

April 3, 2014

# Linux/Unix Basics for HPC

Jay Alameda, NCSA

Mahidhar Tatineni, SDSC

Extended Support for Training, Education and Outreach

# XSEDE

Extreme Science and Engineering  
Discovery Environment

# XSEDE Training Survey

- **Make sure you sign the sign-in sheet!**
- At the end of the module, we will ask you to please complete a short on-line survey about this module. The survey can be found at <http://bit.ly/CSUXSEDE>
- This URL will also be available on the last slide.



San Bernardino Peak – 10,649 ft.  
Photo Credit: [www.summitpost.org](http://www.summitpost.org)



# Goals of tutorial

- Motivate need to know linux basics
- Basic linux editing
- Linux filesystem capabilities
- Basic linux diagnostics
- Basic profiling and debugging
- Data issues
- Networking
- File operations

# Credit, original material

- From Galen Arnold's tutorial at Florida International University
  - [http://www.ncsa.illinois.edu/~arnoldg/FIU\\_intro\\_linux.html](http://www.ncsa.illinois.edu/~arnoldg/FIU_intro_linux.html)
- NB: I use a subset of the capabilities presented (ie, pick what you find useful)

# Why Linux?


- Unix could not just get along : Solaris, HP-UX, Irix, AIX, Mach, Unicos, Ultrix...
  - Each vendor extended and customized Unix based on either BSD, SystemV, or a combination of both
  - It was not free.
    - Updates and bug fixes were sporadic and access to them varied by your ability to pay
    - Moving an application from one vendor's Unix to another could be tricky (vendor lock-in)
- Community loves free stuff!

# Where is Linux?

- All your phones belong to us:
  - Android
  - MacOS (mostly)
- Windows (!) (check out command prompt...)
- SOHO routers, networking equipment
- Sony Blu-ray player
- Refrigerator

## Special Apps and Built-in LCD Screen

Keep your kitchen and family organized with special apps made for your refrigerator. Leave notes for your loved ones. Display photos from your Picasa library, mobile phone or SD card. Stay up to date with all your family activities with Google Calendar. Access hundreds of recipes from Epicurious. Plus, get the latest weather and news via Weather Bug and Associated Press. All through Samsung's brilliant, WiFi-enabled 8" LCD screen.



**3D Blu-ray Disc Player with Super Wi-Fi**  
BDP-S5100  
★★★★★ (17)

**Model Features:** super Wi-Fi®, premium design, access to Sony Entertainment Network, Full HD 1080p Blu-ray Disc™ playback, 2D to 3D conversion, Sideview™ smartphone/tablet remote app

**COMPARE**

**WiFi CERTIFIED** **3D WORLD**

**In Stock**

~~\$100.99~~ **\$119.99**

# XSEDE



# Proceed with caution

- Recommend using a password safe, eg, KeePass (<http://keepass.info/>)



The image shows the official KeePass website and a screenshot of the KeePass application. The website has a header with the KeePass logo and the text "KeePass Password Safe". Below the header, there is a navigation menu on the left with links to Home, Forums, Feature List, Screenshots, Downloads, Translations, Plugins / Ext., Help, FAQ, Security, Awards, Links, Search, and Donate. The main content area features a "Latest News" section with several announcements about new releases of KeePass (2.23, 1.26, 2.22, and 2.21) with their respective dates and "Read More" links. There is also a "News Archive" link. A sidebar on the right contains links to "Awards" and "RSS Feed". Below the news section, there is a "What is KeePass?" section with a paragraph explaining the need for a password manager. At the bottom, there is a brief description of KeePass as a free open source password manager. The screenshot of the KeePass application shows a window titled "MyDatabase.kdb - KeePass" with a menu bar (File, Edit, View, Tools, Help) and a toolbar. The main area displays a tree view of the database structure, including a "Groups" list and a "Entries" list. The "Entries" list shows a table with columns for Title, User Name, Password, URL, and Icon. The table contains several sample entries. A context menu is open over the table, showing options like Copy User Name, Copy Password, URL(s), Perform Auto-Type, Add Entry, Edit/Move Entry, Duplicate Entry, Delete Entry, Select All, Select All, Clipboard, and Reconnect.

**KeePass Password Safe**

OSI certified

This is the official website of KeePass, the free, open source, light-weight and easy-to-use password manager. [\[Awards\]](#) [\[RSS Feed\]](#)

**Latest News**

- KeePass 2.23 released**  
2013-07-20 13:26. [Read More](#) \*
- KeePass 1.26 released**  
2013-07-12 13:32. [Read More](#) \*
- KeePass 2.22 released**  
2013-04-05 14:28. [Read More](#) \*
- KeePass 2.21 released**  
2013-02-03 14:16. [Read More](#) \*

[\[News Archive\]](#)

**What is KeePass?**  
Today you need to remember many passwords. You need a password for the Windows network logon, your e-mail account, your website's FTP password, online passwords (like website member account), etc. etc. etc. The list is endless. Also, you should use different passwords for each account. Because if you use only one password everywhere and someone gets this password you have a problem... A serious problem. The thief would have access to your e-mail account, website, etc. Unimaginable.

KeePass is a free open source password manager, which helps you to manage your passwords in a secure way. You can put all your passwords in one database, which is locked with one master key or a key file. So you only

# XSEDE HPC Resources

- All run linux

XSEDE User Portal | User Guides

https://portal.xsede.org/user-guides

MY XSEDE RESOURCES DOCUMENTATION ALLOCATIONS TRAINING USER FORUMS HELP

Get Started Access Resources Manage Data User Guides News Usage Policy Knowledge Base File Management

Below are links to each resource's user guide. Each guide provides information and instructions on system access, computing env to that resource. Resources are listed alphabetically within each resource type: [High Performance Computing](#), [High Throughput Co systems](#), [Special Purpose systems](#), [Testbeds](#) and [Software](#)

XSEDE is committed to providing quality, useful documentation to its users. Please feel free to leave your suggestions and comm guide.

