

Linux Cheat Sheet

by Charlie Zender
University of California at Irvine

Department of Earth System Science
University of California
Irvine, CA 92697-3100

zender@uci.edu
Voice: (949) 824-2987
Fax: (949) 824-3256

Contents

Contents	i
List of Tables	1
1 CCSM: Community Climate System Model	2
1.1 CAM: Community Atmosphere Model	2
1.2 CLM: Community Land Model	2
2 ESS 200B: Earth System Physics	4
2.1 iPod	8
2.2 Cameras	8
2.2.1 CompactFlash	8
2.2.2 Memory Stick	8
3 X and XFree86	9
3.1 X and XFree86: Configuration	9
3.2 XFree86: NVidia Graphics drivers	13
3.3 Performance Profiling	14
3.4 Virtual Private Network: VPN	16
3.5 XFree86: Security	17
3.6 XFree86: Multiple sessions	17
3.7 LCD Projectors	18
3.8 Desktop	18
3.9 JPEG	18
4 Disk Management	19
4.1 RAID	19
4.1.1 Hardware	19
4.1.2 Getting Promise cards to work	19
4.1.3 Software	19
4.2 Servers: Sand	20
4.3 Formatting Disks	21

4.4	Creating LVM	21
4.5	LVM on ashes	22
4.6	Knoppix	23
4.7	Mirroring Disks	24
4.8	Garbage collection	24
4.9	Disk corruption	25
5	Software	25
5.1	netCDF Browsers	25
5.1.1	ncBrowse	26
5.1.2	panoply	26
5.1.3	ncview	26
5.1.4	NCVweb	26
5.2	Python	26
5.3	Mail	26
5.4	Macintosh OS X	28
5.4.1	Fink	28
6	Programming	29
6.1	Operator precedence	29
6.2	Regular expressions	29
6.3	Bash	32
6.4	Internationalization I18N	33
7	Files	34
7.1	Date conventions	34
8	System Administration	36
8.1	Backups	36
8.2	Mounting devices	37
9	rsync	37
9.1	Slink	38
9.2	Documentation	38
9.3	PPP	38
9.4	Batch Queues	39
9.5	Remote shell service rcp, rsh, ssh, telnet...	39
9.6	Root	39
9.7	L ^A T _E X	39
9.8	Library	40
9.9	Info	40
9.10	Networking	40
9.11	C development	40
9.12	Time and TimeZone	40
9.13	Perl	41

9.14	Links	41
9.15	Install /etc/initscript to boost stacksize so CCM can run	42
9.16	Math Libraries	42
9.17	Hardware description of Zender group computers:	42
9.17.1	Firefox	59
9.18	Create krein disk structure	59
9.19	NFS export	59
9.20	NFS mount these disks on local machine:	60
9.21	Install sudo	62
9.22	Install IDL and PGI	63
9.22.1	PGI on sand	65
9.23	PCMCIA	65
9.24	Building new kernel	66
9.25	LILO configuration	67
9.26	Restarting daemons	68
9.27	Monitor Kernel Startup	68
9.28	Strip downloaded text files of DOS linefeeds	68
9.29	Installing Debian	69
9.30	Recreating Server Files	74
9.31	i8500s	76
9.32	Images	76
9.32.1	Gimp	76
9.32.2	KView	77
9.32.3	Kuickshow	77
9.33	CDs	77
9.34	DVDs	77
9.34.1	MPlayer	77
9.34.2	dvdrip	78
9.34.3	Xine	78
9.35	Network	79
9.36	Kernel Modules	84
9.37	Change network topology	85
9.38	LGGE	88
9.39	Router	89
9.40	Sound	89
9.41	dakine = msw, lanina = linux	90
9.42	DHCP	90
9.43	Wireless	91
9.44	ARP Address Resolution Protocol	92
9.45	Hibernate	93
9.46	PCMCIA	94
9.47	core dumps	94
9.48	Printing	95
9.48.1	CUPS	95
9.49	Virtual Memory	96

9.50	tramp.el	97
9.51	RPM packages in /usr that were installed (and may need to be uninstalled or reinstalled) by hand:	97
9.51.1	IPCC and PBS configuration	97
9.51.2	GPG: GNU Privacy Guard	99
9.51.3	Building RPMs	100
9.51.4	RPM commands	100
9.51.5	APT commands	101
9.52	Required software	103
9.53	Packages installed in /usr/local (some RPM, some *.tgz, some proprietary):	104
9.54	NCO and patches	104
9.55	Skype	104
9.56	Installing RedHat Linux	105
9.56.1	Updating RedHat Linux with latest patches	105
9.56.2	Upgrading RedHat Linux	106
9.57	Debugging	106
9.58	Valgrind	107
9.59	DDD	108
9.60	Mailman	109
9.61	Web Servers	109
9.61.1	Group Web Server	110
9.61.2	ESS Web Server	110
9.62	FTP and firewalls	110
9.63	Accounts	111
9.64	NCAR	111
9.65	Autotools	111
9.66	SSH	112
9.67	Security	113
9.68	GCC	113
9.69	Groups	114
9.70	WINE	114
9.71	Partitioning	114

Bibliography	115
---------------------	------------

Index	116
--------------	------------

List of Tables

1	Funky Keystrokes	18
2	PCD Commands	30
3	GDB Commands	108
4	Partitions	114

1 CCSM: Community Climate System Model

```
# Preliminary system modification
sudo ln -s /usr/bin/make /usr/bin/gmake
aptitude install netcdf mpich mpi-doc # Debian GNU/Linux
```

1.1 CAM: Community Atmosphere Model

1.2 CLM: Community Land Model

Make Linux workstation capable of running CLM offline:

```
cd ~
# Download and unpack CLM 2.1 source code into your home directory
http://www.cgd.ucar.edu:8080/accept/license?action=fillOut&file_id=7
# Use CVS to check-out model if you have an NCAR login:
cvs -d :ext:zender@goldhill.cgd.ucar.edu:/fs/cgd/csm/models/CVS.REPOS \
co -r clm2_deva_52 -d clm2_deva_52 clm2

# Create model run space
sudo mkdir /ptmp
sudo chmod 777 /ptmp
if [ -n "${LOGNAME}" ]; then export LOGNAME=${USER}; fi
mkdir /ptmp/${LOGNAME}

# Create and populate boundary condition directories
export CSMDATA_NCAR=/fs/cgd/csm/inputdata/lnd/clm2
export CSMDATA_UCI=/datashare/csm/inputdata/lnd/clm2
export CSMDATA_UCI=/ptmp/${USER}/inputdata/lnd/clm2
export CSMDATA_UCI=${HOME}/tmp/inputdata/lnd/clm2
if [ ! -d ${CSMDATA_UCI} ]; then sudo mkdir -p ${CSMDATA_UCI}; fi
data_drc_fll_lst="NCEPDATA" # Directories to be copied in full
# Directories to selectively copy
data_drc_prt_lst="inidata_2.1/cam srfddata/cam pftdata rawdata rtmdata"
# Create all boundary condition directories
for drc in ${data_drc_fll_lst} ${data_drc_prt_lst}; do
  mkdir -p ${CSMDATA_UCI}/${drc}
done
# Populate entire directories
for drc in ${data_drc_fll_lst}; do
  scp -r "babyblue.ucar.edu:${CSMDATA_NCAR}/${drc}/*" ${CSMDATA_UCI}/${drc}
done # done populating directories
# Populate selected directories with selected files
for fl in \
inidata_2.1/cam/clmi_0000-09-01_64x128_T42_c021125.nc \
```

```
pftdata/pft-physiology \  
rawdata/mksrf_glacier.nc \  
rawdata/mksrf_lai.nc \  
rawdata/mksrf_lanwat.nc \  
rawdata/mksrf_navyoro_20min.nc \  
rawdata/mksrf_pft.nc \  
rawdata/mksrf_soicol_clm2.nc \  
rawdata/mksrf_soitex.10level.nc \  
rawdata/mksrf_urban.nc \  
rtmdata/rdirc.05 \  
srfdata/cam/clms_64x128_T42_c020514.nc \  
; do  
if [ ! -f ${CSMDATA_UCI}/${f1} ] ; then  
    scp babyblue.ucar.edu:${CSMDATA_NCAR}/${f1} ${CSMDATA_UCI}/${f1}  
fi # endif  
done  
  
export USER_FC=1f95  
cd ~/clm;  
clm.sh
```

2 ESS 200B: Earth System Physics

Atmospheric Physics and Radiation

Course Description:

Physical processes which mediate the transformation of energy, heat, and momentum in the Climate system. Topics include hydrostatics, convection, radiative forcing and climate feedbacks.

Here are topics for the 5 Week ESS Atm. Phys/Radiation mini-course. Each bullet is one week, and is subdivided into two parts:

1. Hydrostatics, Thermodynamics
2. Buoyancy/Convection, Radiance/Absorption
3. Flux/Emission, Scattering/Actinic Flux
4. Clouds/Direct Effects, Aerosols/Indirect Effects
5. Earth's Energy Budget, T \leftrightarrow Climate interactions

The text has been chosen to dovetail with other courses, ``Global Physical Climatology'' by Hartmann.

Each day's topic (Tues./Thurs. structure) comprises the key points enumerated below, which are to be illustrated by theoretical and practical examples from the course text and from current research.

Hydrostatics:

1. Mass, density, pressure
2. Ideal gas law
3. Hydrostatic equilibrium
4. Dry lapse rate
5. Realistic T profiles

Thermodynamics:

1. Work, Expansion compression
2. Heat, heat capacity
3. Adiabatic vs. Diabatic
4. Saturation
5. Condensation, Evaporation
6. Latent heat
7. Moist lapse rate

Earth's Energy Budget:

1. Planetary Radiative Equilibrium
2. Energy Transformations
3. Dry, Moist Static Energy
4. Gravitational Potential Energy

5. Available Energy
6. Global Mean Energy Storage

Climate Models:

1. Planetary Radiative Equilibrium
2. Boltzmann's Law, Effective Temperature
3. Solar and Infrared Radiative Budgets
4. Greenhouse Effect
5. Forcing, Response, Sensitivity
6. 0-D, 1-D Climate models
7. Multiple Climate Equilibria

Buoyancy/Convection:

1. Dry/Moist static energy
2. Stability/Instability
3. Convectively Available Potential Energy
4. Simple convection: Adiabatic adjustment
5. Convective redistribution of heat, vapor, tracers

Flux, Emission:

1. Solid angles/Radiance/Isotropy
2. Blackbody radiation
3. Radiance-Flux relationships
4. Boltzmann's law
5. Radiative equilibrium T structure of grey atmosphere

Absorption/Extinction:

1. Molecular cross-sections/spectra
2. Absorption: Vibration, rotation bands
3. Extinction
4. Beer's law
5. Explaining the observed spectrum of the clear atmosphere
6. GHGs and the atmospheric window

Scattering/actinic flux:

1. Reflection
2. Direct vs. diffuse radiation
3. Particle-light interaction
4. Size parameters, Mie theory
5. Photochemical radiation
6. PAR, NDVI definitions

Clouds:

1. Formation, Lifecycle
2. Cloud climatology (cloud classification taught in Hydrology)

3. Particle size distributions
4. Radiative properties
5. Liquid/ice distinctions
6. Direct radiative forcing: LW, SW, Net

Aerosols:

1. Formation, Lifecycle
2. Global sources
3. Sinks: scavenging, dry deposition
4. Global distribution & radiative forcing
7. Vertical flux/heating profiles for realistic atmospheres

Indirect Effects:

1. Daisyworld?
2. Indirect forcing via clouds
3. Surface/Atmosphere partitioning of absorption
4. Global energy budget, including latent, sensible heat
5. Meridional heat transport (atm, ocn)
6. Nutrient transport?

T<-->Climate interactions, Physical Climate Change

1. Coupling of troposphere/stratosphere change
2. Coupling of troposphere/land surface change
3. Water vapor feedback
4. Effects of altered lapse rates
5. Thermal expansion of oceans
6. Estimating RF of GHG increases
7. Observed indirect effects

Notes on course design:

Topics which might belong in other courses which

I assume will be covered here (and thus might be duplicated):

1. Buoyancy/convection (maybe in Boundary layer)

Topics which might belong here but which I assume

will be covered in other courses (and thus might fall through the cracks):

1. Ice/albedo feedback (Hydrology cryosphere)
2. Mass continuity equation (Dynamics)
3. Advection (Dynamics or Geosci. Model)
4. Diffusion (Dynamics or Geosci. Model)
5. Box models, feedbacks, residence times (Hydrology)
6. Aerosol composition, nutrients, pH (Chemistry)
7. Gibbs/Helmholtz energy (Terr BGC)

Topics which might belong here but which I assume will not be covered in any first year unit due to time constraints. These topics should be available in an advanced graduate course.

1. Formulation of RT Equation
2. Solution for pure absorbing atmosphere with reflecting lower boundary
3. Definition of single scattering, asymmetry
4. Radiative-convective equilibrium?

2.1 iPod

The primary reference for using the Apple iPod with Linux is <http://www.cavecanen.org/linux/ipod>. Two Linux applications for managing iPod music are GNUpod and gtkpod.

2.2 Cameras

If your Linux uses modules, load the following modules: Support for USB. Also select either UHCI (Intel PIIX4, VIA, ...) support, UHCI Alternate Driver (JE) support or OHCI-HCD (Compaq, iMacs, OPTi, SiS, ALi, ...) support.

```
usbcore.o
usb-uhci.o, uhci.o or usb-ohci.o
```

In particular, CompactFlash relies on the USB driver.

2.2.1 CompactFlash

To transfer files from your Camera's CompactFlash to disk, first mount the CompactFlash card as a `vfat` file system:

```
mount -t vfat /dev/hde1 /cf # CompactFlash, root-only
```

This CompactFlash is enumerated as an IDE drive here, because it is inserted into the system on the PCMCIA bus (in a PCMCIA adaptor). CompactFlash to USB adaptors are at least as common. The USB sub-system uses SCSI drivers, so CompactFlash drives will be enumerated as SCSI drives (e.g., `/dev/sda1`) when mounted through a USB port.

By default, CompactFlash devices require root privileges—they are not user-writable unless opened with the appropriate mask:

```
mount -t vfat -o umask=022 /dev/hde1 /cf # CompactFlash, users R/W
```

In this mode users may read and write files but can not actually own files or directories.

The `vfat` filesystem does not support multiple user identities and privileges. However, the device may be mounted with a particular UID and GID which will allow full user read-write access.

```
sudo mount -t vfat -o umask=022,uid=3555,gid=2400 /dev/hde1 /cf # CompactFlash
```

This is the recommended way to mount CompactFlash used for backup and file transfer purposes.

2.2.2 Memory Stick

A Memory Stick, aka a USB Flash Drive, is very similar to CompactFlash device. As always the Flash Memory HOWTO is very useful.

To quickly backup and copy material to my black attache memory stick, use

```
sudo mkdir -p /media/USB20FD/${HOME} /media/USB20FD/${DATA}
sudo chmod 755 /media/USB20FD/${HOME} /media/USB20FD/${DATA}
sudo chown zender /media/USB20FD/${HOME} /media/USB20FD/${DATA}
sudo chgrp cgdcsm /media/USB20FD/${HOME} /media/USB20FD/${DATA}
```

To quickly backup and copy material from *virga* to my grey Attache memory stick, use

```
sudo mkdir -p /media/disk/${HOME} /media/disk/${DATA}
sudo chmod 755 /media/disk/${HOME} /media/disk/${DATA}
sudo chown zender /media/disk/${HOME} /media/disk/${DATA}
sudo chgrp cgdcsm /media/disk/${HOME} /media/disk/${DATA}
```

In the old days, to mount a Memory Stick one did this:

```
# Install necessary sysfs utilities
aptitude install hotplug sysfsutils udev
# Create a mountpoint for sysfs
cd /
mkdir sys
mount -t sysfs none /sys
# Mount USB Memory Stick
sudo mkdir -m 777 /memstick
sudo mount -t vfat -o umask=022,uid=3555,gid=2400 /dev/sda1 /memstick # Ash
sudo mount -t vfat -o umask=022,uid=3555,gid=2400 /dev/sdb1 /memstick # Vir
```

For Linux 2.6 kernels, make sure to install *udev*.

The Memory stick with a USB adaptor appears automatically in the */sys* filesystem as a SCSI drive:

```
zender@elnino:~$ more /sys/block/sda/device/model
IntelligentStick
zender@elnino:~$ more /sys/block/sda/device/vendor
I-Stick2
```

It must still be mounted. See the *rsync* description in Section 9 for examples of how to back up to CompactFlash and Memory Stick media. I bought a PNY Technologies “Attache” 4 GB USB 2.0 Flash drive on 20060729. I bought a second PNY Technologies “Attache” 4 GB USB 2.0 Flash drive on 20071006.

3 X and XFree86

X tends to have large memory leaks. Sometimes the problems are due to the X server, and sometimes the problems are due to programs which never free graphics space requested by the X server. The *xrestop* monitors resources consumed by the X server.

3.1 X and XFree86: Configuration

A configuration tool has been built directly into the XFree86 server accessible with `XFree86 -configure`. `XFree86 -configure` generated a perfectly working `XF86Config` file for *lanina*, and is the method I now recommend. XFree86 4.0.1 also has a different and new configuration tool, `xf86cfg`. `xf86cfg` did not generate a working/optimal configuration file for *lanina*,

so I do not recommend this method. Both these methods should generate fairly good configuration files which must be stored as `XF86Config-4`, e.g., `/etc/X11/XF86Config-4`. Note the extra "-4", which indicates suitability for `XFree86-4.x`. Version 4 uses dynamically loaded modules for each particular chipset so there is no need to symbolically link the X Window System server (e.g., `/etc/X11/X`) to anything.

XFree86 version 3.3.x uses the file `/etc/X11/XF86Config`, which may be generated by Xconfigurator:

```
Xconfigurator
cp /etc/X11/XF86Config ${HOME}/linux
cp ${HOME}/linux/XF86Config /etc/X11
```

In XFree86 version 3 the correct X server should be dynamically linked to the generic server file, e.g.,

```
ln -sf /usr/X11R6/bin/XF86_I128 /etc/X11/X
```

The Google `googleearth` application requires correctly functioning OpenGL drivers. The `xorg.conf` file instructs X to load these drivers with

```
Section "Module"
    Load "glx"
    Load "GLcore"
EndSection
```

It is important that `glx` precede `GLcore`. The `glxinfo` command reports the OpenGL driver status.

Use `import` to save an X window to an image format

```
import san_antonio.gif
import -quality 100 san_antonio.png
```

Tweaking the default `XF86Config` file is often necessary for special performance. The following tweaks to `swcursor` and `XkbOptions` are useful

1. Software cursor to allow large cursors. In Section "Device" add

```
# csz++
# 19991005: "swcursor" required for 96 x 94 pixel ~/.lightning cursor
Option "swcursor"
# csz--
```

2. Swap positions of capslock and control keys. In the `InputDevice` section for the keyboard add

```
#csz++
# Swap positions of capslock and control keys
XkbOptions "ctrl:swapcaps"
#csz++
```

In the `InputDevice` section for the keyboard add

```
#csz++
# Make pressing both touchpad buttons at same time emulate middle
# button on 3 button mice
Option      "Emulate3Buttons" "on"
#csz++
```

Adjust the speed of the touchpad mouse using the `MinSpeed`, `MaxSpeed`, and `AccelFactor`, options in the `Touchpad` section

```
Section "InputDevice"
Identifier "Synaptics Touchpad"
Driver "synaptics"
Option "SendCoreEvents" "true"
Option "Device" "/dev/psaux"
Option "Protocol" "auto-dev"
Option "HorizScrollDelta" "0"
Option "MinSpeed" "0.5"
Option "MaxSpeed" "0.7"
Option "AccelFactor" "0.0350"
EndSection
```

Another way to do this is to add

```
Option "SHMConfig" "on"
```

and then

```
aptitude install ksynaptics
```

Make sure the 100-dpi fonts are installed or else many `xterm` and `emacs` fonts will appear chunky.

Font management and installation is described by the Linux Documentation Project (LDP) at <http://www.linuxdoc.org/HOWTO/Font-HOWTO.html>. X can be made to use any font desired, but the required procedures to do so are arcane. Here is a skeleton outline:

```
ps ax | grep xfs # Check for running font server xfs
xset -q # Check font path
xset fp+ unix/:port_number # Add font server xfs port to font path
xset fp rehash # Rehash font path
sudo /etc/rc.d/init.d/xfs restart # Restart font server xfs
end{verbatim}
```

None of this should be necessary on a `\trmid{x}{RedHat}` GNU/Linux machine, since RedHat has an ```automatic``` way of serving fonts.

Instead, what should be done on all machines, is to list in order of preference the valid font directories in the `\verb'Files'` section of the `\flidx{XF86Config-4}` file.

```
\begin{verbatim}
Section "Files"
```

```

FontPath "/usr/X11R6/lib/X11/fonts/local/"
FontPath "/usr/X11R6/lib/X11/fonts/misc/"
FontPath "/usr/X11R6/lib/X11/fonts/75dpi:unscaled"
FontPath "/usr/X11R6/lib/X11/fonts/100dpi:unscaled"
FontPath "/usr/X11R6/lib/X11/fonts/Type1/"
FontPath "/usr/X11R6/lib/X11/fonts/CID/"
FontPath "/usr/X11R6/lib/X11/fonts/Speedo/"
FontPath "/usr/X11R6/lib/X11/fonts/75dpi/"
FontPath "/usr/X11R6/lib/X11/fonts/100dpi/"
FontPath "/usr/share/AbiSuite/fonts/"
FontPath "/wnd/windows/fonts/"

```

```
EndSection
```

Also, make sure the X-server loads the modules to handle freetype and Type 1 (Adobe) fonts. This is accomplished by having

```

Section "Module"
    Load "freetype"
    Load "type1"
EndSection

```

Notice that the font server can use the MS Windows fonts directly from the Windows fonts directory, `/wnd/windows/fonts`. However, the fonts in this directory first must be prepared for use by running a few commands in that directory:

```

cd /wnd/windows/fonts
ttmkfdir -o fonts.scale
mkfontdir

```

Also, make sure all the fonts are world-readable.

When fooling with X configurations it is convenient to have a clean way to shut down and restart all X processes. This can be done by initializing the system to a new runlevel. Most Unices start multi-user mode and all network applications *except* the X server in runlevel 3. Runlevel 5 starts the same processes as runlevel 3 and starts the X server as well. Thus initializing a system to a different runlevel (usually 3 or 5) is a clean way of starting and stopping X processes. The `telinit` command serves this purpose. Shut down all X processes with `telinit 3`. Restart all X processes with `telinit 5`. Set the default system runlevel in `/etc/inittab`. This is what determines whether X automatically starts on reboots.

Some computers have terrible default X settings which may be corrected by use of `xset`, usually done in `~/xinitrc`. For example, lanina has DPMS turned off by default, and very slow key repeat settings so its `~/xinitrc` contains

```

xset q # Show current settings
xset +dpms # Enable DPMS
xset dpms 300 600 1200 # Seconds until Standby, Suspend, Off
xset mouse 4 4 # Mouse acceleration and threshold

```

The keyboard repeat rate may be changed with the `kbdrate` command. The `-r` options sets the repeat rate in characters per second. The maximum allowed rate, 30 cps, gives a nice response.

```
kbdrate -r 30 # Set keyboard repeat rate to 30 cps
```

The default mouse movement rate is also susceptible to fubaration.

`xvidtune` generates modelines which center the picture on the monitor.

The `XAPPLRESDIR` contains X defaults for various programs.

```
mkdir ${DATA}/app-defaults
export XAPPLRESDIR="${X11}/lib/X11/app-defaults:${DATA}/app-defaults"
```

These defaults may be over-ridden on a per-user basis by implementing X resources in the `.Xdefaults` file.

3.2 XFree86: NVidia Graphics drivers

Source code drivers that support all of the advanced features of the ubiquitous NVidia graphics cards are unavailable. Two type of drivers are available. First, the XFree86 project `nv` driver is open source and works well with all 2D drawing commands. Second, NVidia supplies binary-only `nvidia` drivers which support full 3D acceleration and OpenGL graphics. The `nvidia` kernel modules are closed source, prone to errors, and must be re-installed after the kernel is recompiled. Since there is no easy mechanism to retain multiple `nvidia` modules for multiple kernels on one machine (e.g., for testing), the constraints of the closed-source modules are bothersome.

Since these drivers are modules, installing them requires kernel recompilation which, in turn, requires kernel headers.

```
dpkg --get-architecture '*686*'
apt-cache search kernel | grep 2.6.15 | grep 686
```

```
# Ubuntu Breezy Badger
uname_r=`uname -r`
aptitude install linux-image-${uname_r} linux-restricted-modules-${uname_r}
aptitude remove linux-image-${uname_r} linux-kernel-headers-${uname_r} linux-
```

Note that recent Debian kernel packages require the initial ram-disk (`initrd`) to be specified in the `GRUB menu.lst`.

```
aptitude install libncurses5-dev
cd ${DATA}/zender/tmp
svn checkout http://powertop.googlecode.com/svn/trunk/ powertop
cd ${DATA}/tmp/powertop
make
sudo make install
```

Linux experts re-compile their own kernels frequently. GNU/Linux will not boot into X until the correct X server module is available for the current kernel. Download NVidia drivers from <http://www.nvidia.com/object/linux.html>. Integrate the NVidia drivers into these kernels by hand:


```
wget http://download.nvidia.com/XFree86/Linux-x86/1.0-7174/NVIDIA-Linux-x86-
cd /usr/src; sh NVIDIA-Linux-x86-1.0-4496-pkg2.run
cd /usr/src; sh NVIDIA-Linux-x86-1.0-5336-pkg2.run
cd /usr/src; sh NVIDIA-Linux-x86-1.0-6111-pkg1.run
cd /usr/src; sh NVIDIA-Linux-x86-1.0-7174-pkg1.run
```

When using the closed source NVidia module driver, make sure that XF86Config-4 references `nvidia` not the open-source XFree86 driver `nv`.

The X.org and XFree86 X servers write output from the initialization process to `/var/log/Xorg.0.log` and `/var/log/XFree86.0.log`, respectively.

Ubuntu Breezy Badger packages for NVIDIA cards:

```
sudo aptitude install linux-restricted-modules-2.6.15-16-686 # Non-free Linu
sudo aptitude install nvidia-glx nvidia-settings avm-fritz-firmware # NVIDIA
```

3.3 Performance Profiling

`clay` is set up to do performance profiling by both `oprofile` and PAPI with HPCToolkit. This required a kernel patch and recompile. There is one module for each approach: `oprofile` for `oprofile` and `perfctr` for any software that requires the PAPI modifications (specifically, `hpcrun` from the HPCToolkit). There are other performance tools that use the PAPI API (`Tau`, for one), but so far we only have HPCToolkit installed. As a side note, most Ubuntu kernels have the `oprofile` module available, but we are not aware of any with the more exciting `perfctr` patch applied and the `perfctr` module available.

The Oprofile homepage is <http://oprofile.sourceforge.net/docs> Oprofile first requires

```
sudo modprobe oprofile
```

Second, initialize the `oprofiled` daemon and start it collecting info. This command depends on your exact hardware/software configuration.

```
opcontrol --vmlinux=/path/to/vmlinux # Normal generic
opcontrol --no-vmlinux # No vmlinux exists
```

The path must lead to the *uncompressed* linux ELF executable, not to typical `vmlinuz` compressed boot sector that is installed in the `/boot` directory. For `clay.ess.uci.edu`:

```
sudo opcontrol --vmlinux=/usr/src/linux-2.6.11/vmlinux
# Explicitly say one CPU:
sudo opcontrol --separate=none --vmlinux=/usr/src/linux-2.6.11/vmlinux
# --separate=cpu reports counts on both CPUs
sudo opcontrol --separate=cpu --vmlinux=/usr/src/linux-2.6.11/vmlinux
# Once separated, must explicitly shut off for succeeding runs
sudo opcontrol --start
# Next line kills daemon, removes small amount of overhead:
sudo opcontrol --shutdown
# Remove profile module (needed for hpcrun)
opcontrol --deinit
```

In order for `hpcrun` to work, `perfctr` module must be `modprobe`-loaded and `/dev/perfctr` must be mode 644. When ready to collect info, do a `sudo ls` to initialize the timeout on the `sudo` command so later commands do not ask for passwords. For `ncbo`, assuming `ncbo` has been compiled with the `-g` option,

```
sudo opcontrol --reset # Reset counters
ncbo -O --op_typ='- ' -p ${DATA}/nco_bm \
gcm_T85.nc gcm_T85_00.nc ${DATA}/nco_bm/ipcc.diff.nc
oproport --exclude-dependent --demangle=smart --symbols > \
oprofile.report.full.ncbo
```

The output is a text file that gives the time spent in each function. The `poll_idle` time is that time which the CPU(s) has spent doing nothing, i.e., idling. For a lightly loaded dual-CPU machine, you would expect to obtain about 50% in `poll_idle` running a single serial job.

To use `HPCToolkit`, make sure that `oprofile` is not loaded, and load the `perfctr` module.

```
lsmod | grep oprofile
sudo opcontrol --deinit # Unload oprofile if loaded
modprobe perfctr # Load perfctr module
```

The `PAPI` API has access to necessary hardware counters once the `perfctr` has been loaded. After this, profiling is relatively straightforward. To profile a command, prefix it with `hpcrun`, e.g.,

```
hpcrun [options] -- ncbo -O --op_typ='- ' -p ${DATA}/nco_bm \
gcm_T85.nc gcm_T85_00.nc ${DATA}/nco_bm/ipcc.diff.nc
```

The `hpcrun` options are typically a set of hardware counters you want to access during the run.

```
% hpcrun -L | grep Yes # List available hpcrun options
PAPI_L2_DCM      Yes      Level 2 data cache misses ()
PAPI_L2_ICM      Yes      Level 2 instruction cache misses ()
PAPI_FPU_IDL     Yes      Cycles floating point units are idle ()
PAPI_TLB_DM      Yes      Data translation lookaside buffer misses ()
PAPI_TLB_IM      Yes      Instruction translation lookaside buffer misses ()
PAPI_L1_LDM      Yes      Level 1 load misses ()
PAPI_L1_STM      Yes      Level 1 store misses ()
PAPI_L2_LDM      Yes      Level 2 load misses ()
PAPI_L2_STM      Yes      Level 2 store misses ()
PAPI_STL_ICY     Yes      Cycles with no instruction issue ()
PAPI_HW_INT      Yes      Hardware interrupts ()
PAPI_BR_TKN      Yes      Conditional branch instructions taken ()
PAPI_BR_MSP      Yes      Conditional branch instructions mispredicted ()
PAPI_TOT_INS     Yes      Instructions completed ()
PAPI_FP_INS      Yes      Floating point instructions ()
PAPI_BR_INS      Yes      Branch instructions ()
PAPI_VEC_INS     Yes      Vector/SIMD instructions ()
```

PAPI_RES_STL	Yes	Cycles stalled on any resource ()
PAPI_TOT_CYC	Yes	Total cycles ()
PAPI_L2_DCH	Yes	Level 2 data cache hits ()
PAPI_L1_DCA	Yes	Level 1 data cache accesses ()
PAPI_L2_DCR	Yes	Level 2 data cache reads ()
PAPI_L2_DCW	Yes	Level 2 data cache writes ()
PAPI_L2_ICH	Yes	Level 2 instruction cache hits ()
PAPI_L1_ICA	Yes	Level 1 instruction cache accesses ()
PAPI_L1_ICR	Yes	Level 1 instruction cache reads ()
PAPI_FML_INS	Yes	Floating point multiply instructions ()
PAPI_FAD_INS	Yes	Floating point add instructions ()
PAPI_FP_OPS	Yes	Floating point operations ()

These options are requested as follows

```
hpcrun -e PAPI_TOT_CYC:32767 -e PAPI_FP_OPS:32767 -e PAPI_FP_INS:32767 \
-e PAPI_HW_INT:32767 -e PAPI_L2_DCM:32767 -- <command_to_profile>
```

hpcrun profiles *everything* that results. For example, `command_to_profile` is a shell command, then hpcrun profiles every subcommand in the shell, and gives each its own output file in the form of: `app_name.PAPI_TOT_CYC.clay.ess.uci.edu.10137.0`.

Process hpcrun output files into something usable with hpcquick. hpcquick is a Perl script that calls some other hpc tools to generate the XML database (in its own subdirectory) that hpcviewer needs.

```
# src_location          hpct DB file to process
hpcquick -I src/nco -P ncwa.PAPI_TOT_CYC.clay.ess.uci.edu.10137.0
# View results with Java hpcviewer
hpcviewer # Open the './hpcquick.dbxxx/hpcquick.hpcviewer' file
```

To use oprofile on clay requires a re-compiled kernel with the profiling switches enabled. Once re-compiled, the kernel may require a new Nvidia driver. This may require un-installing and purging packages, e.g.,

```
aptitude remove --purge nvidia-glx
```

If the package is not purged, then `/etc/init.d/nvidia-glx`, which runs at each startup, may wipe out the TLS links that the new Nvidia driver needs.

3.4 Virtual Private Network: VPN

Use a Virtual Private Network (VPN) to obtain a UCI net address from an off-campus computer.

```
sudo /etc/init.d/vpnclient_init start # Load kernel module
sudo vpnclient connect UCI # Start VPN connection
sudo vpnclient connect UCIFull # Start VPN connection--full mode
```

NACS uses allows browsers to request a VPN by logging into the UCI VPN Proxy at <https://vpn.nacs.uci.edu>.

3.5 XFree86: Security

X supports a number of security measures, not all of them helpful. By default, the Debian distribution prohibits X connections from any processes not started by the current desktop owner. Apparently this is enforced through so-called TCP-forwarding. TCP-forwarding may be prohibited with the `-nolisten tcp` argument to the X server. To permit other users access to your desktop display, remove this argument from appropriate file(s). Debian starts the X server from `xserverrc (/etc/X11/xinit/xserverrc)`. The `gdm.conf (/etc/gdm/gdm.conf)` may also need to be modified.

```
/usr/bin/X11/X -dpi 100 -nolisten tcp
```

Another way of allowing access to your X desktop is to authorize access from all clients using the `xhost` command:

```
xhost + # Allow connections from any server
```

This may be risky unless the desktop sits behind a firewall or is offline.

3.6 XFree86: Multiple sessions

Default X display is accessible as Ctl-Alt-F7 on RedHat Linux. For a different number of bitplanes on default display, use, e.g., `startx`

```
startx -- -depth 24
```

To create a second display at, say, 24 bits-per-pixel (), accessible as Ctl-Alt-F8, use

```
startx -- -depth 24 :1
```

To create a second display using a test XF86Config file, use

```
startx -- -xf86config linux/etc/X11/XF86Config-4.elnino :1
```

For security reasons, the path to the XF86Config file must be a relative path, not an absolute path.

