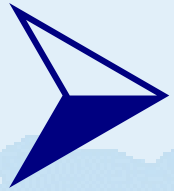

An Introduction to Linux

Jonathan Quick, Hartebeesthoek Radio Astronomy Observatory

Basic Concepts
Top Ten Commands & Procedures
Finding More Information

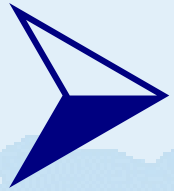

Goals

- Give you basic "survival skills"
- Show you where to find more information when you need it
- Increase your familiarity with Unix/Linux tools and the "Unix way of thinking"
- Show how to browse and manipulate files

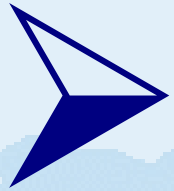

Introduction and Background

- Linux is a free implementation of the traditional Unix operating system
  - multi-process, multi-user, time-sharing
  - paging virtual memory
  - "free" in the sense of "freely distributable"; GPL
- Turns an ordinary PC into a real computer :-)
  - Our 'bootes.hartrao.ac.za' had been up 48 days when I was modifying these transparencies...

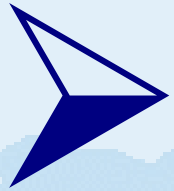

Basic Linux Concepts

- Linux Kernel
  - Base monolithic kernel + loadable modules
- Linux System / "Distribution"
  - Kernel + lots of software
- Process
- Command / Executable
- Shell / Command Line Interpreter
- X Window System

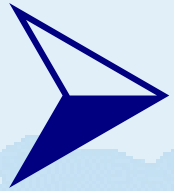

Everything is a File

- find . -type x
  - Simple files, directories
    - f (regular file), d (directory)
  - Devices
    - b block (buffered) special
      - hard disks, disk partitions, floppies
    - c character (unbuffered) special
      - serial ports, tape drives
  - Others
    - p (named pipe, FIFO), l (symbolic link), s (named network socket)

What's there or the Directory Tree

- /
  - "the root directory", "top"
- /usr/doc/bash/README.Debian.gz
  - a "pathname"
- . (current directory), .. (up one level)
- ~ (home directory), ~user
- tree -d .

The Root Directory Level

- /bin, /sbin -- basic commands
- /boot -- boot files
- /dev -- device special files
- /etc -- configuration files
- /home -- standard user home directories
- /lib -- run-time libraries

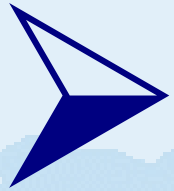

The Root Directory Level 2

- /mnt, /floppy, /dosa -- temporary mount points
- /proc -- process information
- /root -- home directory of 'root' user
- /tmp -- temporary files
- /usr -- "user level software" - see next slide
- /var -- "variable" run-time files

/usr, /var

- /usr
  - /usr/bin, /usr/sbin -- "non-essential" commands
  - /usr/lib -- run-time libraries
  - /usr/X11R6 -- the X Window System
  - /usr/doc, include, local, share, src
- /var
  - /var/log -- run-time log files
  - /var/spool -- queued files (e.g. printer)
  - /var/lock, run, tmp

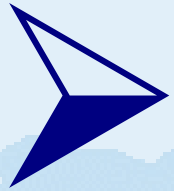

Logging in and out

- Locally:
  - Into X Window System via 'xdm' / 'kdm' / 'gdm'
  - Into a "virtual console" 'login:' prompt
- Remotely:
  - Using 'ssh' or 'telnet'
- The user name and password checked against /etc/passwd --> UID, GID
- Logging out
  - logout, exit, Ctrl-D

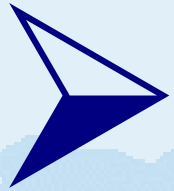

The Shell

- When you have logged in you will be dropped into a shell (or "command line interface", CLI)
- There are many alternative shells:
  - tcsh, used by FS 'oper' and 'prog' accounts
  - Bash (default), used by the 'root' account
  - others (ash, sh ,ksh, zsh,...)
- Using the shell command line you can invoke commands (executables and scripts)

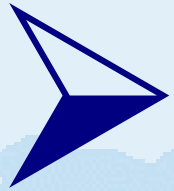

Top Ten Commands

- ls -al, locate, find . -name '*.c'
- cd, pwd, pushd, popd
- less, more, most; zless README.gz
- fgrep, grep, egrep
  - zgrep broken /usr/doc/*/README.gz
  - rgrep `hostname` /etc
- mv -i, cp -av, rm -rf
- emacs . &

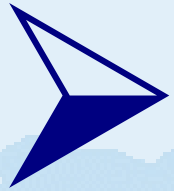

Help! I Need Help!