<b>High Performance Computing</b>	<b>Scientific Visualization</b>
<a href="#">Blacklight (PSC)</a>	<a href="#">Longhorn (TACC)</a> <i>decommissioning April, 2014</i>
<a href="#">Gordon (SDSC)</a>	<a href="#">Maverick (TACC)</a> <i>New!!</i>
<a href="#">Gordon ION (SDSC)</a>	
<a href="#">Keeneland (Georgia Tech)</a>	<b>Storage Systems</b>
<a href="#">Kraken (NICS)</a> <i>decommissioning May, 2014</i>	<a href="#">Data Supercell (PSC)</a>
<a href="#">Mason (IU)</a>	<a href="#">HPSS (NICS)</a>
<a href="#">Lonestar (TACC)</a> <i>offline July, 2014</i>	<a href="#">Data Oasis (SDSC)</a>
<a href="#">Stampede (TACC)</a>	<a href="#">Ranch (TACC)</a>
<a href="#">SuperMIC (LSU)</a> <i>coming April, 2014</i>	<a href="#">XSEDE Wide File System (XSEDE)</a>
<a href="#">Trestles (SDSC)</a>	<b>Special Purpose Systems</b>
<b>High Throughput Computing</b>	<a href="#">Quarry (IU Gateway Web Services Hosting System)</a>
<a href="#">Open Science Grid</a>	<b>Testbeds</b>
	<a href="#">FutureGrid (distributed)</a>

XSEDE



# Login to XSEDE SDSC trestles cluster

- We will “hop” to trestles via the XSEDE single signon (SSO) hub.
- Use training logins provided today
- `ssh -l username login.xsede.org`
- Once on the SSO hub,
  - `gsissh trestles.sdsc.edu`
- We'll use the interactive session we establish to try out linux commands.

# Snapshots illustrating login process

```
mahidhar — mahidhar@trestles-login2:~ — ssh — 82x39

client64-166:~ mahidhar$ ssh -l mahidhar login.xsede.org
mahidhar@login.xsede.org's password:
Last login: Wed Apr  2 16:22:19 2014 from client64-166.sdsc.edu

Welcome to the XSEDE Single Sign-On (SSO) Hub!

Your storage on this machine is limited to 100MB.

You may connect from here to any XSEDE resource on which you have an account.

To view a list of sites where you actually have an account, visit:
https://portal.xsede.org/group/xup/accounts

Here are the login commands for common XSEDE resources:

Blacklight: gsissh blacklight.psc.xsede.org
Gordon Compute Cluster: gsissh gordon.sdsc.xsede.org
Gordon ION: gsissh gordon.sdsc.xsede.org
Keeneland: gsissh gsissh.keeneland.gatech.xsede.org
Kraken: gsissh kraken-gsi.nics.xsede.org
Lonestar: gsissh lonestar.tacc.xsede.org
Mason: gsissh mason.iu.xsede.org
Maverick: gsissh -p 2222 maverick.tacc.xsede.org
Open Science Grid: gsissh submit-1.osg.xsede.org
Stampede: gsissh -p 2222 stampede.tacc.xsede.org
Trestles: gsissh trestles.sdsc.xsede.org

Contact help@xsede.org for any assistance that may be needed.

[mahidhar@gw69 ~]$ gsissh trestles.sdsc.xsede.org
Last login: Wed Apr  2 13:24:39 2014 from gw69.iu.xsede.org
Rocks 6.1 (Emerald Boa)
Profile built 14:25 08-Dec-2013

Kickstarted 15:06 08-Dec-2013
Trestles Login Node
*****
[mahidhar@trestles-login2 ~]$
```

# Learning more about (most) commands

- These command line options often work to help shed light on a command:

`-v -help --v --help`

- Log into `trestles.sdsc.edu`, and try some commands and see:

`stat --help, ls -v, df -help, etc`

Often, if one option does not succeed, suggest another option that may likely work:

```
[jalameda@submit ~]$ whoami --help
whoami: invalid option -- 'h'
Try `whoami --help' for more information.
```

The XSEDE logo is displayed in a large, bold, white sans-serif font against a dark blue background with a grid pattern. The background of the entire slide features a stylized image of a planet's horizon with a blue sky and white clouds, and a grid of glowing blue squares.



# Sample output (stat, ls, df)

```
mahidhar — xdtr1@trestles-login2:~ — ssh — 80x24
[xdtr1@trestles-login2 ~]$ stat ptp_job.o2011011
  File: `ptp_job.o2011011'
  Size: 4285          Blocks: 16          IO Block: 32768   regular file
Device: 1fh/31d Inode: 455262226   Links: 1
Access: (0600/-rw-----)  Uid: (506339/   xdtr1)   Gid: ( 6426/   uic157)
Access: 2014-04-01 06:07:02.000000000 -0700
Modify: 2014-04-01 06:04:36.000000000 -0700
Change: 2014-04-01 06:04:37.000000000 -0700
[xdtr1@trestles-login2 ~]$ ls -l ptp_job.o2011011
-rw----- 1 xdtr1 uic157 4285 Apr  1 06:04 ptp_job.o2011011
[xdtr1@trestles-login2 ~]$ df /home/xdtr1
Filesystem            1K-blocks      Used Available Use% Mounted on
10.2.254.176:/export/home2/xdtr1
                        13456290464 7488135904 5284614080   59% /home/xdtr1
[xdtr1@trestles-login2 ~]$ df /oasis/scratch/trestles
Filesystem            1K-blocks      Used Available Use% Mounted on
172.25.33.252@tcp:172.25.33.124@tcp:/puma
                        437223093248 341585779344 73756871220   83% /oasis/scratch/t
restles
[xdtr1@trestles-login2 ~]$
```

# Good on-line resources:

- Introduction to Linux (Cornell Virtual Workshop)
  - <https://www.cac.cornell.edu/VW/Linux/>
- Also nice list at NCSA:
  - <http://www.ncsa.illinois.edu/UserInfo/GettingStarted/tutorials.html>

# Linux editors

- Ubiquitous: vi (or vim) – hard not to find this editor in your path
  - One nice resource:  
<http://www.cs.fsu.edu/general/vimmanual.html>
- Not so ubiquitous: pico or nano (nano is on trestles)
  - <http://www.nano-editor.org/>