Table 1 summarizes the keystroke commands related to X Windows.

Accessing virtual consoles and starting multiple sessions is the same in XFree86 4.x as XFree86 3.x. However, the specification of the color depth has changed from `-bpp` to `-depth`. To start X with 24 bit color the command is

```
startx -- -depth 24
```

Display managers such as `xdm`, `kdm`, and `gdm` start X themselves, and require that non-default options be passed directly to the X server in the configuration file, e.g., `/etc/X11/gdm/gdm.conf`, or `/etc/kde/kdm/kdmrc`. Before configuring display managers individually, however, it is wise to consider setting many parameters in the system-wide X configuration file, XF86Config-4. For example, adding `DefaultColorDepth 24` to Section "Screen" should cause the X server to always run with 24 bit-planes.

Table 1: **Funky Keystrokes**^a

Key	Description
Ct1-Alt-F1--F5	Virtual consoles
Ct1-Alt-F7	X display :0.0
Ct1-Alt-F8	X display :0.1
Ct1-Alt-+ ^b	Change X Resolution
Ct1-Alt-- ^c	Change X Resolution

^aUse `em or lose `em

^bThis *must* be the plus key (+) on the numeric keypad. On most laptops this requires a *four key combination*, since the numeric keypad requires a key to activate.

^cThis *must* be the minus key (-) on the numeric keypad

3.7 LCD Projectors

LCD Projectors interact with X in strange ways. The ESS Department projector works best at resolution fxm. The IGPP projector works best at resolution fxm.

3.8 Desktop

The X server automatically starts at the specified system runlevel. The default runlevel is specified in `/etc/inittab`, and is OS-dependent. In a given runlevel (RedHat uses runlevel 5 for this, Debian uses runlevel 2), X automatically starts the display manager specified in `/etc/desktop`. The most common choices are GNOME, KDE, and XDM. These choices invoke the display managers `gdm`, `kdm`, and `xm`, respectively. These display managers are capable of starting any desktop on the system (as determined in their configuration settings). Naturally they default to their appropriate desktops, e.g., `kdm` starts the K desktop environment. This default may be over-ridden by `/etc/X11/default-display-manager`

```
$ cat /etc/X11/default-display-manager
/usr/bin/kdm
```

3.9 JPEG

Processing JPEG files produced by IDL.

```
for fl_stb in `ls *.jpg | perl -p -e 's/$*\.\.jpg/$1/g;'` ; do
    echo "Converting ${fl_stb}.jpg to ${fl_stb}.pnm..."
    jpegtopnm ${fl_stb}.jpg > ${fl_stb}.pnm
    echo "Converting ${fl_stb}.pnm to ${fl_stb}.png..."
    pnmtopng ${fl_stb}.pnm > ${fl_stb}.png
    echo "Converting ${fl_stb}.pnm to ${fl_stb}.png..."
    pnmto ppm ${fl_stb}.pnm > ${fl_stb}.ppm
done
```

4 Disk Management

4.1 RAID

RAID is Redundant Array of Independent Disks. Reference: <http://en.tldp.org/HOWTO/Software-RAID-HOWTO.html>

4.1.1 Hardware

(Will Kitto helped with this setup)

The RAID on biogenic consists of 4 × 150 GB Maxtor disks. There are two Promise disk controller cards installed in biogenic, with two disks connected to each card. The jumpers on each disk were set appropriately. All disks are "master" not "slave", so the disks have the following IDs:

```
Card which drive
-----
1      m1      hde
      s1      hdf (does not exist)
      m2      hdg
      s2      hdh (does not exist)
2      m1      hdi
      s1      hdj (does not exist)
      m2      hdk
      s2      hdl (does not exist)
```

4.1.2 Getting Promise cards to work

The Promise cards were not supported by the kernel (that comes with RH7.2). A patch was found to support the Promise chipset, and a new kernel was made, with all the RAID modules included. I do not know if RH8.0 has support for the Promise cards, but I would guess yes.

4.1.3 Software

The devices were setup to run software RAID-5 (i.e. not RAID done in hardware). This means that of the 4 disks, only the capacity of 3 (i.e., 450 GB) is available for users (the other 150 GB are for redundancy). I followed the instructions that you'd see in any typical RAID howto (e.g. as per the reference). The `raidtab` was setup

```
% more /etc/raidtab
raiddev /dev/md0
raid-level          5
nr-raid-disks      4
nr-spare-disks     0
persistent-superblock 1
parity-algorithm   left-symmetric
```

```

chunk-size          32
device              /dev/hde1
raid-disk           0
device              /dev/hdg1
raid-disk           1
device              /dev/hdi1
raid-disk           2
device              /dev/hdk1
raid-disk           3

```

```

% mkraid /dev/md0
% mke2fs -v -j -b 4096 -R stride=8 /dev/md0

```

```
/dev/md0 /raid ext3 defaults 1 2
```

To check the status of the RAID, view `/proc/mdstat`:

```

% more /proc/mdstat
Personalities : [linear] [raid0] [raid1] [raid5]
read_ahead 1024 sectors
md0 : active raid5 hde1[0] (F) hdk1[3] hdi1[2] hdg1[1]
      480238656 blocks level 5, 32k chunk, algorithm 2 [4/3] [_UUU]
unused devices: <none>

```

I have had a disk failure a couple of times (a RAID failure fixable with `raidhotadd`, not an actual problem with the disk). Issuing the `raidhotadd` command fixed it (first view `/proc/mdstat` to see which disk failed):

```
% raidhotadd /dev/md0 /dev/hde1
```

Bootup and shutdown: It does not appear that any of the RAID commands (`raidstart` or `raidstop`) appear in any of the startup or shutdown scripts, but it all seems to come up and down properly.

RAID commands: `raid0run`, `raidhotgenerateerror`, `raidstart`, `raidhotadd`, `raidhotremove`, `raidstop`

4.2 Servers: Sand

Modifications 3.29.05 by hjm

```

# 200 gb disk from lanina moved to sand on IDE bus 1
# new hoary ubuntu distro loaded on 200gb disk as::

```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/hda1	10482044	2048572	8433472	20%	/
tmpfs	1786416	0	1786416	0%	/dev/shm
/dev/hda6	105294788	5897208	99397580	6%	/data
/dev/hda3	73278252	2128912	71149340	3%	/home

```

/dev          10482044    2048572    8433472    20% /.dev
none         5120        2836        2284    56% /dev

# hjm - 6.7.05 - 200GB IDE disk moved from lanina failed SMART test, replaced
# new spare on same IDE bus; Kubuntu (ubuntu + KDE) installed.
# there is currently no floppy on sand
#/dev/fd0      /media/floppy0  auto    rw,user,noauto  0      0

```

4.3 Formatting Disks

Occasionally disks are added or need to be replaced. There is a lot one can know about formatting disks. Fortunately, a few commands suffice for most situations. The venerable `fdisk` family of commands is important. The recommended formatting command is `cdisk`, a curses-based disk partition table manipulator.

The command to build a Linux filesystem is `mkfs`.

```

cdisk /dev/hdc # Build partition table on /dev/hdc
mkfs -t ext3 /dev/hdc1 # Build an ext3 filesystem on /dev/hdc

```

Once the disk is partitioned (with `cdisk`) and formatted with (`mkfs`), add an entry to the `/etc/fstab`. The entry should be for, e.g., `/dev/hdc1` (the partition) not `/dev/hdc` (the disk). Make sure the mount point exists as a directory, e.g., `/data`, before attempting to mount it.

In May 2007, I formatted a 100 GB replacement disk for `elnino` as `ext3` using the partition name `hdc1`. This seemed to work fine for a few reboots until one day in June the `hdc1` device could not be found. I found the other disk devices were using `sda1`-style terminology presumably because they are using SATA drivers. After changing the `/etc/fstab` entry from `hdc1` to `sdb1` the new drive mounted again. There was no data loss.

4.4 Creating LVM

LVM is Logical Volume Management. Reference: <http://tldp.org/HOWTO/LVM-HOWTO/>

```

fdisk /dev/hde1
fdisk /dev/hdi1
vgscan -v
pvcreate -v /dev/hde1
pvcreate -v /dev/hdi1
vgcreate -v -s 16 dust_vg /dev/hde1 /dev/hdi1
lvcreate -L 305000 -n dust_lv dust_vg # Creates /dev/dust_vg/dust_lv
mke2fs -j /dev/dust_vg/dust_lv
vgscan # Creates /etc/lvmtab and /etc/lvmtab.d/* and /etc/lvmconf/*

```

vg commands:

```

vg_annotate  vgcfgbackup  vgck  vgexport  vgmerge  vgrename  vgsplit
vg_cachegen  vgcfgrestore  vgcreate  vgextend  vgmknodes  vgrename
vga_reset   vgchange  vgdisplay  vgimport  vgreduce  vgscan

```


Entry in `/etc/fstab`:

```
/dev/dust_vg/dust_lv /data ext3 exec,dev,suid,rw 1 2
```

Bootup and shutdown: It does not appear that any of the LVM commands appear in any of the startup or shutdown scripts. However, `dust` seems to have a perpetual problem of shutting down, because it hangs on the `umount` command (recall that all the computers mount `/home` from `dust`, so it is tricky for this disk not to be busy). Typically upon shutdown, I am forced to just power-off at the point that `dust` gets stuck. Upon boot-up, everything usually comes up fine after the disk is checked.

4.5 LVM on ashes

We installed LVM managing some partitions on `ashes.ess.uci.edu`.

```
lvscan
lvremove /dev/data/lvol1
lvcreate
lvcreate -L 5g /dev/data # Create space for /usr/local
lvcreate -L 5g /dev/data # Create space for /home
lvcreate -L 42g /dev/data # Failed due to not enough space
lvcreate -L 39g /dev/data # Create space for /data
mkdir /mnt/home
/sbin/telinit 1
mkdir /mnt/data
mkdir /mnt/local
cd /mnt
mkreiserfs /dev/data/lvol1
mkreiserfs /dev/data/lvol2
mkreiserfs /dev/data/lvol3
ls /usr/local
mount /dev/data/lvol1 /mnt/home
mv /home/* /mnt/home
ls /mnt/home
ls /home
mount /dev/data/lvol2 /mnt/local
mv /usr/local/* /mnt/local
ls /mnt/local
ls /usr/local
umount home local
vi /etc/fstab
mkdir /data
mount /data
more /etc/fstab
tune2fs -j /dev/hda1 # Turn ext2 partition into ext3
tune2fs -j /dev/hda6 # Turn ext2 partition into ext3
```

```
tune2fs -j /dev/hda7 # Turn ext2 partition into ext3
more /etc/fstab
tune2fs -j /dev/hda5 # Turn ext2 partition into ext3
emacs -nw /etc/fstab # Turn /home, /usr/local, /boot, /data into reiserfs
shutdown -r now
```

4.6 Knoppix

Knoppix is a Linux distribution useful for previewing hardware and for rescuing downed computers. Knoppix offers many boot time options. Specify these at the `boot :` prompt.

```
knoppix lang=us # Change language from default (German) to English
knoppix 2 # Boot to run-level 2
knoppix screen=1600x1200 # Specify screen resolution
```

Knoppix is very useful for restoring corrupted systems. Typically this requires mounting the disks on the filesystem. Knoppix does this automatically by double-clicking on the disk icon (alternatively, one may use the `mount` command). Then the disk must be set to enable-writing. Knoppix makes this available through the mouse menu. Finally it is usually necessary to change from user `knoppix` to user `root` before doing any serious repairs. The command sequence to fix a bad sector on a reiserfs partition on `ashes` was:

```
su root
reiserfsck /dev/hda1
```

Assuming corruption is found and the disk, may want to unmount the disk and attempt to rebuild the filesystem tree

```
umount /dev/hda1
reiserfsck --rebuild-tree /dev/hda1
```

Many ATA-3 and later ATA, IDE and SCSI-3 hard drives include Self-Monitoring, Analysis and Reporting Technology (SMART) utilities. The `smartmontools` package uses these capabilities to monitor and tests for disk problems. The command line program is `smartctl`, and the daemon is `/etc/init.d/smartmontools` which uses the configuration in

```
/etc/init.d/smartmontools start
/usr/sbin/smartctl --all /dev/hda1
smartctl -i /dev/hda # Print disk information
smartctl -H /dev/hda # Print health status
smartctl -l error /dev/hda # Print error log from test
smartctl -t offline /dev/hda # Perform offline test (then check error log)
smartctl -t short /dev/hda # Perform short disk test
smartctl -t long /dev/hda # Perform long disk test
smartctl -l selftest /dev/hda # Check selftest error log
```

The `lm-sensors` package complements the SMART utilities.

```

sudo aptitude install lm-sensors # Install lm-sensors
# Create devices on the local system
cat > ~/mkdev.sh << EOF
#!/bin/bash
# number of devices to create (max: 256)
NUMBER=32
# Device owner and group
OUSER=root
OGROUP=root
# Set device mode
MODE=600
# Do not run script if devfs is used
if [ -r /proc/mounts ] ; then
if grep -q "/dev devfs" /proc/mounts ; then
echo "You do not need to run this script as your system uses devfs."
exit;
fi
fi

i=0;
while [ $i -lt $NUMBER ] ; do
echo /dev/i2c-$i
mknod -m $MODE /dev/i2c-$i c 89 $i || exit
chown "$OUSER:$OGROUP" /dev/i2c-$i || exit
i=$((i + 1))
done
EOF

chmod 755 ~/mkdev.sh # Make mkdev.sh executable
sudo ~/mkdev.sh

```

4.7 Mirroring Disks

Chad Cantwell used hardware RAID 1 to join two Maxtor 150 GB disks on `dust` as `/home` on `dirt.ess.uci.edu`. Mirroring disks to other disks is a smart strategy for creating high-availability storage. The RAID status is in `/proc/mdstat`. `mdstat` contains a percentage done while setting up, and then the final RAID status afterwards. This is configured via `/etc/raidtab`.

4.8 Garbage collection

Through time most disks accumulate unwanted files such as `core` dumps, intermediate files, and obsolete files. Unless care is taken, a substantial portion of free disk space may be used by unwanted files. The following short script uses `find` to locate and sort in order of descending size the largest 100 files within the `$(drc)` hierarchy.

```
drc=${DATA}
drc=${HOME}
find ${drc} -type f -exec ls -l {} \; | sort -r -n -k 5 | head -100 > \
/tmp/bgf_foo.txt
```

The results are stored in `/tmp/bgf_foo.txt`. Find files that contain `foo` in their names:

```
find . -name '*html'
/tmp/bgf_foo.txt
```

4.9 Disk corruption

Disks die occasionally and it is prudent to gain familiarity with disk recovery strategies (although these are no substitute for making regular off-site backups!). Disk blocks that are unwritable are known as *bad blocks*. When disks develop bad blocks they may still function for quite a while, but restarting the system may become tedious as manual `fsck`'s become mandatory. Specific disk repair commands are filesystem-dependent. The default Linux filesystem is *ext2*, which is supported by a package of utilities known as `e2fsprogs`. The `dumpe2fs` command displays useful information about a device, such as the block size, chunk size, etc. The `e2fsck` command is used to monitor and repair *ext2* filesystems.

```
dumpe2fs -b /dev/hda7 # Print bad block list
umount /dev/hda7
e2fsck -c /dev/hda7 # Write bad blocks to list
e2fsck -p /dev/hda7 # Automatically preen filesystem without prompting
```

Ext2 filesystems support the notion of a bad block list, i.e., a list of blocks it will not attempt to write to. The recommended procedure to create such a list is to first identify the device, e.g., `/dev/hda7`. Then logout of any windows accessing that device and unmount it. This may be tricky if the partition is used by the `umount` itself (or a dynamic library upon which the executable depends, `libc`, for example). If this occurs, one can presumably boot from a rescue floppy and execute `umount` from there. Then run `e2fsck -c` on the device.

A highly recommended rescue floppy known as “Tom’s root boot” is available from <http://www.toms.net/rb/home.html>.

`hjm++ 1.23.06 updatedb` (on sand). ‘`locate`’ is a very useful tool on many **ix* systems, however it can overwhelm you with output, especially if it is used on a system that has filesystems and backup directories. It is useful to `grep` the output to for the lead path or remove backup directories from being included in the `locate` db, by entering them in the `/etc/updatedb.conf`. This has been done on sand to avoid cluttering the output with repeated entries from the backup mirrors. `hjm-`

5 Software

5.1 netCDF Browsers

There are at least four freely available netCDF file browsers: `ncview`, `ncBrowse`, `NCVweb`, and `panoply`.

5.1.1 ncBrowse

[ncBrowse](#) `ncBrowse` is a Java application.

5.1.2 panoply

[panoply](#) Its functionality is somewhat limited in that logarithmic axes are not allowed.

5.1.3 ncview

[ncview](#) Dave Pierson's excellent `ncview` software must be installed by hand. Since it relies on UDUNITS, it is important to build that as well.

```
sudo scp /usr/local/bin/ncview dust.ess.uci.edu:/usr/local/bin
sudo scp -r /usr/local/lib/ncview dust.ess.uci.edu:/usr/local/lib
```

`ncview` allows logarithmic transformation of coordinates.

5.1.4 NCVweb

The [NCVweb](#) package is designed specifically for viewing netCDF files produced by the Atmospheric Radiation Measurement (ARM) Program. NCVweb is for web-based visualization. It uses NCO for aggregating data.

fxm

5.2 Python

The netCDF interface to Python has at least 2 actively maintained netCDF interfaces: <http://met-www.cit.cornell.edu/noon/ncmodule.html> by Bill Noon, and <http://starship.python.net/crew/hinsen/netcdf.html> by Konrad Hinsen.

5.3 Mail

My Linux servers run `sendmail`, or the `postfix` interface to `sendmail`. This program, like other daemons, can be restarted using `killall -HUP sendmail`. To start the program, use the RedHat `control-panel`. The outgoing SMTP server is set in the file `/etc/sendmail.cf` or `/etc/mail/sendmail.cf` in the line beginning with `DS`. Mail that has not yet been delivered is stored in the `/var/spool/mqueue` directory. Incoming mailbox (e.g., `/var/spool/mail/zender`) must be known to Emacs.

There is a good example of how to set up `postfix` for Ubuntu machines at <http://www.ubuntulinux.org/wiki/DialupEmailHowto>. This example works for `ashes.ess.uci.edu`.

All too often, mail does not get delivered. Thus it is very important to keep a file copy of all messages before entrusting them to the Mail Transfer Agent, (MTA). At the very least, CC yourself on all mail you send. If it does not reach you, it did not reach its intended recipient. Undelivered mail is stored in `/var/spool/mqueue-client`. Undelivered messages have non-zero sizes.

Occasionally it is necessary to move mail between machines. Mail folders which are in standard format can be simply joined together using the `cat` command. As of September 2002, the correct UCI POP server to use for incoming mail is `pop.uci.edu`. The SMTP server to use for outgoing mail depends on the Internet address of the client. In general, setting SMTP server to `localhost.localdomain` works fine. This requires correctly configuring a Mail Transport Agent (MTA), such as Postfix, first. When configuring the SMTP client (e.g., `thunderbird`) delivery protocol, a good option to select is (TLS), in particular, the "TLS, if available" option.

Clients directly connected to UCI may use `smtp.uci.edu`. See <http://www.nacs.uci.edu/computing/e4e.html> for additional UCI details. Clients at remote locations have two basic options for for sending mail from the remote location, port forwarding or changing mail servers. Port forwarding works from any remote location. Being portable, it is recommended for all short-term trips where there is no advantage to being on a local mail server. The goal of port-forwarding is to redirect traffic on the local port to the hostport on the remote host. This forwarding is done via encrypted tunnel so the `remote_client` must have trusted access to the remote server. This command forwards 2025 traffic on `localhost` to port 25 on the SMTP server. The tunneling is done on a trusted remoted client.

```
ssh -L localport:remote_server:server_port remote_client
ssh -L 2025:smtp.uci.edu:25 sand.ess.uci.edu
# Access PBS cluster firewalled httpd
# First ssh to pbs, then open local browser to http://localhost:2026
ssh -L 2026:pbs.ess.uci.edu:80 pbs.ess.uci.edu
```

The tunnel must remain open for this port-forwarding to work (do not close the window).

The second option is to login to an authorized SMTP server. All non-UCI SMTP servers must use port 587. Clients directly connected to `wsu.edu` should use `mail.wsu.edu` (WSU blocks mail routed through `localhost.localdomain` SMTP servers). Clients directly connected to `Cox.net` should use `smtp.west.cox.net`. Clients directly connected to `cgd.ucar.edu`, or to `wireless.ucar.edu` should use `mailhub.cgd.ucar.edu`. Clients directly connected to `greenspeedisp.net` should use `mail.greenspeedisp.net`.

Setting the SMTP server in Mozilla is non-intuitive. Select the Outgoing Server (SMTP) setting at the *bottom* of the accounts menu presented by Edit | Mail and Newsgroups Account Settings.

Emacs `rmail` mode provides an excellent environment for editing and sending mail. `rmail` understands your `.mailrc` file and, moreover, works with all features of Emacs. Remember to change the Emacs SMTP server when roaming.

```
(setq smtpmail-default-smtp-server "smtp.uci.edu")
```

Sometimes it is useful to combine or juggle `rmail` and Mozilla mail folders:

```
cd ${HOME}
tar cvzf ${DATA}/tmp/mail.tar.gz ./mail
scp ${DATA}/tmp/mail.tar.gz dust.ess.uci.edu:
tar xvzf mail.tar.gz
rm ${HOME}/mail/*
```

```
# Synchronize mail backup repository (sand) to main repository (virga)
rsync /home/zender/.mozilla-thunderbird/ktwn0dgv.default/Mail/pop.uci.edu sa
rsync /home/zender/.mozilla-thunderbird/ktwn0dgv.default/Mail/Local\ Folder

# Synchronize current client (ashes,elnino,sand) with main repository (virga)
pth_src='virga.ess.uci.edu:/home/zender/.mozilla-thunderbird/ktwn0dgv.default
pth_src='sand.ess.uci.edu:/home/zender/.mozilla-thunderbird/1h6xplh6.default
pth_dst=${HOME}/.mozilla-thunderbird/prznnk5x.zender # elnino
pth_dst=${HOME}/.mozilla-thunderbird/lsymr6b9.default # ashes
pth_dst=${HOME}/.mozilla-thunderbird/1h6xplh6.default # sand
rsync ${pth_src}/Mail/pop.uci.edu ${pth_dst}/Mail
rsync ${pth_src}'/Mail/Local\ Folders' ${pth_dst}/Mail
```

My user preferences file, `prefs.js`, was corrupted on 20060222. Upgrading thunderbird while it is running is usually safe. However, this major upgrade (to version 1.5) appeared to corrupt it. See <http://www.mozilla.org/unix/customizing.html#prefs> Create a new profile with `mozilla-thunderbird -profilemanager`. To see all your firefox settings, view the URL `about:config` in a Firefox browser.

```
act_src='prznnk5x.zender' # elnino
act_dst='ktwn0dgv.default' # virga
pth_src=${HOME}/.mozilla-thunderbird/${act_src}
pth_dst=${HOME}/.mozilla-thunderbird/${act_dst}
/bin/cp -p -f ${pth_src}/*.mab ${pth_dst}
/bin/cp -p -f ${pth_src}/cert8.db ${pth_dst}
/bin/cp -p -f ${pth_src}/key3.db ${pth_dst}
/bin/cp -p -f ${pth_src}/63760765.s ${pth_dst}
/bin/cp -p -f ${pth_src}/mimeTypes.rdf ${pth_dst}
/bin/cp -p -f ${pth_src}/training.dat ${pth_dst}
/bin/cp -p -f ${pth_src}/virtualFolders.dat ${pth_dst}
/bin/cp -p -f ${pth_src}/Mail/pop.uci.edu/* ${pth_dst}/Mail/pop.uci.edu
/bin/cp -p -f ${pth_src}/Mail/Local\ Folders/* ${pth_dst}/Mail/Local\ Folders
```

5.4 Macintosh OS X

OS X is based on FreeBSD. The OS X equivalent to `ldd` is `otool`

```
otool -L `which ncks`
```

5.4.1 Fink

Fink is a Debian-based system for installing packages on OS X. Fink comes as a “disk image” file with a `.dmg` suffix. Clicking with the pointer on a disk image file causes a sequence of actions to occur: the file is mounted, the contents appear as icons, etc. The same effect may be achieved from the command line with the `hdiutil` and `installer` commands.

```

man hdiutil
man installer
# OS X VNC URL: http://netmath.math.uiuc.edu/VNC-osx.htm, http://www.realvnc.com
hdiutil attach OSXvnc-0.6.dmg # Mount disk image file
cd '/Volumes/OSXvnc 0.6' # Look at mounted image
cp -r OSXvnc.app /Applications # Copy package to software directory
cd '/Volumes/OSXvnc 0.6/OSXvnc.app/Contents/MacOS' # Directory with executable
./OSXvnc & # Run VNC
# Window pops up on Macintosh, prompts to start VNC server
# Recommended password is grape
hdiutil unmount '/Volumes/OSXvnc 0.6' # Unmount disk image file
hdiutil unmount /dev/disk2 # Unmount disk image file
# To view from Linux, use
vncviewer c-67-169-127-154.client.comcast.net:1 # IP:display

# Fink URL: http://fink.sf.net
hdiutil attach Fink-0.6.1-Installer.dmg # Mount disk image file
installer -pkginfo -pkg '/Volumes/Fink 0.6.1 Installer/Fink 0.6.1 \
Installer.pkg' -target /sw # Needs a .pkg package # Query package info
sudo installer -pkg '/Volumes/Fink 0.6.1 Installer/Fink 0.6.1 Installer.pkg'
-target /sw # Install package, requires sudo password
hdiutil unmount '/Volumes/Fink 0.6.1 Installer' # Unmount disk image file
hdiutil unmount /dev/disk3 # Unmount disk image file

```

6 Programming

6.1 Operator precedence

The operator precedence of most languages follows that of the C language. Table 2 summarizes operator precedence of C. Precedence decrease from top to bottom (the first line has highest precedence).

6.2 Regular expressions

The alphanumeric pattern specifying a group of strings is called a regular expression. Special characters are \$, ^, ., *, +, ?, [,], and \.

- “.” Matches any character except newline
- “*” Match smallest possible preceding regular expression as many times as possible
- “+” Match preceding regular expression at least once
- “?” Match preceding regular expression once or not at all
- “[. . .]” Character set

Table 2: C/C++ Operator Precedence^a

Operator	Description	Associativity
::a	Unary scope resolution	Left to right
a::b	Binary scope resolution	Left to right
a++	Post-increment	Left to right
a--	Post-decrement	Left to right
a()	Function call	Left to right
a[b]	Array element	Left to right
a->b	Pointer to structure member	Left to right
.	Structure or Union member	Left to right
++a	Pre-increment	Right to left
--a	Pre-decrement	Right to left
!a	Logical NOT	Right to left
~a	Bitwise NOT	Right to left
-a	Unary minus	Right to left
+a	Unary plus	Right to left
&a	Address	Right to left
*a	Indirection	Right to left
sizeof(a)	Size in bytes	Right to left
new	Allocate memory	Left to right
delete	De-allocate memory	Left to right
(type)	Typecast & all C++ cast operators	Left to right
.*	Pointer to member (objects)	Left to right
->*	Pointer to member (pointers)	Left to right
*	Multiply	Left to right
/	Divide	Left to right
%	Remainder	Left to right
+	Add	Left to right
-	Subtract	Left to right
<<	Left shift	Left to right
>>	Right shift	Left to right
<	Less than	Left to right
<=	Less than or equal to	Left to right
>	Greater than	Left to right
>=	Greater than or equal to	Left to right
==	Equal to	Left to right
!=	Not equal to	Left to right
&	Bitwise AND	Left to right
^	Bitwise XOR	Left to right
	Bitwise OR	Left to right
&&	Logical AND	Left to right
	Logical OR	Left to right
?:	Conditional	Right to left
=	Assignment	Right to left
+= -= *= /= %= &= ^= = <<= >>=	Compound assignment	Right to left
,	Comma	Left to right

- “[^ . . .]” Complemented character set
- “^” Match beginning of line
- “\$” Match end of line
- “\” Quote special characters
- “\w” Matches any word constituent character
- “\W” Matches any non-word constituent character

In the replace expression, `&` stands for the match found for the whole regular expression, and “\N” stands for the match to the Nth occurrence of the “\ (. . . \)” pair. A floating point number match is [*Friedl, 1997*, p. 128]

```
-? ([0-9]+ (\.[0-9]*)? |\.[0-9]+)
```

The “?” makes the negative sign optional. The regular expressions used for lexically recognizing doubles in `ncap` and in `ncgen`, respectively, are

```
[0-9]*\.[0-9]*([eE][+-]?[0-9]+)?[LlDd]?|[0-9]*([eE][+-]?[0-9]+)[LlDd]?
[+-]?[0-9]*\.[0-9]*{exp}?[LlDd]?|[+-]?[0-9]*{exp}[LlDd]?
```

Using regular expressions in Emacs is both a pleasure and a problem. Some example regular expression search and replaces that work, and what they do

```
; Replace "!foo" by "! foo"
; (query-replace-regexp "!\\(\\w\\)" "! \\1" nil nil nil)
; Replace "character foo*80" by "character(80)::foo"
; (query-replace-regexp "character \\(\\.\\*\\)\\*\\([0-9]*\\)" "character(\\2)::\\1" nil)
; Replace "end do      !" by "end do !"
; (query-replace-regexp "end \\(do\\|if\\) *!" "end \\1 !" nil)
; Replace "foo      !" by "foo !" as long as "!" is not followed by "="
; (query-replace-regexp "\\([!^!\\t\\n\\) + +!" "\\1 !" nil)
; Replace floating point number X.Y with floating point number X.Yf, e.g.,
; (query-replace-regexp "\\([+-]?[0-9]*\\. [0-9]*[LlDd]?\\)" "\\1f" nil)
; Replace floating point number X.Y with floating point number $X.Y$, e.g.,
; (query-replace-regexp "\\([+-]?[0-9]*\\. [0-9]*[LlDd]?\\)" "$\\1$" nil)
; Replace "command{\\foo}{bar}" by "command{\\foo}{\\ensuremath{bar}}"
; (query-replace-regexp "command{\\\\.\\*}{\\.\\*}" "command{\\1}{\\ensuremath{\\2}}")
```

One often composes text in an Emacs buffer, then copies and pastes that buffer into another program, e.g., a browser or mail program. Browser and mail windows often have column-width limits imposed, so it is helpful to compose with the same column-width in Emacs. The command to set the maximum buffer width is `set-fill-column`. The argument to this command is the number of columns. Entering the argument to Emacs commands is non-trivial. The example of setting the number of columns to 80 show this: ESC 80 ESC x set-fill-column RET. This command might be written in Emacs short-hand as

```
M-x set-fill-column
```

Bash supports a wide variety of pattern operators for shell filename expansion, aka globbing. These globbing operations filter all files and directories present through a filter including one or more wildcard characters.

```
for fl in `ls dst25_8589_??.*nc` ; do
  mv $fl ${fl/8589_/clm} # Change '8589_' to 'clm' in filenames
done
for fl in `ls *khus*.jpg` ; do
  mv $fl ${fl/khus/kuhs} # Change 'khus' to 'kuhs' in filenames
done
for fl in `ls *Johnson,_Jack*` ; do
  mv $fl ${fl/Johnson,_Jack/Jack_Johnson}
done
# On machine with temporary mail, e.g., ashes:
cd ~/mail
for fl in `bin/ls` ; do mv $fl ${fl}.foo ; done
# On machine with archived mail, e.g., elnino:
cd ~/mail
scp 'ashes.ess.uci.edu:mail/*.foo' ~/mail
for fl in `bin/ls *.foo` ; do
  cat ${fl}/.foo/ ${fl} > ${fl}/.foo/.new
  mv -f ${fl}/.foo/.new $fl
done
```

I also have a custom Perl script, `fl_rnm.pl`, for renaming files.

The most common shell pattern matching operators, such as `*` and `?`, have minimal functionality as regular expressions. The Bash shell supports extended regular expressions in filename globbing via the shell option `extglob`. Activate this functionality using the `shopt` command:

```
shopt -s extglob
```

6.3 Bash

Use `hash` to rehash commands in Bash, e.g., `hash -r`. This is equivalent to `rehash` or `reset` in `cs`. Systemwide defaults are set in `/etc/ssh/ssh_config`.

The Bash shell supports a wide variety of built-in commands, command line editing, job control, and history features.

```
!-1:s/foo/bar/ # Repeat previous command, substitute bar for foo
!!:s/foo/bar/  # Repeat previous command, substitute bar for foo
^foo^bar^     # Repeat previous command, substitute bar for foo
!foo:s/foo/bar/ # Substitute bar for foo in most recent command containing
```

6.4 Internationalization I18N

Create a local directory structure to hold *.mo (“machine object”) files created from *.po (“portable object”) by `gettext` machinery. One directory is needed for each language.

```
mkdir -p ${HOME}/share/locale/es/LC_MESSAGES
```

System-wide translation database is under `/usr/share/locale`. The installation of `gettext()` is its own documentation. Examine `/usr/share/gettext/intl` to see how its done.

Emacs `po-mode.el` should be loaded whenever *.po files are loaded.

7 Files

7.1 Date conventions

This section describes conventions for naming files from geophysical models. We adopt the usual convention that `DD` is a two digit sequence to indicate the day of month $DD \in [01, 02, \dots, 31]$, `MM` is a two digit sequence to indicate the month of year $MM \in [01, 02, \dots, 12]$, and `YYYY` is the four digit Common Era year. It is often useful to select files based on their date convention. In such cases it is useful to have regular expressions (cf. §6.2) for each date convention: Date components like `MM` and `YYYY` are not just keyboard inputs, they are also variables, since useful information may be derived from them. For example, the number of years N in a file containing data from years `YYYY` and `ZZZZ` is $N = ZZZZ - YYYY + 1$.

Valid replacements for `DD` are `[0123][0-9]`. Valid replacements for `MM` are `[01][0-9]`. Valid replacements for `YYYY` are `[0-9][0-9][0-9][0-9]`.

