll let you know along with a temporary password. You should immediately change this password using change your password, using the *yppasswd* command; you will need to enter your old password and your new one twice.

2.2 Logging in

Normally you will log into a machine directly from a keyboard and monitor directly attached to the machine. If you want to log in from a remote location like home or another institution, note that we only allow external logins via the SSH protocol. Most Unix machines have the ssh command installed, and SSH software is available for Windows and Mac machines as well.

By default, however, direct access to all of the Linux machines except the server, astronomy, is closed to all outside machines except for the Apache Point subnet; this is done to help to improve security. You can always log into astronomy from other locations, and if necessary, log into other cluster machines from astronomy. If you will be away for an extended period of time, and there are other machines that you would like direct access to, it is straightforward to set this up – just ask.

2.2.1 Logging in using SSH keys instead of passwords

You can always log in using your password. However, if you prefer, you can also log in using an SSH key, which has certain advantages. The SSH protocol allows the use of key pairs for login authentication. There are at least two flavors of SSH, which use different formats for the key pairs. Since mid-2004, we have been using OpenSSH. (As of late 2004, APO is using the other flavor).

To generate a key pair, use the command:

ssh-keygen -t dsa

which will create two files in the .ssh subdirectory under your home directory called id_dsa and id_dsa.pub; it will also prompt you to enter a passphrase to use with the key pair. The .pub file is public and can be copied to other machines; the other file should *never* be shared.

If you wish to log into another machine using the SSH key:

- If the remote machine is using OpenSSH, append the contents of the id_dsa.pub file to the .ssh/authorized_keys2 file on the target machine
- If the remote machine is using the other SSH flavor, you first need to convert your public key, using:

  ssh-keygen -e -f id_dsa.pub ¿ *newname*.pub

  Then copy the converted key to the .ssh2 directory of the remote machine. You will also need to add the name of the .pub file into the .ssh2/authorization file on the remote machine.

Once this is done, when you use ssh to connect to the machine, it will prompt you for the passphrase associated with your key.

If you want to come from a machine running the ohter flavor of SSH into our machines, you can convert that SSH format key into our format using ssh-keygen -i -f *name*.pub on the OpenSSH machine (ours), and then use the converted key.

You can register your passphrase for a session to avoid having to repeat typing it by using ssh-add, which will prompt you for your passphrase. If you are doing this from a remote log in session (as opposed to directly logged onto your workstation, you will need to first start a shell with the SSH agent using: ssh-agent *shell* (where *shell* is the name of your preferred shell, e.g., tcsh or bash). After you register your passphrase, any SSH operation which would have required you to enter your passphrase will automatically use the one you entered with ssh-add. This can make multiple uses of SSH very convenient, but beware, it also would allow someone else who happened to come into your office the capability of logging into some other machine without knowing a password!

2.3 Shells

As with all UNIX systems, you can customize your workspace. Linux provides several shells which you can work in. By default, all users are set up using tcsh as the login shell. You can start any other shell by typing its name (e.g., bash, sh). If you wish to change your default shell, use the *chsh* command.

You can define your own aliases, set environment variables using startup files which are read every time the shell is started up. tcsh uses the .cshrc file in your login directory, while bash uses the .bashrc file. Your account is setup up with a default .cshrc file; the default file has an entry which sources a systemwide file which resides in /home/local/etc/cshrc (or /home/local/etc/bashrc for bash). This file sets up a default path so that you will find all of the locally installed applications and packages, and sets some variables which may be used by these packages. You are welcome to modify your .cshrc file as you see fit, but if you want to use the locally installed packages, we strongly recommend you keep the source /home/local/etc/cshrc line in your .cshrc, and also, be careful about making modifications to the command path.

2.4 Window system/ managers

As will all X11 display system, Linux provide the possibility of using a varienty of different window managers which control the appearance of your display. The default window manager we are using is KDE, but gnome is also pre-installed, among others. If you want to use another window manager, you can set this up using menu entries on the main login screen. Different window managers can generally be customized to provide different resources in your menus, taskbars, mouse buttons, etc. For KDE, use the KDE control center application to customize things.

2.5 Email

By default, outgoing mail has a return address *username*@nmsu.edu. If people reply to this, it gets received at the NMSU, not the department, level. It is possible to read your NMSU email using the my.nmsu.edu web interface. You can also set up a personal machine as a mail client to the main NMSU mail server.

Alternatively, many people in the department chooses to have the NMSU email forwarded either to a personal email account or to the Astronomy cluster. This is accomplished by going to http://my.nmsu.edu, go into the email preferences configuration, and have all of your email forwarded. To forward it to the astronomy cluster, use *username*@astronomy.nmsu.edu.