# Brief vi demo

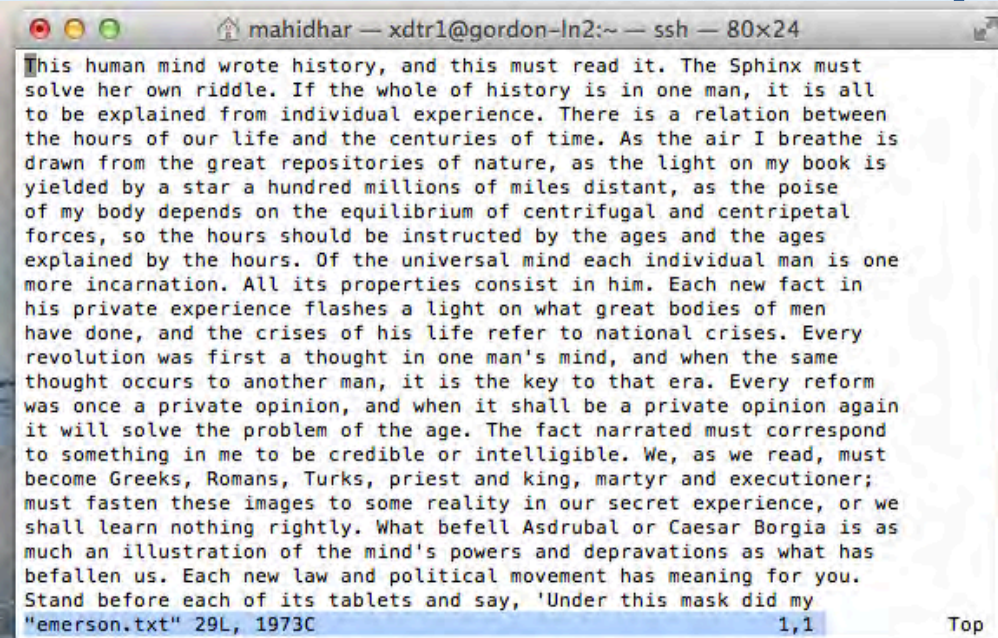
- Open file: `vi filename`
  - Two modes: command (where you start) and input
- Some commands:
  - `dd` (delete current line)
  - `x` (delete current character), can pre-pend with number
- `hjkl`: arrow keys (most reliable!)
  - `h` left, `l` right, `j` down, `k` up

# More commands

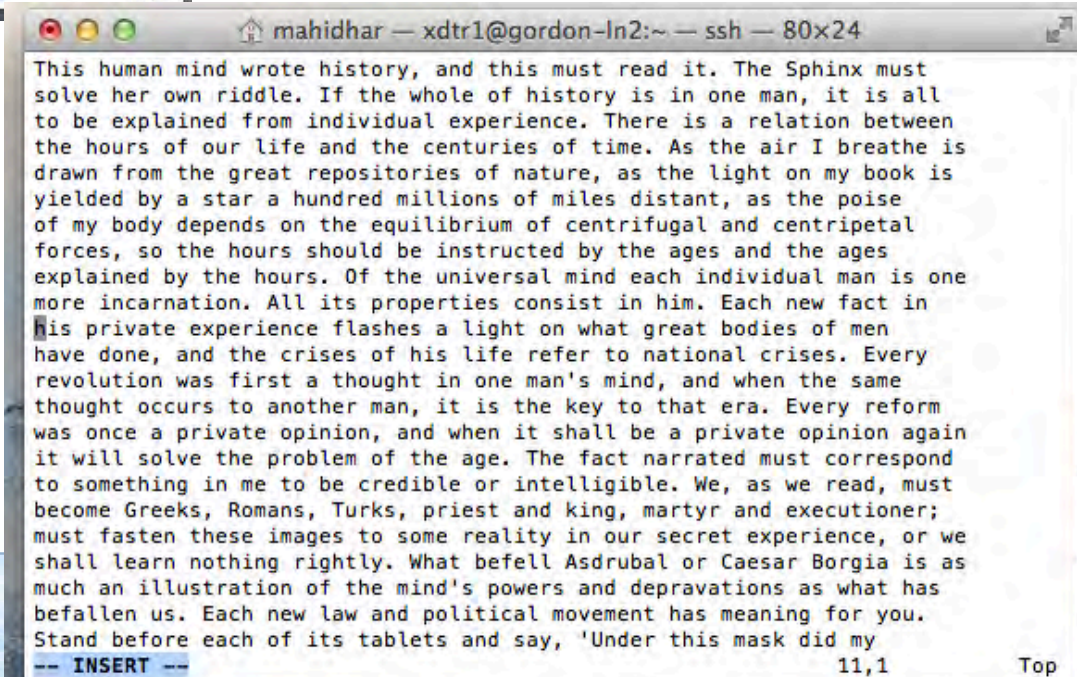
- Shift-G – go to last line in file
- 0 – go to beginning of line, \$ go to end of line
- Shifting to insert mode:
  - Insert text before cursor: I
  - Insert text after cursor: a
  - “o” open blank line below current line (o)
  - “O” open blank line above current line (O)
- “esc” ends insert mode



# vi : command and input modes (screenshots)



```
mahidhar — xdtr1@gordon-ln2:~ — ssh — 80x24
This human mind wrote history, and this must read it. The Sphinx must
solve her own riddle. If the whole of history is in one man, it is all
to be explained from individual experience. There is a relation between
the hours of our life and the centuries of time. As the air I breathe is
drawn from the great repositories of nature, as the light on my book is
yielded by a star a hundred millions of miles distant, as the poise
of my body depends on the equilibrium of centrifugal and centripetal
forces, so the hours should be instructed by the ages and the ages
explained by the hours. Of the universal mind each individual man is one
more incarnation. All its properties consist in him. Each new fact in
his private experience flashes a light on what great bodies of men
have done, and the crises of his life refer to national crises. Every
revolution was first a thought in one man's mind, and when the same
thought occurs to another man, it is the key to that era. Every reform
was once a private opinion, and when it shall be a private opinion again
it will solve the problem of the age. The fact narrated must correspond
to something in me to be credible or intelligible. We, as we read, must
become Greeks, Romans, Turks, priest and king, martyr and executioner;
must fasten these images to some reality in our secret experience, or we
shall learn nothing rightly. What befell Asdrubal or Caesar Borgia is as
much an illustration of the mind's powers and depravations as what has
befallen us. Each new law and political movement has meaning for you.
Stand before each of its tablets and say, 'Under this mask did my
"emerson.txt" 29L, 1973C 1,1 Top
```



```
mahidhar — xdtr1@gordon-ln2:~ — ssh — 80x24
This human mind wrote history, and this must read it. The Sphinx must
solve her own riddle. If the whole of history is in one man, it is all
to be explained from individual experience. There is a relation between
the hours of our life and the centuries of time. As the air I breathe is
drawn from the great repositories of nature, as the light on my book is
yielded by a star a hundred millions of miles distant, as the poise
of my body depends on the equilibrium of centrifugal and centripetal
forces, so the hours should be instructed by the ages and the ages
explained by the hours. Of the universal mind each individual man is one
more incarnation. All its properties consist in him. Each new fact in
His private experience flashes a light on what great bodies of men
have done, and the crises of his life refer to national crises. Every
revolution was first a thought in one man's mind, and when the same
thought occurs to another man, it is the key to that era. Every reform
was once a private opinion, and when it shall be a private opinion again
it will solve the problem of the age. The fact narrated must correspond
to something in me to be credible or intelligible. We, as we read, must
become Greeks, Romans, Turks, priest and king, martyr and executioner;
must fasten these images to some reality in our secret experience, or we
shall learn nothing rightly. What befell Asdrubal or Caesar Borgia is as
much an illustration of the mind's powers and depravations as what has
befallen us. Each new law and political movement has meaning for you.
Stand before each of its tablets and say, 'Under this mask did my
-- INSERT -- 11,1 Top
```