Climatological values are abbreviated `clm`, which is best interpreted as “all of the available data present when the command was run”. Syntactically, `clm` works equivalently to `YYYY`. However, `clm` is only three characters, while `YYYY` is four characters. This distinction helps reduce errors when commands use globbing to do the right thing, e.g.,

```
ncra caseid_????01.nc caseid_clm01.nc # OK globbing
ncra caseid_[0-9][0-9][0-9][0-9]01.nc caseid_clm01.nc # Better
ncra caseid_[0-9]{4}01.nc caseid_clm01.nc # Preferred
```

1. `${caseid}_YYYY` Annual mean
2. `${caseid}_YYYYMM` Monthly mean
3. `${caseid}_YYYYMMDD` Daily mean
4. `${caseid}_YYYYMMDDHH` Hourly mean
5. `${caseid}_YYYYMMDDHHMM` Minute mean
6. `${caseid}_YYYYMMDDHHMMSS` Second mean
7. `${caseid}_clm` Climatological mean
8. `${caseid}_clmMM` Climatological monthly mean
9. `${caseid}_YYYYZZZZ_MM` Ensemble mean of month *MM* from years *YYYY* through *ZZZZ* (one record)
10. `${caseid}_YYYYZZZZ_0112` Ensemble mean of seasonal cycle sampled from years *YYYY* through *ZZZZ* (twelve records)
11. `${caseid}_YYYY_ZZZZ_MM` Timeseries of month *MM* from years *YYYY* through *ZZZZ* ($ZZZZ - YYYY + 1$ records)
12. `${caseid}_YYYY_ZZZZ_MMNN` Timeseries of month *MM* through month *NN* means from years *YYYY* through *ZZZZ* ($ZZZZ - YYYY + 1$ records)

13. $\{\text{caseid}\}_{\text{YYYY_ZZZZ_0305}}$ Timeseries of Springtime means from years *YYYY* through *ZZZZ* ($\text{ZZZZ} - \text{YYYY} + 1$ records)
14. $\{\text{caseid}\}_{\text{clm_0112}}$ Twelve month seasonal cycle
15. $\{\text{caseid}\}_{\text{YYYY_0112}}$ Annual mean seasonal cycle (12 records). Note this is an exception to the general rule. It is unambiguous, however, because the mean of all twelve months is simple representable by the $\{\text{caseid}\}_{\text{YYYY}}$ convention
16. $\{\text{caseid}\}_{\text{YYYY_MMNN}}$ Mean of months *MM* through *NN* (1 record)
17. $\{\text{caseid}\}_{\text{YYYY_0305}}$ Springtime mean (1 record)
18. $\{\text{caseid}\}_{\text{YYYY_ZZZZ}}$ Annual mean timeseries (multiple records)
19. $\{\text{caseid}\}_{\text{YYYYMM_ZZZZNN}}$ Monthly mean timeseries from *YYYYMM* to *ZZZZNN*, inclusive (multiple records)
20. $\{\text{caseid}\}_{\text{YYYYMMDD_ZZZZNNEE}}$ Daily mean timeseries from *YYYYMMDD* to *ZZZZNNEE*, inclusive (multiple records)
21. $\{\text{caseid}\}_{\text{YYYY_ZZZZ_t}}$ Mean of annual mean timeseries from *YYYY* to *ZZZZ*, inclusive (single record)
22. $\{\text{caseid}\}_{\text{YYYYMM_ZZZZNN_t}}$ Mean of monthly mean timeseries from *YYYYMM* to *ZZZZNN*, inclusive (single record)
23. $\{\text{caseid}\}_{\text{YYYYMMDD_ZZZZNNEE_t}}$ Mean of daily mean timeseries from *YYYYMMDD* to *ZZZZNNEE*, inclusive (single record)

8 System Administration

8.1 Backups

Simple but effective system backups are made each night. The Cron system performs specified actions at regular intervals. The `cron` program checks for actions to perform once per minute. Currently, all backups are handled by the `crontab` file of the root user. Use the `crontab` program to alter any user's `crontab`.

```
crontab -e
sudo crontab -u root -e
```

The `crontab` program stores each user's `crontab` in `/var/spool/cron`. For the root user, this is `/var/spool/cron/root` and `/var/spool/cron/crontabs/root` for RedHat and Debian GNU/Linux, respectively. Currently this system backup `crontab` file is

```
15 2 * * * /usr/local/bin/bck_home.sh zender
45 2 * * * /usr/local/bin/bck_home.sh bian
15 3 * * * /usr/local/bin/bck_home.sh alfgr
45 3 * * * /usr/local/bin/bck_home.sh mflanner
15 4 * * * /usr/local/bin/bck_home.sh kwon
45 4 * * * /usr/local/bin/bck_home.sh strombrg
15 5 * * * /usr/local/bin/bck_home.sh lopez
45 5 * * * /usr/local/bin/bck_home.sh jtalaman
15 6 * * * /usr/local/bin/bck_home.sh okin
45 6 * * * /usr/local/bin/bck_home.sh pajarola
30 5 * * * /usr/local/bin/bck_data.sh
35 5 * * * /usr/local/bin/bck_etc.sh
40 5 * * * /usr/local/bin/bck_var.sh
37 3 * * * /usr/bin/rdate -s time.nist.gov
38 3 * * * /sbin/hwclock --systohc
```

The first column is the minute of each hour for an action. The second column is the hour of each day for an action. The third, fourth, and fifth columns are the day of month, month of year, and year for the corresponding actions. Asterisks denote all values for the corresponding field. Always backup data to at least two physically separate locations!

My personal, user-specific `crontab` file is

```
30 2 * * * /home/zender/sh/bck_zender.sh dly
```

Note that the backup commands themselves are only valid when specified with fully qualified paths. This security feature of Cron helps prevent malicious files from being inadvertently executed.

Backup binaries from `/usr/local/bin` (e.g., `ddd`, `netscape`). The script `bck.pl` is devoted to this, and works for ZIP disks and LS120 drives too. “Dot” files (e.g., `.cshrc`, `.mailrc`, `.netscape` directory) should be relatively safe as they are frequently archived by CVS. The most important files to backup thus become files in the mail directories.

8.2 Mounting devices

Make sure CDROM is in `fstab`, mounting it as `/dev/cdrom` is OK, mounting it as `/dev/hdc` is also common.

```
ln -s /mnt/cdrom /cdrom
ln -s /mnt/floppy /flp
```

Occasionally, such as when rescuing lost systems, it is helpful to mount disk partitions from a command line shell, e.g.,

```
mount /dev/hda7 /home
mount /dev/hda8 /data
```

9 rsync

`rsync` is *the* program to synchronize non-archived files among remote machines. Files and directories which are not controlled by a source code control system are notoriously hard to keep synchronized. Examples are my directory of PDF journal papers, \LaTeX class files, and web directories. The general syntax of `rsync` is `rsync src dst`. A common mistake is to use the same path depth in `src` and `dst` arguments when normally the `src` path should be one level deeper than the `dst` path,

```
rsync source:/drcl/drc2/drc3 destination:/drcl/drc2
rsync source:/drcl/drc2/drc3/ destination:/drcl/drc2/drc3
```

The previous two commands are equivalent and the first form is preferred.

Specifying `-av` is usually recommended for recursive synchronization without modifying file attributes. Hence the typical commands to keep machines in sync are

```
# Rsync synchronization methods
# -a (equivalent to -rlptgoD): Copy recursively, preserve file modes & times
# -v: Verbose
rsync -av /data/zender/cccac dust.ess.uci.edu:/data/zender
rsync -av /data/zender/ppr dust.ess.uci.edu:/data/zender
```

Often backups are made to CompactFlash or Memory Stick devices, described in Sections [2.2.1](#) and [2.2.2](#).

```
# Backup home directory to compact flash
/bin/cp -r /home/zender /cf # Backup for dummies
# Rsync backup methods
# -a (equivalent to -rlptgoD): Copy recursively, preserve file modes & times
# -v: Verbose
rsync -av /home/zender /cf
# Exclude backup files from being copied
rsync -av --exclude '*~' /home/zender /cf
```


There are different types of back-ups. Often a primary source location contains all important information, and is the first repository to remove deprecated files. Back-ups of this primary repository, i.e., secondary repositories, should delete these deprecated files.

```
# Delete files on receiving side that are not on sending side
rsync -av --delete-excluded /data/zender/snd dust.ess.uci.edu:/data/zender :
rsync -av --delete-excluded dust.ess.uci.edu:/data/zender/snd /data/zender :
rsync -av --delete-excluded --cvs-exclude --exclude '*~' /home/zender /cf #
# Backup to memory stick
rsync -av --delete-excluded --cvs-exclude --exclude='*~' /home/zender /memst
rsync -av --delete-excluded --cvs-exclude \
--exclude='*~' \
--exclude='*.o' \
--exclude='*.so' \
--exclude='*.a' \
--exclude='*.mod' \
--exclude='.[a-zA-Z]*' \
/home/zender /media/USB20FD/ # Backup
```

9.1 Slink

NCAR CGD uses a tool called `slink` which makes numerous separate software installations appear as one whole tree.

```
cd /data/zender/gsl-0.7
./configure --prefix=/contrib/gsl-0.7
make
make install
cd /contrib
newgrp contrib
co -l slink.conf
ci slink.conf
./slink
exit
```

9.2 Documentation

GNU/Linux documentation is a little scattered. RedHat and Debian GNU/Linux install program-specific documentation in `/usr/share/doc`.

9.3 PPP

Modem is `/dev/ttyS0`, which should be linked to `/dev/modem`. `/dev/modem` must be usable by all. This should be set with control-panel. Must get correct permissions and ownerships on various files and directories: Line speeds, etc. have not changed from RH5 defaults Permissions are a major security issue!

```
cd ${HOME}/linux/usr/local/bin
sudo cp ppp-go ppp-off ppp-on-dialer /usr/local/bin # Or /usr/local/sbin
cd ${HOME}/linux/etc/ppp
sudo cp options chat* *secrets /etc/ppp
```

Some security configurations may require that `pppd` be run with the `setuid=root` bit set, and some distributions ship `pppd` with `mode=644`. In this case, use, e.g., `chmod +s /usr/sbin/pppd`.

9.4 Batch Queues

Two batch queue handlers are in use at NCAR and UCI, the NQS Network Queueing System and the AIX LoadLeveler system. LoadLeveler is described at <http://www.scd.ucar.edu/docs/blackforest/batch.html>. The commands to submit jobs in these systems are `qsub` and `llsubmit`, respectively. The commands to query jobs in these systems are `qstat` and `llq`, respectively. The commands to cancel jobs in these systems are `qdel -k` and `llcancel`, respectively.

```
qstat -a -h ute # Check batch queues on ute from utefe
qdel -h ute 12978.ute # Delete jobs on ute from utefe
```

9.5 Remote shell service rcp, rsh, ssh, telnet...

Make sure `.rhosts` is installed and NOT group/other writable! Turn on `sshd`, `telnetd`, `httpd` in, at least, run levels 3–5. Turn on services in runlevel 2 if they should be running even when the X-server is not.

9.6 Root

Change shell to `tcsh` using control-panel

```
chsh -s /bin/tcsh root
```

Install abbreviations to root's `.cshrc` file, e.g.,

```
alias m 'less'
alias h 'history'
alias csrc 'source ~/.tcshrc'
alias cd 'cd \!*; set prompt=${cwd}" ROOT"#" "'
alias dir 'ls -lga'
setenv PATH "/usr/sbin:/sbin:/bin:/usr/bin:/usr/local/bin:$PATH"
```

9.7 L^AT_EX

All L^AT_EX notes are contained in `ltx.tex` and `ltx.ps`.

9.8 Library

Often unresolved external links are reported by a loader and the required library must be located. A useful procedure to follow is to change to the library directory and use the `nm` command to search for the missing subroutine

```
sbr=foo
cd /usr/lib
for lib in lib*.a ; do
#   printf "Searching ${lib} for ${sbr}...\n"
    nm ${lib} | grep ${sbr}
done # end loop over lib
```

9.9 Info

After installing packages in `/usr/local` which install info in `/usr/local/info`, you need to update `/usr/local/info/dir`

9.10 Networking

`nslookup` returns information about the machine on the Internet with a given name or IP address. Use control-panel to set network parameters such as hostname in `/etc/hosts`, e.g.,

```
128.117.91.216 lanina.ppp.ucar.edu lanina
```

and the nameserver in `/etc/resolv.conf`:

```
nameserver 128.117.24.2 bearmtn-e0
nameserver 128.117.24.3 greenmtn-e0
nameserver 128.117.64.22 NCAR PPP router
nameserver 128.200.1.201 csi.ns.nts.uci.edu
nameserver 128.200.192.202 cpl2.ns.nts.uci.edu
```

9.11 C development

`kernel-headers` and `glibc-devel` are required to develop C language programs, but not installed by default

```
rpm -ivh /mnt/cdrom/RedHat/RPMS/glibc-devel-2.0.5c-10.i386.rpm
rpm -ivh /mnt/cdrom/RedHat/RPMS/kernel-headers-2.0.32-2.i386.rpm
```

9.12 Time and TimeZone

Use `--utc` option to indicate that HW clock is kept in coordinated universal time or UTC, which is virtually identical to Greenwich mean time or GMT. This option is set during the installation of GNU/Linux.

```
/sbin/hwclock --debug
/sbin/hwclock --test --set --date="08/13/98 04:38:00" # Tests the command
/sbin/hwclock --set --date="04/09/00 15:41:00"
```

A better option seems to be using `rdate` to set the system clock and then `hwclock` to set the hardware clock to the system clock.

```
sudo /usr/sbin/rdate -s time.nist.gov
sudo /usr/sbin/rdate -s ntp.ucsd.edu
sudo /sbin/hwclock --systohc
```

Note that `time.nist.gov` may refuse connections, so consider alternatives like `ntp.ucsd.edu`. Inserting this command in `/etc/rc.d/rc.local` ensures time is set correctly on each reboot (assuming machine is on network at boot time). I think this resets the hardware clock, and not just the system time.

The preferred solution is to use NTP, the network time protocol. The NTP homepage is <http://www.eecis.udel.edu/~ntp>, and the NTP FAQ is <http://www.eecis.udel.edu/~ntp/ntpfaq/NTP-a-faq.htm>. As of about 1999, few Stratum 1 timeservers will accept synchronization requests from Stratum 3 machines. Stratum 3 machines should synchronize with any publically accessible Stratum 2 server, listed, for example, at <http://www.eecis.udel.edu/~mills/ntp/clock2.htm>. I chose server 132.239.254.49 = `ntp.ucsd.edu`. Alternatively, the `ntpdate` command works just like `rdate`, and can be used when the NTP daemon itself is unavailable.

Changing the system timezone is described at <http://www.wikihow.com/Change-the-Timezone->

```
mv /etc/localtime /etc/localtime-old # Backup old TZ file
ln -s /usr/share/zoneinfo/Europe/Amsterdam /etc/localtime # Set new TZ
/usr/bin/rdate -s time.nist.gov # Update current system time
# Set ZONE entry (e.g. "America/Los_Angeles") in /etc/sysconfig/clock
/sbin/hwclock --systohc # Set hardware clock
```

9.13 Perl

```
ln -s /usr/bin/perl /usr/local/bin/perl
```

9.14 Links

Whenever `/home` is re-installed, many links must be recreated. Links may be listed with

```
ls --recursive --classify | grep "@"
```

For machines at NCAR, it is best to store source code in `/fs/cgd/home0/zender` and create links to `${HOME}` so that the larger object files and executables will not consume the expensive space on the central fileserver.

```
cd ${HOME}
mkdir -p lib/${PVM_ARCH}
```

```

mkdir -p obj/${PVM_ARCH}
mkdir -p bin/${PVM_ARCH}
mkdir include
ln -s /fs/cgd/home0/zender/sh bin/sh
ln -s /fs/cgd/data0/zender/match match
ln -s /fs/cgd/data0/zender/match_dst/dst aer
ln -s /fs/cgd/home0/zender/aca aca
ln -s /fs/cgd/home0/zender/c c
ln -s /fs/cgd/home0/zender/c++ c++
ln -s /fs/cgd/home0/zender/ck ck
ln -s /fs/cgd/home0/zender/dot dot
ln -s /fs/cgd/home0/zender/f f
ln -s /fs/cgd/home0/zender/fsf fsf
ln -s /fs/cgd/home0/zender/idx_rfr idx_rfr
ln -s /fs/cgd/home0/zender/mk mk
ln -s /fs/cgd/home0/zender/nco nco
ln -s /fs/cgd/home0/zender/perl perl

```

9.15 Install /etc/initscript to boost stacksize so CCM can run

9.16 Math Libraries

Compile `libspecfun.a`. Normally, the double precision version of `libspecfun.a`, located in `src.dp` should be built and used. However, on Alpha chips using `f90`, this results in `gamma_` being undefined. In this case, build the single precision source (and hence function names) using double precision flags:

```

cd specfun/src.sp
setenv FFLAGS "-r8 -i4"
make -e -f Makefile.unix
cp ../libspecfun.a $MY_LIB_DIR

```

9.17 Hardware description of Zender group computers:

Dell Support

1-800-624-9896 (general)

1-800-234-1490 x69080 (general)

<http://premiersupport.dell.com>

For all machines: Netmask = 255.255.255.0

Default gateway (subnet 14) = 128.200.14.1

Default gateway (subnet 24) = 128.200.24.1

Broadcast (subnet 14) = 128.200.14.255

Primary nameserver (UCI) = 128.200.1.201

Secondary nameserver (UCI) = 128.200.192.202

UCI News server = news.service.uci.edu

UCI POP server (newer, should work) = pop.uci.edu
UCI IMAP server = imap.uci.edu
UCI SMTP server = smtp.uci.edu

ashes.ess.uci.edu = 128.200.14.90

Dell Inspiron 8500

Arrived 20030408

Service Tag C2PNM21

Express Service Code 26285412457

UCI Property #: 039003448

Pentium IV 2.4 GHz 512 KB L2 Cache

Wireless card is TrueMobile1400

MAC address: 00:90:4B:B2:09:86 (wireless internal chip TrueMobile1400)

MAC address: 00:0b:db:17:83:0d (wired)

aptitude install bcmwl5driverloader

Broadcom wireless card: http://www.linuxant.com/drivers_bcmwl/bcmwl5/downloads-license.php

Installed driver Broadcom 01/09/2003, 3.10.39.0

ndiswrapper driver for Broadcom wireless chip installed 20061216 as a per <http://ubuntuforums.org/showthread.php>

/cdrom/Setup/I8500/bcmwl5.inf /cdrom/Setup/bcmwl5.sys

License Key: 87-A7-39-A4-18-C0

email address: zender@uci.edu

Registered as eth1

hda: FUJITSU MHS2060AT, ATA DISK drive 60 GB

hdc: HL-DT-STCD-RW/DVD-ROM GCC-4240N, ATAPI CD/DVD-ROM drive

hdc: ATAPI 24X DVD-ROM CD-R/RW drive, 2048kB Cache, UDMA(33) Intel 810 + AC97 Audio, version 0.24, 04:35:38 May 6 2003

Broadcom 4401 Ethernet Driver bcm4400 ver. 2.0.0 (03/25/03)

eth0: Broadcom BCM4401 100Base-T found at mem faffe000, IRQ 11, node addr 000bdb17830d

0: nvidia: loading NVIDIA Linux x86 nvidia.o Kernel Module 1.0-4349 Thu Mar 27 19:00:02 PST 2003

AC97 modem device forced to iobase_0=0xb400, iobase_1=0xb080, irq=11

http://www.linuxvoodoo.com/store/index.php/cPath/45_66 Bought netgear card from Fry's on 20041211

for \$45-\$25=\$20 NetGear WG511 802.11B/G D-Link Wireless Cardbus NIC 802.11 g 108Mbps

MAC address: 00:09:5B:E8:C4:E1 (NetGear WG511 802.11B/G)

20050818: Bad internal disk, problem report filed with Dell <https://wiki.ubuntu.com/HardwareSupportMachinesL>

suggests pci=noacpi,acpi=noirq

biogenic.ess.uci.edu = 128.200.14.73

Dell Precision 530

Arrived Jan 10, 2002

Serial number: HJKZ411

Service Tag HJKZ411

Express Service Code 38189387557

UCI Property #: 019003703

Intel(R) Xeon(TM) CPU 1.50GHz

hp5850.ess.uci.edu = 128.200.15.157:

Hewlett Packard 5850 Color inkjet printer in Croul 1101

silt.ess.uci.edu = 128.200.14.156:

clay.ess.uci.edu = 128.200.14.158:

UCI property number (for both as Los Alamos Cluster:) 059000194

Los Alamos Computers (LAC) order #014214 Both are dual opterons on a Tyan Tyan S2885ANRF motherboard (onboard firewire) Dual AMD Opteron 246 (2.0 GHz, 1M L2 cache) 2G PC3200 registered ECC DDR RAM (1 has 2GB of PC2100, one has 4GB of PC3200) nVidia Corporation NV34 [GeForce FX 5200] 128MB RAM w/ video w/DVI, 3x 250GB WD SD series SATA disks WDC WD2500SD-01K hdc: SONY DVD RW DW-D26A, ATAPI CD/DVD-ROM drive 3.5 inch floppy drive Onboard gigabit NIC Multimedia audio controller: Advanced Micro Devices [AMD] AMD-8111 AC97 Audio (rev 03) Four port USB 2.0 PCI card Logitech Z560 400W speakers (4 satellites + sub) Chenbro SR10403 enclosure (3 case fans) Enermax 460W power supply (quiet and dependable) disks are mostly in raid5 config: Filesystem 1K-blocks Used Available Use/dev/md1 15496084 1864536 12844376 13tmpfs 1028532 0 1028532 0/dev/md0 100954 14318 81424 15/dev/md2 462259168 34468 438743308 1none 5120 2708 2412 53where the md devices are setup like this: DEVICE partitions ARRAY /dev/md2 level=raid5 num-devices=3 UUID=b190d39b:cad75d67:7abb3ee1:8c71f882 devices=/dev/sda8,/dev/sdb8,/dev/sdc8 ARRAY /dev/md1 level=raid5 num-devices=3 UUID=ba4fb7ff:85a95d9c:988a6647:1d9d8f8c devices=/dev/sda6,/dev/sdb6 ARRAY /dev/md0 level=raid1 num-devices=2 UUID=c8d86633:d20e14e3:f9025448:de67a792 devices=/dev/sda5,/dev/sdb5 Monitors are: Viewsonic VP201b 20" LCD: Model number VLCDS26064-2W Serial numbers: A21050401846, A21050401861 Connections to Cisco 3550 first floor switch are 1 Gb s-1 to jacks 75 (sand), 48 (silt), and 56 (clay). These occupy ports 1-3 of the total 10 jacks. In switch closet, blue panel goes to station, yellow goes to switch.

dirt.ess.uci.edu = 128.200.14.25:

Dell Precision 610

Shipped 19990902 from Dell Computer on UCI PO

System Service Tag 4R5EJ

Express Service Code 7985179

Dual 500 MHz Pentium III Xeon with 512 KB Level 2 cache

1 GB RAM

Primary SCSI controller for hard disks: Adaptec AIC-7890 Ultra2/Wide LVD controller (Adaptec 2940 UW-equivalent)

Two 36 GB SCSI disks: QUANTUM Model: ATLAS 10K 36WLS

Secondary SCSI controller for CDROM: Adaptec AIC-7880 internal Ultra/Narrow and external Ultra/Wide (Adaptec 2940 UW-equivalent)

SCSI CDROM NEC Model: CD-ROM DRIVE:466

Audio: System-board-integrated 16-bit Crystal CS4237B audio controller chip which emulates Sound Blaster Pro card from Creative Laboratories, Inc. NIC: 3Com 3C905b-TX Wakeup On LAN-capable (uses a 3Com 3C918v2 ASIC)

24" Dell UltraScan P1690

1920x1200 @ 60 Hz, 75.0 kHz hsync

See <http://support.dell.com/oti/monitors/P1690/En/specs.htm>

Dell warranties monitor syncs at 1920x1200 resolution with hsync = 95 kHz, vsync = 76 Hz, dot-

clock = 245.5 MHz, horizontal/vertical sync polarity = -/-

Video controller: Diamond Viper 770D AGP PCI video adapter with 16Mb of SGRAM

This controller is based on the NVidia RIVA TNT2 chipset and uses the XF86_SVGA driver

IOmega 250 MB ZIP drive

elnino.ess.uci.edu = 128.200.14.97

Dell Precision Mobile Workstation M50 \$4400

Arrived 20030129

Service Tag 4RPK921

Express Service Code 103-826-513-53

UCI Property #: 039003413

1 GB RAM

15" UXGA

IEEE 1394

elnino inherited haze's second battery, which is a 66 Whr JP-04M778-42016-2CR-1656

24x CD R/W 8x DVD ROM

nVidia, Quadro?4 500 GoGL, 64MB, VGA

Intel Mobile Intel(R) Pentium(R) 4 - M CPU 2.20GHz stepping 07 512B L2 Cache

hda: IC25T060ATCS05-0, ATA DISK drive

hdb: Samsung CD-RW/DVD-ROM SN-324B, ATAPI CD/DVD-ROM drive

NB: hdb requires ide-scsi driver hdb mounts as sr0 hdc: HITACHI_DK23EB-40, ATA DISK drive

hde: SanDisk SDCFB-1024, CFA DISK drive Floppy drive(s): fd0 is 1.44M

Attached scsi CD-ROM sr0 at scsi0, channel 0, id 0, lun 0

sr0: scsi3-mmc drive: 0x/24x writer cd/rw xa/form2 cdda tray

agpgart: Detected Intel i845 chipset Intel 810 + AC97 Audio, version 0.21, 09:15:48 Aug 14 2002

eth0: 3c59x eth1: Looks like a Lucent/Agere firmware version 8.10

Uses hermes, orinoco, and orinoco_cs modules MAC address: 00:08:74:E4:EC:3F (wired)

MAC address: 00:02:2D:85:5C:3E (wireless)

Firewire IEEE 1394 uses ohci1394: Linksys Wireless G Broadband Router:

S/N CDF80E406886 MAC 0013107D321C Ownership ID 4HFFS4BW Device ID 5ZRJG2FV

ZyXel AG200 Wireless 802.11a/b/g USB adaptor (Based on Atheros chipset) S/N S510500149

MAC 00A0C5B810C7 <http://www.zyxel.com/product> Cleaned fans 20070917 using on-line ser-

vice manual Found with search for "Dell Precision M50 Service Manual" <http://support.dell.com/support/edocs/sy>

elnino's internal 60 GB disk drive died 20080120 This was an IBM travelstar Model: IC25T060ATCS05-

0 5400 RPM Received replacement 80 GB disk drive died in 20080205 This was an Samsung

esmf.ess.uci.edu = 128.200.197.165

Sony DSCF828 Digital Camera

SanDisk ImageMate USB 2.0 Reader/Writer for CompactFlash Type I & II

Model SDDR-91 Part Number 20-90-00091 1 year warranty SanDisk CompactFlash 1.0 GB www.sandisk.com/reg

5 year warranty Sony Camera Case LCS-VA3

haze.ess.uci.edu: 128.200.14.62

Dell Inspiron 8000

Dell order #: 609532437

UCI is Dell customer #: 6774301

Laptop is UCI PO #: 0119L03200561
 Shipped on 2001/07/31
 System Service Tag JGP4R01
 Express Service Code 423-685-352-17
 Touchpad is PS/2 compatible
 Two 59 WHR LI-ION batteries
 Pentium III 1 GHz
 512 MB SDRAM
 48 GB Hard drive
 Fixed Internal CD-RW/DVD combination drive
 100 MB ZIP drive, interchangeable with 3.5" floppy drive
 15" UXGA display
 Video controller: NVIDIA GeForce 2 Go Video, w/ 32 MB DDR, AGP 4X
 Grey Microsoft USB mouse with 3 buttons
 Audio controller: ESS Maestro 3
 Xircom RealPort 10/100 + 56K Ethernet/Modem Combo PC Card, Type 3
 PCMCIA tulip_cb driver:

```
/etc/pcmcia/config
device "tulip"
  class "network" module "cb_enabler", "tulip"
card "Xircom RealPort2 10/100 Fast Ethernet"
  version "Xircom", "*", "R2E-100BTX"
```

Bought extra battery from <http://www.computergiants.com> Attached to printer

hp5180.zendernet = fxm

HP Photosmart C5180 All-in-One Printer-Scanner-Copier Purchased from Carrefour Echirolles
 20070905 Takes 6-pack ink cartridges type part number 363 Hangs off zendernet router MAC
 address #: 001a4b954ae8

Serial Number: MY6CRQ217F04MK

Firmware Version: R0631MxNxxN0

Service ID: 17249

Problem printing is: Network host '192.168.1.2' is busy; will retry in 30 seconds

Netgear wireless router

WGR614v7 Serial #: 1JS2767T02DC9

MAC address #: 001B2F5B6BFC

Default access: <http://www.routerlogin.net>

ESSID: zendernet Router settings date-stamped and saved to virga:

```
\${DATA}/tmp/netgear_wgr614v7_20070916.cfg
```

```
\${DATA}/tmp/netgear_wgr614v7_20070923.cfg
```

```
cp /home/zender/Desktop/netgear.cfg ${DATA}/tmp/netgear_wgr614v7_20070923.c
```

Came with WG111v2 USB 802.11g card: MAC address #: 00184DFF1D2C

S/N: 1AC174BP09583

Comprehensive Wireless MAC list for zendernet router:

```

airlink_awlc3026_pcmcia 00032F36D551
ashes_trumobile1400 00904BB20986
elnino 00022D855C3E
linksys_dlink_pcmcia fxm
netgear_wg111v2_usb 00184DFF1D2C
netgear_wg511_pcmcia 00095BE8C4E1
orinoco_silver_pcmcia 00022D0982E2
virga_ipw3945 0013020A7CE5
zyxel_a200_pcmcia 00A0C5B810C7
neige_ipw4965 001DE0289E6D
hp5180_printer 001a4b954ae8 # wired

```

zendernet router reserved IPs:

```

192.168.1.2=hp5180
192.168.1.3=virga
192.168.1.4=elnino
192.168.1.5=neige (ipw4965)
192.168.1.6=ashes (trumobile1400 broadcom)
192.168.1.7=ashes (netgear-wg511_pcmcia)
192.168.1.8=neige (broadcom netXtreme 57xx Gigabit Controller)

```

hp4600.ess.uci.edu = 128.200.14.123

COLOR LASERJET 4600DN 17PPM 96MB PAR ENET 2-EIO PS3 PCL6/5C

Purchased from GST. Inc. 17707 Valley View Ave. Cerritos CA 90703-7004 Arrived Jan 23, 2003

Serial number: JPBKB18664

UCI Property #: 039003414

17 ppm clr

600x600 DPI

96 MB RAM

10 GB disk

81.0180 EIO 1 ERROR messages: If you have a 615N/J6057A card and you get EIO errors, open a ticket for your free replacement. You have a bad card. 1-800-HPINVENT.