- command --help; man something
  - man 5 printcap
    - 1) commands, 2/3)programming, 4) /dev, 5) conf-files, 7) concepts 8) sysadmin
- apropos something or man -k something
- locate something
- cd /usr/share/doc/something;zless README.gz
- info something (or 'C-h i' in Emacs)

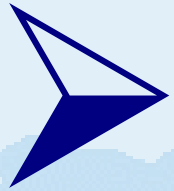

Navigating Directories

- ls -altd
  - -a all files (also "hidden" '.xyz' files)
  - -l long format (dates, protection)
  - -t most recently changed files first
  - -d list directories (not their contents)
- cd /usr2/fs; pushd /var
  - cd (back to home directory)
  - cd .. (up one level, "parent" directory)
  - popd

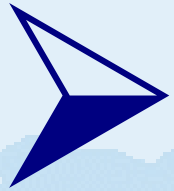

Browsing and Searching

- less file1 file2
  - SPACE forward, 'b' backward
    - 'Page Up' and 'Page Down' keys work, too
  - ESC > end, ESC < beginning
    - Like in Emacs!
  - / searchstring RET
    - n search for next occurrence
  - F wait for more lines to be appended to the file
  - q quit

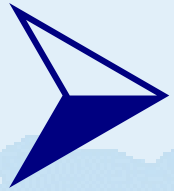

Browsing and Searching 2

- fgrep string files
  - search for a fixed string
  - -i ignore upper/lowercase
- grep, egrep
  - '^beginning', 'end$'
  - 'some.*thing'
  - 'special[0-9][0-9]'

Moving, Copying, Deleting Files

- cp -avi from to
- cp -avi file1 file2 file3 todirectory
  - -a copy all file properties & dates
  - -v verbose; show file names copies
  - -i interactive warning against overwriting
- mv -vi (just like cp)
  - both renaming and moving between directories
- rm -rfi
  - recursively, force, interactive

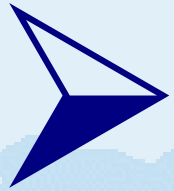

Mastering the Command Line

- TAB Completion
  - of both command and file names
- Ctrl-D Alternatives
  - (in tcsh; double TAB in bash)
- Up-Arrow for previous commands
  - "Command Line History"
  - Editing with Left/Right-Arrow, Backspace, Ctrl-D, Ctrl-A/E
- Mouse copy/paste; Shift-Insert

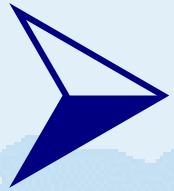

Manipulating Input/Output

- When a command executes it has 3 I/O streams:
  - standard input, output, and error output
  - by default these are bound to the "terminal" /dev/tty
- "Redirecting":
  - < input, > output, >> append output
  - error output / tcsh: >&, >>& append
  - error output / bash: 2>, 2>> append

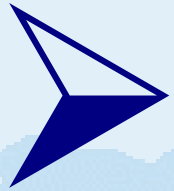

Combining Commands

- "Pipeline": 'fgrep bbc01 *.log | less'
  - Feed the output of the first command to the input of the second
  - Runs the two commands in parallel!
- Grouping commands with parentheses
  - (cd /usr/doc; zgrep -i doc *.gz) | less
- Running commands in the background
  - &, Ctrl-Z, bg, fg, jobs, %1

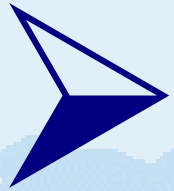

Multiple Jobs

- Any command line can be run in the back-ground by just appending '&'
- A command you have already started can be suspended with Ctrl-Z ("zap" it :-)
- You can resume a suspended job
  - With 'bg': send it to run in the background
  - With 'fg': bring it back to foreground (to your shell)
- Type 'jobs' to see what you have in background
- kill -9 %1

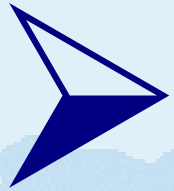

Combination Examples

- How many log files do you have?
  - cd /usr2/log; ls *.log | wc
- Look at valid pointing offset results
  - fgrep '#offset' point.log | grep '1 1' | less
- Where is RealTek mentioned in kernel source code?
  - cd /usr/src/kernel-source-2.0.36 (remember TAB!)
  - fgrep -i realtek `find . -name '*.c'` | less
- updatedb &

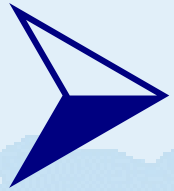

Wildcards

- The shell interprets wildcards and substitutes the resulting list of local file names on the command line
- '*' -- anything (zero or more characters)
  - Can be anywhere in the pathname!
    - /usr/doc/*/README*.gz
- '?' -- any one character
- '[0-9]', '[a-z]' -- any listed character