# Something extra: Eclipse Parallel Tools

- We've been using an ssh terminal view available in Eclipse Parallel Tools Platform
- Full Integrated Development Environment for parallel application development
- Available at [www.eclipse.org/PTP](http://www.eclipse.org/PTP)
- Tutorials (including how to add SSO capability, see NCSA module) at <http://wiki.eclipse.org/PTP/tutorials>; see XSEDE13 tutorial for latest content!

# Saving (or not) and exiting

- In command mode, type colon
  - Cursor drops to status line
  - :w -- writes file out
  - :q -- quits vi
  - :q! – quits vi without saving
  - :wq writes file, then quits

# Getting to know your filesystem

- `df -h` (volume), `df -i` (inodes)
  - (“display filesystem”)

```
[jalameda@submit ~]$ df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda1        20G   5.8G   13G   31% /
tmpfs            7.9G   80K   7.9G    1% /dev/shm
/dev/sda6        74G   6.5G   64G   10% /opt
/dev/sda3        30G   474M   28G    2% /tmp
/dev/sda5        9.7G   733M   8.5G    8% /var
panfs://storage.local:global
                234T  169T   66T   73% /panfs
export-60-14.storage.hpc.fsu.edu:/lustre
                344T  147T  198T   43% /lustre

[jalameda@submit ~]$ df -i
Filesystem      Inodes   IUsed   IFree IUse% Mounted on
/dev/sda1     1310720  178197 1132523   14% /
tmpfs         2049189    36 2049153    1% /dev/shm
/dev/sda6     4874240  146509 4727731    4% /opt
/dev/sda3     1966080    881 1965199    1% /tmp
/dev/sda5     640848   10545 630303    2% /var
panfs://storage.local:global
                244876955 176559489 68317466   73% /panfs
export-60-14.storage.hpc.fsu.edu:/lustre
                1309881124 30189927 1279691197    3% /lustre

[jalameda@submit ~]$
```





# “display usage”

- du – try this on your own
- du –s
  - Summary of usage in directory
- Size can be confusing – be careful
  - (from du –help)

Display values are in units of the first available SIZE from --block-size, and the DU\_BLOCK\_SIZE, BLOCK\_SIZE and BLOCKSIZE environment variables. Otherwise, units default to 1024 bytes (or 512 if POSIXLY\_CORRECT is set).

# File listing and status

- Listing: ls
  - ls -a (show all files)
  - ls -l (long listing)
  - ls -t (sort by time stamps)
- Status: stat – detailed file status

```
[jalameda@submit shallow]$ stat shallow
  File: `shallow'
  Size: 49903          Blocks: 320          IO Block: 16384   regular file
Device: 12h/18d Inode: 128654194   Links: 1
Access: (0775/-rwxrwxr-x)  Uid: ( 2654/jalameda)   Gid: ( 2654/jalameda)
Access: 2013-09-25 20:30:14.697511669 -0400
Modify: 2013-09-25 20:30:14.694712000 -0400
Change: 2013-09-25 20:30:14.697511669 -0400
[jalameda@submit shallow]$
```

# Symbolic links

- Can be convenient to jump from one filesystem to another

```
[jalameda@submit ~]$ ls -la
total 1376
drwxr-xr-x  7 jalameda jalameda 4096 Sep 26 16:25 .
drwxr-xr-x 457 root      root    65536 Sep 23 04:13 ..
-rw-----  1 jalameda jalameda 2542 Sep 26 18:58 .bash_history
-rw-r--r--  1 jalameda jalameda  18 Apr 23 2012 .bash_logout
-rw-r--r--  1 jalameda jalameda 176 Apr 23 2012 .bash_profile
-rw-r--r--  1 jalameda jalameda 124 Apr 23 2012 .bashrc
-rw-r--r--  1 jalameda jalameda 500 Feb 27 2012 .emacs
drwxr-xr-x  2 jalameda jalameda 4096 Jul 14 2010 .gnome2
-rw-----  1 jalameda jalameda 548 Sep 26 18:58 .history
drwxr-xr-x  4 jalameda jalameda 4096 Jun  2 2011 .mozilla
-rw-----  1 jalameda jalameda  36 Sep 23 04:13 .mpd.conf
lrwxrwxrwx  1 jalameda jalameda  37 Sep 23 04:13 scratch -> /panfs/storage.local/scratch/jalameda
drwxrwxr-x  4 jalameda jalameda 4096 Sep 26 09:21 shallow
-rw-rw-r--  1 jalameda jalameda 237 Sep 26 09:22 shallow-batch.sh
-rw-----  1 jalameda jalameda 252 Sep 25 23:22 shallow-test.o6661105
-rw-----  1 jalameda jalameda 4340 Sep 26 09:23 shallow-test.o6661462
-rw-----  1 jalameda jalameda 4340 Sep 26 16:25 shallow-test.o6663422
drwx-----  2 jalameda jalameda 4096 Sep 25 20:23 .ssh
drwxr-xr-x  4 jalameda jalameda 4096 Sep 26 15:14 trainingSC12\_C
-rw-----  1 jalameda jalameda 1940 Sep 26 09:24 .viminfo
[jalameda@submit ~]$
```