<https://forums1.itrc.hp.com/service/forums/questionanswer.do?threadId=71839>

My card is an HP615N/J6057A card

S/N: SG2B303C8E

Case #1: 3213917944

Case #2: 3213977696

CSO#: 2689174068000

Called second time on 20060223 switched from port 2 to port 3

Green Laser Pointer Originally with three Vinnic L1154 batteries <http://www.batterymart.com/battery.mv?p=B>
 LR44 has a picture of these batteries and gives their technical specifications as Volts: 1.5, mAh: 120., Chemistry: Alkaline Dimensions: 11.6 X 5.4 (Dia x Ht. MM) The bumpy (negative) side

faces into the pointer, against the spring The flat (positive) side faces out of the pointer, toward the chain I believe the pointer came with three L1154 batteries The pointer does not work with three Rayovac 303/357 batteries The pointer does work with four Rayovac 303/357 batteries, but the lid will not shut Hence, the problem seems to be that the aggregate power is not enough The Rayovac is a Silver Oxide chemistry, like the Vinnic S1154 (but not the L1154).

ipcc.ess.uci.edu = ipcc.calit2.uci.edu = 128.195.185.75
 pbs.ess.uci.edu = pbs.calit2.uci.edu = 128.195.185.76
 The Wiki for PBS is at <http://tephra.ess.uci.edu/PBSWiki>
 tephra.ess.uci.edu = 128.200.14.171

neige.ess.uci.edu = 128.200.14.fxm
 Dell Precision M6300 64-bit architecture uses EM64T instruction set Received at UCI/SMU 20071221/20071228
 Dell Service Tag: JB1S8F1 Complete Care until 20111213 Gold support until 20111213 Express
 Service Code: 42026989213 4 year limited warranty plus 4 year NBD on-site service and complete
 care Gold technical support Intel Core2 Extreme CPU X7900 2.8 GHz 800 MHz FSB
 17" UltraSharp Wide Screen WUXGA
 24x CD-RW/DVD burner, 8x DVD+/-RW, DVD-ROM, Blu-Ray Re-writable
 4 GB RAM NVidia Quadro FX 1600M 512 MB Windows Vista Business Product ID: 89576-
 OEM-7332141-00054 Intel Wireless WiFi 4965AGN Broadcom NetXtreme 57xx Gigabit Con-
 troller MAC address: 001C231F2730 (wired broadcom)
 MAC address: 001DE0289E6D (wireless ipw4965AGN)
 Sound controller problems:

https://wiki.ubuntu.com/Gutsy_Intel_HD_Audio_Controller

Get video working: [nvidia-glx-new UPEK Fingerprint Reader](#)

<http://thinkfinger.sourceforge.net>
<http://packages.ubuntu.com/hardy/admin/libpam-thinkfinger>

```
# Dell site for BIOS updates
http://linux.dell.com/wiki/index.php/Repository/firmware
sudo aptitude install firmware-tools firmware-addon-dell
deb http://linux.dell.com/repo gutsy dell-software
http://direct2dell.com/one2one/archive/category/1021.aspx
# Four steps to firmware updates:
sudo wget -q -O - http://linux.dell.com/repo/firmware/bootstrap.cgi | bash
sudo aptitude install firmware-addon-dell
sudo aptitude install $(bootstrap_firmware -a)
sudo update_firmware
```

virga.ess.uci.edu = 128.200.14.189
 Dell Inspiron 9400
 Received at UCI 20060221
 Dell Service Tag: F1N0J91

Express Service Code: 327-508-573-33

4 year limited warranty plus 4 year NBD on-site service and complete care

Gold technical support

Intel Core Duo Processor 2 GHz/667 MHz FSB

17" UltraSharp Wide Screen UXGA WUXGA

8x CD/DVD burner (DVD+/-RW) with double layer DVD+R write capability

2 GB RAM Logitech MX1000 Laser Cordless mouse: 29.1 ROM GP Logitech Resolution 800

DPI, 5.8 MP s-1 Logitech S/N: LZB533350679 MAC address: 00:14:22:EF:61:8E (wired)

MAC address: 00:13:02:0A:7C:E5 (wireless ipw3945)

Bluetooth Dell Wireless 350 Bluetooth Internal (2.0+ enhanced data rate)

Bluespoon AX2 Logitech QuickCam Pro for notebooks IEEE 1394

scsi0 : ata_piix Vendor: ATA Model: Hitachi HTS72101 Rev: MCZO scsi1 : ata_piix Vendor:

_NEC Model: DVD+-RW ND-6650A Rev: 102C SCSI device sda: 192426570 512-byte hdwr

sectors (98522 MB) 1 PCI Express card slot Zero pcmcia slots Video: 256 MB Nvidia GeForce

Go 7800 eth0: Broadcom 4400 10/100BaseT Ethernet 00:14:22:ef:61:8e Intel PRO/Wireless 3945

80211a/g minicard

Sound: <http://wiki.ubuntu.com/DebuggingSoundProblems> # Turn off internal speakers when headphones plugged in # echo options snd-hda-intel model=ref — sudo tee -a /etc/modprobe.d/alsa-base

Ubuntu wireless cards:

<https://wiki.ubuntu.com/HardwareSupportComponentsWirelessNetworkCards>

This helpful site:

http://www.thinkwiki.org/wiki/Intel_PRO/Wireless_3945ABG_Mini-PCI_Express_A

mentions that the correct driver is available from sourceforge at

<http://ipw3945.sourceforge.net/>

The upshot is that the ipw3945 driver is expected to be available Q1 2006

When complete, Intel/Linux wireless drivers are available at

<http://support.intel.com/support/notebook/sb/CS-006408.htm>

Thinkwiki used to (20060215) recommend using the ipw2200 driver

Here's what happens when you try that:

http://www.thinkwiki.org/wiki/Ipw2200#Installation_on_Debian

```
aptitude remove linux-headers-2.6.15-18-686 linux-headers-2.6.15-18
```

```
aptitude install linux-headers-2.6.15-20-686
```

```
sudo aptitude install ieee80211-source
```

```
sudo module-assistant -t build ieee80211-source
```

```
sudo aptitude install ipw2200-source
```

```
sudo module-assistant -t build ipw2200-source
```

Install firmware

```
cd ${DATA}/tmp
```

```
sudo tar xzvf ipw3945-linux-1.0.0.tgz
```

```
cd ${DATA}/tmp/intel-ipw3945-1.0.0/
```

```
-C /lib/hotplug/firmware
```

```
modprobe ipw2200
```

```
iwconfig
```

```
lsmod | grep ipw
# Module loads but eth1 does not show up in ifconfig
```

flyash.ess.uci.edu (originally named lanina):
 usually dhcp-14116.ess.uci.edu, dhcp-14118.ess.uci.edu (wired) or dhcp-086159.mobile.uci.edu,dhcp-086183.mobile.uci.edu (wireless)
 Dell Inspiron 7500R
 Dell order #: 320036478
 UCI is Dell customer #: 6774301
 Laptop is UCI PO #: 0120L03002256
 Shipped on 2000/01/11
 System Service Tag Y20Z8
 Express Service Code 572-015-24
 Touchpad is PS/2 compatible
 Pentium mobile III 500 MHz
 512 MB 100 MHz RAM
 25 GB Hard disk
 120 MB Super disk, aka LS120 (mounted as /mnt/lS120 = /dev/hdc)
 Display is 15.0" SXGA+ active matrix color
 Video controller: ATI Mobility P, 64 bit, AGP 2X w/ 8 MB RAM (ati X.org server)
 Controller supports 32 bit color but LCD display limited to 18 bit
 Infrared IrDA 1.1 port
 Logitech first mouse 3 button PS/2 compatible
 Audio controller: ESS Technology Maestro-2e (Sound Blaster Pro-compatible)
 udev:DEVPATH=/bus/pci/drivers/ES1968 (ESS Maestro) udev:UDEV [1170546089.233339] add@/module/snd.
 udev:PHYSDEVDRIVER=ES1968 (ESS Maestro) <http://occy.net/taxonomy/term/6?from=20>: "In order to get sound working on my Mom's Dell Inspiron 8200, using Ubuntu Linux, I had to do the following:"

```
# edit /etc/hotplug/blacklist
add snd_intel8x0m
edit /boot/grub/menu.lst
add # kopt=root=/dev/hda1 ro acpi_irq_isa=7
# make SURE you leave in the # above
run update-grub
```

Wireless: Lucent technologies Orinoco silver card Orinoco MAC address: 00:02:2D:09:82:E2 Xircom 32bit Cardbus Ethernet 10/100 + Modem 56 (aka CBEM56G 1.03)
 Xircom MAC address: 00:10:A4:08:12:31 PCMCIA tulip_cb driver:

```
/etc/pcmcia/config
device "tulip_cb"
    class "network" module "cb_enabler", "tulip_cb"
card "Xircom CardBus 10/100 Ethernet + 56K Modem"
    version "Xircom", "*", "CBEM56G"
    bind "tulip_cb" to 0, "serial_cb" to 1
```

CBEM56G ethernet cable: CABLE-ETH122 CBEM56G modem cable: CABLE-MOD444 PCI controller: Texas Instruments PCI 1225 Cardbus

Airlink AWLC3026 802.11b/g card purchased 20061216 from Fry's: MAC 00:03:2F:36:D5:51

```
hda: IBM-DARA-225000, ATA DISK drive\\
hdc: LS-120 SLIM3 00 UHD Floppy, ATAPI FLOPPY drive\\
hdd: TorISAN DVD-ROM DRD-U624, ATAPI CDROM drive\\
hda: IBM-DARA-225000, 24207MB w/418kB Cache, CHS=3278/240/63\\
hdd: ATAPI 0X DVD-ROM drive, 256kB Cache\\
```

Video modes supported:

```
720x400 16 colors @ 70 Hz text mode
1280x1024 8,16,24,32 bit @ 100 Hz
1400x1050 8,16,24 bit @ 100 Hz
1400x1050 32 bit @ 85 Hz
1600x1200 8,16,24 bit @ 100 Hz
1600x1200 32 bit @ 85 Hz
```

Do horizontal and vertical refresh rates have any meaning for TFT/LCD displays? Installation programs ask for these variables and the following settings are reported to work for the i7500 at 1400x1050: hsync range 31.5 -82, vrefresh 40-110. For haze the horizontal refresh should be set at 28-90, while the vertical should be set at 40-110. **NB: ATI RAGE card at 1400x1050 resolution does not work on Fedora/Ubuntu unless magic option "vga=792" is added to kernel boot line, e.g., in grub.conf or menu.lst**

lanina.zender.org = 128.200.14.80

Compaq Presario 5240

Purchased 19990312 from CompUSA in Boulder \$1100

AMD K6-2 3d Processor at 400 MHz

Serial number 1X91CFDH8662

10 GB Hard disk

hjm:03-03-05 added 200GB disk, added ProMEPIS (Debian) OS, in following partitions: hda1-10GB -OS, hda2-1GB swap, hda3- /home (78GB) hda4 -/data(spare) 64 MB RAM

hjm: 03-03-05 upgraded to 256MB (with simms from home) 120 MB Super disk, aka LS120 (mounted as /mnt/l120 = /dev/hdd)

```
hda: WDC AC310000B, 9541MB w/512kB Cache, CHS=1292/240/63
hda: WDC AC310000B, ATA DISK drive
hdc: LTN382, ATAPI CDROM drive
hdc: ATAPI 40X CD-ROM drive, 120kB Cache
hdd: LS-120 COSM 05 UHD Floppy, ATAPI FLOPPY drive
```

Floppy drive(s): fd0 is 1.44M

Rockwell HCF 56 kbps DataFax modem on PCI bus COM2 (/dev/ttyS1)

3COM-US Robotics 56 kbps winmodem on ISA bus COM1 (/dev/ttyS0) IRQ4

NIC: PCI Fast ethernet DEC 21143-based controller
 Graphics: Rage LT PRO AGP 2X (XF86_Mach64 server)
 USB
 PCI disk controllers

```
# Graphics: Rage LT PRO AGP 2X
# 19" Monitor purchased separately 19990901 from CompUSA ~\$350:
# Horizontal scan frequency 30--95 KHz
# Vertical scan frequency 47--150 Hz
# Supports 1024x768 @ 68.6 KHz / 85 Hz non-interlaced
# Supports 1280x1024 @ 80 KHz / 75 Hz non-interlaced
# Supports 1600x1200 @ 93.7 KHz / 75 Hz non-interlaced

# Computed following modeline settings using above info and
# http://www.dkfz-heidelberg.de/spec/linux/modeline/
# V-freq: 85.00 Hz // h-freq: 68.79 KHz
Modeline "1024x768" 97.40 1024 1072 1192 1416 768 768 771 809
# V-freq: 75.00 Hz // h-freq: 80.42 KHz
Modeline "1280x1024" 151.83 1280 1360 1544 1888 1024 1024 1027 1072
# V-freq: 75.00 Hz // h-freq: 94.24 KHz
Modeline "1600x1200" 242.01 1600 1728 2024 2568 1200 1200 1204 1256
```

sand.ess.uci.edu = 128.200.14.132

Western Scientific \$4000

Chassis Serial number 1012003577

Arrived 20040221

Service Tag fxm

Express Service Code fxm

UCI Property #: 049003617

2 GB RAM

Two AMD Opteron(tm) Processor 244s at 1.8GHz

Bought without monitor

Now using Dell UltraSharp monitor same as biogenic

Network card sk98lin Logitech cordless mouse/keyboard: Navigator Duo Graphics Card: ASUS V9520 Series CPU Graphics card powered by NVIDIA GeForce FX 5200 GPU. Supports AGP 8x Motherboard: IWILL DK8X: Eight sockets for up to 16 GB RAM Use 184-pin Registered PC2700/2100 ECC DDR memory modules Two 32-bit/33 MHz PCI slots One 64 bit/66 MHz PCI slot Two PCI-X slots One AGP Pro 8X slot hda: DVD: Plextor Model PX-708A Internal 40X ATAPI DVD+-R/RW CD-R/RW drive hdc: DV-516D 0106 ATAPI 48X CD/DVD-ROM drive Tech Support (800) 443-6699 Hank Vu (800) 443-6699 x. 211 hank@wsm.com

2.28.05 mods 3.29.05 by hjm 200 gb disk from lanina moved to sand on IDE bus 1 new hoary ubuntu distro loaded on 200gb disk as:: Filesystem 1K-blocks Used Available Use/dev/hda1 10482044 2048572 8433472 20tmpfs 1786416 0 1786416 0/dev/hda6 105294788 5897208 99397580 6/dev/hda3 73278252 2128912 71149340 3/dev 10482044 2048572 8433472 20none 5120 2836 2284 56

sand was upgraded to 4GB ram, of which 3.5 is available to the system - a little mentioned weirdness of Opteron systems is that the 0.5 GB of RAM just below 4 GB disappears into the memory-mapping black hole that AMD inherited from INtel.

After the SW raid experiment, sand has 3ware Escalade HW raid card driving eight identical Western Digital 2500jd drives in a RAID 5 config yielding 1.7TB usable. It looks to the system like a single SCSI disk.

5.13.05 - disk failure on sand's RAID5 partition - detected on reboot in dmesg, not by log or email as expected. At this point, the data was still intact and SHOULD have been backed up to another system, but since it was HW RAID5 AND it was 3ware controller (known for reliability (HA!) and robustness (HA!)) AND this was acting as the backup for other systems (which were still OK) AND the data was 200GB at this point, I thought it was ok to go ahead. MISTAKE!

First thing was to find the problem of why we hadn't been informed of the failure beforehand. The controller (a 3ware Escalade 86506-8port driving 8x250GB identical WD disks) has, like most such cards, a BIOS-based utility for setting up the RAID which actually worked pretty well, except that unlike SW RAID, you can't use the raid immediately in degraded mode (while it's building the checksum info across the raid) - you have to let it sit there for hours (it's a 1.6 TB array) while it checksums the entire array (even tho there's nothing on it to begin with). That done, it looks like a giant scsi disk to the OS - so far so great.

3ware also comes with a web-frontend utility called 3dm and a commandline utility called tw_cli. When I had installed the 3dm, I had gone thru the installation script, checked that there were no error messages, checked that I got an email verification and then forgot about it - altho thinking about it - I must have gotten the email from the script, not the app. I did not check that the web server interface was working as I didn't think I'd ever use it. MISTAKE.

Now I DID need to talk to the controller and the 3dm/tw_cli were the only things that could while the OS was running. THIS is one of the downsides of a hardware RAID - you're stuck with the tools that the vendor gives you. Since I was running on a 64-bit SMP Linux (Ubuntu), dual opteron, the installation bash script ran fine, but the monitoring daemon silently failed (32 bit code and I was running a 64-bit-only OS). So nothing was hearing the controller screaming that a disk had died and the RAID was now running in degraded mode. (as noted above, the only thing that let us know this was an entry in dmesg on a reboot.)

After verifying that this software was in fact incompatible with the OS, I tried to find an upgrade that WOULD let me talk to the controller. I figured that 3ware being a vendor of high-end hardware, my kind of machine would be among their main targets. And I was right - BUT ... Trying to find the software that was compatible with my system was an exercise in frustration - 3ware's web site is walled off from google's bots (like almost all corporate sites) and since 3ware is relatively high end hardware, there are not a lot of messages on the linux BBSs about such failures and how to deal with them. So after a couple hours of browsing I had to go back to the 3ware site and deal with their oh-so cool web design that doesn't show URLs in any way different than regular text. The text only shows up as a hyperlink if you mouse over it. I noticed this and then had to mouse over entire pages of text, line by line to search for likely hyperlinks.

The one that finally took me to the page I needed was buried in a paragraph that I almost overlooked. Turns out the SW does exist, but is NOT specified for the controller I have (8506-8) but the 9000-series controller (which is noted in the fine print as being backwards compatible with the 8000-series). ALso, it's not 'released SW', it's being 'In Engineering Phase'. To make a long story shorter, I ended up downloading and trying several versions of software until I finally stumbled

over the right software - the 64bit versions of the 3dm2 and cli for the 9000-series controllers. This installed OK and apparently ran. The web interface software however, while it started up and presented an optimistic login screen, gave no indication of what the passwords should be or where to go to set them. After looking in the config file (/etc/3dm2/3dm2.conf) only to find encrypted passwords, I then wandered around the 3ware web site trying to find documentation about how to set or even find out what the initial passwords were. There were no docs or help files or README's with the software (it's '3ware' for those of you going thru the same hell; you can change them via the web interface when you finally get in.)

Re the passwords - nothing - or at least nothing I could find in about an hour's searching. I finally decided to look in the installation script - bingo. The passwords are set and encrypted into the config file from there. SO after setting them to what I wanted, FINALLY I was able to log into the web interface and talk to the controller. And in fact after being able to log in, the help file DOES tell you what the password is and how to change it.

Actually the `tw_cli` app also works, but it's pretty ugly (altho give them credit - they DID make 2 linux-specific clients). The one that I needed was the 3ware 9000 series 3DM2 Linux64-bit one - helpfully, on the web page I eventually found: (<http://3ware.com/support/downloadpageeng.asp?SNO=4>), both the 32 bit and 64 bit one are named the SAME THING.

So here I am, talking to the 3ware controller via the web interface and while it's not fantastic, it's really not bad. And one disk has been marked bad. So now I have to replace the bad disk. I'm just about to bring the system down to do this, when I realize the disks are sitting in the expensive hotswap cages we bought for this specific purpose, so (after unmounting the filesystem) I take a deep breath, and pop the offending disk. nothing happens - the system doesn't freeze or explode or anything - it looks like it has actually worked - and the 3dm2 interface shows that the bad disk is now gone. GREAT! I quickly replace the disk with a spare and slam it back in again - and there it is on the web interface. Now isolated all by itself.

Now - how to go about adding it back to the RAID? The Web interface is a bit dodgy on how to go about adding this disk back into the array. And the help pages are not particularly helpful; the Maintenance help page sort of obliquely refers to this scenario, but certainly doesn't give any specific step-by-step instructions. You'd have thought that since one of the primary reasons for buying such an expensive controller is to be able to replace a RAID5 disk on the fly, they might have a specific mention of such an eventuality. The way I did it is to add the disk to a new 'UNIT' and then add that UNIT to the previously defined RAID5 UNIT and request that the new combined unit be rebuilt. That seemed to work and the controller went about integrating the new disk into the raid 5 array. Again, it was not possible to mount the array and use it while it was being integrated, like you can do with SW RAID under linux. This took several hours, and in the end, it FAILED. That was the just about the last straw. After spending \$ and time (= \$) on this escapade (that's what the Escalade series SHOULD be named), the thing fails to rebuild the array. (But at least it now reports via email that it has failed.) So now what???

The filesystem was a reiserfs to begin with. As a last resort, I try to rescue the thing with a `fsck.reiserfs`. After reading the dire warning about this being the last thing you should try, I give it the `-rebuild-tree` option and go home. This being 2TB of disk, it takes a while. Later that night I see it's completed and try to mount it. To my astonishment it mounts. I do a 'df' - hmm - that's not good - only 3started (the raid had only been running a short while). I'm not at all happy to see that it the only directory on the partition is ... `lost+found`. This dir contains the rubble of what used to be about 200GB of expensive and carefully groomed earth-sensing and atmospheric data.

So go ahead - ask me - Am I happy that I spent the extra *to buy a hardware raid card rather than two 30 4port sata controllers and using SW RAID?*

I probably couldn't have done all this disk hot-swapping with a non-HW RAID card, but the cost of a reboot for most of us is not that big a deal. That said, I'm not sure of the total complexity that doing such a thing under SW RAID would have entailed. To do this with 2x 4 port controllers would have required additional complexity and I'm not sure it can be done easily with mdadm. And it is possible that I did something wrong in the 3ware rebuild - I'll be sending this narrative back to them as well.

As a postscript to this, I should also mention that while most Linux server vendors sell 3ware cards, at least one (Los Alamos Computers) suggests SW raid as being both significantly cheaper and faster. They suggest the Promise SATA TX4 for about 70. *Newegg has the supported - in - kernel Silicon Image chip setted Syba 4port card for 30.* If you remember my previous posts, I was surprised to find SW RAID to be a bit (10-20 that I probably should have tried the SW RAID on a full 2TB array.

Well, you makes your choices and you takes your chances.

Currently, the 3dm2 SW is running and should be restarted with a reboot (commands to re-start are in /etc/init.d/local. Will test to be sure. The 3dm2 web interface can be gotten at sand:888 and the password for user (read-only) is '3ware'. The admin password is about a coyote-resistant neighbor's cat.

6.7.05 - noticed that the 3dm2 daemon was reverting to original settings on each reboot. It hadn't written a config dir (/etc/3dm2.3dm2.conf) and so was losing the config. It also wasn't complaining that it couldn't do it, to stdout/err, to dmesg nor to syslog.

rsync backups: There is an rsync cron job in /etc/crontab/rsync.home2data that will rsync the /home dir (where the web site is, the local dir tree, an all user dirs, as well as some additional stuff) to the /data/home dir. It's stored uncompressed, so we can squeeze a bit more data out of it if we want. The rsync is initialized and tested and will write an entry into the syslog. excluded dirs/files can be entered in the file: /etc/rsync/sand.home2data.exclude, one per line. 6.7.05 - also rsyncing the web site to soot so it can be used as a backup server in case sand explodes again. Note that it is being owned as 'hjm' so it may have to be recursively chowned on soot, tho I think it will work fine

There is a /etc/init.d/local file that is responsible for starting a variety of local services - the license manager, the nco-bm benchmark server, making sure apache starts, mounting the /data dir, etc. It should be updated on a change with:

```
$ update-rc.d local defaults 80
```

```
hjm++ 9.20.05
```

```
sand's syslog and dmesg ring was filling up with ethernet errors. While you
degree of control by editing /etc/syslog.conf, the easiest way to stop the v
of such errors like:
```

```
Sep 20 10:14:56 sand kernel: DROPPED IN=eth0 OUT= MAC=ff:ff:ff:ff:ff:ff:00:
Sep 20 10:14:56 sand kernel: DROPPED IN=eth0 OUT= MAC=ff:ff:ff:ff:ff:ff:08:
Sep 20 10:14:57 sand kernel: DROPPED IN=eth0 OUT= MAC=ff:ff:ff:ff:ff:ff:00:
```

```
is to stop the logging using guarddog (with the "Logging" panel, surprising
```

That stops 99% of the lines.

hjm--

flexlm license manager running on sand

If the license manager screws up (infrequent, but not unknown):

test the status:

```
% /usr/local/rsi/idl/bin/lmstat
```

if it needs to be restarted, bring it down smoothly with:

```
% /usr/local/rsi/idl/bin/lmdown -q
```

(you can't just kill the flexlm daemon and restart it as it sets up l

files all over the place)

#After it comes down, restart it via:

```
su -c "/usr/local/rsi/idl/bin/lmgd" hjm
```

or

```
su -c "/usr/local/rsi/idl/bin/lmgd" cluo
```

(the manager has to run as a regular user, not as root)

installed gnome as well for Dan. Just as well - kde has some issues with the kernel 2.6.11-1 SMP installed. extras include: libxml2-dev libxml2-doc libxml2-utils bonnie++ xosview tkdiff kompare kdesdk-doc-html ssh apcupsd guarddog vncserver synaptic ddd ddd-doc pydb glibc-doc gnuplot libqt3-dev tree nedit joe gnome

sand is now sitting behind a pretty restrictive firewall so if things don't work, that may be the reason. In fact, it was preventing the nco-bm server from getting data on udp port 29659 until guarddog was changed to address that issue.

also running ubuntu kernel 2.6.11-1 smp sort of successfully, but it kills the usb connectivity so apcupsd doesn't see teh ups. And then it did... ANd then it didn't. Then it did. Then it didn't. Still to be resolved.

07-13-05 hjm - Now runnning 2.6.10-5-smp-k8 more or less successfully, een with the USB also upgraded all the KDE stuff, so now it appears to be running more stably. printing hint for remote use: kcmshell printmgr will bring up the KDE print manager in full.

DODS server on sand. The DODS server on sand is not a server at all, but just a series of cgi scripts that are active as soon as the apache server comes up. DODS urls are constructed as:

```
http://sand.ess.uci.edu/cgi-bin/dods/nph-dods/dodsdata/in.nc.dds?
```

server name	^^^	^^^^^^^^	^^^^^^^^	^^^^^	^^^
					file name ext :
name of the subdir* in the cgi -+					to define acti
tree where the DODS cgi are					DODS
kept					
					data file name

```

|         subdir* rooted from htmlroot
|         actual files live
name of the cgi script that is called

```

*this is the same across zender DODS servers

seasalt.ess.uci.edu = 128.200.14.39:

Dell Dimension T Minitower 800 MHz PIII

Arrived 20000615

System Service Tag DS8J10B

Express Service Code 300 055 479 95

Intel Pentium III (Coppermine) 800 MHz

256 MB RAM

40 GB Ultra ATA 7200 RPM with ATA 66 controller

NIC: 3Com EtherLink 10/100 PCI PCI For Complete PC Management (3c905c-TX)

Microsoft PS/2 Mouse (Intellipoint)

Promise Technology Inc. Ultra66 IDE Controller

Intel 82371AB/EB PCI Bus Master IDE Controller

120 MB Super disk, aka LS120

Sony 8X/4X/32X CD-RW drive

```
# Graphics: NVIDIA TNT2 M64 4xAGP with 32MB RAM
```

```
07-12-05 hjm
```

```
modified the xorg.conf to read:
```

```
Section "Device"
```

```
    Identifier      "NVIDIA Corporation NV5M64 [RIVA TNT2 Model 64/Mode
```

```
    Driver          "nvidia"
```

```
    BusID          "PCI:1:0:0"
```

```
    Option "BackingStore" "True" <--- allows covered windows to refresh
```

```
    Option "SaveUnders" "True" <--- ditto
```

```
    Option "RenderAccel" "true" <--- allows better HW acceleration.
```

```
EndSection
```

```
# 19" Dell M990 purchased with computer
```

```
# See http://support.dell.com/docs/monitors/m990/En/Index.htm
```

```
# Video controller: NVIDIA TNT2 M64 4xAGP with 32MB RAM
```

```
# Uses XF86_SVGA driver
```

```
# Horizontal scan frequency 30--96 kHz (automatic)
```

```
# Vertical scan frequency 50--160 Hz (automatic)
```

```
# Optimal preset resolution 1024x768 at 75 Hz or 85 Hz
```

```
# Highest preset resolution 1600x1200 at 75 Hz
```

```
# Highest addressable resolution 1600x1200 at 75 Hz
```

```
# Preset modes:
```

```
# Supports 1024x768 @ 68.7 kHz 85 Hz +/+
```

```
# Supports 1280x1024 @ 80.0 kHz 75 Hz +/+
```

```
# Supports 1280x1024 @ 91.1 kHz 85 Hz +/+
```

```
# Supports 1600x1200 @ 93.7 kHz 75 Hz +/-

# Computed following modeline settings using above info and
# http://www.dkfz-heidelberg.de/spec/linux/modeline/
# V-freq: 85.00 Hz // h-freq: 68.79 KHz
Modeline "1024x768" 97.40 1024 1072 1192 1416 768 768 771 809 +HSync
# V-freq: 80.00 Hz // h-freq: 86.05 KHz
Modeline "1280x1024" 167.97 1280 1368 1576 1952 1024 1024 1027 1075 +HSync
# V-freq: 85.00 Hz // h-freq: 91.72 KHz
Modeline "1280x1024" 185.64 1280 1376 1600 2024 1024 1024 1028 1079 +HSync
# V-freq: 75.00 Hz // h-freq: 94.24 KHz
Modeline "1600x1200" 242.01 1600 1728 2024 2568 1200 1200 1204 1256 +HSync
```

Linux 2.2.x kernels do not recognize the ATA66 controller, but patches and workarounds are available. HHPT366 HOWTO at <http://www.csie.ntu.edu.tw/%7Eb6506063/hpt366/> UDMA Mini-HOWTO <http://www.linuxdoc.org> My solution was to plug the hard drive directly into the IDE controller on the motherboard and then install linux, i.e., I bypassed the ATA66 controller.

soot.ess.uci.edu = 128.200.14.98

Dell Precision Workstation 650n Minitower

Arrived Jan 23, 2003

Service Tag GC9L921

Express Service Code 35570219545

UCI Property #: 039003412

Dual Xeon 2.8GHz

2 GB RAM

hjm: 4.11.05 - added "/etc/init.d/httpd restart" to /etc/rc.local to restart web server on reboot.

```
# hjm - new 200 GB disk added (4.14.05)-on same IDE bus as 1st; suboptimal,
# but physically tricky to take over other IDE bus. If very slow, will
# try to move it.
```

```
# 6.1.05 - moved 2nd disk to other controller channel - much faster.
```

```
1024 $ df
```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/mapper/VolGroup00-LogVol100	113177264	94615204	12812916	89%	/
/dev/hda2	101105	30766	65118	33%	/boot
none	1037348	0	1037348	0%	/dev/shm
/dev/hdc2	191440744	73245868	118194876	39%	/data3

```
6.7.05 - installed Intel icc compiler/debugger on soot in /opt.
```

```
% end soot.ess.uci.
```

```
\subsection{SuSE Linux}\label{sxn:suse}
```

```
The Zender group server, \mchidx{sand.ess.uci.edu}, runs
```

```
\trmidx{SuSE Linux} Professional version~9.0.
```

SuSE offers phone support at 1-510-628-3385 Monday--Friday from 9:00~AM--3:00~PM PST.

SuSE email support is at `\url{support@suse.com}`.

Our registration logyn is csz and the password is pw_lw.

Our SuSE software serial number is 48205940904172.

Yast Online Update location should be `ftp://ftp.suse.com/pub/suse`.

Configuring printers with SuSE is non-trivial, but good technical details are given at `\mchidx{portal.suse.com}`.

Search the support database for `\trmidx{CUPS}`.

```
\subsection{Browsers}\label{sx:moz}
```

UCI libraries make many `\trmidx{journal subscriptions}` available for free on the campus network or from remote locations provided that browser `\trmidx{proxy}` settings are correctly configured.

```
\begin{verbatim}
```

```
http://www.lib.uci.edu/services/how/connect.html # Describes proxies
```

```
http://www.nacs.uci.edu/help/proxy/ # UCI proxy instructions
```

Change Preferences->Advanced->Proxies->Automatic proxy configuration to

```
http://proxy.uci.edu/proxy_config.pac # Use for automatic proxy configuration
```

9.17.1 Firefox

Disable popup windows by adding the following line to `user.js`:

```
user_pref("dom.disable_open_during_load", true);
```

Do not modify `prefs.js` directly since it is automatically overwritten on shutdown and thus custom changes are not persistent.

9.18 Create krein disk structure

```
mkdir -p /data/zender/ZENDER
```

```
mkdir -p /dks9d2s0/BIAN/match
```

```
mkdir -p /dks9d3s0/ZENDER/match
```

```
mkdir -p /dks9d4s0/DSS
```

```
mkdir -p /dks9d5s0/ZENDER/csm
```

```
ln -s /dks9d2s0/BIAN/match /data/zender/BIAN/match
```

```
ln -s /dks9d3s0/ZENDER/match /data/zender/ZENDER/match
```

```
ln -s /dks9d4s0/DSS /data/zender/DSS
```

```
ln -s /dks9d5s0/ZENDER/csm /data/zender/ZENDER/csm
```

9.19 NFS export

One may use the network file system (NFS) to mount or export directories across the network. To export local disks to a host, edit the `exports` file, e.g., `/etc/exports`.

If remote host needs root access (root on remote machine has root privilege on the disks), add the hostname behind the `root=` chain. Otherwise add the host behind the `-access=` chain. If not sure which one you want then add the system to `-access=`. Note that the host separator is `:.`

Details of managing setup are described in excellent NFS-HOWTO <http://www.linux.org/docs/ldp/howto/NFS-HOWTO>. It is easy to export the disks on an already running filesystem. First make sure that the NFS daemon, `nfsd`, is running. Then make sure the NFS lock daemon, `nfslock`, is running. `exportfs`:

```
sudo /etc/init.d/nfs restart
/usr/sbin/exportfs -rv
```

By default, NFS exports filesystems with a feature called root-squashing enabled. With root-squashing enabled, processes do not maintain root privileges when writing to an NFS-exported partition on a different machine. This is a useful security measure, but plays hell with poor man's cron scripts that backup files across machines. To override root-squashing, use the `no_root_squash` option in the `exports` file.

To export NFS filesystems securely, update `hosts.allow` and `hosts.deny` as appropriate. In particular, `hosts.allow` must allow access to all machines that request INET services from a machine with NFS-mounted home directories. To restrict INET services to specific hosts use, e.g.,

```
$ cat /etc/hosts.allow
portmap: 128.200.14.25
```

To allow INET services to all hosts use, e.g.,

```
$ cat /etc/hosts.allow
portmap:ALL
```

To restrict INET services to all hosts use, e.g.,

```
$ cat /etc/hosts.allow
portmap: 128.200.14.25
portmap: 128.200.14.62
portmap: 128.200.14.39
portmap: 128.200.86.234
```

This may prevent errors revealed by `gconf-sanity-check-1`:

```
$ gconf-sanity-check-1
```

9.20 NFS mount these disks on local machine:

Here are the mountpoints to use on `lanina` when mounting the central `fileserv` partitions at NCAR:

```
sudo mkdir -p /ncar/fs/cgd/home0
sudo mkdir -p /ncar/fs/cgd/data0
sudo mkdir -p /ncar/fs/cgd/csm
```

Once these mountpoints exist, the machine is on the NCAR network, and the mount options are listed in the `/etc/fstab` file, the partitions may be mounted with the following commands:

```
sudo mount fileserver-n8.ucar.edu:/fs/cgd/home0
sudo mount fileserver-n8.ucar.edu:/fs/cgd/data0
sudo mount fileserver-n8.ucar.edu:/fs/cgd/csm
```

Note that the partitions should not be automatically mounted at boot time since usually the laptop is not directly on the NCAR network. The `noauto` option to `mount` accomplishes this. Thus the appropriate `fstab` entries are

```
fileserver-n8.ucar.edu:/fs/cgd/home0    /ncar/fs/cgd/home0    nfs    exec, de
fileserver-n8.ucar.edu:/fs/cgd/data0    /ncar/fs/cgd/data0    nfs    exec, de
fileserver-n8.ucar.edu:/fs/cgd/csm      /ncar/fs/cgd/csm      nfs    exec, de
```

When disconnecting the laptop from the network, it is best to `umount` these partitions so that the NFS-related daemons do not waste time looking for them when the laptop is turned on again somewhere off the network.

```
sudo umount fileserver-n8.ucar.edu:/fs/cgd/home0
sudo umount fileserver-n8.ucar.edu:/fs/cgd/data0
sudo umount fileserver-n8.ucar.edu:/fs/cgd/csm
```

Here are the mountpoints on `krein`:

```
mount krein.math.uci.edu:/dks9d2s0 at mountpoint /dks9d2s0
mount krein.math.uci.edu:/dks9d3s0 at mountpoint /dks9d3s0
mount krein.math.uci.edu:/dks9d4s0 at mountpoint /dks9d4s0
mount krein.math.uci.edu:/dks9d5s0 at mountpoint /dks9d5s0
```

Create mountpoints on local machine (e.g., `dust.ess.uci.edu`):

```
mkdir /dks9d2s0
mkdir /dks9d3s0
mkdir /dks9d4s0
mkdir /dks9d5s0
```

Add mount commands to `/etc/fstab`:

```
krein.math.uci.edu:/dks9d2s0    /dks9d2s0    nfs    exec, dev, suid, rw 1
krein.math.uci.edu:/dks9d3s0    /dks9d3s0    nfs    exec, dev, suid, rw 1
krein.math.uci.edu:/dks9d4s0    /dks9d4s0    nfs    exec, dev, suid, rw 1
krein.math.uci.edu:/dks9d5s0    /dks9d5s0    nfs    exec, dev, suid, rw 1
```

Mount these partitions interactively the first time:

```
mount krein.math.uci.edu:/dks9d2s0 /dks9d2s0
mount krein.math.uci.edu:/dks9d3s0 /dks9d3s0
mount krein.math.uci.edu:/dks9d4s0 /dks9d4s0
mount krein.math.uci.edu:/dks9d5s0 /dks9d5s0
```


Create softlinks on local machine (e.g., `dust.ess.uci.edu`) so these directories may be accessed identically on `krein` and on local machine

1. Disk `dks9d2s0` is for match runs `/BIAN/match`,
2. Disk `dks9d3s0` is for match runs `/ZENDER/match`
3. Disk `dks9d4s0` is for NCEP data `/DSS`
4. Disk `dks9d5s0` is for boundary data `$/DATA/data`, diagnostics `$/DATA/dgn`, `$/DATA/aca`, `$/DATA/map`, and csm runs `/ZENDER/csm` and all run directories `/data/zender/match`, `/data/zender/csm`

```
mkdir -p /data/zender/ZENDER
mkdir -p /dks9d2s0/BIAN/match
mkdir -p /dks9d3s0/ZENDER/match
mkdir -p /dks9d4s0/DSS
mkdir -p /dks9d5s0/ZENDER/csm
mkdir -p /dks9d5s0/aca
mkdir -p /dks9d5s0/csm
mkdir -p /dks9d5s0/data
mkdir -p /dks9d5s0/dgn
mkdir -p /dks9d5s0/match
mkdir -p /dks9d5s0/tmp
ln -s /dks9d2s0/BIAN/match /data/zender/BIAN/match
ln -s /dks9d3s0/ZENDER/match /data/zender/ZENDER/match
ln -s /dks9d4s0/DSS /data/zender/DSS
ln -s /dks9d5s0/ZENDER/csm /data/zender/ZENDER/csm
ln -s /dks9d5s0/aca /data/zender/aca
ln -s /dks9d5s0/csm /data/zender/csm
ln -s /dks9d5s0/data /data/zender/data
ln -s /dks9d5s0/dgn /data/zender/dgn
ln -s /dks9d5s0/match /data/zender/match
ln -s /dks9d5s0/tmp /data/zender/tmp
```

9.21 Install sudo

As root, install `sudo`

```
scp dust.ess.uci.edu:/etc/sudoers /etc/sudoers
scp dust.ess.uci.edu:/usr/bin/sudo /usr/bin/sudo
```

The `sudo` command has some subtle problems. First, `sudo` does not alter `HOME`. To execute commands with `HOME` changed to the target user, i.e., root, invoke `sudo` with `-H /root` but does not change

```
sudo -H foo # Change ${HOME} to /root
```

Second, `sudo` carries the user's environment with it. Often it is desirable to see a command executed exactly as it would be if it were being executed from a root login shell.

```
# Initialize $HOME, $SHELL, $USER, $LOGNAME, and $PATH
# Un-set all other variables
sudo -i foo # Run command as in initial login shell
```

Some GNU/Linux distributions, such as Ubuntu ¹, do not enable the root account by default. Instead, Ubuntu expects the first user created to use `sudo` to perform administrative tasks. In such situations, it may be useful to create a shell with root privileges

```
sudo -s # Create shell with root priveleges
```

Another option is to enable the root account

```
sudo passwd root # Enable root account
sudo passwd -l root # Lock root account after enabling it
```

hjm - 08-18-05 I've had a continuing problem with sand periodically complaining about my `.ICEauthority` and `.Xauthority` files being screwed up and then refusing to let me log in on the console. This may not affect you too much as you probably never use GUI tools to admin the system, but I do and this is the reason.

