

CSN09101

Networked Services

Week 4 : Basic Administration Concepts

Module Leader: Dr Gordon Russell

Lecturers: G. Russell

This lecture

- Disks
- The boot process
- User Management
- Discussions

Disks and Partitions

Disks

- /dev/hda – primary disk on first IDE controller
- /dev/hdb – slave on first IDE controller
- /dev/hdc – primary disk on second IDE controller
- /dev/hdd – slave on second IDE controller
- /dev/sda – lowest numbered SCSI device
- /dev/sdb – next lowest SCSI device
- ...

Partitions

- Rather than use the whole disk for one purpose...
- Split disk up into chunks.
- The chunks are known as partitions.
- Partitions can be primary or secondary.
- This is partially a hang-over from when DOS could only handle 4 partitions...

```
> sfdisk -l /dev/sda
```

```
Disk /dev/sda: 19449 cylinders, 255 heads, 63 sectors/track
```

```
Units = cylinders of 8225280 bytes, blocks of 1024 bytes, counting from 0
```

Device	Boot	Start	End	#cyls	#blocks	Id	System
/dev/sda1	*	0+	1274-	1275-	10240000	83	Linux
/dev/sda2		1274+	3824-	2550-	20480000	82	Linux swap / Solaris
/dev/sda3		3824+	19449-	15625-	125506560	83	Linux
/dev/sda4		0	-	0	0	0	Empty

> cat /etc/fstab

- When the system boots the fstab file tells the kernel what filesystems to load

```

UUID=d40d9bef-1306-491f-bcba-61990e1bf886 /                               ext4
    defaults          1 1
UUID=f9d23007-9414-498f-9cc6-553eeb685213 /home                          ext4
    defaults          1 2
UUID=5501a6af-ee7f-4c73-81a7-cf5c75cb8661 swap                          swap
    defaults          0 0
# /dev/sda2                swap                swap    defaults    0 0
tmpfs                    /dev/shm          tmpfs    defaults    0 0
devpts                   /dev/pts          devpts   gid=5,mode=620 0 0
sysfs                    /sys              sysfs    defaults    0 0
proc                    /proc             proc     defaults    0 0

```

Partition identifiers

> **blkid**

```
/dev/sda1: UUID="d40d9bef-1306-491f-bcba-61990e1bf886"
```

```
TYPE="ext4"
```

```
/dev/sda2: UUID="5501a6af-ee7f-4c73-81a7-cf5c75cb8661"
```

```
TYPE="swap"
```

```
/dev/sda3: UUID="f9d23007-9414-498f-9cc6-553eeb685213"
```

```
TYPE="ext4"
```


> df

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/sda2	10080520	3142968	6425484	33%	/
/dev/sda1	101086	9665	86202	11%	/boot
none	1038660	0	1038660	0%	/dev/shm
/dev/sda6	56340828	3853984	49624868	8%	/home

- “df -h” is also useful, translating bytes in MB or GB as appropriate...

* UML

- In UML, there are no IDE or SCSI drives.
- The disks are called `/dev/ubd/n` where `n` is a number
- They are actually implemented by files in the host operating system, but this is hidden from you.
 - `/dev/ubd/0` is `/`
 - `/dev/ubd/1` is swap
- However if you are using the standard QEMU/KVM virtualisation then the drives are the normal `/dev/sda1` style devices.

Disk Usage

- If you want to find out how much disk space a directory is using, the “du” command does this easily.

```
$ du -s /usr/lib
```

```
477464 /usr/lib
```

```
$ du -sh /usr/lib
```

```
467M /usr/lib
```

- “-s” is useful, otherwise it tells you about all subdirectories too.
- “-h” puts it into human readable form.

Linux Boot Process

Booting to kernel

- From switch-on:
 - PC BIOS selects a boot disk
 - BIOS loads the boot block and executes it.
 - This loads a stage 1 boot loader.
 - Stage 1 loads stage 2 loader.
 - Linux loader (e.g. Grub, lilo) runs
 - Operator selects from loader menu
 - Kernel loaded with device ramdisk

GRUB version 0.90 (638K lower / 261128K upper memory)

Red Hat Enterprise Linux ES (2.4.9-e.12smp)
Red Hat Enterprise Linux ES-up (2.4.9-e.12)

Use the ↑ and ↓ keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, or 'c' for a command-line.



The highlighted entry will be booted automatically in 7 seconds.