You can also set forwarding to some alternate personal address if you prefer. If you receive your email somewhere other than the astronomy cluster, then you should also set things up so that any email that happens to be sent directly to the astronomy cluster gets forwarded to your preferred address. This can be accomplished by placing a .forward file in your home directory which contains a single line with the email address to which you wish to have your local mail forwarded (note that if you use a .forward file, both this file and the home directory must **not** allow world-write priviledge).

In any case, it is wise to set things up so that all of your email ends up in one place, so you will be sure not to miss anything!

If you set up your mail to be delivered to our server, mail can be seen from *any* of the Linux machines; you do not need to log onto the Linux server to read it.

Many mail clients are available; some of the most commonly used ones are netscape/mozilla/thunderbird (various versions of a similar package), pine (text-based), the MH suite (text and graphical interfaces), plus others. If you don't have a strong preference, we recommend you use netscape/mozilla/thunderbird. We also have a web interface for reading mail (http://astronomy.nmsu.edu/webmail).

If you use a mail client, the server (astronomy) is set up to be either an IMAP or a POP server; we recommend using IMAP as it is more powerful (it allows you to manipulated mail directly on the server instead of just locally on your machine). The server name, both for incoming and outgoing mail, is astronomy.nmsu.edu

You do need to be careful about only using one tool to access mail at a given time; if two different tools are being run (e.g., webmail and a mail client), you may not see all of your messages on each.

Whatever mail tool you use, we prefer that you create folders for filing your mail; if you leave everything in your inbox, it clutters up our mail spooling disk.

2.5.1 Useful mail aliases

Current email lists for subsets of the Astronomy department, as well as the entire department are kept fairly up to date. You can use these by sending email to *listname*@astronomy.nmsu.edu, where *listname* is one of the following:

- astrotenure : tenure track faculty
- astrofac : all faculty
- astroother : research associates
- astrostudents : all graduate students
- astrostudents15, astrostudents14, ... astrostudents02: graduate students starting in academic year 2015/14, 2014/13, ...2002/03, etc.
- astrostaff : office staff
- astroall : all of the above
- astroacad : all of the above except astrostaff
- announce : for, e.g. colloquium announcements, includes some people external to astronomy (e.g. apostaff, physics contact, etc.).

The exact contents of the list can be viewed here. Although some effort is made to keep the lists up to date, please contact Jon if modifications/additions/deletions need to be made to the list.

2.6 Web pages

A Web server is run on the Linux server, astronomy. By default, users will have a Web page that is accessed at http://astronomy.nmsu.edu/*username*.

We have set aside a specific partition, /home/httpd, to host individual Web pages. You will have a directory, /home/httpd/html/*username*, where you can place HTML files that can be seen on the WWW. This is where you would make a personal web page.

2.6.1 Department web pages

Individual person webpages

A set of research-oriented webpages for department members are available for your perusal on the server. To view them, simply load the webpages

http://astronomy.nmsu.edu/dept/html/directory.faculty.shtml

http://astronomy.nmsu.edu/dept/html/directory.rstaff.shtml

http://astronomy.nmsu.edu/dept/html/directory.students.shtml

and click on the appropriate sublinks in the top navigation bar, the left hand menu (under Directory), or on individual names in the primary window.

The contact information on the top of each page is created automatically from a text database, while the text on research interests, publications, and the like is stored in an include file on astronomy. You own this file for your webpage, and have the freedom to update its contents at any time.

You can modify the information on your own page by changing the file /home/httpd/html/dept/directory/[faculty,rstaff,grads]/*username*.include (for faculty, research staff, and graduate students, respectively).

Due to the use of dynamic html, any changes that made to the include file are instantly reflected online.

The *username*.include file, and parallel versions for everyone else (named according to your login ID), are owned by each individual user. If for some reason you find that this is not the case and find it impossible to edit your own file, contact Jon to fix the permissions.

There are a few details about the format of the include files worth mentioning, if you decide to update them.

1. The extra line breaks (¡br¿) at the end of the file extend the length of the primary window so that the page footer lies below the bottom of the left hand menubar. If you have a short research entry and then add in a few paragraphs as your interests develop, some of the line breaks might be removed to compensate (but this depends upon the width of the viewer's browser window, and the browser in use, so it is always safer to have extra line breaks rather than too few).

2. If you have a personal webpage, feel free to place a link to it within the include file text (as has been done on Jon's page, for example). Please use the following format to construct the link.