# Symbolic links

- Try following the link
  - cd scratch
  - cd .. (return to parent directory)
- Make another symbolic link: In –s
  - ln-s /home/ux400689/shallow-SR1\_kepler\_try2/shallow shallow\_works

```
-bash-4.1$ ls -lat
total 60
drwxr-xr-x 37 root      root    0 Mar 31 10:17 ..
-rw-----  1 ux400689 sdsc  4619 Mar 31 10:13 .bash_history
drwxr-xr-x 10 ux400689 sdsc  4096 Mar 31 10:13 .
lrwxrwxrwx  1 ux400689 sdsc   46 Mar 31 10:13 shallow_works -> /home/ux400689/shallow-SR1_kepler_try2/shallow
drwxr-xr-x  4 ux400689 sdsc  4096 Mar 31 10:10 shallow-SR1_kepler_try2
-rw-----  1 ux400689 sdsc   5043 Mar 31 10:10 .viminfo
drwxr-xr-x 14 ux400689 sdsc  4096 Mar  6 09:24 .eclipsesettings
drwxrwxr-x  4 ux400689 sdsc  4096 Mar  6 06:44 shallow
drwxr-xr-x  4 ux400689 sdsc  4096 Jan 23 11:39 shallow-SR2-RC1-23jan2014
drwxr-xr-x  3 ux400689 sdsc  4096 Jan 23 10:50 shallow-sr1-kepler-jan2014
-rw-r--r--  1 ux400689 sdsc   410 Dec 20 14:02 .bash_profile
-rw-----  1 ux400689 sdsc    70 Dec 13 06:43 .Xauthority
drwxr-xr-x  3 ux400689 sdsc  4096 Sep 27  2013 .visit
drwx-----  2 ux400689 sdsc  4096 Sep 26  2013 .ssh
drwxr-xr-x  4 ux400689 sdsc  4096 Jul 21  2013 trainingSC12_C
-bash-4.1$ ^C
-bash-4.1$
```

# Symbolic links pitfalls

- Be careful when
  - Navigating
  - Deleting
  - Not so fast all the time!

# Special files

- `/dev/null`: infinite black hole for data
  - Can redirect unwanted IO to `/dev/null`
- Eg, `cat filename > /dev/null`
- `/dev/zero`: source of null characters, eg, can use to fill file with zeros
  - Example: `dd if=/dev/zero of=filename count=1024 bs=1024` (create 1 MB file filled with zeros)
- `/dev/random`: source of psuedo-random numbers



# Checking out your environment

- What is my login?:
  - whoami
- What is my environment?
  - env
- What options do I have to augment my environment with (modules system):
  - module avail
- What is actually in my environment?
  - module list

# Environment (screenshots)

```
mahidhar — xdtr1@gordon-ln2:~ — ssh — 80x24

[xdtr1@gordon-ln2 ~]$ whoami
xdtr1
[xdtr1@gordon-ln2 ~]$ env
PDSHROOT=/opt/pdsh
MANPATH=/opt/intel/composer_xe_2013.1.117/man:/opt/torque/man:
HOSTNAME=gordon-ln2.sdsc.edu
IPPROOT=/opt/intel/composer_xe_2013.1.117/ipp
INTEL_LICENSE_FILE=/opt/flexlm/license/license-intel.txt
SHELL=/bin/bash
TERM=xterm-256color
HISTSIZE=1000
LIBRARY_PATH=/opt/intel/composer_xe_2013.1.117/compiler/lib/intel64:/opt/intel/c
omposer_xe_2013.1.117/ipp/./compiler/lib/intel64:/opt/intel/composer_xe_2013.1.
117/ipp/lib/intel64:/opt/intel/composer_xe_2013.1.117/compiler/lib/intel64:/opt/
intel/composer_xe_2013.1.117/mkl/lib/intel64:/opt/intel/composer_xe_2013.1.117/t
bb/lib/intel64//cc4.1.0_libc2.4_kernel2.6.16.21:/opt/gridengine/lib/lx26-amd64:/
opt/intel/composer_xe_2013.1.117/debugger/lib/intel64:/opt/intel/composer_xe_201
3.1.117/mpirt/lib/intel64
FPATH=/opt/intel/composer_xe_2013.1.117/mkl/include
QTDIR=/usr/lib64/qt-3.3
QTINC=/usr/lib64/qt-3.3/include
ROCKSROOT=/opt/rocks/share/devel
ANT_HOME=/opt/rocks
USER=xdtr1
```



# Environment (screenshots)

```
mahidhar — xdtr1@gordon-ln2:~ — ssh — 80x24

[xdtr1@gordon-ln2 ~]$ module li
Currently Loaded Modulefiles:
  1) intel/2013.1.117   2) mvapich2_ib/1.9   3) gnubase/1.0
[xdtr1@gordon-ln2 ~]$ module av

----- /opt/modulefiles/mpi/.intel -----
mpich2_ib/1.5(default)      mvapich2_ib/1.9(default)
mv2profile_ib/2.0a(default) openmpi_ib/1.6.5(default)

----- /opt/modulefiles/applications/.intel -----
eigen/3.1.3(default)      lapack/3.4.2(default)      petsc/3.3.p5(default)
fftw/2.1.5                mxml/2.7(default)         scalapack/2.0.2(default)
fftw/3.3.3(default)      netcdf/3.6.2              sprng/2.0b(default)
gsl/1.15(default)         netcdf/4.3.0(default)     sundials/2.5.0(default)
hdf4/2.9(default)        papi/5.0.1(default)       superlu/3.2(default)
hdf5/1.8.11(default)     parmetis/4.0.2(default)   tau/2.22.p1(default)
ipm/2.0.0(default)       pdt/3.18.1(default)       trilinos/11.0.3(default)

----- /opt/modulefiles/applications -----
R/3.0.1(default)          garli/2.01(default)       namd/2.10
amber/12(default)         gaussian/09.D.01(default) namd/2.9(default)
apbs/1.3(default)         gdal/1.9.2(default)       nwchem/6.3(default)
beagle/1.1(default)       geos/3.3.8(default)       octave/3.6.2(default)
beagle/2.1                globus/5.2.5              phylobayes/1.4e(default)
```



# And helpful environment variables

- What is my home directory?
  - \$HOME
- What is my login?
  - \$USER
- What is my current working directory?
  - \$PWD
- Can inspect these with the echo command, eg
  - echo \$HOME
  - Try it out!

