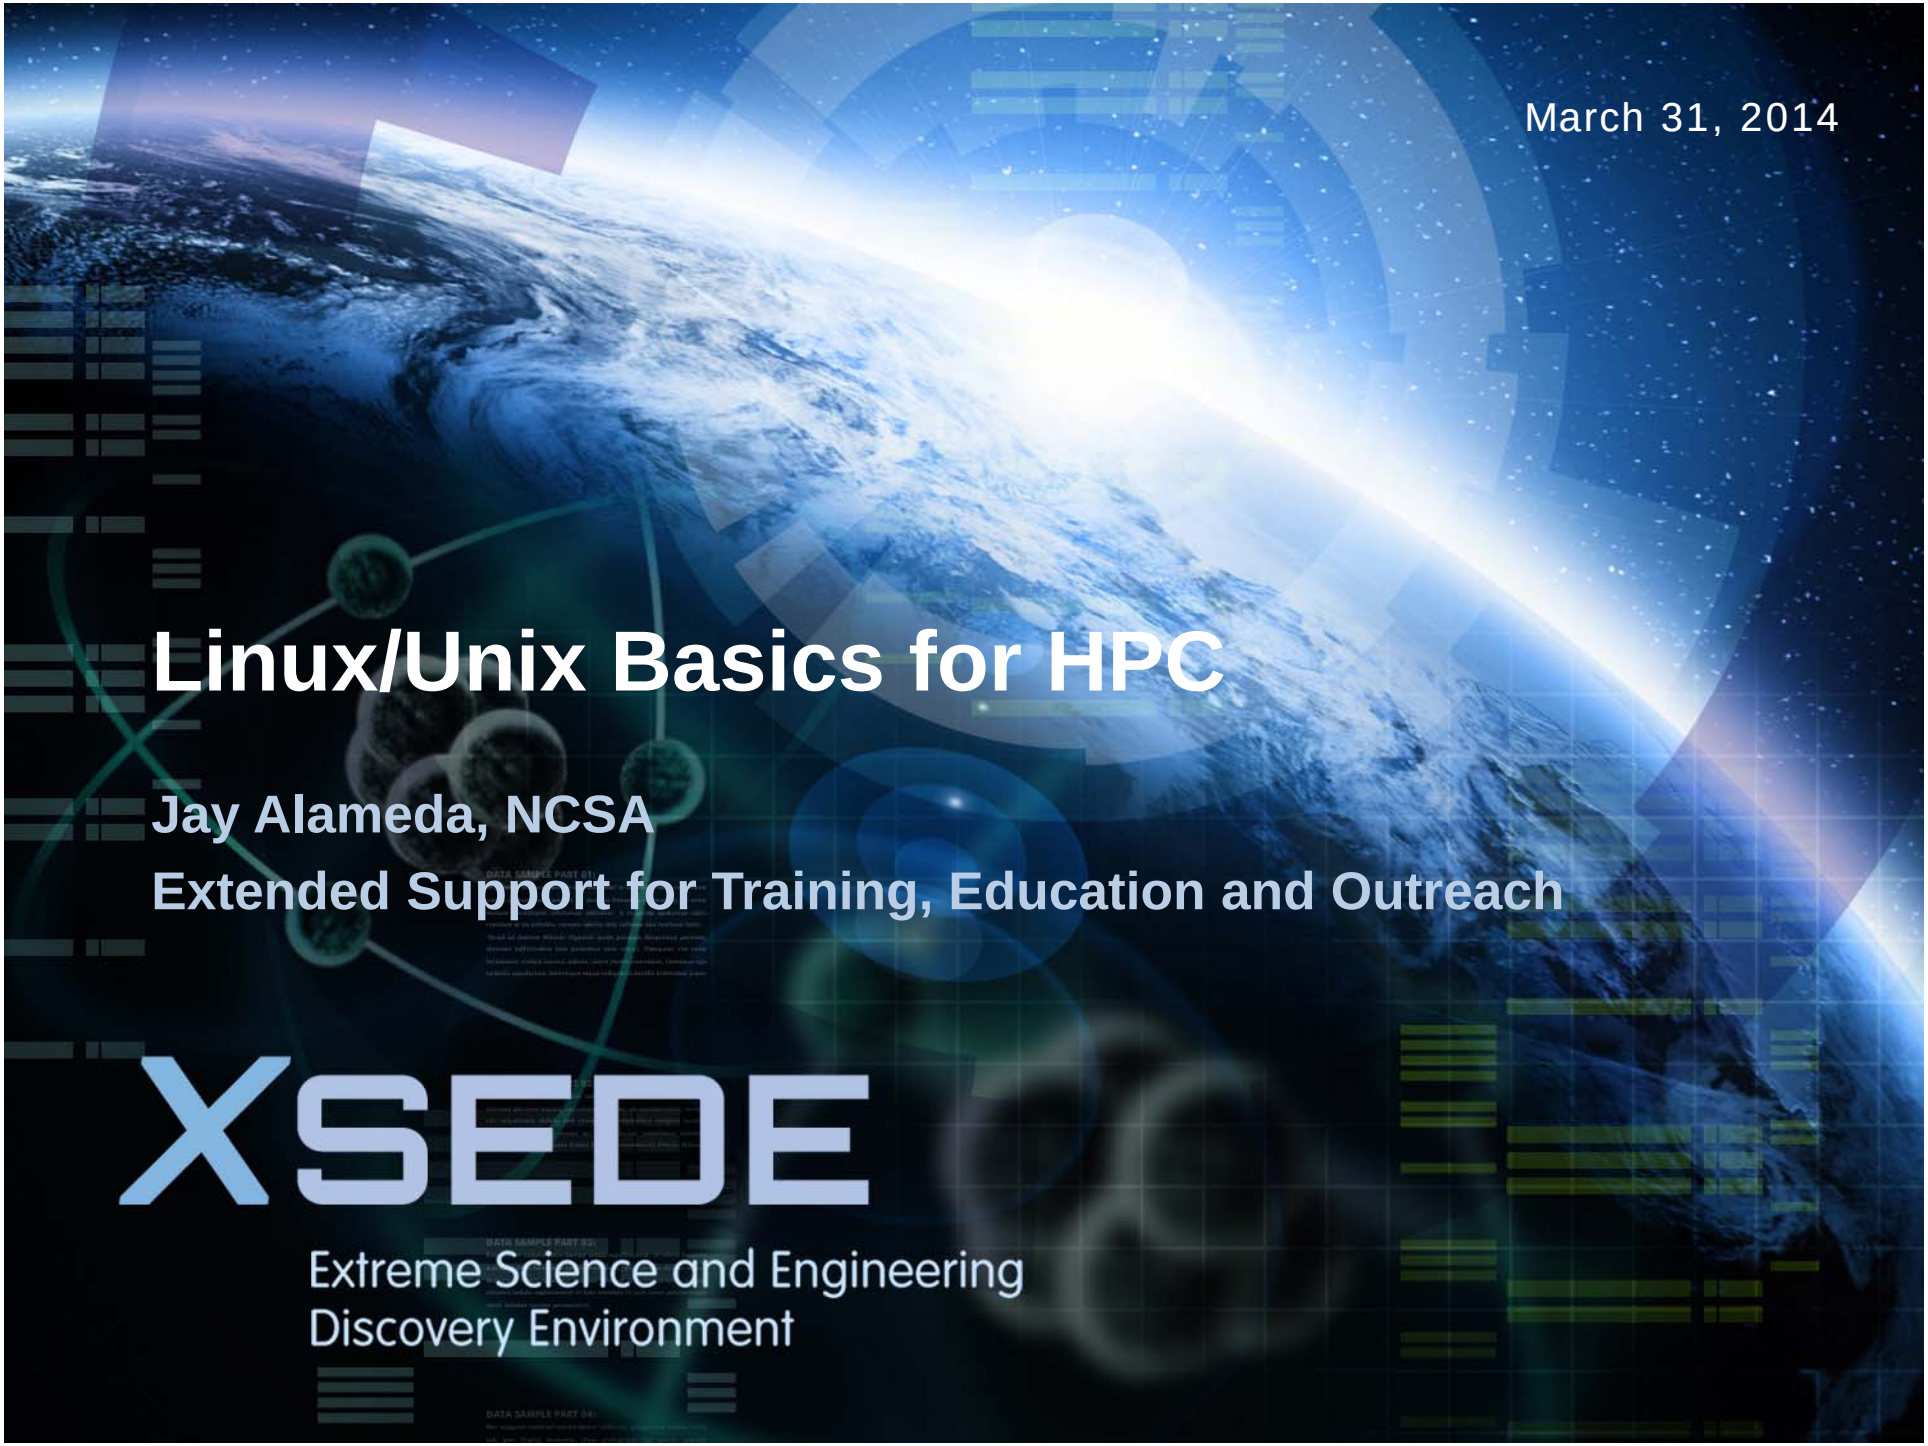
March 31, 2014
Linux/Unix Basics for HPC
Jay Alameda, NCSA
Extended Support for Training, Education and Outreach
XSEDE
Extreme Science and Engineering
Discovery Environment