¡p¿ ¡a href=http://astronomy.nmsu.edu/holtz target="_public"¿¡b¿Home Page¡/b¿¡/a¿

The purpose of the _public target item within the ¡a href=¿ flag is to insure that someone who is browsing through a set of our webpages keeps sending them to the same external page (to avoid cluttering up a viewer's screen with many spawned pages, if they are browsing through our entire directory tree).

3. If you publish a new paper, I hope that you will enjoy placing a link to the journal article within your include file. Please use the following format to construct the link.

¡p¿ ¡a href=http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=1997ApJ...479L.121V&db_key=AST&high=3f5d06d35119202 target ="_research"¿Optical Rotation Curves of Distant Field Galaxies: Sub-L* Systems¡/a¿ ¡br¿ Vogt, N.P., Phillips, A.C., Faber, S.M., Gallego, J., Gronwall, C., Guzman, R., Illingworth, G.D., Koo, D.C., & Lowenthal, J.D. 1997, ApJ, 479, L121

The purpose of the _research target item within the ¡a href=¿ flag is to send external journal paper pages to the same spawned page every time, again to avoid clutter for the viewer.

4. If you put in a link to a close collaborator's page, or some other research link, please use the _research target flag again. We are trying to give all spawned webpages one of these two flags throughout the entire set of department webpages, to distinguish broadly between research information and public (which includes academic) information.

5. Please keep the ¡/font¿ flag at the end of the file, so that the default fonts for the menu bars can come into effect after your prose.

The individual photos are stored in a parallel directory, at
/home/httpd/html/dept/directory/pictures

If you would like to replace your current photo, please let us know. If you have a digitized version of a photo of your own, just send it to Jon.

If the information at the top of your webpage needs updating (for example, if you change offices or you receive a M.S. degree), please let Jon know and we will update your database entry. Your shtml file is overwritten automatically every time the files are created, so do not try to put updates in there (they will disappear).

If you would like to make your research-oriented webpage your public homepage on astronomy, with an easy path for people to load, then place the following redirect file inside your owned directory at /home/httpd/html/*username*. The redirect file should be called index.html, and should contain the following text (replacing faculty.nicole with rstaff.sallam or grads.carrande, etc.).

¡html¿ ¡meta http-equiv=”Refresh” content=”0; URL=http://astronomy.nmsu.edu/dept/html/directory.faculty.nicole.shtml”¿ ¡body¿¡/body¿ ¡/html¿

For example, Nicole placed this file within /home/httpd/html/nicole and now if you load
http://astronomy.nmsu.edu/nicole
into your browser, you will be redirected to the page at
http://astronomy.nmsu.edu/dept/html/directory.faculty.nicole.shtml

Research group webpages

We also have a set of research group web pages, at:
http://astronomy.nmsu.edu/dept/html/research.top.shtm

The content for these research group webpages can be found in files: /home/httpd/html/dept/include/research.[extragalactic,galactic,planetary,solar,stellar].inc

These files are associated with the faculty group, of which all faculty are members, so any faculty can modify the content in these.

2.7 File transfer

The preferred method for file transfer is to use encryped file transfer via the scp or sftp commands.

Please note that standard FTP service (not sftp) to individual accounts is turned off for all of the Linux machines. This is because username and password information used when accessing FTP service is passed over the net unsecured, and can be stolen. **Please do not attempt to ftp into your account from another machine; it will not succeed, but can still compromise security!**

The only FTP access that is allowed is anonymous FTP access to the Linux server, astronomy. The anonymous ftp area that is accessible can be seen directly from any of the Linux machines, in /home/ftp. Upon account creation, there will be a directory /home/ftp/pub/*username* created; files placed here will be available to external users accessing astronomy via anonymous ftp. If you wish to allow external users to put files onto the FTP server, external users will need to cd to incoming/username after logging into the ftp server. You will be able to see files they have downloaded in /home/ftp/incoming/*username* from any of the Linux machines. *Files uploaded into the /home/ftp/incoming/ area will be owned by the user ftp, but can be deleted by anyone; copy these files into your area so that you will own them, and then please delete them from the incoming/ftp area.* **ALL LARGE FILES IN THE INCOMING AREA OLDER THAN 6 MONTHS MAY BE AUTOMATICALLY DELETED TO KEEP THE AREA FROM FILLING UP**.

2.8 Network usage

Please be aware of the distinction between CPU usage/location and disk usage/location. When you log into a machine, you are using its CPU. However, each CPU can access either local disks or disks over the network. The locations /home/users, /home/httpd, and /home/ftp are all located on the astronomy server. Thus, if you log into your machine but access files in one of these directories, you are creating network traffic. This is not a problem so long as the volume is relatively small, but if you were to start using or writing programs that do a lot of transfer to and from disk, it can cause poorer performance.