Since sand has no root user, it manipulates the user's `.Xauthority` to enable X to open displays. That changes the permission of the file and thus prohibits a console login because `xdm` needs to write to that file.

So if you use an X app as root, you will need to explicitly change the permissions on these files before you will be able to log in.

I've added the following into my `.alias` file which fixes (but does not solve) the problem:
`alias chice="sudo chown hjm.cgdcsm /home/hjm/.ICEauthority /home/hjm/.Xauthority"`

9.22 Install IDL and PGI

Intel recently purchased Kai software and is now distributing their compilers for Linux.

```
http://www.intel.com/software/products/compilers/flin/noncom.htm
http://www.coyotegulch.com/hpc/intel_1st_look.html
```

Select the "non-commercial unsupported software" for Linux, not the free evaluation software. The Intel compilers are OpenMP-compliant. The Fortran and C+ compiler work fine on RedHat 7.2. Only the OpenMP features of the C+ compiler seems not to work. Serial numbers of `icc` and `ifc` are 1110-19809410 and 1130-70846464, respectively. Intel compiler support questions go to pto.support@intel.com.

We bought Lahey 6.1 Pro fortran compiler serial number LP072528. Technical support is support@lahey.com.

My PGI PINs are 109584 (dust), 502421 (seasalt), and 508553 (sand). Download releases from PGI website at <http://www.pgroup.com/downloads>. Generate permanent licenses keys using personalized account on the PGI website at <http://www.pgroup.com/pgilogin.htm> Username `pn109584` and Password `Ma11!21` (dust). Username `pn502421` and Password `Teen:28` (seasalt), Username `pn508553` and Password `Jane&40` (sand). Updating PGI license keys. To learn the FLEXlm-style hostid of the system, execute

¹From the Ubuntu website: "Ubuntu" is an ancient African word, meaning "humanity to others". Ubuntu also means "I am what I am because of who we all are". Ubuntu means "towards humanity" in Bemba.

```

${PGI}/linux86/bin/lmutil lmhostid
${PGI}/linux86-64/6.0/bin/lmutil lmhostid

```

For lanina, this results in 0010a4081231 For seasalt, this results in 0001031c5c7f For seasalt, this results in 00d0680399a4 PGI-style hostids do not require a license daemon, but are locked to the username that installed the compilers, and to the hardware configuration in existence when the compilers were installed. The PGI-style hostid can be found by running the command

```

${PGI}/linux86/bin/pghostid
unset LM_LICENSE_FILE

```

The host ID must match that in the license file.

To copy files from dust to other machines, e.g., lanina:

```

tar cvzf /data/zender/tmp/pgi.tar.gz ./pgi
tar cvzf /data/zender/tmp/rsi.tar.gz ./rsi
sudo scp dust.ess.uci.edu:/data/zender/tmp/pgi.tar.gz /usr/local
sudo scp dust.ess.uci.edu:/data/zender/tmp/rsi.tar.gz /usr/local
cd /usr/local
sudo tar xvzf pgi.tar.gz
sudo tar xvzf rsi.tar.gz
sudo ln -s ./rsi/idl_5.6 ./rsi/idl
/bin/rm ./rsi/idl_linux.tar.gz

```

Add the following to `.bashrc`:

```

# Licenses for IDL and PGI are handled by FlexLM
# FlexLM is started in /etc/rc.d/rc.local
# Running lmgrd as root is not recommended but is expedient
# IDL lmgrd starts vendor-specific license daemons for both IDL and PGI
# This is because contents of PGI license file are in IDL license file
# Users must have LM_LICENSE_FILE environment variable point to both licenses
case ${HOST} in
    dust* )
        export LM_LICENSE_FILE="${IDL_DIR}/../license/license.dat:${PGI}/li
        ;; # endif using licenses directly on license server
    dakine* | lanina* | seasalt* | dhcp* | ras* | cgd85* )
# Randerson's server
        export LM_LICENSE_FILE="${PGI}/license.dat:1700@ess1.ess.uci.edu"
# Zender's server
        export LM_LICENSE_FILE="${PGI}/license.dat:1700@dust.ess.uci.edu"
        ;; # endif using network licenses
esac # endcase ${HOST}

```

PGI Fortran may not run on lanina unless `LM_LICENSE_FILE` is undefined with `unset LM_LICENSE_FILE`. This is because any licenses specified in `LM_LICENSE_FILE` appear to override searching for

PGI-style node-locked licenses first. Since `LM_LICENSE_FILE` must be defined to allow network access to the IDL software, this means it will be difficult if not impossible to get PGI and IDL both working on `lanina` at the same time.

To remove a checked out license use the `lmremove` option to the `lmutil` command.

```
lmutil lmremove [-c license_file_list] feature user user_host display
lmutil lmremove idl zender lanina.ess.uci.edu /dev/pts/0
```

The argument to `lmremove` were obtained from the `lmstat -a` command, which is now obsolete. Instead give sub-commands arguments directly to `lmutil`, e.g.,

```
lmutil lmstat
```

9.22.1 PGI on sand

So that any system user may run the software, we create a FlexLM license daemon `/etc/init.d/pgi-lmgrd`. A corresponding user, `flexlm`, runs this daemon.

```
/usr/local/pgi-5.1-3 # Western Scientific installation
/usr/local/src/ # Source tarball
/etc/init.d/pgi-lmgrd # daemon to start FlexLM
/usr/local/pgi-5.1-3/license.dat # License
/var/tmp/flexlm.log # License daemon error messages
```

The server's hostname configuration is very important. When the hostname changes then update the license file accordingly.

9.23 PCMCIA

```
cd /usr/src
tar xvzf pcmcia-cs-3_x_x.tar.gz
cd /usr/src/linux/pcmcia-cs-3.x.x
make config
make all
make install
```

Customize in `/etc/pcmcia` if necessary. For RedHat systems, read Section 2.5.2 (p. 12) of PCMCIA HOWTO

As recommended, delete (or comment out) contents of default `/etc/pcmcia/network.opts`

and replace with script given in HOWTO

This fix enables ethernet connections on power-up

Alternatively, I modified `/etc/pcmcia/network` with two suggestions from PCMCIA Xircom Hypermail list, but, as stated in the HOWTO,

RedHat may not actually run this script.

JWZ has similiar card and uses three commands

```
usernetctl eth0 down
/etc/rc.d/init.d/pcmcia restart
usernetctl eth0 up
```

9.24 Building new kernel

Before building kernel make sure default compiler is known to build stable kernels. For RedHat 7.0 systems use `kgcc` not `gcc`. New kernel (2.2.18+) Makefiles do this automatically, but just in case do this as root `export CC=/usr/bin/kgcc`. For guaranteed kernel stability, GCC version 2.95 is recommended. `export CC=/usr/bin/gcc-2.95`. Finally, the `System.map` file must be copied along with the kernel.

```
cd /usr/src/linux
# Choose one of the following:
cp cnf_fl /usr/src/linux/.config # Use archived config file
make clean (remove old object files)
make oldconfig # Use .config file, or
make config # Create/modify .config file (text), or
make menuconfig # Create/modify .config file (curses), or
make xconfig # Create/modify .config file (Xwindows)
make mrproper # Remove .config (CAREFUL!!!), clean up disk
# Enable: APM,ACPI,ISA,LVM,UDF,USB
# DHCP requires CONFIG_PACKET and CONFIG_FILTER
# USB requires fxm
# DVD and CD-ROM requires UDF filesystem
# Following steps are mandatory
make dep # Dependency check
make # Create bzImage (NB: bzImage, not zImage)
make modules
make modules_install
make install # Copy bzImage to boot. Rename it intelligently.
# This automagically does the following three steps
cp /usr/src/kernel-source-2.4.23/arch/i386/boot/bzImage /boot/vmlinuz-2.4.23
cp /usr/src/linux/arch/i386/boot/bzImage /boot/vmlinuz-2.4.23
cp /usr/src/linux/System.map /boot/System.map-2.4.23
make clean (remove old object files)
rm /boot/vmlinuz
rm /boot/System.map
ln -s /boot/vmlinuz-2.4.23 /boot/vmlinuz
ln -s /boot/System.map-2.4.23 /boot/System.map
/sbin/lilo
shutdown -r now
# Exit X the rebuild NVidia driver
cd /usr/src;sh NVIDIA-Linux-x86-1.0-4496-pkg2.run
# Get 2.6 kernel patches for NVidia drivers
```

`http://www.minion.de/files`

I like the `make xconfig` method. It is very clean and allows storing and retrieving configuration files. By default, the configuration file is saved as `.config` in the top-level make directory. Save a visible version of the configuration file as, e.g., `/home/zender/linux/usr/src/linux/config.lanina` or `config.lanina.2.4.23.20010322` and then `cp /usr/src/linux/config.lanina.2.4.23`. Specific configuration options which are not the default must be set.

1. Block Devices: `CONFIG_BLK_DEV_IDEFLOPPY` for LS-120/Iomega Zip support.
2. Sound: `CONFIG_SOUND_MAESTRO` for Lanina soundcard
3. Networking Options: for IEEE 1394/Firewire
4. IrDA (infrared) support: `CONFIG_IRDA`
5. USB support: Various. Important for future goodies like mouse, joystick...
6. Character devices: `CONFIG_PRINTER`
7. Filesystems: `CONFIG_FAT_FS`, `CONFIG_VFAT_FS` for MS Windows filesystem support

The `sysctl` command allows one to print (and set) kernel parameters at runtime. The `Procfs` file system which displays `/proc/sys` is required for `sysctl` to work.

```
sysctl -a # Print all kernel configuration values in /proc/sys
```

9.25 LILO configuration

The LILO (linux loader) system is a venerable method of loading Linux which has lately been superseded by GRUB. One problem with the LILO method is that the `lilo` command must be run *after* installing a kernel and *prior to* rebooting the machine. If it is necessary to rescue a machine that boots with LILO, one must construct rather complex `lilo` commands to synchronize the system.

```
# Find major, minor device numbers on other RedHat machine
ls -l /dev/sda1
ls -l /dev/sda3
# Create devices with correct numbers
mknod /dev/sda1 b 8 1
mknod /dev/sda3 b 8 3
# Create mountpoint directory
mkdir /mnt/root
# Mount disks with root and kernel images
mount /dev/sda3 /mnt/root
mount /dev/sda1 /mnt/root/boot
# Change root directory to simplify paths for lilo
chroot /mnt/root # Change root from RedHat rescue to disk
```

```
chroot /mnt/sdb3 /sh # Change root from Knoppix CDROM to disk
cd /;ls # Verify we've changed root
/sbin/lilo # Simplest form
# Unless chroot succeeded, lilo may need following arguments
# -b: Location of boot device
/sbin/lilo -b /mnt/sdb1 #
# -C: Location of configuration file (default /etc/lilo.conf)
/sbin/lilo -C /mnt/sdb3/etc/lilo.conf
# -i: Location of bootloader (default /boot/boot.b)
/sbin/lilo -i /mnt/sdb1/boot.b
```

When updating the kernel, make the old, working image available as a backup.

```
image=/boot/vmlinuz-2.4.23
    label=linux
image=/boot/vmlinuz-2.2.16-22
    label=linux-2.2.16-22
```

`lilo.conf` must specify the `linear` keyword to boot off a SCSI disk (e.g., `dirt.ess.uci.edu`).

9.26 Restarting daemons

Send HUP (hangup) signal to process:

```
kill -HUP pid
```

Restart daemon manually:

```
/etc/rc.d/init.d/inet restart
/etc/rc.d/init.d/network restart
```

Use alias:

```
restart inet
```

9.27 Monitor Kernel Startup

Trace with the kernel startup with `dmesg`. This command prints `/var/log/dmesg`.

```
dmesg > foo
```

9.28 Strip downloaded text files of DOS linefeeds

There are three different systems used for representing the end of the line in text files. MS Windows uses ASCII CR/LF, Macintosh uses CR, and Unix uses LF as end of line characters. To convert from Windows to Unix, use the `tr` command to strip the file of the excess carriage returns:

```
# Strip excess carriage returns from Windows file
tr -d '\015' < fl_in > fl_out
```

To convert from Unix to Windows, insert carriage returns in front of line feeds:

```
tr ??? < fl_in > fl_out
```

FTP servers attempt to handle these translations automatically when text mode is selected.

The end-of-line convention is part of what is known in Emacs as the encoding. Files written with the DOS encoding convention may be reset to the Unix convention within Emacs using

```
C-x RET f undecided-unix
```

9.29 Installing Debian

Useful mailing list debian-user@lists.debian.org <http://www.debian.org/MailingLists>

A few packages cause many warnings when upgrading Debian. These include `kbuildsyscoa` and `kio`. What causes these warnings?

```
kio (KService*): WARNING: The desktop entry file Utilities/kfilereplacepart
k
```

```
kbuildsyscoa: WARNING: 'kfile_koffice.desktop' specifies undefined mimetype.
```

```
# Following procedure was tested on ashes.ess.uci.edu 20030508, neige 20080
# Bring linux.tex over first to revise procedure as needed
scp dust.ess.uci.edu:/home/zender/linux/linux.tex ~
# Upgrade base installation
aptitude update; aptitude dist-upgrade
# Install packages which make installations easier
aptitude install apt-file cvs sudo subversion openssh-server
aptitude install kubuntu-desktop
scp dust.ess.uci.edu:/etc/sudoers /etc/sudoers
sudo groupadd -g 2400 cgdcsm # Create group cgdcsm
sudo useradd -D -g cgdcsm # Make group cgdcsm default
sudo usermod -g cgdcsm -u 3555 zender # Change UID and default GID of user
# If NFS-mounting home directory from dust, do this:
# sudo usermod -d /dhome/zender zender
# Log out then log back in so UID, GID will be consistent
chgrp -R cgdcsm /home/zender
export CVSROOT=:ext:zender@esmf.ess.uci.edu:/u/zender/cvs
export CVSUMASK=002
export CVS_RSH=ssh
sudo slocate -u
cvs co -kk dot elisp linux sh # Retrieve dot files, emacs, Linux Cheatsheet
# Use KDE control center to left-hand mouse, auto-raise
xmodmap ~/dot/xmodmaprc # Swap ctl and caps-lock
export PVM_ARCH=~ /sh/pvmgetarch `
cd ~;mkdir -p bin/${PVM_ARCH} include lib/${PVM_ARCH} obj/${PVM_ARCH} crm ma
# Ubuntu: First activate 'universe' sources in sources.list
```



```

sudo aptitude install auctex metamail preview-latex tramp # Install Emacs and
sudo ln -s /usr/bin/mimencode /usr/local/bin/mmencode # Rmail uses mmencode
emacs & # Edit dot files to contain new system name
cd dot;make dot; . bashrc;cd # Install dot files, load custom bash settings
ssh-keygen -t rsa1 # Generate RSA1 key
ssh-keygen -t rsa # Generate RSA key
ssh-keygen -t dsa # Generate DSA key
cd ~/.dot;cp authorized_keys authorized_keys2 known_hosts known_hosts2 ~/.ssh/
# Edit known_hosts and authorized_keys
# Insert ~/.ssh/id_dsa.pub, ~/.ssh/id_rsa.pub into authorized_keys2
# SSH to remote machine, copy value from ~/.ssh/known_hosts* on remote
# machine into known_hosts*. Or take values directly from /etc/ssh/ssh_host*
cd ~/.dot;make ssh;cd
# Directories for personal machines
cvs -d :ext:zender@esmf.ess.uci.edu:/u/zender/cvs co -kk \
aca aereoce aeronet afgl anl anv arese avhrr bxm c c++ cld clm crr \
dmr dst esmf \
ess ess_acc ess_atm ess_bnd ess_ccc ess_ccp ess_lsp ess_phz ess_prc ess_rdn
f fsf grd hdf hire idea idl idx_rfr igbp igpp improve job jrn lsm ltr \
map match matlab mie mk mny ncep ncl perl phd poetry pr prp rvw sdn \
slr_spc tex time toms uci www
# Directories for computational machines
cvs -d :ext:zender@esmf.ess.uci.edu:/u/zender/cvs co -kk \
aca acia afgl anl c c++ cld clm crr dst f fsf icr idx_rfr map mie mk \
mtrdst ncl perl slr_spc time
# Papers
cvs -d :ext:zender@esmf.ess.uci.edu:/u/zender/cvs co -kk \
ppr_BiZ03 ppr_BiZ04 ppr_GrZ04 ppr_ZeK05 \
ppr_ZeM07 ppr_Zen08 ppr_ZMT04 ppr_ZeT06b ppr_ZeT06 sltsbl
cvs -d :ext:charlesz@hox.uio.no:/mn/hox/dl/alfgr/CVS co -kk ppr_GMZ05
# Mark's papers
cvs -d :ext:biogenic.ess.uci.edu:/home/mflanner/cvs co -kk ppr_FlZ06
cd ~/tex;scp biogenic.ess.uci.edu:/home/mflanner/tex/mflanner.bib .
cd ~/ppr_FlZ05;scp biogenic.ess.uci.edu:/home/mflanner/ppr_FlZ05/fgr*.eps .
cvs -d esmf.ess.uci.edu:/u/zender/cvs co -kk -r prp_itr -d prp_itr prp_arl
cvs -d esmf.ess.uci.edu:/u/zender/cvs co -kk prp_ans prp_ids
# Other Proposals
cvs -d :ext:visa.eng.uci.edu:/home/cvs-user co -kk -d prp_JZK05 SEIII05
cvs -z3 -d :ext:zender@nco.cvs.sourceforge.net:/cvsroot/nco co -kk nco
svn checkout https://swamp.googlecode.com/svn/trunk/ swamp --username charl
cvs -d :ext:zender@goldhill.cgd.ucar.edu:/fs/cgd/csm/models/CVS.REPOS \
co -r ccm_brnch_dst -kk ccm_dst
cvs -d :ext:zender@esmf.ess.uci.edu:/u/zender/cvs co -kk -r match_brnch_dst
cd crm;cvs -d :ext:zender@goldhill.cgd.ucar.edu:/fs/cgd/csm/models/CVS.REPOS
co -r ccm_brnch_crm -kk crm

```

```

sudo mkdir /data;cd /data;sudo mkdir zender;sudo chown zender zender;sudo ch
cd /data/zender;mkdir -p mie ps tmp # Create data directories which may be e
# Hardware monitoring
sudo aptitude install xosview hddtemp gkrellm
# Security
sudo aptitude install chkrootkit clamav rkhunter
sudo chkrootkit
sudo rkhunter -c
# Packages for building NCO .debs
sudo aptitude install antlr autoconf autoconf-doc automake1.9 \
binutils-multiarch bison ccache \
debhelper debootstrap devscripts dput debian-policy dupload \
figlet fakeroot flex gnuplot libantlr-dev libtool linda lintian \
pbuilder reportbug wdiff
sudo aptitude install manpages-dev glibc-doc
sudo aptitude remove gcc-3.3 g++-3.3 libstdc++5-3.3-dev
sudo aptitude remove gcc-4.0 g++-4.0 libstdc++6-4.0-dev gfortran-4.0 libstdc
sudo aptitude install gcc gcc-4.2 gcc-doc gcc-4.2-locales
sudo aptitude install g++ g++-4.2 libstdc++6-4.2-doc
sudo aptitude install gfortran-4.1 # Needed for NCL/NCAR Graphics
sudo aptitude install gfortran gfortran-4.2 gfortran-doc
sudo aptitude install gij gcj
# If necessary, link default executables to appropriate Ubuntu binaries
sudo ln -s -f /usr/bin/gcov-4.2 /usr/bin/gcov
sudo ln -s -f /usr/bin/cpp-4.2 /usr/bin/cpp
sudo ln -s -f /usr/bin/g++-4.2 /usr/bin/g++
sudo ln -s -f /usr/bin/gcc-4.2 /usr/bin/gcc
sudo ln -s -f /usr/bin/gfortran-4.2 /usr/bin/gfortran
# MPI: Choose one MPI distribution: MPICH2 (recommended), MPICH, or LAM-MPI
scp 'dust.ess.uci.edu:/data/zender/tmp/*mpich2*' ${DATA}/tmp
cd ${DATA}/tmp;sudo dpkg --install mpich2-doc_1.0.2-3_i386.deb mpich2_1.0.2
sudo aptitude install lam4c2 lam4-dev lam-mpidoc lam-runtime
sudo aptitude install mpich mpich-common mpi-doc
# Install GSL
sudo yum install rpmlint
sudo yum install tetex tetex-fonts tetex-dvips texinfo-tex
sudo yum install automake autoconf bison curl-devel flex librx librx-devel
sudo yum install antlr antlr-c++-devel antlr-manual gsl gsl-devel netcdf net
sudo yum install libdap libdap-devel libnc-dap libnc-dap-devel
sudo aptitude install gsl-bin libgsl0-dev
scp dust.ess.uci.edu:/data/zender/tmp/lf9562.tar.gz ${DATA}/tmp # Install L
# Install netcdf by hand (uses Fortran90 build, finishes C++ build)
sudo aptitude install netcdf-perl libxml-simple-perl
# Packages Opendap requires:
sudo aptitude install libcurl3-dev libxml2-dev

```

```

# Install nr by hand (needed by fff)
# Install specfun (needed by fff)
# Install udunits (needed by NCO)
cd ~/c;make makdep # Build makdep only
cd ~/nco/bld;make # Build NCO module (needed by mie)
cd ~/nco/src/nco_c++;make -f Makefile.old inc;cd ~ # Build libnco_c++ (needed by mie)
cd ~/c++;make # Build c++ module (including getopt_bsd)
cd ~/c;make # Build c module
cd ~/mie;make # Build mie module
cd ~/f;make # Build f module
cd ~/dead;make # Build dead module
# Copy important data directories wholesale
for drc in aca arese avhrr cccac data dst fgr hitran lsmdata map mny no2 pp2;do
  rsync -av dust.ess.uci.edu:/data/zender/${drc} ${DATA}
  printf "Copied dust.ess.uci.edu:/data/zender/${drc}\n"
done
# Mount /dev/sg0 as /cdrom to enable CD-burning (done in fstab)
# Install packages not on default Debian unstable install
# Printing, web, A/V-related
sudo aptitude install \
cupsys cupsys-client cupsys-driver-gimpprint \
gimp foomatic-filters-ppds foomatic-filters imagemagick \
kaffe pump sane traceroute vncserver xvncviewer xrestop xsane
# Debugging
sudo aptitude install \
ddd ddd-doc gdb graphviz valgrind valgrind-callgrind \
graphviz kcachegrind kcachegrind-converters \
# TeX, document-related
sudo aptitude install bibtool bibtex2html cmatrix-xfont dvipng emacs-intl-french
# Video, sound-related:
sudo aptitude install ubuntu-restricted-extras # includes flashplugin-nonfree
sudo aptitude install regionset kubuntu-restricted-extras
sudo aptitude install libdvdcss2 # this comes from medibuntu
# sudo /usr/share/doc/libdvdread3/install-css.sh # redundant with libdvdcss2
# 20070819: NB: Flash differs from Shockwave. There is no Linux Shockwave plugin
sudo aptitude install flashplugin-nonfree libmad0 totem-xine
# Not available in Ubuntu
sudo aptitude install foiltex latex2html tth
# Install audio player software
sudo aptitude install gtkpod xmms
# Install TightVNC for webcasts
sudo aptitude install tightvnc-java tightvncserver xtightvncviewer
# Install acroread (not necessarily recommended)
sudo aptitude install acroread acroread-plugins
sudo aptitude install xine-ui

```

```

xine # Watch DVDs
# Get LaTeX working
scp -r dust.ess.uci.edu:tex/cls ~/tex
scp -r dust.ess.uci.edu:/usr/share/texmf/tex/latex/ifdraft .;sudo mv ifdraft
scp -r dust.ess.uci.edu:/usr/share/texmf/tex/latex/movie15 .;sudo mv movie15
scp -r dust.ess.uci.edu:/usr/share/texmf/tex/latex/datetime .;sudo mv datet
scp -r dust.ess.uci.edu:/usr/share/texmf/tex/latex/fmtcount .;sudo mv fmtco
scp -r dust.ess.uci.edu:/usr/share/texmf/tex/latex/ifsym .;sudo mv ifsym /u
scp -r dust.ess.uci.edu:/usr/share/texmf/tex/latex/revnum .;sudo mv revnum
scp -r 'dust.ess.uci.edu:/data/zender/ps/erbe_b_sld012d_8589_0[17]_x_[LS]WCI
scp dust.ess.uci.edu:/data/zender/ps/uci_fax.pdf dust.ess.uci.edu:/data/zend
scp -r dust.ess.uci.edu:tex/bst ~/tex;
cd ~/tex;sudo mkdir /usr/share/texmf-texlive/bibtex/bst/csz;sudo cp bst/* /u
# sudo chmod a+w /var/cache/fonts/ls-R
# sudo chmod 666 /dev/dsp # Sound device mode 660 causes permissions errors
# Make sure /cdrom is mount point for /dev/sg0 in /etc/fstab
# Get hardware tools (USB mouse, reiserfs utilities)
sudo aptitude install gpm lm-sensors smartmontools xfonts-artwiz # Get narr
cd ${DATA};scp -r dust.ess.uci.edu:/data/zender/map . # Allow bds to run
# Create links to directories in ${DATA}
sudo mkdir -p /fs/cgd/home0;sudo ln -s /home/zender /fs/cgd/home0/zender
sudo mkdir -p /fs/cgd/data0;sudo ln -s /data/zender /fs/cgd/data0/zender
mkdir -p ${DATA}/csm/inputdata/lnd/clm2 ${DATA}/csm/input
sudo mkdir -p /fs/cgd;sudo ln -s /data/zender/csm /fs/cgd/csm
cd ${DATA}/csm/inputdata/lnd/clm2;scp -r dust.ess.uci.edu:${DATA}/csm/input
cd ${DATA}/csm/input;scp -r dust.ess.uci.edu:${DATA}/csm/input/atm dust.ess
sudo mkdir -p /datashare/inputdata/csm
sudo ln -s ${DATA}/csm/inputdata/lnd /datashare/inputdata/csm/lnd
# Remove automatic gdm/kdm/xdm
cd /etc/init.d;mkdir unused;mv gdm unused
# Install ccache as default compiler
sudo cp /usr/bin/ccache /usr/local/bin
sudo ln -s /usr/local/bin/ccache /usr/local/bin/gcc
sudo ln -s /usr/local/bin/ccache /usr/local/bin/g++
sudo ln -s /usr/local/bin/ccache /usr/local/bin/icc
sudo ln -s /usr/local/bin/ccache /usr/local/bin/icpc

```

Commands to set up a basic scientific computing environment

```

mkdir -p bin
/bin/cp -r ~/zender/sh ~
export PVM_ARCH='~/sh/pvmgetarch`
mkdir -p ~/bin/${PVM_ARCH} ~/include ~/lib/${PVM_ARCH} ~/obj/${PVM_ARCH}
cp ~/zender/bin/${PVM_ARCH}/makdep ~/bin/${PVM_ARCH}
cp ~/zender/sh/pvmgetarch ~/sh

```

```
mv ~/.bashrc ~/.bashrc.old
cp ~zender/.bashrc ~
```

9.30 Recreating Server Files

Server port to leave open when registering over UCInet.

```
29659 # NCO
1700 # PGI, IDL license servers
7143 # PathScale licence server on IPCC
```

When the system goes down and takes the /etc partition with it, the system services must be re-installed from scratch.

```
#sudo mv /home/backup/etc/passwd /etc/passwd
ls /home/backup
```

```
# Reset SSH keys to ease work on remote machine
```

```
# Reinstall sudo permissions
```

```
sudo scp biogenic.ess.uci.edu:/etc/sudoers dust.ess.uci.edu:/etc/sudoers
```

```
# Restore accounts (RedHat Linux: add -n to turn off creation of group name)
```

```
sudo groupadd -g 2400 cgdcsm # Create group cgdcsm
sudo groupadd -g 1965 esmfadm # Create group esmfadm
sudo useradd -c 'Charlie Zender' -d /home/zender -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Mark Flanner' -d /home/mflanner -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Eun Young Kwon' -d /home/kwon -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Huisheng Bian' -d /home/bian -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Alf Grini' -d /home/alfgr -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Dave Newman' -d /home/newman -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Sarah Bortz' -d /home/bortz -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Renato Pajarola' -d /home/pajarola -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Jorge Talamantes' -d /home/jtalaman -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Dan Stromberg' -d /home/strombrg -g esmfadm -G users -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Francisco Lopez' -d /home/lopez -g esmfadm -G users -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Greg Okin' -d /home/okin -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Harry Mangalam' -d /home/hjm -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Test User' -d /home/test -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'ESMF System Administrator' -d /home/esmfadm -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'CVS' -M -g cgdcsm -u 3562 cvspub
sudo useradd -c 'ESS 211 Class Account' -d /home/ess211 -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Hsun-Ying Kao' -d /home/hkao -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Gayathri Venkitachalam' -d /home/gvenkita -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Chao Luo' -d /home/chaoluo -g cgdcsm -m -p '$1$charlie$VJFpvyB'
sudo useradd -c 'Qin Han' -d /home/qhan -g cgdcsm -m -p '$1$charlie$VJFpvyB'
```

```

sudo useradd -c 'Scott Capps' -d /home/scapps -g cgdcsm -m -p '$1$charlie$VJ
# Oops, got to change Robynn's UID
sudo useradd -c 'Robynn Zender' -d /home/robynn -g cgdcsm -m -p '$1$charlie$VJ
sudo useradd -c 'Mike Brown' -d /home/brownmc -g cgdcsm -m -p '$1$charlie$VJ
sudo useradd -c 'Danielle Johnson' -d /home/daniellj -g cgdcsm -m -p '$1$charlie$VJ
sudo useradd -c 'Mike Tosca' -d /home/mtosca -g cgdcsm -m -p '$1$charlie$VJ
sudo useradd -c 'Henry Butowsky' -d /home/hbutowsk -g cgdcsm -m -p '$1$charlie$VJ
sudo useradd -c 'Olivia Zender' -d /home/olivia -g cgdcsm -m -p '$1$charlie$VJ
sudo useradd -c 'Daniel Wang' -d /home/wangd -g cgdcsm -m -p '$1$charlie$VJ

# Create work directories and change permissions
for usr in mtosca; do
    sudo chmod 755 /home/${usr}
    sudo mkdir /data/${usr};cd /data;sudo chown ${usr} ${usr};sudo chgrp cgdcsm ${usr}
    sudo mkdir /var/www/html/${usr};cd /var/www/html;sudo chown ${usr} ${usr}
done

# Patch RPMs

# Restore /usr/local (includes NCAR graphics, Lahey fortran, PGI, RSI)
cd /usr;tar cvzf local.tar.gz ./local
sudo scp local.tar.gz dust.ess.uci.edu:/usr
cd /usr;tar xvzf local.tar.gz

# Restore CVS
sudo mkdir /var/lock/cvs
sudo chmod 777 /var/lock/cvs
sudo scp ~/linux/etc/xinetd.d/cvsserver /etc/xinetd.d
sudo /etc/rc.d/init.d/xinetd restart

# Restore WWW
#sudo /bin/rm -r /var/www/html
#sudo scp -r /var/www/html dust.ess.uci.edu:/var/www
sudo scp ~/linux/etc/apache/httpd.conf dust.ess.uci.edu:/etc/httpd/conf/httpd.conf
sudo chown -R zender *
sudo chgrp -R cgdcsm *
sudo /etc/init.d/httpd restart
# Rebuild contents: doc, ppr, facts, bxm, nco,

# Restore BXM
cd ~/bxm;sudo make bxm;cd -
cd ~/bxm;sudo make bxm_aer;cd -
sudo mkdir -p /var/ftp/dead
sudo chown apache /var/ftp/dead
sudo chgrp apache /var/ftp/dead

```

```
# Restore NFS Crossmounts
sudo /etc/init.d/nfs restart
sudo /usr/sbin/exportfs -rv
sudo mount -a

# Restore Scientific computing
sudo rpm -i gsl-1.4-5mdk.src.rpm
netCDF

# Restore FTP
cd /var/ftp/pub
sudo mkdir zender
sudo chown zender zender
sudo chgrp cgdcsm zender
```

9.31 i8500s

Inspiron 8500

Use the `xev` program to find the keycodes for the volume keys. In my case they turned out to be 174 (lower) and 176 (raise). Ran the commands `xmodmap -e "keycode 174 = XF86AudioLowerVolume"` and `xmodmap -e "keycode 176 = XF86AudioRaiseVolume"` to map the buttons to the right key symbols. Configured the Sawfish window manager to run the command `aumix -c-5` on lower and `aumix -c+5` on raise. Tried it, and it worked. Put the commands into `.xsession` to make this permanent.