XSEDE Training Survey

- **Make sure you sign the sign-in sheet!**
- At the end of the module, I will ask you to please complete a short on-line survey about this module. The survey can be found at http://bit.ly/ASUXSEDE
- I'll post this URL at the end of the course as well.

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
- NB: I use a subset of the capabilities presented (ie, pick what you find useful)

Why Linux?

- Unix could not just get along : Solaris, HPUX, 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 we 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

3D Blu-ray Disc Player with Super Wi-Fi

BDP-S5100

★★★★★ (17)

COMPARE

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

In Stock

~~$139.99~~ $119.99

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.

XSEDE

Proceed with caution

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

XSEDE HPC Resources

- All run linux –

XSEDE
Extreme Science and Engineering Discovery Environment

HOME ABOUT USER SERVICES RESOURCES EDUCATION & OUTREACH TECHNOLOGY D

User Guides

Below are links to each resource's user guide. Each guide provides information and instructions on system access, computing environment and running jobs specific to that resource. Resources are listed alphabetically within each resource type: High Performance Computing, High Throughput Computing, Visualization, Storage 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 comments at the bottom of each user guide.

High Performance Computing

Blacklight (PSC)

Gordon (SDSC)

Gordon ION (SDSC)

Keeneland (Georgia Tech)

Kraken (NICS)

Mason (IU)

Lonestar (TACC)

Stampede (TACC)

Scientific Visualization

Longhorn (TACC)

Nautilus (NICS) *production ends 9/30/2013*

Storage Systems

Data Supercell (PSC)

HPSS (NICS)

Data Oasis (SDSC)

Ranch (TACC)

XSEDE Wide File System (XSEDE) *New!!*

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.

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.
```

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/vimanual.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

Something extra: Eclipse Parallel Tools

- I'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
- 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: ln –s
  - ln-s /home/ux400689/shallow-SR1_kepler_try2/shallowshallow_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 psuedorandom 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

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!

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

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
  - 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
- 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)

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

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

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/ASUXSEDE
- We value your feedback, and will use your feedback to help improve our training offerings.