# Environmental Variables (screenshot)

```
mahidhar — xdtr1@gordon-ln2:~ — ssh — 84x26

[xdtr1@gordon-ln2 ~]$ echo $HOME
/home/xdtr1
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ echo $USER
xdtr1
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ echo $PWD
/home/xdtr1
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ echo $PATH
/opt/gnu/bin:/opt/gnu/gcc/bin:/opt/mvapich2/intel/ib/bin:/opt/intel/composer_xe_2013.1.117/bin/intel64:/usr/lib64/qt-3.3/bin:/usr/local/bin:/bin:/usr/bin:/usr/local/sbin:/usr/sbin:/sbin:/usr/java/latest/bin:/opt/sdsc/bin:/opt/maui/bin:/opt/torque/bin:/opt/torque/sbin:/opt/torque/bin:/opt/torque/sbin:/state/partition1/catalina/bin:/opt/pdsh/bin:/opt/rocks/bin:/opt/rocks/sbin:/home/servers/gordon/bin:/home/xdtr1/bin
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ echo $LD_LIBRARY_PATH
/opt/gnu/gmp/lib:/opt/gnu/mpc/lib:/opt/gnu/gcc/lib64:/opt/gnu/mpfr/lib:/opt/gnu/lib:/opt/gnu/lib64:/opt/mvapich2/intel/ib/lib:/opt/intel/composer_xe_2013.1.117/compiler/lib/intel64:/opt/intel/composer_xe_2013.1.117/ipp/./compiler/lib/intel64:/opt/intel/composer_xe_2013.1.117/ipp/lib/intel64:/opt/intel/composer_xe_2013.1.117/compiler/lib/intel64:/opt/intel/composer_xe_2013.1.117/mkl/lib/intel64:/opt/intel/composer_xe_2013.1.117/tbb/lib/intel64/cc4.1.0_libc2.4_kernel2.6.16.21:/opt/gridengine/lib/lx26-amd64:/opt/intel/composer_xe_2013.1.117/debugger/lib/intel64:/opt/intel/composer_xe_2013.1.117/mpirt/lib/intel64
[xdtr1@gordon-ln2 ~]$
```

# Learning about your system

- /proc directory – with special files (inspect via “more filename”) – `cpuinfo`, `meminfo`
- Inspecting your PCI bus architecture – especially useful with GPGPUS, etc – `lspci` (and check out options for interesting variants)
- How long has the system been up, who is on, etc: `uptime`, `users`, `who`, `w`



# System info - screenshot

```
mahidhar — xdtr1@gordon-ln2:~ — ssh — 84x26
[xdtr1@gordon-ln2 ~]$ more /proc/cpuinfo
processor       : 0
vendor_id      : GenuineIntel
cpu family     : 6
model          : 45
model name     : Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz
stepping       : 6
cpu MHz        : 1200.000
cache size     : 20480 KB
physical id    : 0
siblings       : 8
core id        : 0
cpu cores      : 8
apicid         : 0
initial apicid : 0
fpu            : yes
fpu_exception  : yes
cpuid level    : 13
wp             : yes
flags          : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat
pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm c
onstant_tsc arch_perfmon pebs bts rep_good xtopology nonstop_tsc aperfmperf pni pclm
ulqdq dtes64 monitor ds_cpl vmx smx est tm2 ssse3 cx16 xtpr pdcm pcid dca sse4_1 sse
4_2 x2apic popcnt tsc_deadline_timer aes xsave avx lahf_lm ida arat xsaveopt pln pts
dts tpr_shadow vnmi flexpriority ept vpid
bogomips      : 5187.29
```

# What is my system doing

- Static process listing: `ps`
  - `ps` – private to you
  - `ps -efl` – global view of processes
- Dynamic listing – `top`
  - Typing `h` while displaying shows options
    - Try typing “`m`” to toggle memory view
- Inspect virtual memory – `vmstat`
- Am I running out of memory: `free`
- Do I need to kill processes?
  - `kill PID`, `kill -9 PID`, `killall`: be careful!

# Learning about the process I'm running

- Timing a command:
  - time cmd
    - real (wall clock time)
    - user cpu time (process in user space)
    - system cpu time (process in kernel space)
- Learning about system calls – strace
  - Strace cmd, or attach with strace -p PID
- Learning about shared library calls – ltrace
  - Not in path on trestles system



# Trying to learn about IO

- dd command is helpful
  - Reads and then writes a file with the `-if` and `-of` flags
  - Can change block sizes, do data conversions, do direct IO (bypass buffer cache)
  - See [http://www.gnu.org/software/coreutils/manual/html\\_node/dd-invocation.html](http://www.gnu.org/software/coreutils/manual/html_node/dd-invocation.html)
  - Can learn about IO characteristics of system

# Getting a little tricky on IO

- Want to pass (small amounts) of data quickly between programs: consider /dev/shm

```
[jalameda@submit ~]$ df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/sda1              20642428    6022764   13571088  31% /
tmpfs                  8196756         80    8196676   1% /dev/shm
/dev/sda6              76744752    6752144   66094192  10% /opt
/dev/sda3              30963708    483140    28907704   2% /tmp
/dev/sda5              10079084    750832    8816252   8% /var
panfs://storage.local:global
                      250754002400 181208247024 69545755376  73% /panfs
export-60-14.storage.hpc.fsu.edu:/lustre
                      369089829984 157459695392 211629820512  43% /lustre
[jalameda@submit ~]$
```

- And, to learn about your processors memory bandwidth capabilities: STREAM benchmark  
– <http://www.streambench.org/>

# What about my data

- Globus Online – powerful utility to move (lots) of data – [www.globusonline.org](http://www.globusonline.org)
- Combining files into a single archive
  - tar
  - Can compress on the fly (try some experiments to see if this makes sense)
  - tar –cvf filename (list of items to combine)
    - Create archive “filename”, verbosely



# More on tar

- Listing contents of archive –
  - `tar -tvf`
- Extract archive
  - `tar -xvf`
  - (I like the verbose option)
- Compression
  - `tar -czf` uses gzip compression
  - Experiment – compare `tar -cvf` (twice) against `tar -cf` (use time command)

# More on tar, and data

- Why compression?
  - Limited link speed, limited data space
  - Not an obvious choice always
- Tar alternatives –
  - zip – package and compress archive file
    - `man zip`
  - gzip – compress files
    - `gz` extension is a clue
    - `man gzip`