9.32 Images

Images refers to all still images, including JPEG, tiff, etc. The best program to use for viewing images depends on your intent. For image processing, use `gimp`. For slideshows of raster images, use `gimp`.

9.32.1 Gimp

The `gimp` program is excellent for viewing and editing images. However, it has now slideshow capabilities.

9.32.2 KView

9.32.3 Kuickshow

Kuickshow is a very fine tool with excellent preview and slideshow modes. In preview mode, just click on a filename and keep hitting Pg Dn to see new pictures. The slideshow mode has an appropriate delay between slides. Remember to de-activate power-saving and screen-blanking before expecting the slideshow to continue unattended.

9.33 CDs

Compact Disks (CDs) store up to 700 MB. Formatting and writing data to the CD, known as burning, is not always straightforward with Linux. The program `cdbakeoven` has an intuitive interface. It provides continuous real-time status reports during burning, and sometimes works when K3b fails. On `sand.ess.uci.edu`, K3b tends to fail yet `cdbakeoven` works. On `elnino.ess.uci.edu`, K3b tends to work and `cdbakeoven` fails.

The KDE CD/DVD-burning utility is called K3b. It automatically loads when it detects a blank CD in a writable drive. However, K3b does not have a completion meter and just hangs when it fails to burn CDs (at least under SuSE. This is annoying because one does not know whether the program is working, and how long until completion.

9.34 DVDs

I use `xine` and `mplayer` to play DVDs. Due to legal concerns, few Linux distributions automatically install DeCSS, the pre-requisite library for decrypting DVDs. Multimedia software works with media resource locators (MRLs). MRLs are similar to URLs with media-like protocols. Most audio/video players do not allow direct control of volume, perhaps because the computer speakers are a shared resource. Hence it is wise to start the audio controls before the video player. `Kmix` is a fine audio controller to use.

9.34.1 MPlayer

MPlayer accepts many of the same commands that work with Xine:

```
kmix & # Volume control
mplayer dvd://1 # Quickstart DVD playing
mplayer dvd://1 -dvdangle 2 # Multiangle DVD playing
mplayer http://mplayer.hq/example.avi # Stream from HTTP
mplayer rtsp://server.example.com/streamName # Stream using RTSP
mencoder "mf://*.jpg" -mf fps=25 -o output.avi -ovc divx4 # Encode all *.jpg
mencoder -tv driver=v4l:width=640:height=480 tv:// -o tv.avi -ovc raw # I
```

Mplayer has difficulty with full-screen mode.

9.34.2 dvdrip

```

sudo apt-get install lame
sudo apt-get install dvdrip vcdimager cdrdao subtitleripper
sudo apt-get install rar unrar
sudo ln -fs /usr/bin/rar /usr/bin/rar-2.80
rsync /data/zender/snd dust.ess.uci.edu:/data/zender
cd /data/zender;rsync dust.ess.uci.edu:/data/zender/snd .

```

9.34.3 Xine

Xine accepts many of the same commands that work with Mplayer:

```

kmix & # Volume control
xine dvd://1 # Quickstart DVD playing
xine dvd://1 -dvdangle 2 # Multiangle DVD playing
xine http://mplayer.hq/example.avi # Stream from HTTP
xine rtsp://server.example.com/streamName # Stream using RTSP

```

Xine provides trouble-free operation in full-screen mode.

```

/usr/share/doc/libdvdread2/examples/install-css.sh #
aptitude install xine-ui

```

Make sure the DVD is in the drive an mounted

```

sudo ln -s -f /dev/sr0 /dev/cdrom # ashes
sudo ln -s -f /dev/sr0 /dev/dvd # ashes
sudo mount -t iso9660 -o ro,user,noauto,unhide /dev/cdrom /cdrom
sudo mount -t udf -o ro,user,noauto,unhide /dev/dvd /dvd
sudo mount -t udf -o ro,user,noauto,unhide /dev/dvd /dvd
# Xine mounts DVDs like this on elnino:
sudo mount -t udf -o ro,noexec,nosuid,nodev /dev/scd0 /cdrom

```

xine generates lots of warning messages:

Run xine as root until permissions are straightened out.

```

libdvdcss error: failed opening raw device, continuing
main: probing <aadxr3> video output plugin
load_plugins: failed to load video output plugin <aadxr3>
main: probing <dxxr3> video output plugin
load_plugins: failed to load video output plugin <dxxr3>
main: probing <xv> video output plugin
main: probing <alsa> audio output plugin
load_plugins: failed to load audio output plugin <alsa>
main: probing <oss> audio output plugin
xine_interface: unknown param 10
xine_interface: unknown param 10

```

```
xine_interface: unknown param 10
xine_interface: unknown param 10
vo_scale: invalid ratio, using 4:3
vo_scale: unknown aspect ratio (0) in stream => using 4:3
load_plugins: plugin dxr3-spudec failed to instantiate itself.
libspudec:init_plugin called
liba52:No accelerated IMDCT transform found
```

9.35 Network

Linksys Etherfast 10/100 hub

netstat -rn: BuH98 p. 177

Destination 0.0.0.0 is default destination of all datagrams not explicitly routed elsewhere

Flags: U = Route is Up or Usable, G = Route uses a gateway, H = Route targets a host

Sometimes connectivity to a host or a subnet can be lost when a static route gets stuck in the routing table. This has happened twice in the past when system administrators log into lanina and manipulate the routing table. In such cases, the static route must be manually removed from the routing table as follows. Assuming 128.200.14.0 is unreachable,

```
zender@lanina:~$ /sbin/route -n
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref    Use Ifa
128.200.14.0    0.0.0.0         255.255.255.0  U        0      0      0 eth
zender@lanina:~$ sudo route del -net 128.200.14.0 netmask 255.255.255.0 eth
```

```
zender@dust:~/dot$ netstat -rn:
Kernel IP routing table
Destination      Gateway          Genmask         Flags  MSS Window  irtt Ifa
128.200.14.25    0.0.0.0         255.255.255.255 UH      0  0      0 etl
128.200.14.0     0.0.0.0         255.255.255.0  U        0  0      0 etl
127.0.0.0        0.0.0.0         255.0.0.0      U        0  0      0 lo
0.0.0.0          128.200.14.1    0.0.0.0        UG      0  0      0 etl
```

```
zender@lanina:~$ netstat -rn (lanina at UCI):
Kernel IP routing table
Destination      Gateway          Genmask         Flags  MSS Window  irtt Ifa
128.200.14.0    0.0.0.0         255.255.255.0  U        0  0      0 etl
192.168.82.0    0.0.0.0         255.255.255.0  U        0  0      0 vmi
127.0.0.0        0.0.0.0         255.0.0.0      U        0  0      0 lo
0.0.0.0          128.200.14.1    0.0.0.0        UG      0  0      0 etl
```

```
zender@lanina:~/dot$ netstat -rn (with dakine as PPP gateway to UCI)
```

Kernel IP routing table

Destination	Gateway	Genmask	Flags	MSS	Window	irrtt	Iface
192.168.1.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
127.0.0.0	0.0.0.0	255.0.0.0	U	0	0	0	lo
0.0.0.0	192.168.1.1	0.0.0.0	UG	0	0	0	eth0

zender@dakine:~/dot\$ netstat -rn (with dakine as PPP gateway to UCI)

Kernel IP routing table

Destination	Gateway	Genmask	Flags	MSS	Window	irrtt	Iface
192.168.1.1	0.0.0.0	255.255.255.255	UH	0	0	0	eth0
128.195.187.2	0.0.0.0	255.255.255.255	UH	0	0	0	ppp0
192.168.1.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
127.0.0.0	0.0.0.0	255.0.0.0	U	0	0	0	lo
0.0.0.0	128.195.187.2	0.0.0.0	UG	0	0	0	ppp0

zender@lanina:~/dot\$ netstat -rn (with lanina as PPP gateway to UCI)

Kernel IP routing table

Destination	Gateway	Genmask	Flags	MSS	Window	irrtt	Iface
128.195.186.2	0.0.0.0	255.255.255.255	UH	0	0	0	ppp0
192.168.1.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
127.0.0.0	0.0.0.0	255.0.0.0	U	0	0	0	lo
0.0.0.0	128.195.186.2	0.0.0.0	UG	0	0	0	ppp0

zender@dakine:~/dot\$ netstat -rn (with lanina as PPP gateway to UCI)

Kernel IP routing table

Destination	Gateway	Genmask	Flags	MSS	Window	irrtt	Iface
192.168.1.1	0.0.0.0	255.255.255.255	UH	0	0	0	eth0
192.168.1.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
127.0.0.0	0.0.0.0	255.0.0.0	U	0	0	0	lo
0.0.0.0	192.168.1.2	0.0.0.0	UG	0	0	0	eth0

zender@lanina:~\$ netstat -rn (from NCAR as dhcp25.cgd.ucar.edu)

Kernel IP routing table

Destination	Gateway	Genmask	Flags	MSS	Window	irrtt	Iface
128.117.0.0	0.0.0.0	255.255.0.0	U	0	0	0	eth0
127.0.0.0	0.0.0.0	255.0.0.0	U	0	0	0	lo
0.0.0.0	128.117.22.251	0.0.0.0	UG	0	0	0	eth0

zender@lanina:~/dot\$ netstat -rn (with lanina as PPP gateway to NCAR)

Kernel IP routing table

Destination	Gateway	Genmask	Flags	MSS	Window	irrtt	Iface
128.117.64.227	0.0.0.0	255.255.255.255	UH	0	0	0	ppp0
127.0.0.0	0.0.0.0	255.0.0.0	U	0	0	0	lo
0.0.0.0	128.117.64.227	0.0.0.0	UG	0	0	0	ppp0

ifconfig -a:

```

zender@ashes:~$ ifconfig -a (at UCI, wired working wireless not quite)
dummy0    Link encap:Ethernet  HWaddr 00:00:00:00:00:00
          BROADCAST NOARP  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)

eth0      Link encap:Ethernet  HWaddr 00:0B:DB:17:83:0D
          inet addr:128.200.14.90  Bcast:128.200.14.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:9379086 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1543592 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0
          RX bytes:1840313132 (1.7 GiB)  TX bytes:1113652359 (1.0 GiB)

eth1      Link encap:Ethernet  HWaddr 00:90:4B:B2:09:86
          BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:13954461 errors:0 dropped:0 overruns:0 frame:0
          TX packets:13954461 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0
          RX bytes:1488519751 (1.3 GiB)  TX bytes:1488519751 (1.3 GiB)

zender@dust:~/dot$ ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:C0:4F:60:73:8D
          inet addr:128.200.14.25  Bcast:128.200.14.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:5406688 errors:2103 dropped:0 overruns:0 frame:3179
          TX packets:1398865 errors:0 dropped:0 overruns:0 carrier:24
          collisions:62451 txqueuelen:100
          Interrupt:17 Base address:0xdc00

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:3924  Metric:1
          RX packets:59429 errors:0 dropped:0 overruns:0 frame:0

```

```

TX packets:59429 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
zender@lanina:~$ ifconfig -a (from UCI)
eth0      Link encap:Ethernet  HWaddr 00:10:A4:08:12:31
          inet addr:192.168.1.2  Bcast:192.168.1.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:674 errors:0 dropped:0 overruns:0 frame:0
          TX packets:541 errors:2 dropped:0 overruns:0 carrier:2
          collisions:73 txqueuelen:100
          Interrupt:11 Base address:0x200

eth0:1    Link encap:Ethernet  HWaddr 00:10:A4:08:12:31
          inet addr:128.200.14.80  Bcast:128.200.255.255  Mask:255.255.0.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          Interrupt:11 Base address:0x200

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:3924  Metric:1
          RX packets:45 errors:0 dropped:0 overruns:0 frame:0
          TX packets:45 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0

zender@lanina:~/dot$ ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:10:A4:08:12:31
          inet addr:192.168.1.2  Bcast:192.168.1.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:897 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1029 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          Interrupt:11 Base address:0x200

eth0:1    Link encap:Ethernet  HWaddr 00:10:A4:08:12:31
          inet addr:128.200.14.80  Bcast:128.200.255.255  Mask:255.255.0.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          Interrupt:11 Base address:0x200

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:3924  Metric:1
          RX packets:2982 errors:0 dropped:0 overruns:0 frame:0
          TX packets:2982 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0

zender@dakine:~$ ifconfig -a

```

```

eth0      Link encap:Ethernet  HWaddr 00:40:D0:0B:E9:E4
          inet addr:192.168.1.1  Bcast:192.168.1.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:2290 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1273 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          Interrupt:10 Base address:0x2000

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:3924  Metric:1
          RX packets:222 errors:0 dropped:0 overruns:0 frame:0
          TX packets:222 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0

ppp0      Link encap:Point-to-Point Protocol
          inet addr:128.195.187.37  P-t-P:128.195.187.2  Mask:255.255.255.255
          UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1500  Metric:1
          RX packets:216 errors:1 dropped:0 overruns:0 frame:1
          TX packets:221 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:10

zender@lanina:~$ ifconfig -a (from NCAR as dhcp25.cgd.ucar.edu)
eth0      Link encap:Ethernet  HWaddr 00:10:A4:08:12:31
          inet addr:128.117.22.248  Bcast:128.117.255.255  Mask:255.255.0.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:9198 errors:0 dropped:0 overruns:0 frame:0
          TX packets:529 errors:0 dropped:0 overruns:0 carrier:0
          collisions:28 txqueuelen:100
          Interrupt:11 Base address:0x200

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:3924  Metric:1
          RX packets:468 errors:0 dropped:0 overruns:0 frame:0
          TX packets:468 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0

zender@lanina:~/dot$ ifconfig -a (with lanina as PPP gateway to NCAR)
eth0      Link encap:Ethernet  HWaddr 00:10:A4:08:12:31
          inet addr:192.168.1.2  Bcast:192.168.1.255  Mask:255.255.255.0
          BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:12 dropped:0 overruns:0 carrier:12
          collisions:0 txqueuelen:100

```

```

Interrupt:11 Base address:0x200

lo      Link encap:Local Loopback
        inet addr:127.0.0.1  Mask:255.0.0.0
        UP LOOPBACK RUNNING  MTU:3924  Metric:1
        RX packets:120 errors:0 dropped:0 overruns:0 frame:0
        TX packets:120 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0

ppp0    Link encap:Point-to-Point Protocol
        inet addr:128.117.68.69  P-t-P:128.117.64.227  Mask:255.255.255.255
        UP POINTOPOINT RUNNING NOARP MULTICAST  MTU:1500  Metric:1
        RX packets:621 errors:4 dropped:0 overruns:0 frame:4
        TX packets:650 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:10

zender@haze: ifconfig -a (from UCI)
eth0    Link encap:Ethernet  HWaddr 00:10:A4:A0:9D:9D
        inet addr:128.200.14.62  Bcast:128.200.14.255  Mask:255.255.255.0
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:618084 errors:0 dropped:0 overruns:0 frame:0
        TX packets:416759 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:100
        Interrupt:11 Base address:0xe000

lo      Link encap:Local Loopback
        inet addr:127.0.0.1  Mask:255.0.0.0
        UP LOOPBACK RUNNING  MTU:16436  Metric:1
        RX packets:28 errors:0 dropped:0 overruns:0 frame:0
        TX packets:28 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0

```

9.36 Kernel Modules

Modules may be listed with the `lsmod` command, inserted with `insmod`, probed with `modprobe`, or removed with `rmmmod`.

```

lsmod # List loaded modules
insmod foo # Install loadable kernel module
modprobe ohci1394
modprobe raw1394

```

Occasionally it is necessary to insert standalone kernel modules.

```

# First bring down interface to avoid 'interface busy' errors
ifconfig eth0 down

```

```
# Remove current module from kernel
rmmod tulip
# Insert new module into kernel
insmod /home/zender/tulip.o
# Examine messages
tail -10 /var/log/messages
```

9.37 Change network topology

Linux Network (BuH98) p. 138, Linux Bible p. 374 Following scenarios include possibly superfluous commands to shut down `eth0:1`. If `eth0:1` is not shutdown and is running, route table may be bad. Machine which runs PPP (calls ISP) tries to set default route to `ppp0` interface. If default route (0.0.0.0) on PPP host is preset to `eth0`, PPP fails doing this Hence unset and preset `eth0` default route on PPP host before dialing

On RedHat systems, the key networking information is kept in `/etc/sysconfig/network` and in `/etc/sysconfig/network-scripts/ifcfg-eth0`. Things such as `GATEWAY`, `NETMASK`, must be set correctly in these files. These files are easiest to change by hand.

1. `lanina` is PPP dialin host and gateway for home LAN Execute following commands on `lanina`:

```
# Delete existing Ethernet interface(s), if any
/sbin/ifconfig eth0 down
/sbin/ifconfig eth0:1 down
/sbin/ifconfig eth0 lanina.zender.org # Connect IP address with Ethernet
# Add entry to kernel routing table which directs all datagrams to any l
/sbin/route add -net 192.168.1.0
# Route to gateway host all datagrams bound for hosts outside intranet
# fxm: "metric 1" argument appears necessary
/sbin/route add default gw 192.168.1.1 metric 1
# Execute following commands on dakine
/sbin/ifconfig eth0 down
/sbin/ifconfig eth0:1 down
/sbin/ifconfig eth0 dakine.zender.org # Connect IP address with Ethernet
# Ensure default route interface is not preset to eth0 before dialing
/sbin/route del default gw 192.168.1.1 metric 1
ppp-go
/etc/ppp/chain start # Turn on IP masquerading
```

2. `lanina` is PPP dialin host and gateway for home LAN

```
# Execute following commands on lanina
# Delete existing Ethernet interface(s), if any
hostname lanina.zender.org
hostname -i -v
/sbin/ifconfig eth0 down
```



```

/sbin/ifconfig eth0:1 down
# Ensure default route interface is not preset to eth0 before dialing
/sbin/route del default gw 192.168.1.1 metric 1
ppp-go
/sbin/ifconfig eth0 lanina.zender.org # Connect IP address with Ethernet
/etc/ppp/chain start # Turn on IP masquerading
# Execute following commands on dakine
# Point dakine to lanina gateway
/sbin/ifconfig eth0 down
/sbin/ifconfig eth0 dakine.zender.org
# Route to gateway host all datagrams bound for hosts outside intranet
# fxm: "metric 1" argument appears necessary
/sbin/route add default gw 192.168.1.2 metric 1

```

These commands plus some other goodies (clock synchronization) are executed by the `lanina.sh` script.

3. lanina is fixed IP (lanina.ess.uci.edu) running Linux

```

# Execute following commands on lanina
# This topology setup works, but lanina does not resolve names on local
# ssh dust.ess.uci.edu works but ssh dust does not
hostname lanina.ess.uci.edu
hostname -i -v
/sbin/ifconfig eth0 down
/sbin/ifconfig eth0:1 down
#/sbin/ifconfig eth0 lanina.zender.org # Connect IP address with Ethernet
/sbin/ifconfig eth0 inet 128.200.14.80 # Connect IP address with Ethernet
/sbin/ifconfig eth0 netmask 255.255.0.0 broadcast 128.200.14.255
# Creating eth0:1 puts potentially troublesome entries in route table
# These entries may cause problems when machine is again connected to h
#/sbin/ifconfig eth0:1 lanina.ess.uci.edu # Connect IP address with Eth
# Route to gateway host all datagrams bound for hosts outside intranet
# fxm: "metric 1" argument appears necessary
/sbin/route add default gw 128.200.14.1 metric 1

```

4. elnino is wireless on Harry's Earthlink.net DSL network. elnino uses eth0 for wired connection and eth1 for wireless.

5. lanina is wireless on Harry's Cox DSL network Make sure /etc/resolv.conf contains the cox.net nameserver:

```

nameserver 68.4.16.25
nameserver 68.2.16.30
nameserver 68.5.16.30

```

lanina is 192.168.1.73 on Harry's network. elnino is 192.168.1.74 on Harry's network. ashes is 192.168.1.75 on Harry's network. Harry's router is 192.168.1.1. Following lines switch lanina to work on wireless network:

```

/sbin/ifconfig eth0 down
sleep 1
/sbin/ifconfig eth0 add address 192.168.1.73 add netmask 255.255.255.0
/sbin/route add default gw 192.168.1.1

```

These commands plus some other goodies (clock synchronization) are executed by the `harry.sh` script.

Harry's DSL server, `moon`, is assigned a "permanent" IP address by `cox.net`. Since I masquerade as `moon`, it may be useful to try to display X programs from other machines on my laptop.

```

# As of 20050504:
# Hiperstore = 128.195.169.30
export DISPLAY='ip68-109-196-84.oc.oc.cox.net'
# Cliff's house in Lyons as of 20041221:
65.112.206.3
# Harry's Cox DSL
ssh ip68-109-196-84.oc.oc.cox.net # (to moon from outside)
# LGGE Static DHCP
ssh lgge-pc194.lgge.obs.ujf-grenoble.fr # 192.168.102.194 (to virga static IP)
ssh lgge-pc194.obs.ujf-grenoble.fr # 192.168.102.194 (to virga static IP from outside)
ssh lgge-asa50.obs.ujf-grenoble.fr # (virga appears as this from outside)
# SMU ADSL
ssh 62.147.226.50 # (neige appears as this from outside)
ssh -p 265 62.147.226.50 # Forward to 192.168.1.4 (to neige from outside)
ssh 62.147.226.50 # Forward to 192.168.1.4 (to neige from outside)
# UCI wireless for ashes as of 20050411:
/sbin/route del default gw 192.168.33.1 # Delete gateway
/sbin/route add default gw 192.168.0.1 # Add gateway
kwifimanager # Convenient GUI for WiFi configuration

```

`moon` is visible to the internet, but mine is not since all IP packets are masquerading as Harry's machine. Hence two consecutive SSH's are required unless `moon` is contacted on port 261. This port forwards SSH requests directly to `elnino`'s DHCP address of `192.168.1.226`

```

ssh -p 262 68.109.196.84 # Forward to 192.168.1.73 (to haze from outside)
ssh -p 263 68.109.196.84 # Forward to 192.168.1.74 (to elnino from outside)
ssh -p 264 68.109.196.84 # Forward to 192.168.1.75 (to ashes from outside)
ssh -p 265 68.109.196.84 # Forward to 192.168.1.76 (to virga from outside)
ssh -p 262 192.168.1.1 # Forward to 192.168.1.73 (to haze from inside)
ssh -p 263 192.168.1.1 # Forward to 192.168.1.74 (to elnino from inside)
ssh -p 264 192.168.1.1 # Forward to 192.168.1.75 (to ashes from inside)
ssh -p 265 192.168.1.1 # Forward to 192.168.1.76 (to virga from inside)

```

`bonk`'s network sometimes gets into strange states which dramatically reduce network speed. If you suspect this might be the case, log into `bonk` and do a

```
/sbin/ifconfig eth1
```

If it shows errors, alert Harry. The fix is often simply a quick `ifdown/ifup` cycle. NCAR standard for routers on each subnet is 128.117.xxx.251. CGD router is 128.117.22.251 or 128.117.24.251. The default gateway should be the router The default nameserver (in `resolv.conf`) should be `bearmtn = 128.117.24.2`.

9.38 LGGE

The external web homepage at the Laboratoire de Glaciologie Géophysique de l'Environnement (LGGE) in Grenoble, France is <http://www-lgge.ujf-grenoble.fr/eng>. LGGE is cordoned off by a firewall. The outside-facing LGGE SSH server is `triolet.obs.ujf-grenoble.fr`. The two-step hop to log into `ashes` from outside the LGGE firewall is

```
ssh triolet.obs.ujf-grenoble.fr
ssh lgge-pc194 # Virga static IP
ssh 192.168.102.194 # Virga static IP
ssh 192.168.102.5 # Virga inexplicably drops static and switches to DHCP
ftp ftp # Triolet is also ftp server: login to ftp as zender
ftp ftp # Triolet is also ftp server: login to ftp as anonymous
cd pub/depot/zender
ftp://ftp/pub/depot/zender
```

The automatic proxy configuration file for accessing the web via a browser inside the firewall is

```
http://proxy.ujf-grenoble.fr/auto-proxy.pac
```

The internal LGGE homepage is <http://www> (yes, you read that correctly). The LGGE homepage for Linux is: <http://lgge-pc240> and for printer configuration is <http://lgge-pc240:631/printers>. My USA/UCI printer configuration file is in `/etc/cups/cupsd.conf.bak`.

LGGE supports NFS-mounting the Linux partitions on `lgge-pc240`:

```
ls /mnt
mkdir /mnt/lgge-pc240
mount lgge-pc240:/usr/local /mnt/lgge-pc240
ls -l /mnt/lgge-pc240/bin/ncbo
/mnt/lgge-pc240/bin/ncbo -r
```

Pour imprimer partir du poste linux: imprimante Noir et Blanc HP laserjet 4250 au 1er tage couloir (format papier A4 recto/verso-duplex): 192.168.101.39 Photocopieuse Canon 3570 deuxieme tage salle reprographie: (format A4 et A3 recto/verso-duplex): 192.168.101.36

Imprimante couleur: Dell 3110CN deuxieme tage salle reprographie: (format A4 recto/verso-duplex): 192.168.101.42

9.39 Router

Whistlernet routers: Linksys Wireless-G router WRT54G Install HyperWRT Linux router distribution from <http://www.hyperwrt.org>.

1. Connect ethernet of laptop to router
2. Log into router. Login,password when shipped is admin,admin
3. Destination is http://192.168.1.1
4. Router Name is zender-WRT54G
5. Local Address is 192.168.1.5
6. Internet Address is 192.168.1.6
7. Subnet Mask is 255.255.255.0
8. Enable DHCP server
9. Starting IP address: 192.168.1.160
10. Ending IP address: 192.168.1.179
11. Static DNS: 128.200.1.201, 64.4.16.25,
12. Router password: dakinet
13. Wireless Network Name (SSID): zendernet
14. Setup to use MAC-based filtering
15. Startup script for each router to bond to eachother:

```
/usr/sbin/wl lazywds 0
# Give Mangonet MAC to zendernet
/usr/sbin/wl ws 00:0F:66:8E:B6:D2
```

9.40 Sound

Audio can be one of the most pleasant features of a computer, or its most annoying. Usually Linux does recognize the special keys many keyboard manufacturers have added to automatically control audio features. In public places such as planes, meetings, and libraries, it is often very desirable to turn off not just the sound server, but all audio beeps. For this purpose, the `xset` command does very nicely:

```
xset b 0 # Turn off sound
xset b off # Turn off sound
# Set beep parameters: % maximum volume, pitch (Hz), duration (ms)
xset b 20 440 50
```

<http://aries17.uwaterloo.ca/~dmg/brick> suggests adding

```
alias sound maestro
options sb io=0x220 irq=5 dm9=1
```

to `/etc/conf.modules` for Dell Inspiron 7500.

Sometimes the sound device gets controlled by processes that do not play well with others. One symptom of this is when `amarok` complains that the sound device “is unavailable”. When this occurs, try to kill the errant processes revealed to be controlling the sound device. To find those processes, follow these steps:

1. Open the system monitor: System-Administration-System Monitor
2. Open in the menu: Monitor-Search for open files
3. In the text box, enter: `/dev/snd/*` then click on Find

On 20080325, `kpdf`, of all things, was responsible for stealing the sound outputs.

9.41 dakine = msw, lanina = linux

List `dakine msw` shares accessible via Samba on `lanina lanina`: `smbclient -L dakine Work` interactively in FTP-like environment with `msw` shares `lanina`: `smbclient \\dakine\home` Ability to browse Linux from Windows is determined by `/etc/smb.conf`:

```
hosts allow = 192.168.1. 128.200.14. 127. 128.117.24.
[printers]
  browseable = yes
[]
```

Specifically, you must set up encrypted passwords by using the line “`encrypt passwords = yes`” in your `smb.conf` file, and you must create an appropriate `smbpasswd` file

9.42 DHCP

Dynamic Host Configuration Protocol allows computers to configure as an Internet node without a static IP address. DHCP clients receive their IP address from the DHCP server. The network configuration files must be altered to enable DHCP. On Debian systems, this is configured in `/etc/network/interfaces` with a line for the appropriate interface. For example, to configure `eth0` (usually the wired ethernet connection) for DHCP, use

```
iface eth0 inet dhcp
```

The `pump` command may work when DHCP does not.

```
/sbin/pump -i eth0 #
/sbin/pump -i eth1 --status # Inquire settings
```

Note, however, that `pump` may corrupt the ARP table (Section 9.44).

9.43 Wireless

Wireless is a PITA. Logging into wireless access in Starbuck's cafe Barcelona.

```
# Virga procedure:
# ifdown eth1;iwconfig eth1 essid "UCInet Mobile Access";ifup eth1
# ifdown eth1;iwconfig eth1 essid "mangonet";ifup eth1
# Ashes procedure:
# iwconfig eth2 rate 54M # Increase speed on internal broadcom chip
# iwconfig eth2 rate 5.5M # Increase speed on internal broadcom chip
# Alphabetical by city name
iwlist eth1 scan # List available hotspots
iwconfig eth1 essid any key off # Reset wireless interface
iwconfig eth1 essid 313a2a ap 00:02:2D:31:3A:2A # Albuquerque taylornet
iwconfig eth1 s:13a2a # Albuquerque taylornet
iwconfig eth1 essid sid_cienciaswl # Barcelona UAB ICTA
iwconfig eth2 essid IAMAS-2005 # Beijing IAMAS
iwconfig eth1 essid UCAR # Boulder NCAR/UCAR (heartbeat on wireless.ucar.edu)
iwconfig eth1 essid UCAR ap 00:1B:2A:64:F1:50 # Boulder NCAR Director's Con
iwconfig eth1 essid UCAR ap 00:1B:2A:64:F3:40 # Boulder NCAR Chapman Room
iwconfig eth1 essid UCAR ap 00:1B:2A:64:F3:60 # Boulder CMS offices
iwconfig eth1 ap 00:18:F8:FA:BD:46 # Greenbelt Marriott EOS Aerosol Cloud R
iwconfig eth1 essid COMINN key s:wificomfort38 # Grenoble Comfort Inn
iwconfig eth1 ap 00:0F:24:F1:89:B0 # Irvine 1103 Croul
iwconfig eth1 ap 00:0F:24:F1:89:40 # Irvine 3101 Croul
iwconfig eth1 ap 00:0F:24:F1:87:10 # Irvine 3228 Croul
iwconfig eth1 ap 00:0F:24:F1:86:90 # Irvine 3103 Croul
iwconfig eth1 essid "UCInet Mobile Access" # Irvine UCInet access points
iwconfig eth1 essid panera # Irvine Panera
iwconfig eth1 ap 00:0F:66:8E:B6:D3 # Irvine mangonet
iwconfig eth1 essid mangonet ap 00:0F:66:8E:B6:D3 # Irvine mangonet
iwconfig eth1 essid "Hale_La'_Aloa" key 6efd8c1bee # Kona
iwconfig eth1 essid traveler99 key FEEBED7A49 # Lyons cliffnet
iwconfig eth1 essid linksys ap 00:18:39:87:E1:60 # Lyons cliffnet
iwconfig eth1 essid "Blue Tile Beach House" key s:beach # Maui
iwconfig eth1 essid Guest # New London
iwconfig eth1 essid NICEXPO # Nice EGU
iwconfig eth1 essid EGU2004-1 # Nice EGU 192.168.33.181
iwconfig eth1 ap 00:02:72:59:D9:0A key s:1122334455 # Oslo Carlton Hotel
iwconfig eth1 essid IUGG-R # Perugia IAMAS Garden iuggpg4015, X/cp+a4C
iwconfig eth1 essid 2WIRE467 # Sacramento curtisnet
iwconfig eth1 essid "Free Public WiFi" ap 02:18:DE:04:D0:48 # San Antonio a
iwconfig eth1 essid WiFi4public ap 00:90:0E:00:32:C1 # San Antonio airport
iwconfig eth1 essid "Public Access Wifi" ap 00:0D:67:00:5E:61 # San Antonio
iwconfig eth1 essid linksys ap 00:18:39:87:E1:60 # San Miguel de Allende cl
iwconfig eth1 essid default # Sutter Creek dadnet
```

The Intel `ipw3945` driver does not automatically associate with the best quality network available, e.g., from multiple UCInet access points. You may configure a wireless interface to try access points in a certain order specified in file `/etc/wpa_supplicant.conf`, e.g.,

```
% cat /etc/wpa_supplicant.conf
ctrl_interface=/var/run/wpa_supplicant
network={
    ssid="name"
    psk="pass"
}
% cat /etc/network/interfaces
allow-hotplug eth1
iface eth1 inet dhcp
wpa-conf /etc/wpa_supplicant.conf
```

9.44 ARP Address Resolution Protocol

The *Address Resolution Protocol* (ARP) table maps ethernet addresses to IP addresses. Examine the ARP table contents (BuH98 p. 173) with

```
arp -a
```

Sometimes the ARP table contains duplicate or ill-defined entries. For example, `pump` may corrupt the ARP table. One symptom of this is when `ssh` to a working host fails with, e.g.:

```
ssh: connect to host ashes.ess.uci.edu port 22: No route to host
```

When this occurs, the *first thing to check* is that the `sshd` daemon is installed and running. Many security-conscious operating systems ship by default with very few services running. Another cause of this message is incorrect network configuration. The suspect device is `eth0` or `eth1`. Ensure that only one Ethernet device is configured. The best solution is to remove the offending entry:

```
sudo arp --delete ashes.ess.uci.edu
```

A simple but dangerous fix is to bring down (and then up) the offending device. This does not work remotely since the machine will be off the network.

`ssh` and `scp` will also fail with the error

```
ssh: connect to host ashes.ess.uci.edu port 22: No route to host
```

when the `sshd` server is not running. To check whether `sshd` is running, use `ps`. To check whether port 22 is open for SSH connections, use the `nmap` command:

```
ps -ef | grep sshd
sudo nmap -sS -F localhost
```

9.45 Hibernate

Laptops these days have various power consumption options controlled by the Linux Advanced Configuration and Power Interface (ACPI). The states are standby, suspend, and hibernate. The standby action causes the screen to go blank. The suspend action invokes ACPI state S3, aka suspend-to-RAM. According to linux.com, “everything in the system enters a low-power state except for RAM, which consumes a small amount of power in order to retain its contents, so that upon resuming, everything is loaded back from the memory and all running applications are restored immediately.” According to the article, the following should be in the device section of `xorg.conf`:

```
Option                "VBERestore"        "true"
```

The suspend action may be invoked by, e.g., the `suspend.sh` script in the article.