In addition to being able to access the server disks, each Linux workstation has local disk space. It is important to remember that performance of applications which access disk files depends on whether the file access is local or over the network. In addition, if you are accessing things over the network, you affect the performance of other computer systems in the department as well.

By default, when you get an account, your home directory will be located on the server disk, in /home/users/*username*. Once you start to use significant I/O, you should request a directory on your local machine.

It is important to try to work on local disks whenever possible. Responsible network usage will help everybody in the department.

2.9 Computing resources

There are two department servers, astronomy and astrodisk, that provide basic services. In general, these machines should not be used for any significant processing, as this could affect the cluster performance for all users.

Your desktop is meant to provide basic computer services like editing of files, web access, email access, image display, plotting, and the capability to compile and run basic programs. Many basic programs are installed either on each individual system, or centrally on the server disks, and thus accessible to all machines. However, the desktop computers are probably not optimal for significant computing.

If significant computing resources are needed for your work (e.g. more powerful/faster processing, more memory, etc.), there are several centrally located computers that are available via remote login:

machine	CPUs	Memory (Gby)	Dedicated usage
praesepe	48	132	Large jobs
hyades	48	32	Large jobs, better for parallel code?
virgo	64	64	Large jobs
antigona	16	16	Klypin group
seismo	32	64	Jackiewicz group
solarstorm	48	132	McAteer group
milkyway	64	64	Holtzman group

3 Software/data resources

3.1 RedHat built-in applications

RedHat/CentOS Linux comes with a wide variety of pre-installed applications, which will be found locally on all machines. A subset of these that we know to be useful are (and let us know if you find others on the system which should be included in this summary list):

- vi: A standard UNIX text editor.

- emacs: A powerful text editor.
- tex/latex: A standard text processing system.
- evince/kghostview: Packages for viewing PostScript files (and more)
- xfig: A package for making diagrams. Particularly useful for making large format posters.
- convert: command from the ImageMagick package that allows you to convert images from one format to another.
- pine: A text-based powerful email client
- kmail: A graphical email client for the KDE window system.
- gcalctool: An X11-based calculator.

3.2 Locally installed applications

We have installed a variety of additional packages on the NIS server. These are accessible to all clients over the network through the /home/local directories. A clists of installed packages follows:

- idl. A powerful interpreted language with graphical capabilities. This is a licensed product. Our license allows ten people, anywhere on the cluster, to be running it at a given time.
- IRAF. A standard astronomical image processing package. Set it up before using the first time with the mkiraf command. Start it up with the cl command.
- STSDAS/TABLES. Add-ons to the IRAF package for work related to HST
- AIPS. Another standard astronomical image processing package, with widespread use in the radio community. If you want to use AIPS, you will need to have some special setup on your computer - see a system administrator.
- xvista. Another astronomical image processing package, distributed worldwide by Jon Holtzman of NMSU. Take advantage of the local expertise!
- soffice: link to the OpenOffice suite, which has text processing, presentations, spreadsheet, etc, i.e. an open source alternative to Microsoft Office that runs on the Linux platform (and reads Microsoft files).
- Supermongo. Another plotting package based on the original Mongo but with significant enhancements. Start it up with the sm command.
- Lick Mongo. A plotting package based on the original Mongo by John Tonry, with very similar usage. Start it up with the lmongo command. Also provides a subroutine library which can be linked to using -lmongo
- xemacs. A powerful screen based editor.
- skycalendar/skycalc. Programs for making observing calendars (skycalendar) and computing airmass tables, etc. for astronomical observations (skycalc)
- xv: A package for viewing images in a wide variety of formats, doing some simple image manipulation, grabbing images from the screen, etc

- acroread. Program to read Adobe PDF format files.
- gifmerge. Program to take a set of individual GIF format files and combine them into a single animated GIF file.
- Numerical Recipes library. Provides many numerical subroutines and functions. See the book by Press et al. The routines are provided through an object library. Use -lrecipes_f for f77/g77, -lrecipes_c for c routines.
- slalib. Subroutine library which provides a wide variety of tasks used for making astronomical observations (e.g., airmass, sidereal time, etc.)
- xanim. Package which display movies in a variety of formats
- latex2html. Package which converts LaTeX files into HTML files.

Table of packages, versions, links to sources/docs, update dates, is a TBD item for this manual....