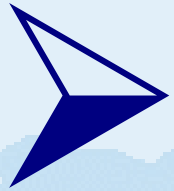

Handling Disks and Floppies

- Formatting
  - Floppies: superformat /dev/fd0
  - Hard disk partitions: mke2fs /dev/hda1
- Using (MS-DOS) floppies directly:
  - mdir a:, mcopy *.txt a:, mcopy 'a:*.vex' .
- Mounting to a mount point (=directory)
  - mount /dev/fd0 /mnt
- Use normally, unmount with 'umount /mnt'

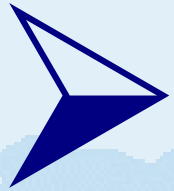

File and Directory Protection

- All files and directories are owned by one user and one group
- All files and directories have three sets of protection "bits"
  - What the owner user can do ('u')
  - What the owner group can do ('g')
  - What all the others ("world") can do ('o')
- Protection "bits" 'r' read, 'w' write, 'x' execute

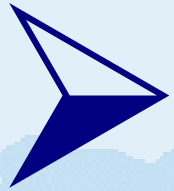

'rwx' for Files

- The protection bits affect the file contents, not the file name!
  - 'r' means the user/group/others can read the contents of this file
  - 'w' means the user/group/others can change the contents of this file
  - 'x' means that execution of this file (script or binary) can be attempted
  - Typical default is '-rwxr-xr-x'

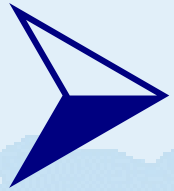

'rwx' for Directories

- The protection bits affect what the user/group/others can do to the file that contains the file names (=the directory itself)
  - 'r' means the user/group/others can read and list the file names in the directory
  - 'w' means the user/group/others can add, delete, and change the file names in the directory
  - 'x' means that this directory can be "used" (e.g. as a part of a pathname)
  - Typical default is 'drwxr-xr-x'

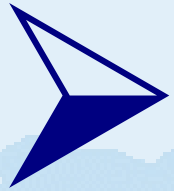

Changing Protection

- Changing protection bits
  - chmod -R ugoa=+-rwx file1 file2
    - u(ser), g(roup), o(thers), a(ll)
    - =(set exactly), +(add), -(remove)
    - r(ead), w(rite), x(ecute)
- Changing ownership
  - chown oper.rtx file1 file2 file3
  - Only 'root' can change file owner

Inodes, Hard Links, Soft Links...

- The same file can be present in multiple directories
  - Same name / different names allowed
- Directories refer to the file with an "inode" number
- Deleting a file actually only removes the name from the directory
  - If this was the last name referring to a given inode, only then the file contents is actually deleted

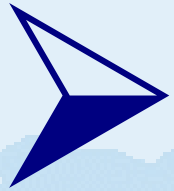

Hard Links, Soft Links

- "Additional" names for the same files are created with 'ln oldfile newfile'
  - This is a "hard link"
  - Notice the order: just like in 'cp' copy command
- A symbolic link is created in the same way 'ln -s oldfile newfile'
  - This is a "soft link"
  - 'newfile' only stores the name of the 'oldfile'

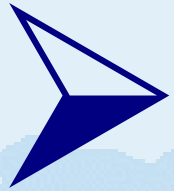

Getting System Information

- ps axf, top; kill, kill -9
- free
- df, mount
- netstat -n
- w, who
- cat /proc/cpuinfo (and others)

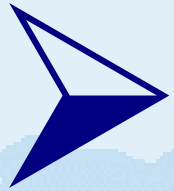

Finding More Information

- The HOWTO collection of documents
  - cd /usr/share/doc/HOWTO
- Linux Documentation Project
  - http://www.tldp.org/
- Debian Bug Tracking System
  - http://www.debian.org/Bugs/
- www.google.com

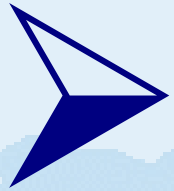

Summary

- What we have covered today:
  - Basic Linux Concepts
  - The Shell
  - Top Ten Commands
  - Help! I Need Help!
  - Mastering the Command Line
  - Manipulating Input/Output
  - Combining Commands, Using Wildcards
  - Combination Examples

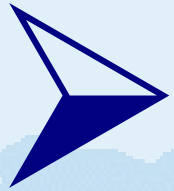

Summary

- What we have covered today:
  - The Directory Tree
  - Everything is a File
  - Handling Disks and Floppies
  - File and Directory Protection
  - Getting System Information
  - Finding More Information