The hibernate action saves the state to disk and powers down. According to linux.com, hibernate is also known as ACPI State S4 or suspend-to-disk. On reboot, the kernel senses that a saved state is available and reloads it.

Experimented with ACPI on Ubuntu 6.06 LTS on 20060604 by connecting the actions to the lid shut button action. Results follow. Suspend on `ashes`: blanks screen and turns off wireless. With `nv` driver, lid-up restores without problems except must manually restart wireless interface. With `nvidia` driver, lid-up get that black melty look and must manually re-boot. Also, both `nv` and `nvidia` driver intermittently (after a few hours) may reboot on `ashes`. This may indicate an `xorg.conf` problem. However `nv` driver seems to have problems with translucency. In fact, translucency could be a problem with `nvidia` as well.

Hibernate on `ashes`: With `nv` driver, saves session to disk and powers down. Session restores correctly on re-boot, except must manually restart wireless interface. With `Nvidia` driver, nothing happens—no shut down at all.

Suspend ... on `virga`: works perfectly with `suspend.sh`! Hibernate ... on `virga` works perfectly after modifying `xorg.conf` to allow session restores. Only checked `virga` with `NVidia` driver, not sure if suspend and hibernate work on `virga` with `nv` driver.

The `i8kutils` package provides a number of Inspiron-specific control commands including `i8kmon`, `i8kctl`, `i8kbuttons`. These appear to work when the kernel has loaded the `i8k` kernel module, e.g., with

```
insmod /lib/modules/2.6.15-23-686/kernel/drivers/char/i8k.ko
```

The `gkrellm` utility will graphically display output from the `i8kutils` package.

```
sudo aptitude install xosview hddtemp gkrellm i8kutils gkrellm-i8k
# To load module automatically add "i8k force=1" to end of /etc/modules
sudo emacs -nw /etc/modules
sudo modprobe i8k # Restart or modprobe to load module
```

See [here](#) for more details.

This thread <https://launchpad.net/distros/ubuntu/+source/util-linux/+bug/66637> describes how to fix broken swap partitions with Ubuntu Edgy.

1. Determine your swap with `fdisk -l`

2. Run `mkswap` on your swap partition and record the output UUID.
3. Put this UUID into `fstab`.
4. Put `RESUME=UUID=<the-swap-partition-uuid-from-vol_ID` into `/etc/initramfs-tools/conf.d/resume`
5. Run `update-initramfs -u`
6. Reboot normally after this finishes
7. Run `swapon -s` to check if your swap is active.
8. Run `ls -la /dev/disk/by-uuid/`

Do not change any symlinks. Just try these instructions.

9.46 PCMCIA

Get rid of annoying beep when PCMCIA is inserted Add to `/etc/sysconfig/pcmcia`:

```
CARDMGR_OPTS = -q
kill -HUP cardmgr_pid
```

Use `cardctl` to notify system of PCMCIA card status changes. To suspend, eject, or resume the card, use

```
cardctl suspend # Shut down, disable power for socket
cardctl eject # Notify client drivers card will be ejected, cut power
cardctl resume # Restore power to socket, re-configure for use
cardctl reset # Send reset signal to socket
cardctl status # Display current socket status flags
cardctl ident # Display card identification information
cardctl config # Display the socket configuration
```

To restart the entire `pcmcia` subsystem, or to remove the `eth0` interface, use

```
/sbin/ifconfig eth0 down
/etc/rc.d/init.d/pcmcia restart
/sbin/ifconfig eth0 up
```

To trace shell execution use, e.g.,

```
sh -x /etc/sysconfig/network-scripts/ifup ifcfg-eth0
```

9.47 core dumps

RH6.1 default `/etc/profile` uses the `ulimit` command to limit size of core files with `ulimit -c 1000000`. This causes an error on telnet logins so comment it out.

9.48 Printing

RedHat causes spool directories in `/var/spool/lpd` to be owned by root with group `lp`. If this is not the case then printing may fail with permission denied errors `comp.os.linux.setup` suggested adding to `/etc/conf.modules`

```
alias parport_lowlevel parport_pc
```

This appears to fix printer queue “permission denied” errors on dakine But this is untested on lanina, which has given parport errors on bootup in the past.

If lanina is booted up unconnected to the printer, and then the printer is connected later and printing is attempted, an error may result such as `2000-12-30-17:13:18.388 Get_local_host: host`. When this occurs printing will fail and restarting the print daemon will also fail. The solution is simply to rename the machine, e.g., `sudo hostname lanina` and then restart the printer daemon and then printing should succeed.

9.48.1 CUPS

CUPS is the Common Unix Printing System. Configure CUPS by pointing a web browser to <http://localhost:631>. The HP4600 network color laser printer is configured to accept IPP as <http://hp4600.ess.uci.edu:631/ipp/port1>. The HP4350 network laser printer is configured to accept IPP <http://192.168.14.2:631/ipp/port1>. The Tek850 color printer is `tek850.ess.uci.edu = 128.200.24.133`. The HP6840 color printer is `hp6840.ess.uci.edu = 128.200.24.134`. The most informative site for printing on Linux is <http://www.linuxprinting.org>. They sponsor a [newsgroup](#) with an email gateway for virtually all types of printer, especially [HP printers](#). [Web Portal](#) to HP printers may be useful.

```
aptitude install cupsomatic-ppd # Install latest CUPS ppd drivers
killall -HUP cupsd # Restart CUPS daemon
/etc/init.d/cupsys restart # Restart CUPS daemon
lp ~/.plan # Print file to default printer
lp ~/.plan http://hp4600.ess.uci.edu:631/ipp/port1 # Specify printer
lp ~/.plan http://xerox.ess.uci.edu:631/ipp/port1 # Specify printer
lp ~/.plan http://tek850.ess.uci.edu:631/ipp/port1 # Specify printer
lp ~/.plan http://hp6840.ess.uci.edu:631/ipp/port1 # Specify printer
# Xerox drivers
# http://www.office.xerox.com/perl-bin/opb_drivers.pl
lp ~/.plan http://xerox.ess.uci.edu # Xerox Phaser 4500 is 128.200.24.127
lp -o docs ~/.plan # Print printer options to printer
lp -dhp4600/hp4600t -oraw # GIMP2 uses this by default
lpr -z InputSlot=Tray1 # CUPS command line option to send to Tray 1
lpr -P XeroxPhaser4500 ~/.plan # Xerox Phaser 4500 is xerox.ess.uci.edu = 128.200.24.127
lpr -P hp4600 ~/.plan # HP4600 is hp4600.ess.uci.edu = 128.200.24.134
```

Enabling extra printers in OpenOffice.org

```
sudo ln -s /usr/lib/openoffice/program/spadmin /usr/local/bin/spadmin
```

```
PATH=${PATH}\:/usr/lib/openoffice/program
/usr/lib/openoffice/program/spadmin
# Set printer to standard input
kprinter --stdin
```

The Xerox Phaser 4500 printer is set up as

```
Xerox/Docuprint 4508
lpd://xerox.ess.uci.edu:515
Add port 515 (LPD port) to /etc/rc.firewall the same way 631 is added?
Add port 9100 (JetDirect port) to /etc/rc.firewall the same way 631 is added?
```

On 20061108, ESS installed a new printer for faculty on a private network. This requires modifying the `interfaces` file to start a new interface solely for this printer.

```
# HP4350 DTN printer in ESS mail room IP 192.168.14.2
iface eth2 inet static
    address 192.168.14.2
    broadcast 192.168.14.255
    netmask 255.255.255.0
```

Printing protocols may be enabled and disabled by firewall controls on particular ports. On Debian-compatible GNU/Linux systems with 2.6 Kernels, the firewall is built with `iptables`. The `iptables` rules are set in `/etc/rc.firewall`.

```
iptables -L # Print IP tables configuration
iptables -F # Flush IP tables
iptables -X # Delete extra chains
iptables -P OUTPUT ACCEPT
iptables -P INPUT ACCEPT
iptables -P FORWARD ACCEPT
ip route flush cache # Flush existing routes
ip route show cache # Show existing routes
```

Under Debian, use the `guarddog` program to configure firewalls. At a minimum, allow these protocols: DNS, FTP, HTTPS, HTTP, IPP, NNTP, NTP, POP3, Ping, rsync, SMTP, SSH, VNC. Systems without these protocols set will not behave well. License daemons on other monstrosities will usually need to be configured by hand. For instance, the IDL license daemon wants to send and receive TCP packets on port 1700. This must be defined as a new “User Defined” protocol under “Advanced Settings” and then manually enabled in all intervening network zones.

9.49 Virtual Memory

When compiling large programs like `mie.cc`, `g++` may fail with an error like “virtual memory exhausted”. This may be due to having too small a swap partition. To see system parameters use `ulimit`, e.g., `ulimit -a`. According to Linus, the size of the swap partition should be twice the amount of RAM, see §9.71. Here is an untested way to use a swapfile instead of repartitioning:

```
# Make a swap file large enough to work around "virtual memory exhausted" e
# $SIZE is number of kilobytes, $SWAP is name of swapfile
$ dd if=/dev/zero of=$SWAP bs=1024 count=$SIZE && chmod 0600 $SWAP && sync &
```

9.50 tramp.el

Tramp stands for “Transparent Remote (file) Access, Multiple Protocol”. It is a most excellent Emacs extension which facilitates editing files on remote machines in a local Emacs session, using ssh, rcp, or any number of other protocols to transfer edits to the remote machine. Checkout tramp from the SourceForge archive:

```
cd ${DATA}
cvs -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/tramp login
cvs -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/tramp co -kk tramp
```

Here are Tramp formats that appear to work:

```
/dust.ess.uci.edu:.plan
/esmf.ess.uci.edu:.plan
/nco.sf.net:/home/groups/n/nc/nco/nco/doc/nco.texi
/scp/zender@krein.math.uci.edu:/home/ess/zender/zender/.plan
/scp/krein.math.uci.edu:/home/ess/zender/zender/.plan
/zender@krein.math.uci.edu:/home/ess/zender/zender/.plan
/krein.math.uci.edu:/home/ess/zender/zender/.plan
/multi:ssh:zender@dust.ess.uci.edu:ssh:zender@krein.math.uci.edu:~/plan
/[multi/ssh:zender@dust.ess.uci.edu/ssh:zender@krein.math.uci.edu]~/plan
```

Yes, Tramp does expand `~` correctly (i.e., on the remote machine). This section falls out of date quickly. Tramp filename conventions have changed at least three times in three years.

9.51 RPM packages in /usr that were installed (and may need to be uninstalled or reinstalled) by hand:

```
opendx-4.1.0-1.i386.rpm
opendx-docs-4.0.10-1.i386.rpm
sgi-opengl-1.2.1-1.i386.rpm
ImageMagick-5.1.1-1.i386.rpm
ghostscript-6.01-1.i386.rpm
ghostscript-fonts-6.0-2.noarch.rpm
```

9.51.1 IPCC and PBS configuration

The IPCC and PBS clusters, `ipcc.ess.uci.edu` and `pbs.ess.uci.edu`, run Rocks. Rocks-based clusters use the Community Enterprise Operating System, CentOS, a re-packaged version of RHEL. CentOS is an RPM-based system. Copies of CentOS are kept on the system so that new RPMs may be easily installed.

```
# Get CentOS bison locally or from the net:
/data/centos/4.3/os/x86_64/CentOS/RPMS/bison-1.875c-2.x86_64.rpm
cd ${DATA}/tmp
wget ftp://ftp.nluug.nl/pub/os/Linux/distr/CentOS/4.3/os/x86_64/CentOS/RPMS/
sudo rpm -ivh bison-1.875c-2.x86_64.rpm
```

Rocks names compute nodes in a 1-based enumerated format so a 26-compute node cluster has nodes:

```
compute-0-1, compute-0-2, compute-0-3, ... compute-0-26
```

Some IPCC compute nodes are loaded by the Medium Performance Cluster (MPC). MPC names compute nodes in a 1-based enumerated format so a 26-compute node cluster has nodes:

```
ipcc-n1, ipcc-n2, ipcc-n3, ... ipcc-n26
```

Each node is, by default, accessible via `ssh`. The `top` command shows the load on all processors in the line labeled “Cpu(s)” (near line 3). This mode presents statistics as a fraction of total available resources for all CPUs so that full usage of 1 CPU on a dual CPU system shows up as $\sim 50\%$ user usage, and $\sim 50\%$ idle. The load average (uppermost line), on the otherhand, is shown as a fraction of maximum load for one CPU.

Pressing `1` while `top` is running toggles the presentation mode between a single-line “Cpu(s)” summary and multiple lines of per-CPU summaries. The load average in the top-most line approaches N when an N -processor node is fully utilized, e.g., about 4.0 for a quad-CPU system. The `free` and `xosview` commands are also useful at monitoring system usage.

The PBS batch queuing system may be interrogated with a number of commands such as `pbsnodes`

```
cluster-ps # Connects to all computes nodes and does 'ps'
free-nodes # Display all free compute nodes (MPC)
job-uptime # Show load averages for all user running jobs on Torque
pbsnodes -a # Display all compute nodes and jobs assigned to them
qdel 161608 # Cancel job number 161608 in queue
qmon & # GUI for qstat-functionality for SGE
qstat -a # Display all jobs running/queued
qstat -r # Display all running jobs
qstat -r | grep ipcc # Display all running jobs
qsub -I -q opteron # Request interactive node from Opteron pool
esmfusers # display processes owned on each ESMF node
```

Running (and killing) MPI jobs can leave clusters with un-released resources such as shared-memory allocations and shared semaphores. This unclean state may block further MPI execution. To clean up the leftover state, use Inter-Process Communication (IPC) management commands such as `ipcs`:

```
# On interactive node:
ipcs -m | awk '/^ *0x/ {print $2 }' | xargs -n 50 ipcrm shm
ipcs -s | awk '/^ *0x/ {print $2 }' | xargs -n 50 ipcrm sem
```

9.51 RPM packages in /usr that were installed (and may need to be uninstalled or reinstalled) by hand: 99

```
# To compute-nodes:
```

```
cluster-fork "ipcs -m | awk '/^ *0x/ {print \\$2 }' | xargs -n 50 ipcrm sh  
cluster-fork "ipcs -s | awk '/^ *0x/ {print \\$2 }' | xargs -n 50 ipcrm ser
```

As its name implies, `cluster-fork` issues its argument as a command to all cluster nodes.

9.51.2 GPG: GNU Privacy Guard

The OpenPGP standard defines the cryptographic authentication protocol that `gpg` implements.

```
gpg --export -a 'Charlie Zender' # -a: Add ASCII armor  
gpg --keyserver pgp.mit.edu --send-keys 6F635D10  
gpg --keyserver pgp.mit.edu --search-keys 'Charlie Zender'  
gpg --list-public-keys  
Type  Bits / KeyID  Date      User ID  
pub   1024D/6F635D10 2006-11-20 Charlie Zender <zender@uci.edu>  
sub   2048g/9148C6AF 2006-11-20
```

```
Key fingerprint = DBD0 E788 E13C 56A2 6C5D 2C62 CB91 49AD 6F63 5D10
```

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
```

```
Version: GnuPG v1.4.3 (GNU/Linux)
```

```
mQGibEVh1bcRBADRCbLu37Qlsu2cOjckAlAgzvHBjMYQMtdQjJOe/rv9+p6StXyM  
SNp6GhpwX3b8j6r0NkT6fOqwaUh/uvxdTRR1FQJF7tHVAq0n9uY1stKFqRQC+cu+  
UP74FgfEvqh5pFL8sI05ayOMO7GZ0lh54rQQAQ/17TK29ttNDIyg2hpyWwCg7t fm  
nQ3gbeqdOhC9N1EikO7jCd0EAMq2ceX8hJv41701dpb405zogKYcto7BHULlUPAY  
v9FYt42MsF1kH6E0XC+HAA/pJ9sfCCrWKzbRtYkZGipJnB8bx19kxUv1RtPHXSdJ  
PHyckPOTseyCAX0/YbyOWMsZYKMzf2t0VileIMLSjTgwpu2OY/zSyi0Lx9kr2K8y  
bUdpBACXoGOSVBZdEX3z/ck5Ogan874UbTbOTxB3yw5Ry58CL52eZXaYP0v9pWL9  
cCe9hAoVlY2sH5fqZnKwQpT0HeNYcCdcY7wBtVnnW/4bjeljv5MsP2tkqsT2ucYg  
S4HDpfshfH0GwpUMHgDgfwXqyOUG4setLd5jfhct9UVptdZ8tLQfQ2hhcmxpZSBa  
ZW5kZXIgpPhplbmRlckB1Y2kuZWR1PohgBBMRAgAgBQJFYdW3AhsDBgsJCAcDAgQV  
AggDBBYCAwECHgECF4AACgkQy5FJrW9jXRC60ACgkWtdilBohVmiAxTwGXgstMSF  
JzUAnApKDwlVOOpTlDUw1CxpCqeAYwRmuQINBEVh1b8QCACyZg28taaSucl0vMSC  
7ZB5KOZgr+SDuuEj19p7Wwu83E6MppsZKQWyhjKJAhmjnAMYAgOexN5xSjgi50Or  
Df+pPrULN49Ec90qu2z/B2O74NZA3237t4ATZnMz8l0saL7R+BjtpEKjh43SDOmd  
lD9Da3aNYP04V2tU78F8Iusq3/k66Ppi+v9XZ+vLsApCgjBt0PJUB0vznt0cUYBS  
M4wvUsMzMB6iffw+CfAo2gaPCsqYEbEvMN1KiMl5jOi7wNFCAOtx+s8sWQ6LVIH8  
vpqMNcEPilUrH95FTuPIklNEC2sk8WFzmCNWqRBTP/iTeXF51zvz58OXg7joVoHp  
oN0rAAMGB/9n/AOrRMLIv9nSzM/R2505VBaiAU/vka+1lSABpN4H/S2F8B+7ehrK  
OWEo2yuStwcu2Hl/F02OwaaNbMzK0sUGMrbgHDGd/venbIKu/F7mEWfXndP8MbXx  
eT9tTjD0Oirm42JSTIQ7SA6RQg5eX0iPapCGbh+BLemv0itRl0E3YBf1t86UElL6  
MzSJUV8xo/lodt8FxFvM8nKSEs/CAXwYfkaUV/8I+Gdmd6ZTgwFcegNSm5V+811K  
BBe2xcNkpokPH2Zfam7W7TeoldSvF7RUXGDoEMorkXvuKGkb6rjwv4k/j8/krbNu  
V/WGP7GerNFkblJDu3fcC2CFqlpvsFfqieKegBECAAKFAkVh1b8CGwwACgkQy5FJ
```

```
rW9jXRAZnwCdHhbbY0H5bszcM7cj3kXnYoEIPoAn3DXHYEmnSXBZ//FJJc0LTG8
WEy6
=YK0f
-----END PGP PUBLIC KEY BLOCK-----
```

9.51.3 Building RPMs

```
# HOWTO build RPMs as a user
# http://my.helia.fi/~karte/linux/doc/rpm-build-as-user.html
sudo yum -y install rpm-build
# /etc/rpm/macros # (c) GPL 2003 Tero.Karvinen at-sign iki.fi
%packager %(echo "$USER")
%_topdir %(echo "$HOME")/rpmbuild
%_rpmtopdir %{_topdir}/%{name}
%_builddir %{_rpmtopdir}/BUILD
%_rpmdir %{_rpmtopdir}
%_sourcedir %{_rpmtopdir}
%_specdir %{_rpmtopdir}
%_srcrpmdir %{_rpmtopdir}
%_tmppath %{_rpmtopdir}/TMP
%_buildroot %{_tmppath}/%{name}-root

# Build rpms in your home directory, without root privileges
# Users can copy this file to $HOME/.rpmmacros to override
# settings, such as packager to "Firstname Lastname email".
# Example of typical directory and file hierarchy for
# rpm building:
# ~/rpmbuild/
# ~/rpmbuild/nano/ # replace nano with rpm name
# ~/rpmbuild/nano/TMP/
# ~/rpmbuild/nano/BUILD/ # user must create this dir
# and files (can be extracted from a source rpm with
# rpm2cpio *.src.rpm |cpio -dvi
# ~/rpmbuild/nano/nano-1.2.tar.gz
# ~/rpmbuild/nano/nano.spec
# Then you can build your rpm with
# rm -rf TMP/* BUILD/* ; rpmbuild -ba *.spec 2>&1 |tee rpmbuild.log
# For more information, see www.iki.fi/karvinen
```

9.51.4 RPM commands

The RedHat Package Manager (RPM) is used to maintain sources and binaries on a wide variety of Linux systems. The packages, so-called RPMs, are manipulated with the `rpm` command.

```
rpm -qa | grep foo # List all installed packages and search for foo
rpm -Uhv *.rpm # Upgrade packages (even if not installed)
```

9.51 RPM packages in /usr that were installed (and may need to be uninstalled or reinstalled) by hand: 101

```
rpm -Fhv *.rpm # Freshen packages (only if already installed)
rpm -q --whatprovides foo # Which installed package provides file foo?
rpm -qf foo # Which installed package provides file foo?
rpm -ql foo | less # Which files does installed package foo own?
rpm -qpl foo.rpm | less # Which files will package foo.rpm install?
rpm -e foo.rpm # Uninstall package
rpm --info foo.rpm # Print descriptive information about package
rpm -qR foo.rpm # Upon what packages does foo depend?
rpm --nodeps -e foo.rpm # Uninstall package regardless of dependencies
rpm -F --replacefiles --nodeps foo*.rpm # Do not check dependencies, allow
```

9.51.5 APT commands

The Debian package manager is called APT, Advanced Package Tool. Debian packages, so-called .deb's ("dot debs"), are manipulated with the apt command.

```
apt-cache search foo # Which packages relate to subject foo?
apt-cache search XML::Simple # Which package contains XML::Simple?
apt-file update # Refresh apt-file database
apt-file search foo # Which (non-installed) package installs/provides file
aptitude --help
aptitude -t unstable install foo # Install package foo from unstable
aptitude autoclean # Remove only package files that can no longer be downloa
aptitude clean # Remove everything except lock files from /var/cache/apt
aptitude dist-upgrade # Smarter version of upgrade
aptitude install `apt-show-versions -u -b | grep unstable` # Upgrade unstabl
aptitude install foo # Locate, download, and install package foo
aptitude install foo http://marillat.free.fr/ # Locate, download, and instal
aptitude install foo=x.y # Locate, download, and install package foo versio
aptitude remove foo # Uninstall foo-x.y.deb
aptitude update # Resynchronize package index files from sources
aptitude upgrade # Install newest versions of all packages currently install
aptitude upgrade foo # Install newest version of package foo
dpkg --configure foo # Configure unpacked package foo
dpkg --contents foo # Examine contents of archive
dpkg --force-help # Help on force actions
dpkg --install --force-overwrite # Install and overwrite files from one pack
dpkg --info foo # Examine archives
dpkg --install foo # Install package foo (Use this not aptitude install)
dpkg --list 'foo*' # Locate all packages named foo*
dpkg --listfiles foo # Which files does installed package foo own?
dpkg --print-avail foo # Print details about package foo
dpkg --purge foo # Remove package foo.deb, including configuration files
dpkg -P foo # Remove package foo.deb, including configuration files
dpkg --remove foo # Remove package foo.deb, leave configuration files intact
```



```
dpkg --search 'foo*' # Which installed package owns file foo?
dpkg -S 'foo*' # Which installed package owns installed file foo?
dpkg --status foo # Print installation status of package foo
dpkg --vextract foo bar # Display filenames contained by package foo in dir
dpkg-reconfigure foo # Re-configure installed package foo
```

9.52 Required software

required software Following is a list of scientific software required for research/teaching. Most of these packages are installed by default with standard and/or “power” GNU/Linux distributions, such as Debian and RedHat. Since these are free (as in no-cost) software packages, it is easiest if they are all installed on all machines, i.e., no differences between servers and clients. This Software should be installed in `/usr/local` by default. This avoids pathname proliferation. Packages may be installed in `/usr` when installation is automated so that the package becomes indistinct from one supplied with the base operating system. This is true of most `.debs` and `RPMs`, for example. When any administrative intervention is required, the package is best installed in `/usr/local`. For this reason, it is wise to back-up `/usr/local` and unnecessary to back-up `/usr`. Packages which have up-to-date pre-compiled binaries for most Linux distributions:

1. $\text{T}_{\text{E}}\text{X}/\text{L}\text{A}\text{T}_{\text{E}}\text{X}$, Autoconf, Autoheader, Automake, Bash, Bison, CVS, DDD, Emacs, Flex, GCC, GDB, GSL, Gettext, Ghostscript, Ghostview Gnuplot, Gzip, Libtool, M4, Make, Octave (Matlab clone), Perl (must be executable as `/usr/bin/perl`) R (S+ clone), Tar, wget, ldd, locate,

Packages which may not have up-to-date pre-compiled binaries for most Linux distributions:

1. Adobe: Acroread
2. Intel Fortran 95 and C++ compilers
3. Java runtime environment (1.4+)
4. NASA: Panoply
5. NCSA: HDF
6. NOAA: Ncbrowse
7. Sourceforge: NCO
8. UCSD: Ncview
9. Unidata: netCDF (must be compiled with `gcc/g++`), netcdf-perl, UDUnits, DODS

9.53 Packages installed in /usr/local (some RPM, some *.tgz, some proprietary):

DODS Distributed Oceanographic Data System dx IBM Data explorer hdf Hierarchical Data Format gnuplot Gnuplot (for DDD) java Java ncBrowse ncBrowse 1.2.1 ncarg NCAR graphics pgi PGI Fortran rsi Research Systems International Interactive Data Language (IDL) udunits Unidata units conversion package

9.54 NCO and patches

Updating NCO with patches: The only difficult thing about patches is remembering the switches used to create and apply them.

```
cd ~/nco # Create patch relative to top level NCO directory
cvs diff -c > feature.patch # Create patch
cd ~/nco # Apply patch from top level NCO directory
patch -p0 < feature.patch # -p0 = Causes patch to use unmodified file names
```

I have simplified remembering the last command by adding `lis patch='patch -p0'` to my `.bashrc` file. When patches go wrong they can be unapplied with

```
patch -p0 -R < feature.patch # -R = Reverse patch
patch -p0 --reverse < feature.patch # --reverse = Reverse patch
```

9.55 Skype

The UCI VoIP FAQ is [here](#).

Skype is an application which allows free IP-based telephony between registered users. The `skype` program package for Debian is at <http://www.skype.com/download/skype/linux>. Once registered, simply start `skype` from a console.

Inspiron 9400 users have found that the audio captured via microphone (plugged into the external jack) is almost inaudible on playback. One solution is use the command-line `alsamixer` program to set the “Capture” recording level to 75/75. Apparently the 9400 uses the microphone input for dual purposes and this may eventually be sorted out by software.

Ekiga is another program that uses VoIP. My `ekiga` address is sip:zender@ekiga.net. Ekiga is supposed to work with cameras like my Logitech QuickCam Pro for Notebooks. This camera works with Ubuntu “out-of-the-box” if one selects the `v4l2` driver. Otherwise, compile and install the driver yourself with:

```
sudo apt-get install linux-headers-`uname -r` linux-restricted-modules-`un
cd ${DATA}/tmp
svn checkout http://svn.berlios.de/svnroot/repos/linux-uvc/
cd linux-uvc/linux-uvc/trunk
make
sudo make install
# Start ekiga with debugging option
```

```

ekiga -d 1 &
# Plug in camera look at dmesg for the device listing, e.g., /dev/video1
# Point your application at that device and see if it works
# If the video device does not work, try restarting the module
sudo rmmod uvcvideo
sudo modprobe uvcvideo trace=15
dmesg | grep uvc

# New method
vncviewer -via tephra.ess.uci.edu -shared localhost:15
# Enter password when asked

# Old method
ssh -L5915:localhost:5915 zender@tephra.ess.uci.edu # Start tunnel to/on tep
vncviewer :15 # Open VNC window on local machine

```

9.56 Installing RedHat Linux

Get following non-default packages: Disk 1: openssh-server ElectricFence X100dpifonts DisXk 2: sudo xsane gimp Powertools: ddd openmotif openmotif-devel acroread dxpc octave blas lapack Other: gsl, RealPlayer, abisuite, HDF Do not install netcdf RPM as it is built with double-underscore versions of all Fortran functions.

9.56.1 Updating RedHat Linux with latest patches

Read UCI instructions at <http://www.dcslib.uci.edu/linux/index.html>. Script is installed as `/usr/local/bin/uci_dcs_lnx_pch.pl`. Replace “7.0” and “i386” with appropriate version and architecture information below, then run as root The kernel packages and C library are stored in CPU-specific directories for maximum performance. Update these directories first, e.g., i686 instead of i386. Then update the rest of the packages.

```

sudo mount ftp.uci.edu:/extra/ftp/mirrors/redhat /mnt
cd /mnt/linux/updates/9/en/os/i686 # CPU-specific, e.g., kernel updates
rpm -Fhv *.rpm
cd /mnt/linux/updates/9/en/os/i386
rpm -Fhv *.rpm
rpm -Fhv /mnt/linux/7.2/en/os/i386/preview/RPMS
rpm -Fhv /mnt/linux/7.2/en/os/powertools/i386/RedHat/RPMS
# Remotely update all RedHat RPMs against UCI Redhat Mirror
rpm -Fhv 'ftp://ftp.uci.edu/mirrors/redhat/linux/updates/7.2/en/os/i386/*.rpm'
# Remotely update all RedHat RPMs against UCI Redhat Mirror
rpm -Fhv 'ftp.uci.edu:mirrors/redhat/linux/updates/7.2/en/os/i386/*.rpm'

```

When a large number of simultaneous updates are required to patch a system, the command `rpm -Fhv *.rpm` may not work. In this case, break the task down into smaller tasks for `rpm`, e.g.,

```
for ltr in a b c d e f g h i j k l m n o p q r s t u v w x y z A B C D \
E F G H I J K L M N O P Q R S T U V W X Y Z; do
rpm -Fhv ${ltr}*.rpm
done
```

Another option is to install the RedHat Rawhide distribution. This distribution is available from <ftp://ftp.redhat.com/pub/redhat/linux/rawhide/i386>. A useful, comprehensive list of distributions is maintained at <http://distrowatch.com>.

9.56.2 Upgrading RedHat Linux

For network upgrades and installs, you must first create a `bootnet.img` floppy to install from, and then specify the network address of the RedHat FTP repository.

```
dd if=bootnet.img of=/dev/fd0 bs=1440k
ftp://ftp.uci.edu/mirrors/redhat/linux/6.2/i386/RedHat/
ftp://ftp.uci.edu/mirrors/redhat/linux/6.2/i386/images/bootnet.img
```

9.57 Debugging

Debugging is an art as much as anything. A list of debugging methods sorted by past efficacy is

1. Compile programs on different platforms. One compiler may notice errors that another compiler does not report. SGI compilers are especially good at finding errors.
2. Compile with bounds checking if possible
3. Run with `MALLOC_CHECK_=1` in environment
4. Link to memory debugging libraries like Electric Fence <ftp://ftp.perens.com/pub/ElectricFence> or `dmalloc` (<http://dmalloc.com>).
5. Be sure that the shell has enough resources (e.g., memory) to run the program or the crash may be especially mysterious. This can involve using the `ulimit` command, e.g., `ulimit -s unlimited`. However, this command may require special permissions to execute.
6. When all else fails, use a symbolic debugger like `DDD/gdb` (§9.59).

The Electric Fence debugger works very well with C-language programs. However, Electric Fence may hinder debugging C++ programs. C++ programs linked to the Electric Fence library (`libefence.a`) may generate obscure errors within `DDD/GDB` such as `libpthread.so.0: cannot load shared object`. If this occurs, simply compile the program without `-lefence` before loading into `GDB`.

Once bugs have been found and identified, consider sending a notice to those who might have been affected by the bug.

1. Identify bug symptoms so users may determine whether they were affected by it. It is also helpful to identify versions and/or dates of the code releases known to be affected by the bug.

2. Identify the bug cause in plain English, i.e., passed wrong pointer, transposed two arguments, error in equation, etc.
3. Describe extent of side-effects which bug *may* cause. Estimates of magnitude of bug, whether results were randomly or systematically biased, which particular regions were more or less affected, etc.
4. Complete the classification of the bug's damage by mentioning what processes, routines, regions, or versions, are *not* affected by the bug.
5. Determine and report on what future changes, if any, should be made to software management to ensure this type of bug does not occur again.
6. Acknowledge whether or not the bug fix is known to be solid
7. Provide actual code patch to fix bug

Linux Journal #87 (July, 2001, p. 82) gives helpful debugging tips. The `MALLOC_CHECK_` environment variable is one such method. Setting this variable and then running a faulty program will cause the program to print some verbose error messages when the fault is triggered. See the man page for `malloc` for more information.

```
MALLOC_CHECK_=0 ccc --tst=mmr # Heap corruption silently ignored
MALLOC_CHECK_=1 ccc --tst=mmr # Heap corruption prints diagnostics to stderr
MALLOC_CHECK_=2 ccc --tst=mmr # Heap corruption causes immediate abort()
```

Many commercial tools are available to find memory leaks. `ccmalloc` is a free tool. Simply link with `ccmalloc g++`, and then run with `MALLOC_CHECK_ = 1` to produce a report

```
ccmalloc g++ -o ccc ccc.o
gcc -o ncks ncks.o /usr/local/lib/ccmalloc-gcc.o -L/usr/local/lib -lccmalloc
g++ -o ccc ccc.o /usr/local/lib/ccmalloc-g++.o -L/usr/local/lib -lccmalloc
MALLOC_CHECK_=1 ccc
```

`ccmalloc` generates a report, and the number of memory leaks in the program is shown in the garbage column. <http://ieee.uow.edu.au/~mjpl6/prog/memleak.html> describes the procedure.

9.58 Valgrind

The most useful memory debugging tool since about 2002 has been `valgrind`². Simply precede the program invocation with `valgrind [options]`.