> cat /etc/grub.conf

```
default=1
```

```
timeout=10
```

```
splashimage=(hd0,0)/grub/splash.xpm.gz
```

```
title Fedora Core (2.6.6-1.435.2.3)
```

```
    root (hd0,0)
```

```
    kernel /vmlinuz-2.6.6-1.435.2.3 ro root=LABEL=/ rhgb quiet
```

```
    initrd /initrd-2.6.6-1.435.2.3.img
```

```
title Fedora Core (2.6.5-1.358)
```

```
    root (hd0,0)
```

```
    kernel /vmlinuz-2.6.5-1.358 ro root=LABEL=/ rhgb quiet
```

```
    initrd /initrd-2.6.5-1.358.img
```

Init.d Startup Commands

- As linux boots, it runs various system scripts.
- Eventually it runs one for your standard “runlevel”.
- The runlevel startup enables the services (like ssh and apache”) which you may want to start.
- Such startup scripts all live in:
/etc/init.d/
- For example, apache is looked after in:
/etc/init.d/httpd

/etc/init.d/*

- The scripts in init.d can:
 - Start
 - Stop
 - Restart
 - Reload
 - + a few others
- You should not call these scripts directly
 - This can interfere with different security models in use (such as SELinux).
 - Instead you need to use the "service" command.
- Example: sending "start" to the "httpd" control script in init.d is performed as follows:
 - > service httpd start

Run levels

- The run level determines what init.d files run.
- As you enter a run level services not running which should run at that run level start.
- As you leave a run level services which should not be running at the new run level stop.
- What start and stop are determined by the soft links found in the /etc/rc?.d directories.
- Usually all we need to know is the standard runlevel is 5.

> ls /etc/rc5.d

K01yum	K35vncserver	K74ypxfrd	S13portmap	S80sendmail
K05saslauthd	K35winbind	K89netplugd	S14nfslock	S90crond
K10dc_server	K45named	K95kudzu	S18rpcgssd	S90xfs
K10psacct	K50netdump	K96init.cssd	S19rpcsvcgssd	S95anacron
K12dc_client	K50snmpd	S00microcode_ctl	S20random	S95atd
K12mysqld	K50snmptrapd	S04readahead_early	S25netfs	S96init.cssd
K20nfs	K50tux	S06cpuspeed	S26apmd	S96readahead
K24irda	K54dovecot	S08iptables	S28autofs	S97messagebus
K25squid	K70aep1000	S09isdn	S44acpid	S97rhnsd
K34dhcrelay	K70bcm5820	S10network	S55sshd	S99local
K34yppasswdd	K74ntpd	S12syslog	S56rawdevices	S99mdmonitor
K35dhcpd	K74ypserv	S13irqbalance	S56xinetd	

S/K priority service-name

- S99mdmonitor :
- ls -l /etc/rc5.d/S99mdmonitor

```
lrwxrwxrwx 1 root root 19 Jul 27 13:00 S99mdmonitor ->  
../init.d/mdmonitor
```

- Starts at priority 99 – runs last
- Start mdmonitor at this runlevel

Link management

- We used to have to create the soft links ourselves to manage the run levels. This is disgusting!
- In Redhat, chkconfig does this job.

```
> chkconfig --list mdmonitor
```

```
mdmonitor 0:off 1:off 2:on 3:on 4:on 5:on 6:off
```

```
> chkconfig --levels 345 mdmonitor off
```

```
> chkconfig --list mdmonitor
```

```
mdmonitor 0:off 1:off 2:on 3:off 4:off 5:off 6:off
```

Newer service managers

- One problem with init.d is that it is linear.
 - Each service is started in turn before the next service starts.
- There are newer service managers available which use tree dependency models and parallel execution.
 - Fedora 15 uses systemd.
 - Still uses init.d for some services
 - More complex to understand but much faster to boot.
 - This is still an area under active development

The syslog

- Clicking on the syslog console output link on the control window finishes with:

Starting system message bus: [OK]

Starting mdmonitor: [OK]

- That confirms mdmonitor was started
- You can also read this from within linux using the “dmesg” command.

The xinetd super-daemon

- Some of the services (e.g. sshd) are processes.
- They start running from an rc script.
- They wait on their own for comms.
- They terminate only when the machine does down.
- Some people say this wastes resources.
- The super-server concept was born.

XINETD

- Xinetd waits for requests from the internet.
- From the requests it works out what program would like to deal with that request.
- It then starts that program running and gives it the waiting requests.
- In this way resources are only used if someone actually requests access to a particular service.

- The approach of xinetd is not as popular as services started explicitly from init.d.
- The virtual machines use xinetd to control telnet.
- If you connect to your VM with telnet, xinetd starts up the telnet daemon for you.
- Services from xinetd usually start “in.”.
- Telnet is “in.telnetd”
- /etc/xinetd.d/ contains all the services it manages.

> cat /etc/xinetd.d/telnet

```
service telnet
{
    flags                = REUSE
    socket_type          = stream
    wait                 = no
    user                 = root
    server                = /usr/sbin/in.telnetd
    log_on_failure       += USERID
    disable              = no
}
```

Terminating a process

- If you know the process id (the PID) of a program you can terminate it quickly and easily
- You send it a message (a signal) to tell it to end.
- The message to end now is called SIGKILL.

```
> ps aux | grep sshd
```

```
root 1796 ..... /usr/sbin/sshd
```

```
> kill -s SIGKILL 1796
```

User Management

User Management

- A wide topic...
 - Adding/Removing/Changing current users
 - Default Scripts
 - Global Scripts

Manual Creation

- User entries in passwd,shadow, group,gshadow.
- Home directory in /home.
- Copy basic .files into their home directory.
- Make new user own their own directory and files.