# Directories, and directory policies

- Take the time to learn your site local policies
  - \$HOME – small files, not purged usually, sometimes (but not always!) backed up
  - \$SCRATCH – large files, large quota, usually purged, never backed up
  - Other directories



# Learning about your network

- Not all of these work everywhere, but if you suspect problems, they are worth a try
  - ping hostname – measures latency to host, packet loss, etc
  - traceroute hostname – shows path to host
  - netstat – learn who is connected, from where, to your system!
  - ifconfig – learn about status of your physical network devices (low-level)

# Network info - screenshot

```
mahidhar — xdtr1@gordon-ln2:~ — ssh — 84x26

[xdtr1@gordon-ln2 ~]$ ping trestles.sdsc.edu
PING trestles.sdsc.edu (198.202.118.30) 56(84) bytes of data.
64 bytes from trestles-login1.sdsc.edu (198.202.118.30): icmp_seq=1 ttl=61 time=0.243 ms
64 bytes from trestles-login1.sdsc.edu (198.202.118.30): icmp_seq=2 ttl=61 time=0.254 ms
^C
--- trestles.sdsc.edu ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1994ms
rtt min/avg/max/mdev = 0.243/0.248/0.254/0.016 ms
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ traceroute stampede.tacc.utexas.edu
traceroute to stampede.tacc.utexas.edu (129.114.62.13), 30 hops max, 60 byte packets
 1  dolphin.sdsc.edu (198.202.104.1)  0.176 ms  0.199 ms  0.180 ms
 2  TACC-SDSC.gw.xsede.org (198.17.196.199)  50.037 ms  50.028 ms  50.008 ms
 3  vl664-ex9214-roc.net.tacc.utexas.edu (129.114.0.166)  50.162 ms  50.091 ms  50.065 ms
 4  login3.stampede.tacc.utexas.edu (129.114.62.13)  49.684 ms  49.702 ms  49.755 ms
[xdtr1@gordon-ln2 ~]$
```

# Higher level networking tools

- ssh – securely connect to another host
- scp -- securely copy files between hosts
  - Gsi variants useful in XSEDE
- wget – retrieve files via URL (http, https, ftp)
  - <http://www.gnu.org/software/wget/>
- curl – retrieve files via URL (lots of protocols!)
  - <http://en.wikipedia.org/wiki/CURL>



# wget - example

- Download “Pride and Prejudice” from Project Gutenberg:

wget <http://www.gutenberg.org/cache/epub/1342/pg1342.txt>



```
mahidhar — xdtr1@gordon-ln2:~ — ssh — 84x26
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ wget http://www.gutenberg.org/cache/epub/1342/pg1342.txt
--2014-04-03 17:50:16-- http://www.gutenberg.org/cache/epub/1342/pg1342.txt
Resolving www.gutenberg.org... 152.19.134.47
Connecting to www.gutenberg.org[152.19.134.47]:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 717569 (701K) [text/plain]
Saving to: "pg1342.txt"

100%[=====>] 717,569      1.06M/s   in 0.6s

2014-04-03 17:50:17 (1.06 MB/s) - "pg1342.txt" saved [717569/717569]

[xdtr1@gordon-ln2 ~]$
```

# File operations

- sort, uniq, wc, grep, cat, cut, paste

- Nice quick reference:

<http://wtuto.com/redhat/tpt.html>

- grep: search for strings in a file

- Simple example: grep string filename

- sort - can use as command (on a file), or have output from other commands piped to it

- Eg – grep bash /etc/passwd | sort
  - Vertical bar pipes output from grep into sort

# More file operations

- `uniq` – find unique lines in a file (many options, including counting duplicates)
  - Eg, `uniq -c file` will give you a list of all the unique lines in a file, with a count prepended for the number of times a line is repeated
- `wc` – “word count” – but counts lines, words, characters and can be restricted to any of those
- `cat` - lists file to stdout (simple, easy!)
- Cut, paste: not as familiar with, cut extracts text from a file based on position in file, paste merges data from one file into another
  - See <http://www.techrepublic.com/article/lesser-known-linux-commands-join-paste-and-sort/>



# Redirection

- Incredibly useful
- Slightly different syntax for different shells
- `Command > file` – redirects stdout of a command into a file
- `Command >& file` – redirects stdout and stderr from command into a file (csh, tcsh)