Package	Installed version	Install date	Source	notes
IRAF and related packages				
IRAF	2.11-3			
CTIO				
MSCRED	3.2.3			
STSDAS	22			
TABLES	22			
STSCI				
APT/VTT	1210			
GMS				
SEA				
Text processing				
LaTeX2html	2K.1beta			
OpenOffice	3.0			
SDSS software				
Plotting software				
gnuplot	3.7			
pgplot	5.2.0			
supermongo	2.4.1			
Astronomical software				
HEASOFT	6.6.3	6/09	GSFC	
cfitsio				
AIPS				
difmap				
ds9				
hstphot				
saoimage	1.24			
skycal	4.1			
tinytim	6.0			
other				
gifmerge				
printpro	4.2.1			
slalib				
xanim	2.80.1			

3.3 Astronomical data/catalogs

Some widely-used astronomical catalogs have been downloaded to a disk on the server, /home/catalogs. If additional large data sets would be of use, we can add them here where they are accessible to everybody.

TBD: list of catalogs, etc.

4 Peripherals: Printers, CDROMs, floppies, and tape drives

4.1 Printers

We are currently have several network printers: the HP laser printer in the downstairs computer room (down), the HP printer in the upstairs computer room (up), the color laser printer in the downstairs computer room (color), the Tektronix color dye-sublimation printer in the upstairs computer room (c1), and the HP Design Jet 500PS poster printer (hp_poster), and Rene's laser printer (rene). The downstairs laser printer is set as the default printer. Printer information is summarized in the following table:

Printer queue name	Location	Default?	Printer model	Interface	Network name	Language	Driver
copier/newcopier	Room 116		Canon	Ethernet	astro-copier	Postscript	CUPS on astronomy
expensive	Room 118	No	Brother MFC-9840CDW	Ethernet	astro-mfc9840cdw	Postscript	CUPS on astronomy
down	Room 116	Yes	HP LaserJet 2400dtn	Ethernet	astro-hp2430	PostScript	CUPS on astronomy
up	Room 217	No	HP LaserJet 4200	Ethernet	astro-hp	PostScript	cups on astronomy
c1	Room 116A	No	Tektronix Phaser 450	Ethernet	astro-color	PostScript	cups on astronomy
hp_poster	Room 107A	No	HP DesignJet 500PS	Ethernet	astro-bigprint	PCL	cups on astronomy
color	Room 116	No	HP Color LaserJet 2605dn	Ethernet	astro-hp2605dn		cups on astronomy
rene	Room 204	no	HPLaserJet6P	Ethernet	prt-rene	PostScript	cups on astronomy
holtz	Room 102	no	Brothter HL4570CDW	holtz	PostScript	cups on astronomy	

For the upstairs and downstaris printers (-Pup, -Pdown, -Pcolor), you can specify options on the command line for varions printing options, e.g.:

- lpr -Pup filename, will print the file with the default double-sided mode
- lpr -Pup/single filename, will print the file in single-sided mode (-Pdown/single for downstairs printer)
- To request manual paper feed input (e.g. for department letterhead, add
  - -o ManualFeed=True for -Pdown or -Pup
  - -o InputSlot=Tray1_Man for -Pcolor
- lpoptions -p up -l, will list a whole slew of possible options...

Because of this capability, it should never be necessary to use the printer control panel to change printer configuration.

You can also configure your print jobs using the graphical xpp command.

Printers are managed on the Linux machines using a commercial software product, PrintPro; this functionality is built-in to RedHat distributions from version 7.3 on. This allows printing via the normall UNIX commands: lp and lpr (on some of the machines, there is an additional graphical printing interface, glp). Documentation is available for the user interface and for the administrators guide. Configuration can be done using /usr/sbin/printers or via the Web at http://astrodisk:631.

The two laser printers and the HP can be accessed directly via embedded web servers.

4.1.1 Making posters with the poster printer

Making posters is relatively straightforward. There is a non-negligible cost to make these, so it is best to get it right the first time and not to have to make too many attempts.

We have three kinds of paper available; in order of expense, these are: lightweight paper, heavyweight paper, and glossy paper. It is very easy to switch between different rolls (you use the menu items on the printer itself to load and unload rolls). The poster printer has 42 inch wide paper; hence you want one of the two dimensions of your poster to be smaller than this size. To avoid needing to make long cuts, you might want to choose to have either the vertical or horizontal dimension be exactly 42 inches.