> **adduser gordon**

- This does all the magic for you.
- It copies the default .files from /etc/skel/

```
> ls -a /etc/skel/
```

```
.bash_logout .bash_profile .bashrc .gtkrc .kde
```

- In bash, .bashrc is executed in non-login shell, and .bash_profile in a login shell.

Skel files

- These files are the basic .files created for a new user.
- Users are free to edit these when they log in.
- This allows them to control their own path, env, and other settings (such as aliases).
- However, if you install a new package which needs something set for each user at login, editing all these copies by hand would be tiresome.

> **ls /etc/profile.d**

```
colorls.csh  gnome-ssh-askpass.csh  krb5.csh  less.csh  vim.csh
colorls.sh   gnome-ssh-askpass.sh   krb5.sh   less.sh   vim.sh
glib2.csh   kde.csh                 lang.csh  qt.csh   which-2.sh
glib2.sh    kde.sh                  lang.sh   qt.sh
```

- If you log in with bash, all the .sh files are executed before your .files
- If you log in with csh, all the .csh files are executed before your .files.

> **cat /etc/profile.d/vim.sh**

```
if [ -n "$BASH_VERSION" -o -n "$KSH_VERSION" -o -n
    "$ZSH_VERSION" ]; then
    # for bash, pdksh and zsh, only if no alias is already set
    alias vi >/dev/null 2>&1 || alias vi=vim
fi
```

- I.e. if this is bash, and you have not set an alias for “vi”, then set one to run “vim” when you type “vi”.

Example

- Create a user jim, in group staff
- But how to set the group?
- You could do:

```
$ man adduser
```

- Usually commands also take the flag “-h”

```
$ adduser -h
```

```
adduser: invalid option -- h
```

```
Usage: useradd [options] LOGIN
```

Options:

-b, --base-dir BASE_DIR base directory for the new user account

...

-g, --gid GROUP force use GROUP for the new user account

...

```
$ adduser jim -g staff  
$ tail -1 /etc/passwd  
jim:x:502:100::/home/jim:/bin/bash  
$ grep 100 /etc/group  
staff:x:100:
```

Moving a uid or gid

```
$ tail -1 /etc/passwd  
jim:x:502:100: /home/andrew:/bin/bash
```

```
$ grep 100 /etc/group  
staff:x:100:
```

```
$ ls -lnd /home/jim  
drwx--x--x. 6 502 100 4096 Mar 27 11:53 /home/jim
```

```
$ ls -lan /home/jim  
drwx--x--x. 6 502 100 4096 Mar 27 11:53 .  
drwxr-xr-x. 6 0 0 4096 Jul 13 2007 ..  
-rw-----. 1 502 100 10553 Apr 8 14:48 .bash_history  
-rw-r--r--. 1 502 100 747954 Dec 20 2007 planner.zip
```

Useful Commands

\$ chown jim.staff filename

\$ chown jim filename

\$ chgrp staff filename

When a User logs in

- When a user logs in the appropriate . files are executed (.login, .cshrc, etc).
- If you want to change to a different user, you could log out and log in again, or you could do
 - su – gordon (change to the gordon user)
 - su – (change to root)
- Without the “-”, you still change users, but the . scripts don’t get executed.
- To go back to the previous user, press CTRL-D

FILE SEARCHING

A file CONTAINING something

- You are looking for a file containing “gordon”
- You think it is in /etc/ something

```
$ grep “gordon” /etc/*
```

```
/etc/group:gordon:x:500:
```

```
/etc/group-:gordon:x:500:
```

```
/etc/gshadow:gordon:!!::
```

```
/etc/gshadow-:gordon:!!::
```

```
/etc/passwd:gordon:x:500:100:Dr Gordon ...
```

A FILENAME containing something

- You know somewhere in /etc there is a filename with the word “host” in it.

```
$ find /etc -name '*host*'
```

```
/etc/hosts.deny
```

```
/etc/ghostscript
```

```
/etc/ssh/ssh_host_dsa_key.pub
```

- Find can find on a range of things, not just names. Other things include sizes, permissions, types, ownership, and combinations of tests.

Find to do something

- Usually find prints the things which match.
- You can get it to execute instead.
- Here you want to find all files called core, and delete them:

```
$ find . -name core -print -exec rm {} \;
```

- I didn't invent the syntax, so don't blame me...

Discussion

- A user keeps getting logged out each time they log in... why?

Discussion

- A user find their ls command is broken... why?

Discussion

- Here are some past exam questions you should now be able to answer:

Question 1

- What is the function of “su”, and what is the difference between “su – gordon” and “su gordon”?

Question 2

- What type of files would you expect to find in /sbin?

Question 3

- Consider the following line:

`gordon:x:44:`

In which file in `/etc` would you expect to see such a line, and what does it mean?

Question 4

- The following commands are typed on a Unix computer.

```
$ mkdir temp  
$ cd temp/  
$ mkdir txt.txt/  
$ cd txt.txt/  
$ touch hello  
$ cd ..  
$ ls *.*
```

What is printed on the screen in response to the last line of the commands?