```
valgrind ccc >foo 2>&1
valgrind --leak-check=yes ccc >foo 2>&1 # Print location of likely leaks
valgrind --tool=memcheck ccc >foo 2>&1
```

²The term is from Norse mythology and is pronounced “val-grinned”

Table 3: GDB Commands

Command	Result
<code>start</code>	Break at beginning of program (all languages)
<code>break rt_cls::var_put</code>	Break at start of function <code>rt_cls::var_put</code>
<code>break nco_var_fll</code>	Break at start of function <code>nco_var_fll</code>
<code>break spc_slr.cc:spc_slr_cls::spc_slr_cls</code>	Break at start of function <code>spc_slr_cls</code> of class
<code>print grd[0]@(grd_nbr-0)</code>	Print first <code>grd_nbr</code> items of array <code>grd</code>
<code>print var_prc[idx]->val.fp[0]</code>	Print first item of array <code>var_prc[idx]->val</code>
<code>watch idx</code>	Set watchpoint on <code>idx</code>
<code>break dstmbl_MP_dst_mbl_</code>	Set breakpoint on G95-compiled module procedure
<code>print dstctl_MP_nstep</code>	Print G95-compiled Module Procedure (MP) varia

Since there are many errors in the default GNU compilers and libraries, these errors may be suppressed using the `valgrind` command. To make it easier to write suppressions, tell `valgrind` to print the suppression command for each error it encounters

```
valgrind --quiet --gen-suppressions=yes ccc # Generate suppression messages
```

In this mode, `valgrind` queries whether to print suppression text for each warning. The default suppressed warnings are contained in `/usr/local/lib/valgrind/default.supp`, and you can store additional suppressions you wish to ignore in a suppressions file, e.g., `valgrind.txt`. Multiple suppressions files may be used by specifying them with the `--suppressions` flag:

```
valgrind -v --suppressions=${HOME}/c++/valgrind.txt ccc >foo 2>&1
```

9.59 DDD

DDD has many obscure capabilities, especially when running in GDB mode. Table 3 summarizes frequently used GDB commands. GDB does not know the size of dynamically allocated arrays. To print the first `grd_nbr` items of array `grd`, use

```
print grd[0]@(grd_nbr-0)
```

The `-0` is necessary to get DDD to recognize that the argument is an integer. This may only be necessary when argument is type `long`. To display this array in the display window, highlight `grd[0]@(grd_nbr-0)` so that it appears in the argument line at the top next to `() : .` Then click the display button. The required specification can be tedious, especially when indirection and class structures are involved, e.g., `tst_obj->flx_slr_frc_in[0]@tst_obj->wvl_nbr_in` or, for, say, $N-3$ elements, the even more complex `(*tst_obj)->flx_slr_frc_in[0]@(tst_obj->wvl_nbr_in-3)`.

Setting conditional breakpoints can be done with the `watch` function. For example, consider the problem of breaking inside of a loop once the value of the counter, `idx`, is 37. The GDB command is `watch idx`, which sets a watchpoint on `idx`. A watchpoint is a breakpoint that is called whenever the expression changes value.

“Stepping” through C+ code is tedious because GDB takes the long route through all the interface files. Instead, set breakpoints at the start of the desired function with, e.g., `break rt_cls::var_put`.

Using `gdb` on G95 code is possible but not pretty. Symbols are case sensitive (use lower case). Dummy arguments are actually pointers. Module variables have a `modulename_MP_` prefix. Module procedures have the same, plus an underscore `_` suffix.

On SGI 64-bit machines, `dbx` and `gdb` do not work and one must use the `cvd` debugger. This debugger is powerful but non-intuitive. To view the stdout stream one must open the “Execution View” window. Also in the Views menu is the “Variable Browser” which does what it says. Clicking on the variable names in the “Variable Browser” will bring up a nice “Array Browser” for arrays. Breakpoints are set in the “Traps” menu in `cvd`.

On AIX, use the `xldb` debugger. Invoke with

```
export DISPLAY=ashes.ess.uci.edu:0.0
xldb `which swnb2` --drc_in=${DATA}/aca -D 1 -E -e 1603 -d foo.nc &
dbx -d 100 `which mie`
```

The Lahey lf95 compiler comes with the `fdb` debugger. Invoke with

```
fdb `which fff` &
```

9.60 Mailman

Mailman is used to create and manage mailing lists. First, run `mailman` to create the “site list”:

```
newlist mailman
```

9.61 Web Servers

The [DEAD box model](#) is interactively served on the web. Doing so requires coordination between model scripts, the HTTP server (Apache), and system permissions. The Apache server keeps its transaction logs in `/var/log/httpd`. Users will get an Internal Server Error if the CGI script fails.

```
tail /var/log/httpd/access_log # RedHat
tail /var/log/httpd/error_log # RedHat
tail /var/log/apache2/access.log # Debian
tail /var/log/apache2/error.log # Debian
```

Track web server usage using `webalizer`.

```
scp ~/linux/etc/webalizer.conf /etc
mkdir -p /var/www/html/usage
chmod -R 755 /var/www/html/usage
webalizer
```


9.61.1 Group Web Server

The group relies on a number of non-default webserver features. Most of these requirements can be met by small modifications of the `apache2.conf` or `httpd.conf` files.

1. `DirectoryIndex` should include `index.html` and `index.shtml`
2. PHP should be enabled (although not currently used).
3. Processing of server-side includes (SSI) should be enabled. This handles processing of `.shtml` files. The `apache2.conf` handles this:

```
# ++hjm to include handling for server-parsed files (.shtml)
<IfModule mod_mime.c>
    AddType text/html .shtml
    AddHandler server-parsed .shtml
</IfModule>
```

4. `DocumentRoot` (top directory) of webpage hierarchy should be `/var/www/html`
5. DODS/OPeNDAP access must work. This may require enabling web-server usage of CGI scripts in the `cgi-bin` directory
6. Webserver must re-start automatically on reboots

9.61.2 ESS Web Server

The ESS webserver is `www.ess.uci.edu`. It does not support remote SSH access. The ESS server, `ess1.ess.uci.edu`, allows remote SSH access. Both the webserver and the department server are virtual hosts which appear to be hosted by the same physical hosts. The physical hosts which support SSH access include `swamis.ps.uci.edu` and `lunada.ps.uci.edu`. Physical hosts which firewall SSH include `mavericks.ps.uci.edu`.

9.62 FTP and firewalls

On Linux, the `ncftp` client may be used to circumvent firewalls. `ncftp` supports most of the commands as the standard `ftp` client, but many more intuitive commands as well. Experience shows that `ncftp` should replace `ftp` in nearly every situation.

```
/bin/rm -r /data/zender/tmp/rpm
mkdir /data/zender/tmp/rpm
cd /data/zender/tmp/rpm
ncftp ftp.uci.edu # This will circumvent NCAR's firewall
cd mirrors/redhat/updates/8.0/en/os/i386
get *rpm
set passive yes # Other firewalls may require this command
bye
sudo /usr/local/bin/uci_dcs_lnx_pch.pl /data/zender/tmp/rpm
```

9.63 Accounts

When GUIs go bad, accounts need to be added by hand. The following commands work to add accounts and appropriately cross-mounted home directories to the cluster:

```
sudo groupmod -g 2400 cgdcsm
sudo useradd -D -g cgdcsm
sudo useradd -d /dhome/mflanner -g cgdcsm -n -u 3563 mflanner
sudo usermod -p '$1$charlie$VJFpvyBLLZzEvLkainaFW/' chaoluo
sudo usermod -L daniellj # Lock user pasSsword
sudo usermod -a -G adm,admin,audio,cdrom,cgdcsm,dialout,dip,floppy,lpadmin,
```

When adding an account to the home directory server itself (dust.ess.uci.edu), the default home directory should be `/home/${USER}` rather than `/dhome/${USER}`.

The `userdel` command is for deleting user accounts.

```
userdel foobar # Delete account
userdel -r foobar # Same and remove all files in ${HOME} and mail spool
```

9.64 NCAR

Over the years NCAR has developed a number of idiosyncratic procedures for managing user accounts. One way to gain access to any NCAR machine is to connect through the gatekeeper machine, gate.ucar.edu

```
ssh gate.ucar.edu
```

This machine will open a proxy to any other machine at UCAR.

9.65 Autotools

The GNU Autotools refers to an integrated set of software development and portability tools including Libtool.

```
cd ${DATA}
cvs -z3 -d :pserver:anoncvs@subversions.gnu.org:/cvsroot/libtool login
cvs -z3 -d :pserver:anoncvs@subversions.gnu.org:/cvsroot/libtool co -r bran
cd ${DATA}/libtool
./bootstrap;./configure --prefix=/usr/local;make;sudo make install
hash -r
cvs -z3 -d :pserver:anoncvs@subversions.gnu.org:/cvsroot/autoconf login
cvs -z3 -d :pserver:anoncvs@subversions.gnu.org:/cvsroot/autoconf co -r AUT
cvs -z3 -d :pserver:anoncvs@subversions.gnu.org:/cvsroot/m4 login
cvs -z3 -d :pserver:anoncvs@subversions.gnu.org:/cvsroot/m4 co -r m4-1_4o m
```

9.66 SSH

Most Unices use OpenSSH from the FreeBSD folks. Systemwide defaults are set in `/etc/ssh/ssh_config`. The directory `.ssh` contains authorization files which make passwordless access possible. An `authorized_keys` file, if present on machine A, contains public keys of users generated on machines B–Z. These users will be allowed to log in to machine A without entering any password. A `known_hosts` file, if present, contains public host keys of known remote machines. These keys were generated on remote machines.

1. `authorized_keys` contains RSA1 public keys of authorized users gathered from the `identity.pub` files on remote machines. These keys employ SSH protocol version 1.
2. `authorized_keys2` contains RSA and DSA public keys of authorized users gathered from the `id_dsa.pub` files on remote machines. These keys employ SSH protocol version 2.
3. `known_hosts` contains RSA1 public keys of authorized machines gathered from the `/etc/ssh/ssh_host_key` files on remote machines. These keys employ SSH protocol version 1.
4. `known_hosts2` contains DSA public keys of authorized machines gathered from the `/etc/ssh/ssh_host_dsa` and `/etc/ssh/ssh_host_rsa_key.pub` files on remote machines. These keys employ SSH protocol version 2.
5. Note that the system-wide configuration files (e.g., `/etc/ssh/ssh_config`) may disallow recognizing user-specific known hosts.

Execute `ssh-keygen -t rsa1` to generate `identity` and `identity.pub` files. Execute `ssh-keygen -t rsa` to generate `id_rsa` and `id_rsa.pub` files. Execute `ssh-keygen -t dsa` to generate `id_dsa` and `id_dsa.pub` files.

When upgrading or renaming a server, one may need to re-generate the server's keys.

```
cd /etc/ssh
ssh-keygen -t rsa1 -f /etc/ssh/ssh_host_key
ssh-keygen -t rsa -f /etc/ssh/ssh_host_rsa_key
ssh-keygen -t dsa -f /etc/ssh/ssh_host_dsa_key
```

To exercise all the various keys, force `ssh` to use the different algorithms:

```
ssh -1 # Use SSH version 1 protocol only
ssh -2 # Use SSH version 2 protocol only
```

Recently, OpenSSH added some useful commands. The command `ssh-copy-id` automates the process of installing the `identity.pub` and `id_[dr]sa.pub` files in a remote machine's `authorized_keys` and `authorized_keys2` files.

```
ssh-copy-id [-i [identity_file]] [user@]machine
```

Some machines have multiple IP addresses. `blackforest.ucar.edu` has four IP addresses. The entries in the `known_hosts` files might conflict with one another and lead to warnings like "POSSIBLE DNS SPOOFING DETECTED!" or "REMOTE HOST IDENTIFICATION HAS CHANGED!". Machines with multiple IP addresses and multiple host keys should be entered in the `known_hosts` files in one of two ways. First, one may omit the optional FQDN from the entry and specify only the IP address(es) associated with the host. When SSH looks at IP addresses only,

it cannot get confused. Second, one may enter the optional FQDN with the IP address, but, in this case, one must make sure that entries for each possible IP address are present. Otherwise the above warnings will result if one tries to use one of the missing addresses.

Misconfiguration of the login shell, e.g., causing `.bashrc` to print, will cause `scp` to fail with a `lost connection error`.

The NCO project uses the Sourceforge shell and CVS servers, `nco.sf.net` and `nco.cvs.sf.net`, respectively. Set SSH keys for these machines by uploading keys to the Sourceforge web interface at <https://sourceforge.net/account>. Uploading keys directly (e.g., with `scp`) to the shell server is allowed, but discouraged. Uploading keys directly to the Sourceforge CVS server is not allowed.

9.67 Security

It happens. Passwords get compromised. When this occurs, the security environment of all computers logically connected to the compromised account should be rebuilt. First, disable password-less logins from the affected account by removing the `authorized_keys` files from all machines. This firewalls the machines while the passwords are changed. Then search the systems for evidence of compromises. Two packages that check for installation of root-kits are `rkhunter` and `chkrootkit`. These may help get discover/eliminate spyware and malware as well. Login individually to all machines and change the passwords with `usermod`:

```
# 1. Remove password-less authorization
mch_lst='biogenic.ess.uci.edu dust.ess.uci.edu esmf.ess.uci.edu goldhill.cgs.uci.edu'
for mch in ${mch_lst}; do
    printf "Removing authorized_keys from ${mch}..."
    ssh ${mch} '/bin/rm ${HOME}/.ssh/authorized_keys*'
    printf "done\n"
done
# 2. Change passwords
for mch in ${mch_lst}; do
    ssh ${mch} sudo usermod -p '$1$salt$hashedpassword' zender
done
# 3. Check for compromises
rkhunter
# 4. Change keys
```

9.68 GCC

The GNU Compiler Collection, GCC, is the default compiler on Linux systems. To report a GCC bug, send following compile command

```
g++ -v -save-temps OPTIONS PROGRAM
```

and its text output in plain text to `bug-gcc@gnu.org`. Attach (with MIME) the resulting fully preprocessed file (`*.i*`). It is OK to compress the file before attaching it. Do not attach the assembly language file (`*.s*`).

Table 4: Partition sizes used on various computers

Computer	/wnd GB	/ GB	/usr GB	/usr/local GB	/home GB	/tmp GB	swap GB	/data GB	/var GB	/boot GB
seasalt	5	0.5	2	2	2	0.5	0.5	27		
dust		10	20		2			33		
lanina	5	0.1	2	2	3	0.2	0.4	11.5		
dakine										
ashes										
biogenic										
elnino										
haze										
soot										

9.69 Groups

Make default group `cgdcsn` with GID 2400. It is also helpful to have a consistent user ID or UID across all systems. My CGD UID is the same as my NCAR scientist number, 3555, as can be seen on Solaris with the `ypcat passwd | grep zender`.

9.70 WINE

```
cd /wnd/Program\ Files/Microsoft\ Office/Office
wine excel.exe
```

9.71 Partitioning

Linus Torvalds recommends setting swap partition size equal to twice the amount of RAM. Table 4 shows the partitioning schemes used on various computers.

1. `/usr`: 4 GB may not be enough on Debian. 10 GB would be better.
2. `/boot`: 100 MB is insufficient as having up to five kernels is reasonable (i.e., UP and SMP versions of old and new kernels plus one extra for testing).
3. `/var`: 1 GB is too small. Web server may grow to 5 GB soon. Use `aptitude clean` to free space used by Debian package installation.
4. `/data`: All remaining space.
5. `/home`: 10 GB seems adequate

Bibliography

Friedl, J. E. F., *Mastering Regular Expressions*, O'Reilly, Sebastopol, CA, 1997. [6.2](#)

Index

.deb, 103
-bpp, 17
-depth, 17
-nolisten tcp, 17
.dmg, 28
/etc/init.d/smartmontools, 23
AccelFactor, 11
Bash, 32
CR/LF, 68
CR, 68
DefaultColorDepth 24, 17
DirectoryIndex, 110
DocumentRoot, 110
Emacs, 26, 31, 69
GATEWAY, 85
GNUpod, 8
HOME, 62
InputDevice, 10
LF, 68
LoadLeveler, 39
MALLOC_CHECK_, 107
Macintosh, 68
MaxSpeed, 11
MinSpeed, 11
NCVweb, 25
NETMASK, 85
NQS, 39
PPP, 85
Tau, 14
XAPPLRESDIR, 13
XFree86-4.x, 10
X server, 12
Xconfigurator, 10
XkbOptions, 10
alsamixer, 104
amarok, 90
aptitude clean, 114
apt, 101
blackforest.ucar.edu, 112
cardctl, 94
cat, 27
ccmalloc g++, 107
ccmalloc, 107
cdbakeoven, 77
cfdisk, 21
cgdcsm, 114
chkrootkit, 113
cluster-fork, 99
control-panel, 26
crontab, 36
cron, 36, 60
csh, 32
cvd, 109
ddd, 36
dmalloc, 106
dmesg, 68
dumpe2fs, 25
e2fsck, 25
e2fsprogs, 25
ekiga, 104
emacs, 11, 97
eth0, 90
exportfs, 60
extglob, 32
fdb, 109
fdisk -l, 93
fdisk, 21
find, 24
firefox, 28
flexlm, 65
free, 98
fsck, 25
ftp, 110
g++, 96
gcc, 66
gconf-sanity-check-1, 60
gdm, 17, 18
gimp, 76
gkrellm, 93
glibc-devel, 40
glxinfo, 10
googleearth, 10
gpg, 99
gtkpod, 8
guarddog, 96
harry.sh, 87

- hash, 32
- hdiutil, 28
- hpcquick, 16
- hpcrun, 14, 15
- hpcviewer, 16
- httpd, 39
- hwclock, 41
- i8kbuttons, 93
- i8kctl, 93
- i8kmon, 93
- i8kutils, 93
- i8k, 93
- icc, 63
- ifc, 63
- ifdown, 88
- ifup, 88
- import, 10
- insmod, 84
- installer, 28
- ipcs, 98
- iptables, 96
- ipw3945, 92
- kbdrate, 13
- kdm, 17, 18
- kernel-headers, 40
- kgcc, 66
- kpdf, 90
- krein, 62
- ldd, 28
- lf95, 109
- lilo, 67
- llcancel, 39
- llq, 39
- llsubmit, 39
- lm-sensors, 23
- lmremove, 65
- lmstat -a, 65
- lmutil, 65
- lp, 95
- ls -la /dev/disk/by-uuid/, 94
- lsmod, 84
- mailman, 109
- malloc, 107
- man, 107
- mkfs, 21
- mkswap, 94
- modprobe, 15, 84
- mount, 23, 61
- mozilla-thunderbird -profilemanager, 28
- mplayer, 77
- ncBrowse, 25
- ncap, 31
- ncbo, 15
- ncftp, 110
- ncgen, 31
- ncview, 25
- netscape, 36
- nfslock, 60
- nmap, 92
- nm, 40
- no-root_squash, 60
- noauto, 61
- nslookup, 40
- ntpdate, 41
- nvidia, 13
- nv, 13
- oprofiled, 14
- oprofile, 14
- otool, 28
- panoply, 25
- pbsnodes, 98
- perfctr, 15
- postfix, 26
- pppd, 39
- ps, 92
- pump, 90, 92
- qdel -k, 39
- qstat, 39
- qsub, 39
- raid0run, 20
- raidhotadd, 20
- raidhotgenerateerror, 20
- raidhotremove, 20
- raidstart, 20
- raidstop, 20
- rcp, 97
- rdate, 41
- rehash, 32
- reset, 32

rkhunter, 113
rmmmod, 84
root, 23
rpm, 100, 105
rsync, 9, 37
scp, 92, 113
sendmail, 26
set-fill-column, 31
shopt, 32
skype, 104
smartctl, 23
smartmontools, 23
ssh-copy-id, 112
sshd, 39, 92
ssh, 92, 97, 98, 112
startx, 17
sudo ls, 15
sudo, 62
suspend.sh, 93
swapon -s, 94
swcursor, 10
sysctl, 67
telinit, 12
telnetd, 39
thunderbird, 27, 28
top, 98
touchpad, 11
tramp, 97
tr, 68
ulimit, 94, 96, 106
umount, 22, 25, 61
update-initramfs -u, 94
userdel, 111
usermod, 113
v4l2, 104
valgrind, 107
vfat, 8
webalizer, 109
xdm, 17, 18
xf86cfg, 9
xhost, 17
xine, 77
xldb, 109
xosview, 98
xrestop, 9
xset, 12, 89
xterm, 11
xvidtune, 13
.Xdefaults, 13
.bashrc, 64
.config, 67
.deb, 101
.mailrc, 27
.rhosts, 39
.shtml, 110
.ssh, 112
/boot, 14, 114
/data, 21, 114
/dev/modem, 38
/dev/perfctr, 15
/dev/sda1, 8
/dev/ttyS0, 38
/etc/X11/default-display-manager, 18
/etc/X11/gdm/gdm.conf, 17
/etc/X11/xinit/xserverrc, 17
/etc/cups/cupsd.conf.bak, 88
/etc/desktop, 18
/etc/fstab, 21, 22, 61
/etc/gdm/gdm.conf, 17
/etc/hosts, 40
/etc/init.d/nvidia-glx, 16
/etc/init.d/pgi-lmgrd, 65
/etc/initramfs-tools/conf.d/resume, 94
/etc/inittab, 12, 18
/etc/kde/kdm/kdmrc, 17
/etc/network/interfaces, 90
/etc/pcmcia/network.opts, 65
/etc/pcmcia/network, 65
/etc/pcmcia, 65
/etc/raidtab, 24
/etc/rc.firewall, 96
/etc/resolv.conf, 40, 86
/etc/ssh/ssh_config, 112
/etc/sysconfig/network-scripts/ifcfg-eth0, 85
/etc/sysconfig/network, 85
/etc/wpa_supplicant.conf, 92
/etc, 74
/home, 22, 114
/proc/mdstat, 20, 24
/proc/sys, 67

- [/root](#), 62
- [/sys](#), 9
- [/usr/bin/perl](#), 103
- [/usr/local/lib/valgrind/default.supp](#), 108
- [/usr](#), 114
- [/var/log/XFree86.0.log](#), 14
- [/var/log/Xorg.0.log](#), 14
- [/var/log/dmesg](#), 68
- [/var/log/httpd](#), 109
- [/var/spool/cron/crontabs/root](#), 36
- [/var/spool/cron/root](#), 36
- [/var/spool/cron](#), 36
- [/var/spool/mqueue-client](#), 26
- [/var/www/html](#), 110
- [/var](#), 114
- [XF86Config-4](#), 10, 14, 17
- [XF86Config](#), 9, 10, 17
- [apache2.conf](#), 110
- [authorized_keys2](#), 112
- [authorized_keys](#), 112, 113
- [bck.pl](#), 36
- [cgi-bin](#), 110
- [core](#), 24
- [crontab](#), 36
- [eth0:1](#), 85
- [eth0](#), 86, 92
- [eth1](#), 86, 92
- [exports](#), 59, 60
- [fl_rnm.pl](#), 32
- [fstab](#), 61, 94
- [gdm.conf](#), 17
- [hdc1](#), 21
- [hosts.allow](#), 60
- [hosts.deny](#), 60
- [httpd.conf](#), 110
- [id_\[dr\]sa.pub](#), 112
- [identity.pub](#), 112
- [index.html](#), 110
- [index.shtml](#), 110
- [initrd](#), 13
- [interfaces](#), 96
- [known_hosts](#), 112
- [lanina](#), 9
- [libefence.a](#), 106
- [libspecfun.a](#), 42
- [lilo.conf](#), 68
- [ltx.ps](#), 39
- [ltx.tex](#), 39
- [mail.greenspeedisp.net](#), 27
- [mail.wsu.edu](#), 27
- [mailhub.cgd.ucar.edu](#), 27
- [mdstat](#), 24
- [menu.lst](#), 13
- [ncview](#), 26
- [nvidia](#), 14
- [nv](#), 14
- [oprofile](#), 14
- [perfctr](#), 14
- [pop.uci.edu](#), 27
- [prefs.js](#), 28, 59
- [raidtab](#), 19
- [resolv.conf](#), 88
- [sda1](#), 21
- [sdb1](#), 21
- [sk98lin](#), 52
- [smtp.uci.edu](#), 27
- [smtp.west.cox.net](#), 27
- [suspend.sh](#), 93
- [user.js](#), 59
- [valgrind.txt](#), 108
- [vmlinuz](#), 14
- [xorg.conf](#), 10, 93
- [xserverrc](#), 17
- [~/xinitrc](#), 12
- [*](#), 32
- [+](#), 18
- [-g](#), 15
- [-](#), 18
- [1](#), 98
- [80](#), 31
- [?](#), 32
- [DD](#), 34
- [ESC](#), 31
- [MM](#), 34
- [Pg Dn](#), 77
- [RET](#), 31
- [YYYY](#), 34
- [set-fill-column](#), 31
- [x](#), 31
- [128.117.22.251](#), 88

- 128.117.24.251, [88](#)
- 128.117.24.2, [88](#)
- 128.117.xxx.251, [88](#)
- ashes.ess.uci.edu, [22](#)
- ashes, [23](#), [86](#), [88](#), [93](#)
- bearmtn, [88](#)
- biogenic, [19](#)
- clay, [14](#)
- dirt.ess.uci.edu, [24](#), [68](#)
- elnino.ess.uci.edu, [77](#)
- elnino, [86](#)
- ess1.ess.uci.edu, [110](#)
- gate.ucar.edu, [111](#)
- hp6840.ess.uci.edu, [95](#)
- https://vpn.nacs.uci.edu, [16](#)
- ipcc.ess.uci.edu, [97](#)
- lanina, [86](#)
- lgge-pc240, [88](#)
- localhost.localdomain, [27](#)
- localhost, [27](#)
- lunada.ps.uci.edu, [110](#)
- mavericks.ps.uci.edu, [110](#)
- moon, [87](#)
- nco.cvs.sf.net, [113](#)
- nco.sf.net, [113](#)
- ntp.ucsd.edu, [41](#)
- pbs.ess.uci.edu, [97](#)
- sand.ess.uci.edu, [77](#)
- swamis.ps.uci.edu, [110](#)
- tek850.ess.uci.edu, [95](#)
- time.nist.gov, [41](#)
- triolet.obs.ujf-grenoble.fr, [88](#)
- virga, [93](#)
- www.ess.uci.edu, [110](#)
- YYYY, [34](#)
- ZZZZ, [34](#)
- 2D, [13](#)
- 3D, [13](#)

- absolute path, [17](#)
- ACPI, [93](#)
- Acroread, [103](#)
- Address Resolution Protocol, [92](#)
- Advanced Configuration and Power Interface, [93](#)
- Apache, [109](#)

- APT, [101](#)
- ARM, [26](#)
- ARP, [90](#), [92](#)
- ATA, [23](#)
- Atmospheric Radiation Measurement, [26](#)
- Autoconf, [103](#)
- Autoheader, [103](#)
- Automake, [103](#)

- backups, [36](#)
- bad blocks, [25](#)
- Bash, [32](#), [103](#)
- Bison, [103](#)
- bpp, [17](#)
- burning, [77](#)

- C, [29](#)
- C language, [40](#)
- C library, [105](#)
- CDs, [77](#)
- CentOS, [97](#)
- cgd.ucar.edu, [27](#)
- CGI, [109](#), [110](#)
- command line, [28](#)
- Compact Disks, [77](#)
- CompactFlash, [8](#), [37](#)
- coordinated universal time, [40](#)
- Cox.net, [27](#)
- Cron, [36](#)
- CUPS, [95](#)
- curses, [21](#)
- CVS, [103](#), [113](#)

- daemon, [23](#), [68](#), [92](#)
- daemons, [26](#)
- DDD, [103](#), [106](#)
- Debian, [17](#), [18](#), [28](#), [36](#), [38](#), [90](#), [96](#), [101](#), [103](#)
- DeCSS, [77](#)
- deleting user accounts, [111](#)
- DHCP, [90](#)
- display manager, [18](#)
- DNS, [96](#)
- DODS, [103](#)
- DSA, [112](#)
- DVDs, [77](#)
- Dynamic Host Configuration Protocol, [90](#)

- Electric Fence, 106
- Emacs, 27, 103
- email gateway, 95
- encoding, 69
- Ethernet, 92
- ethernet, 90
- ext2, 25

- filename expansion, 32
- filesystem, 8
- Fink, 28
- firewall, 96
- Flex, 103
- FlexLM, 65
- formatting disks, 21
- FQDN, 112
- FreeBSD, 28, 112
- FTP, 69, 96

- G95, 109
- GCC, 66, 103, 113
- GDB, 103, 106
- Gettext, 103
- Ghostscript, 103
- Ghostview, 103
- GID, 8, 114
- GLcore, 10
- globbing, 32
- glx, 10
- GMT, 40
- GNOME, 18
- GNU, 111
- GNU/Linux, 36
- Gnuplot, 103
- Google, 10
- greenspeedisp.net, 27
- Greenwich mean time, 40
- GRUB, 67
- GSL, 103
- Gzip, 103

- HDF, 103
- hibernate, 93
- high-availability storage, 24
- hostname, 40, 65
- HP printers, 95

- HPCToolkit, 14, 15
- HTTP, 96, 109
- HTTPS, 96
- HyperWRT, 89

- IDE, 8, 23
- IDL, 18, 96
- INET, 60
- initial ram-disk, 13
- Intel, 63
- Inter-Process Communication, 98
- Internal Server Error, 109
- IP address, 40, 87, 90, 92, 112, 113
- IPC, 98
- IPCC, 97
- iPod, 8
- IPP, 96

- Java, 26
- JPEG, 18, 76

- K3b, 77
- K desktop environment, 18
- Kai, 63
- kbuildsycoca, 69
- KDE, 18, 77
- kernal packages, 105
- kernel, 68
- kernel headers, 13
- kio, 69
- Kmix, 77
- Knoppix, 23

- Lahey, 63, 109
- LCD Projectors, 18
- ldd, 103
- LDP, 11
- LGGE, 88
- Libtool, 103, 111
- LILO, 67
- linear, 68
- Linksys Wireless-G router WRT54G, 89
- Linus Torvalds, 114
- Linux Documentation Project, 11
- locate, 103
- LS120 drive, 36

- LVM, [21](#)
- M4, [103](#)
- Mail Transfer Agent, [26](#)
- Mail Transport Agent, [27](#)
- Mailman, [109](#)
- Make, [103](#)
- malware, [113](#)
- Matlab, [103](#)
- Maxtor, [19](#)
- media resource locators, [77](#)
- memory leaks, [107](#)
- Memory Stick, [9](#), [37](#)
- Mirroring disks, [24](#)
- Mozilla, [27](#)
- MPC, [98](#)
- MPI, [98](#)
- MPlayer, [77](#)
- Mplayer, [78](#)
- MRL, [77](#)
- MTA, [26](#), [27](#)

- NACS, [16](#)
- nameserver, [40](#)
- NCAR, [111](#)
- Ncbrowse, [103](#)
- NCO, [26](#), [103](#), [113](#)
- Ncview, [103](#)
- NCVweb, [26](#)
- netCDF, [25](#), [26](#), [103](#)
- netcdf-perl, [103](#)
- network file system, [59](#)
- network time protocol, [41](#)
- NFS daemon, [60](#)
- NFS lock daemon, [60](#)
- nfsd, [60](#)
- NNTP, [96](#)
- NTP, [41](#), [96](#)
- nv, [93](#)
- NVidia, [13](#), [93](#)
- Nvidia, [16](#), [93](#)
- nvidia, [93](#)

- Octave, [103](#)
- open source, [13](#)
- OpenGL, [10](#), [13](#)

- OpenMP, [63](#)
- OpenOffice.org, [95](#)
- OpenPGP, [99](#)
- OpenSSH, [112](#)
- operator precedence, [29](#)
- oprofile, [16](#)
- OS X, [28](#)

- Panoply, [103](#)
- PAPI, [14](#), [15](#)
- PBS, [97](#), [98](#)
- PCMCIA, [8](#)
- PDF, [37](#)
- Perl, [16](#), [32](#), [103](#)
- PHP, [110](#)
- Ping, [96](#)
- poll_idle, [15](#)
- POP3, [96](#)
- port, [87](#), [96](#)
- port 22, [92](#)
- Postfix, [27](#)
- PPP, [ii](#), [38](#)
- printing, [95](#)
- Procfs, [67](#)
- Promise cards, [19](#)
- proxy, [88](#)
- purging, [16](#)
- Python, [26](#)

- R, [103](#)
- RAID, [19](#)
- RAID-5, [19](#)
- RAM, [114](#)
- Rawhide, [106](#)
- RedHat, [17](#), [18](#), [26](#), [36](#), [38](#), [63](#), [65](#), [66](#), [85](#), [95](#),
[103](#), [106](#)
- regular expression, [29](#)
- rehash, [32](#)
- reiserfs, [23](#)
- relative path, [17](#)
- required software, [103](#)
- RHEL, [97](#)
- rmail, [27](#)
- Rocks, [97](#)
- root, [60](#), [78](#)

- root user, [36](#)
- root-kits, [113](#)
- root-squashing, [60](#)
- RPM, [97](#), [100](#), [103](#)
- RSA, [112](#)
- RSA1, [112](#)
- rsync, [96](#)
- runlevel, [12](#), [18](#)

- S+, [103](#)
- SCSI, [9](#), [23](#), [68](#)
- Skype, [104](#)
- SMART, [23](#)
- SMP, [114](#)
- SMTP, [26](#), [27](#), [96](#)
- SMTP server, [27](#)
- SourceForge, [97](#)
- Sourceforge, [113](#)
- spyware, [113](#)
- SSH, [87](#), [92](#), [96](#), [110](#)
- ssh, [92](#)
- SSI, [110](#)
- standby, [93](#)
- suppressions file, [108](#)
- SuSE, [77](#)
- suspend, [93](#)
- suspend-to-disk, [93](#)
- suspend-to-RAM, [93](#)

- Tar, [103](#)
- TCP, [96](#)
- TCP-forwarding, [17](#)
- T_EX/L_AT_EX, [103](#)
- tiff, [76](#)
- TLS, [16](#), [27](#)
- Tramp, [97](#)

- Ubuntu, [14](#), [26](#), [63](#)
- UCAR, [111](#)
- udev, [9](#)
- UDUnits, [103](#)
- UID, [8](#), [114](#)
- Unix, [68](#)
- unmount, [23](#)
- UP, [114](#)
- URL, [77](#)

- USB, [8](#)
- USB Flash Drive, [8](#)
- UTC, [40](#)
- UUID, [94](#)

- Virtual Private Network, [16](#)
- VNC, [96](#)
- VoIP, [104](#)
- VPN, [16](#)

- wget, [103](#)
- wildcard characters, [32](#)
- Windows, [68](#)
- wireless.ucar.edu, [27](#)
- wsu.edu, [27](#)

- X, [9](#), [12](#), [13](#), [17](#), [39](#), [87](#)
- X Window System, [10](#)
- X.org, [14](#)
- X server, [10](#), [18](#)
- XDM, [18](#)
- XFree86, [9](#), [14](#)
- Xine, [77](#), [78](#)

- ZIP disk, [36](#)