There are numerous ways one might make a poster for the poster printer. Here I mention a few possibilities:

- If you have a Postscript or PDF file for your poster, you should be able to print it from any of the Linux machines. The only issue here is that you will need to set the paper size appropriate for the poster that you have created. This is probably most easily done by submitting the print job using the xpp command, and using the graphical interface to specify the correct page size (select the printer hp_poster, hit the "Options" button and set the correct paper size under the "Basic" tab, you can use the "Custom" media size to set an arbitrary size). Alternatively, try: lpr -Php_poster -o pagesize=Custom.XPTSxYPTS file.ps (where XPTS and YPTS are the size of the poster in points, 72 points/inch.)

- If you use a Mac or Windows machine to create your poster (e.g. with Powerpoint, Keynote, etc.), you can probably print to the poster printer directly from your machine: try adding a printer with address astro-bigprint and have it search for the most recent driver.

- Create a poster using the xfig software available on the Linux cluster. This can be used to draw lines, boxes, text, incorporate images, etc. etc. It is straightforward to use: type xfig to start it up; use the z/Z key to zoom out/in. You can get a template for a 42 by 48 inch poster or for a 36 by 42 inch poster; these templates contain border boxes, and some sample headers (in the big one). For significant amounts of text, xfig is not recommended because each line of text will become its own object, so if you need to edit one line, you may need to fix all subsequent lines by hand if it changes the amount of text in a column. As a result, the use of LaTeX or some other program to make text into some format that can be directly incorporated into xfig is suggested. You might find it convenient to include figures in the LaTeX source; alternatively, you can import them directly into xfig and just use LaTeX for text. You will want to use a large font in the text producing program, as you do not want to use a smaller font and then scale it up to a large size using the scaling feature in xfig. If you use LaTeX, good results are obtained using the 12pt Huge font; you can get a sample LaTeX template that works fairly well and also shows how you can include a figure directly in the LaTeX text here. Convert the .dvi file into Encapsulated PostScript using dvips -E *filename*, and then include this into xfig (use the "Picture" button). After you have included the PostScript, you should use the "Use original size" button in xfig to insure that the text is not scaled (probably it can be safely scaled down, but you probably don't want to scale it up).

  To include graphs, photos, illustrations, etc., you can use the "Picture" button to include Postscript, GIF, JPEG, and several other formats. The figure will be resized to fit into the box that you create; in most cases, you will probably want to adjust the height/width ratio to its original value using the "Shrink to orig" or "Expand to orig" buttons in the Picture menu. Once the H/W ratio is correct, you can resize it preserving this ratio using the "Scale" button by pulling on one of the corners of the object.

Once you have a nice looking poster from xfig, save it into a *filename*.fig file, using "Save As" from the "File" button.

There is a simple script available that will convert this file to a large PostScript file and submit it to the printer with the correct page size; you can run this using:

fig2poster *filename*

where *filename* is the name of the file without any extension. After converting the poster to PostScript, this script will ask you to confirm that you really want to send it to the poster printer! Make sure you are happy with the poster and that the appropriate type of paper is loaded in the printer before printing.

You can make a miniature (8.5x11) version of the poster using, e.g.:

fig2dev -L ps -m 0.2 *filename*.fig *filename*mini.ps

You can adjust the size by changing the magnification factor (0.2 in above example). This can then be printed to one of the laser printers or to the Epson color inkjet printer (lpr -Pepson *filename*mini.ps)

- Use LaTeX to make the entire poster. Some nice LaTeX style stuff has been created for this by several groups; a particularly nice example can be found at http://www.cstr.ed.ac.uk/ robert/posters.html. This example produces A0 size posters which are smaller than our printer is capabile of printing.

  I need to customize the style files, etc., to get this to work nicely with our printer.

  More later....

- Related to use of xfig above, note that you can just use xfig to incorporate a single encapsulated Postscript, JPEG, or GIF file, so if you have any software that creates one of these, you can just use xfig to convert it to a printable format. At the same time, you can also use xfig to scale your input figure to a new size (but beware of scaling things up very much, as you will probably see the resolution limits of your original graphic).

- Use any software available on the PC poseidon, which has printer drivers for the poster printer. At this point, I don't know of good Windows packages, but I'm sure they exist; if someone wanted to recommend something, we might be able to obtain it.

- If anyone develops other methods, let me know....

Note that if you want to include an NMSU logo on your poster, you can get one here (also see the official university site)

4.1.2 Printer supplies

- Epson PhotoEX printer. Color ink cartridges: Epson part S020110. BW ink cartridges: Epson part S020093. Standard inkjet paper is adequate for general use, but is relatively poor for printing color images: photographic or glossy paper is preferred for this (but it is more expensive). All of these items have been available at OfficeMax; they also have been available (cheaper) at Insight and buy.com.