# Peeking at files

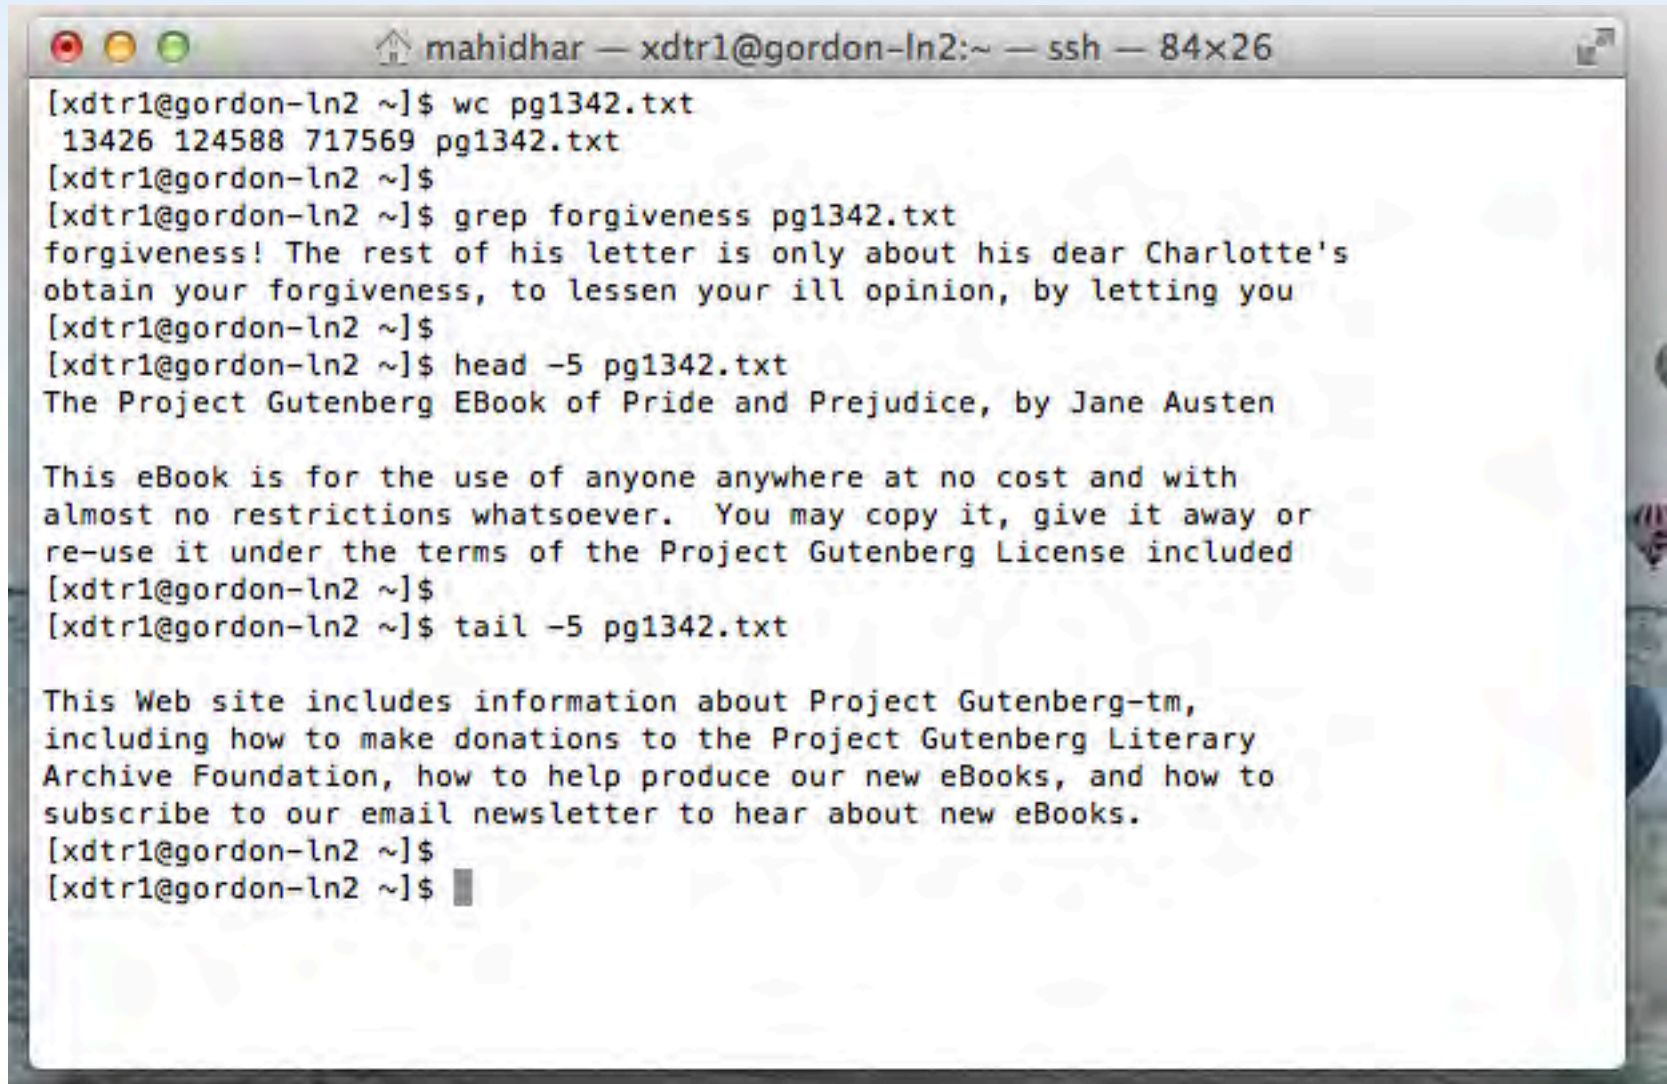
- head – show beginning of file
- tail -- show end of file
  - tail –f – list file dynamically as it grows (really useful)
  - Command > file &
    - (ampersand stuffs command into the background)
  - tail –f file
- more – browse file, page by page
  - more file or cat file | more

# Stream editing

- sed – powerful, but underused by me (!)
- Can get a glimpse of its capabilities at the end of <http://wtuto.com/redhat/tpt.html>



# File operations - screenshot



A screenshot of a terminal window titled "mahidhar — xdtr1@gordon-ln2:~ — ssh — 84x26". The terminal shows a series of commands and their outputs:

```
[xdtr1@gordon-ln2 ~]$ wc pg1342.txt
13426 124588 717569 pg1342.txt
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ grep forgiveness pg1342.txt
forgiveness! The rest of his letter is only about his dear Charlotte's
obtain your forgiveness, to lessen your ill opinion, by letting you
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ head -5 pg1342.txt
The Project Gutenberg EBook of Pride and Prejudice, by Jane Austen

This eBook is for the use of anyone anywhere at no cost and with
almost no restrictions whatsoever. You may copy it, give it away or
re-use it under the terms of the Project Gutenberg License included
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$ tail -5 pg1342.txt

This Web site includes information about Project Gutenberg-tm,
including how to make donations to the Project Gutenberg Literary
Archive Foundation, how to help produce our new eBooks, and how to
subscribe to our email newsletter to hear about new eBooks.
[xdtr1@gordon-ln2 ~]$
[xdtr1@gordon-ln2 ~]$
```

# So you have a binary file...

- Some ways to gain insight
  - `objdump -S` – disassemble, display source with assembly
  - `size` – learn about size of various components of a file
  - `nm -l` -- list line numbers and symbols in a file
  - `od` – octal dump of file (if you need it!)
  - `strings` -- list ascii strings in a binary file
  - `ldd` (useful!) – displays dynamic link libraries

## Other nice utilities

- tee – takes stdin, and redirects it to stdout and files
- script -- makes typescript of terminal session – useful for capturing complex worksession
- screen –s – allows you to make detachable (and reattachable) interactive session



# Beyond the scope of this tutorial

- Shell scripting – utility programming for use on linux systems
- Different syntax for your shell of choice –
  - Popular shells – csh (tcsh), sh (bash)
  - Some critical differences and similarities in their syntax
  - One tutorial to get started –
    - <http://www.freeos.com/guides/lsst/ch02sec01.html>

# XSEDE Training Survey

- Please complete a short on-line survey about this module. The survey can be found at <http://bit.ly/CSUXSEDE>
- We value your feedback, and will use your feedback to help improve our training offerings.