

CST8177 – Linux II

Disks, File systems, Booting

Todd Kelley

kelleyt@algonquincollege.com

Today's Topics

- ▶ bind mounts
- ▶ quotas
- ▶ Installation Disk rescue mode
- ▶ more booting

Bind mounts

- ▶ A bind mount is used to mount a directory onto a mount point: `man mount`
- ▶ use the “bind” option for the mount command
 - # `mount -o bind /some/dir /anotherdir`
 - now `/some/dir` and `/anotherdir` are the same directory
- ▶ Be careful with bind mounts, because they make it possible to form cycles in the file system
- ▶ e.g. dangerous: `"mount -o bind /home /home/user/dir"`
 - serious repercussions for
 - `rm -rf /home/user` # will remove all of `/home`
 - `find /home/user` # will never stop
 - any program that recursively descends directories

Bind mount examples

- ▶ make an inaccessible directory accessible:
 - `mount -o bind /home/user/private/public /public`
- ▶ make disk space in one file system available in another file system
 - suppose you have a large separate file system with lots of free space on `/var`, and root file system with `/home` is nearly full:
 - `mkdir /var/local/home/{user1,user2}`
 - move contents of `/home/{user1,user2,...}` to `/var/local/home`
 - `mount -o bind /var/local/home /home`
 - beware: new `/home` has same mount options as `/var`

Bind mount examples (cont'd)

- ▶ share directories across chroot environments
 - `mount -o bind /dev /home/user/myroot/dev`
 - `chroot /home/user/myroot/dev`
 - in the chroot-ed environment, `/dev` will be the same as the un-chroot-ed `/dev`

Quotas

- ▶ https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/ch-disk-quotas.html
- ▶ Example: enabling quotas on /home
- ▶ /etc/fstab: usrquota,grpquota mount options for file system containing /home
- ▶ quotacheck -cug /home
 - c: don't read quota files, create new quota files
 - u: do user quotas
 - g: do group quotas
- ▶ edquota username or setquota -u user soft hard isoft ihard fs
- ▶ edquota -t # edit grace period
- ▶ quotaon -vaug # turn quotas on
- ▶ repquota -a # report on quotas
- ▶ quotaoff -vaug; quotacheck -vaug; quotaon -vaug #single user mode

Installation DVD for rescue mode / Live CD

- ▶ There are dangers associated with doing file system operations on "system directories" that might be used in system operation.
- ▶ For example, many programs will use the shared libraries in `/usr/lib`, which disappear if we move `/usr`
- ▶ Also, there may come a time when the system won't boot properly: MBR corrupted, bad entry in `/etc/fstab`, inconsistent / file system

linux rescue

- ▶ To boot into rescue mode
 - ensure BIOS boot order is set for booting from CD/DVD before Hard Drive (even in VMware – F2 to enter setup)
 - insert the installation DVD into drive (or the iso image into the virtual DVD drive)
 - boot the system
 - type "linux rescue" at the prompt
 - Linux will run "from" the DVD (Live CD), not from your file systems (your system is not running)
 - It will offer to search for and mount your Linux file systems on `/mnt/sysimage`

linux rescue (cont'd)

- ▶ The Live CD Linux system can see your hard drives, and this is how you can repair or alter what is on those hard drives
- ▶ You need to remember that a Live CD Linux system is running from its own root filesystem (like dual boot?), so this means
 - the users are different /etc/passwd /etc/shadow, etc (or should I say all of /etc) are different
 - the services running, firewalling, and so on, are different

Rescue mode / Live CD

ramdisk

	/		
etc/ passwd shadow		bin/ ls bash	dev/ sda VolGroup00/ LogVol00

/dev/VolGroup00/LogVol00

	/		
etc/ fstab passwd		home/ idallen/ donnelr	dev/ VolGroup00/ LogVol00

linux rescue example 1

- ▶ Fix /etc/fstab
 - mount /dev/VolGroup00/LogVol00 /mnt/sysimage (if it isn't already mounted)
 - vi /mnt/sysimage/etc/fstab
 - fix the problem
 - save and quit
 - exit

linux rescue example 2

▶ fix MBR

- # our root file system is mounted on /mnt/sysimage
- chroot /mnt/sysimage
- # now / is our root file system!
- # our boot filesystem is mounted on /boot
- grub-install /dev/sda

▶ Whoa! That chroot thing was neat

- chroot runs a program or interactive shell using the named directory as the root directory
- Default program is `${SHELL} -i`
- This simulates running off our system's root file system without going through its boot process

Growing a filesystem

- ▶ That LVM tutorial link again:
 - http://www.howtoforge.com/linux_lvm
- ▶ Because Red Hat's installer used Disk Druid to set up LVM and installed the root file system on a Logical Volume, we can
 - add a hard disk
 - create a partition on that hard disk
 - # or, maybe we already had an unused partition, such as a reclaimed Windows partition
 - set up that partition as a physical volume
 - add that physical volume to our Volume Group
 - grow the Logical Volume on the Volume Group
 - grow the file system on that Logical Volume

Growing a file system (cont'd)

- ▶ set up our "new" or "spare" partition as a physical volume for LVM:
 - `pvcreate /dev/sdb1`
- ▶ Add this new physical volume to a volume group (in this case VolGroup00):
 - `vgextend VolGroup00 /dev/sdb1`
- ▶ See how many free extents (Free PE) are available in this volume group (VolGroup00)
 - `vgdisplay`

Growing a file system (cont'd)

- ▶ Suppose the previous "vgdisplay" command showed that VolGroup00 had 319 free extents ("Free PE") and we use them all:
 - `lvextend -l+319 /dev/VolGroup00/LogVol00`
- ▶ Now LogVol00, which contains our root file system, is bigger, but the filesystem is still the same size.
- ▶ Grow the filesystem (ext3) to fill the added space:
 - `resize2fs /dev/VolGroup00/LogVol00`
- ▶ Use `df` command so see we have bigger root file system now!

Booting

- ▶ http://teaching.idallen.com/cst8207/13w/notes/750_booting_and_grub.html
- ▶ page numbers for Fifth Edition Sobell:
 - Chapter 11: 424–431
 - Chapter 15: 551–552

Booting Sequence (CentOS)

- ▶ Power button pressed
- ▶ BIOS
- ▶ POST
- ▶ MBR : contains grub stage 1
- ▶ grub stage 1 : to find grub stage 2
- ▶ grub stage 2 : to launch kernel
- ▶ kernel running
- ▶ init process (PID 1) : consults inittab
- ▶ /etc/inittab
- ▶ /etc/init.d/rc.sysinit
- ▶ /etc/rc.d/rc 3 : assuming default runlevel 3

SysVinit

- ▶ `/etc/init.d/*`
 - these are scripts for starting, stopping, restarting services
- ▶ `/etc/rc.d/rc.N.d/*` #where N is a runlevel
 - these are symbolic links to service's script
 - begins with K means service should not be running in that runlevel: call it with "stop" argument
 - begins with S means service should be running in that runlevel: call it with "start" argument
- ▶ `chkconfig` maintains these scripts