- HP DesignJet 500PS. See http://www.hp.com/cposupport/plotters/support_doc/bpp03084.html. These are available direct from HP, or (cheaper) at Insight and buy.com, although not every item has always been available at the latter two places. Even better prices may be available from http://www.thenerds.net.

- Ink cartridges: C4844A, C4911A, C4912A, C4913A; ballpark $40 each. You can get status of how much ink is left using the control panel on the printer.
- Printheads: C4810A, C4811A, C4812A, C4813A. These do not need to be changed frequently!
- Paper:
  * Basic High Gloss 42in by 100ft (~ $117) C8317A,
  * Basic Heavyweight Coated 42 x 100 (~ 45.99) C7925A,
  * Basic Bond 42x150 (~ 21.99) C7922A.

4.2 USB drives

TBD.

4.3 CDROM

Most computers come with an internal CDROM drive. After inserting a disk, the CDROM can be immediately accessed through the directory /misc/cdrom, as the device is auto-mounted. It should automatically dismount if it is not accessed for a short period of time. You can manually force a dismount (e.g. if you want to eject the disk), using: umount /misc/cdrom.

4.3.1 Writing CDROMs

Many, but not all, of the individual machines have CD writers. The easiest way to burn things to the CDROM is probably to find the menu items available through your window manager.

If this is not available, create a directory hierarchy that you want to be copied to the CDROM; the capacity of a CDROM disk is 650 Mbytes. Issue the command:

mkcd *directory*

to copy all of the files under the specified directory to the CDROM. This will take up to half an hour. The final step in the process compares the files on disk to those on the CDROM; if you see any error messages here, there *may* have been problems with the write (although sometimes you get errors here when everything is still fine).

4.4 DVDROM

Several machines have DVDROM readers. These are accessed just like the CDROMs, and in fact, the DVD readers also read regular CDROMs. On io, there is both a fast CDROM reader as well as a DVD reader (both located in the computer case). The CDROM is accessed through /misc/cdrom; the DVD is accessed through /misc/dvd. Other machines with DVDROM drives are aphrodite and maritimus.

4.5 floppies

Linux can read floppy disks, including those formatted by DOS/Windows. If you have a floppy drive, after inserting a disk, it can be immediately accessed through the directory /misc/floppy.

4.6 Tape drives

Linux supports many SCSI tape drives. Currently, we have one 8mm Exabyte drive and two 4mm DAT (one DDS3 and one DDS) drive installed on the machine astrotape, which is located in the downstairs computer room. It can be accessed through using normal UNIX tape commands (tar, mt, etc.). The Exabyte drive has a capacity of $\sim$ 8 Gby; the DAT DDS has a capacity of $\sim$ 4 Gby. The DAT DDS3 drive on astrotape has a 24 Gby capacity, and an even higher capacity DAT DDS4 drive on maritimus. Tape drive availability is summarized in the following table.

Tape drive	Host computer	Device name
Exabyte 8mm	astrotape	/dev/st? /dev/nst?
DAT 4mm (DDS)	astrotape	/dev/st? /dev/nst?
DAT 4mm (DDS3)	astrotape	/dev/st? /dev/nst?
DAT 4mm (DDS4)	maritimus	/dev/st0 /dev/nst0

The /dev/st$n$ devices will automatically rewind after each access; the /dev/nst$n$ will not rewind.

Occasionally, one may get I/O errors when trying to read tapes. One thing to check if this occurs is whether the block size is incorrectly set on the tape drive. In general, the tape drive block size should be set to 0, which means it should be able to adapt to whatever block size the tape was written with. You can check the tape drive setting using the command: mt -f *device* status, and you can change the blocksize (to 0) using: mt -f *device* setblk 0.

5 Backups

We perform two different types of system backups: a daily backup of small files **only**, and a weekly backup of all files on selected disks. In both cases, the backups are simply to other disks in the department; **we do not have any tape backups**, so there is an inherent risk in our backup systems, as any event that were to wipe out files on multiple disks in the department could wipe out both original and backup versions! If possible, all users are advised to backup critical files themselves on some external backup medium.

5.1 Daily small file backups

For the daily backups, we do not have sufficient space/bandwidth to backup large files; *note that most image data files, e.g. FITS files, will not fall within the size limits*. The partial disk backups currently copy *all* files that are smaller than some size ($I \sim 200$ kbytes), *some* files that fall in the range ($I < size < E \sim 500$ kbytes), and *no* files larger than $E \sim 500$ kbytes). For the intermediate sizes between I and E, the files are backed up if they have any of the following extensions: .c, .h, .l, .y, .sh, .csh, .rc, .awk, .nawk, .sed, .pl, .cl, .pro, .m, .mgo, .sm, .plt, .f, .for, .F, .FOR, .html, .hdl, .tex, .txt, .man, .me, .ms, or .wp; any of these extensions that are compressed (.gz or .Z) will also be included.

The disk backups are made daily and are saved for 7 days. This provides the possibility of recovering files that you might have accidentally corrupted within the past 7 days. However, this does **not** provide long-term archival storage.

Backups can be found in appropriately named directories under /home/backup/base/*machinename*.

5.2 Weekly backups

Weekly full backups are done on many, but not all, of the disks on the department computers. Machines with very large disks are not backed up, since we do not have sufficient space. In addition, it is possible to restrict some files on the disks from being backed up; typically, we do not backup files from people who have left the department.

The weekly backups can be found in /home/backup/astrobackup/*machinename*. Check to see if your files are being backed up!

6 Machines

Machine	OS	Kernel	No. processors	CPU	CPU speed	Memory
ajax	CentOS release 5.9 (Final)	2.6.18-308.24.1.el5	4	Intel(R) Xeon(R) CPU W3550 @ 3.07GHz	3066.885	3086820 kB
algiers	CentOS release 5.9 (Final)	2.6.18-308.13.1.el5	4	Intel(R) Xeon(R) CPU W3570 @ 3.20GHz	3200.237	6103840 kB
antigona	CentOS release 5.9 (Final)	2.6.18-308.8.2.el5	16	Dual-Core AMD Opteron(tm) Processor 8218	2612.041	16437392 kB
denebola	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	2	AMD Athlon(tm) 64 X2 Dual Core Processor 3800+	2004.229	2057548 kB
astronomy	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	8	Dual-Core AMD Opteron(tm) Processor 8218	2612.081	16438128 kB
gera	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	1	Intel(R) Pentium(R) 4 CPU 3.40GHz	3391.577	2057200 kB
hyades	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	48	AMD Opteron(tm) Processor 6168	800.000	32961560 kB
maritimus	CentOS release 5.9 (Final)	2.6.18-308.13.1.el5	4	Intel(R) Core(TM) i5 CPU 750 @ 2.67GHz	1200.000	8132912 kB
milkyway	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	64	AMD Opteron(TM) Processor 6272	1400.000	66007032 kB
nadir	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	2	Intel(R) Core(TM)2 CPU 6600 @ 2.40GHz	2394.000	3983200 kB
orion	CentOS release 5.9 (Final)	2.6.18-308.4.1.el5	4	Intel(R) Xeon(R) CPU W3520 @ 2.67GHz	2666.837	6110200 kB
astrodisk	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	4	Dual-Core AMD Opteron(tm) Processor 2218	1000.000	7930708 kB
praesepe	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	48	AMD Opteron(tm) Processor 6174	800.000	132100252 kB
seiff	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	2	Intel(R) Core(TM)2 CPU 6600 @ 2.40GHz	2394.002	2056552 kB
seismo	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	32	Quad-Core AMD Opteron(tm) Processor 8354	1100.000	66005272 kB
solarstorm	CentOS release 5.9 (Final)	2.6.18-308.24.1.el5.centos.plus	48	AMD Opteron(tm) Processor 6180 SE	800.000	132096164 kB
virgo	CentOS release 5.9 (Final)	2.6.18-348.3.1.el5	64	AMD Opteron(TM) Processor 6272	1400.000	66007032 kB
achernar	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	1600.000	5860316 kB
aldebaran	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2000.000	5860316 kB
altair	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2700.000	5860316 kB
antares	CentOS release 6.4 (Final)	2.6.32-358.2.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	1600.000	3794400 kB
arcturus	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	1600.000	5860316 kB
babylon	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2500.000	5860316 kB
capella	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	1600.000	5860316 kB
castor	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	1600.000	5860316 kB
kermodei	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2200.000	5860316 kB
matrix	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2800.000	5860316 kB
mimosa	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2400.000	5860316 kB
mintaka	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2100.000	5860316 kB
phoebe	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	1600.000	5860316 kB
polaris	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2900.000	5860316 kB
regulus	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2900.000	5860316 kB
rigel	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2600.000	5860316 kB
spica	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	1600.000	5860316 kB
tequila	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2900.000	5860316 kB
vega	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_6x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	1600.000	5860316 kB
zenith	CentOS release 6.4 (Final)	2.6.32-358.11.1.el6.x86_64	2	Intel(R) Pentium(R) CPU G2020 @ 2.90GHz	2900.000	5860316 kB