

CMDS

- [root] export USER=<username>; useradd \${USER} ; passwd \${USER} ; /usr/sbin/setquota -u \${USER} 20000000 30000000 0 0 /export ; # usermod -G vault \${USER}
- force update using “411” line below / or “rock sync users”, but that messes up home dirs

TROUBLESHOOTING

- If nodes won't properly boot (usually caused by unclean shutdown)
 - Remove 'old' boot image on master, for this node:
/tftpboot/pxelinux/pxelinux.cfg/C0A800F? where ? = EDCBA9 for nodes 012345
 - Manually reboot node (force PXE boot with F12 if needed)
 - Prompts for “language”? Change master root dir permissions: `chmod a+rx ~root`
- If network is slow (dmesg | grep eth says “Link is up at 10 Mbps, half duplex”) force with ethtool:
 - `/sbin/ethtool -s eth1 speed 1000` (if fails, set `/sbin/ethtool -s eth1 autoneg off` then on)
 - for eth0 `ifdown eth0; ifup eth0`
- Keys: [ctrl-ctrl] = bring up display server screen ; [ctrl-alt-f2/f3/f4] = emergy sh, dmesg, etc.

DONE

- ROCKS user guide: <http://www.rocksclusters.org/roll-documentation/base/5.1/>
- BACKUP
 - Run `~leeuwen/bin/backup-vault-to-sara` (tar + sftp to /archive/joerivl/drop/)
- Changed aliases c0 .. c5 ; drop ; f for frontend
 - [root@drop ~]# for i in `seq 0 5`; do rocks add host alias compute-0-\$i c\$i; done
- RAID: set up raid check + e-mail
 - Smart Array P600
 - 'hpacucli' rpm (e-mail frank) + RPMs compat-lidstdc*
 - per <http://www.mulder.franken.de/workstuff/> (changed to incorporate “Bay Numbers” vs HW drives)
 - `/etc/cron.hourly/raidcheck` + `/root/bin/check-hp-raid-status.pl`
 - front-end
 - `e2label /dev/sda5 /scratch ..etc`
 - edited `/etc/fstab` + run `sudo mount -a` ; /export now on 4TB raid, /data on 7TB raid; scratch=local disk
 - turned off automount ; `/etc/auto.master` now empty ; `/etc/exports` now `/data/` and `/exports/`
 - using `fdisk` “type fd” and `mdadm -create` etc, added 2×750 software RAID0
 - `/sbin/mdadm -Ac partitions -m 0 /dev/md0` (to bring up after reboot)
 - `/sbin/fsck.ext4 /dev/md0`
 - `mount -t ext4 /dev/md0 /scratch2`
- RAID INITIAL BUILD with `/usr/sbin/hpacucli` ([Guide, p44+](#))
 - `ctrl slot=1 pd all show`
 - `ctrl slot=1 ld all delete`
 - `ctrl slot=1 create type=ld`
`drives=1E:1:1,1E:1:2,1E:1:3,1E:1:4,1E:1:5,1E:1:6,1E:1:7 raid=6`
 - `ctrl slot=1 create type=ld`
`drives=1E:1:8,1E:1:9,1E:1:10,1E:1:11,1E:1:12 raid=0`
 - `parted /dev/cciss/c0d1` (after some fiddling to delete old partitions)

- `mklabel gpt`
 - `mkpartfs primary ext3 0 -0`
- RAID0 REBUILD
 - Kill processes (`/sbin/fuser -m /dev/cciss/c0d1p1`, then kill with `-k`), then unmount
 - Blink LED (`ctrl slot=1 pd 1E:1:10 modify led=on`), Replace disk,
 - Re-enable logical drive (`ctrl slot=1 ld 2 modify reenable forced`) (from [Cheat sheet](#))
 - Make new file system, `/sbin/mkfs.ext3 -L /data /dev/cciss/c0d1p1`
- RAID PHYS CONFIG
 - Drive 1 = left, top; Count seems to be 1-3 = left column, top to bottom. RAID0=two lowest in second-to-right column + entire right column
- CROSSMOUNTS
 - removed automounter. changed `/etc/exports` and node script to hard mount. Homedirs are `/export/home/` (make sure this is correct in `/etc/passwd`)
 - `make -C /var/411 clean; make -C /var/411; make -C /var/411 force; cluster-fork 411get --all`
 - `/etc/rc.d/init.d/nfs restart; /etc/rc.d/init.d/nfs restart; make -C /var/411; service autofs reload; exportfs -a`
 - `mv /usr/local /usr/local.rocks`
 - Write speeds: RAID6 10MB/s, RAID0 40MB/s, local disk 40MB/s (10*2GB file)
 - after `/usr/sbin/hpacucli ctrl slot=1 modify drivewritecache=enable` (`..disable`)
 - Write speeds: RAID6 70MB/s, RAID0 130MB/s, local disk 40MB/s (10*2GB file)
 - Read speeds: RAID0 160MB/s, RAID5/local 140MB/s
 - Increased number of NFS threads from 8 to 32 (link [12](#))
- VNC/ETC
 - RealVNC server
 - FreeNX (yum install nx freenx ; <http://wiki.centos.org/HowTos/FreeNX>)
- GRID ENGINE
 - Some HOWTO links: [1](#), [2](#), [removing state 'E'](#)
 - `qconf -mq all.q` to reduce numbers of slots on nodes
- NODE CONFIG
 - [replace-partition.xml](#); `extend-compute.xml`
 - `rocks remove host partition compute-0-0`
 - `cd /export/rocks/install ; rocks create distro ; ssh c0 "/boot/kickstart/cluster-kickstart-pxe" ; #OR; ssh c0 "/boot/kickstart/cluster-kickstart"`
 - `rocks remove host partition compute-0-1 #etc; cluster-fork -n 'c%d:1-5' 'rm /.rocks-release; /boot/kickstart/cluster-kickstart-pxe' ;`
 - (removed `/tftpboot/pxelinux/pxelinux.cfg/C0A800FE`)
 - `fftw compile + /usr/local; made fftw(l)(f) wisdom, added custom paths, /etc/hostname`
 - tempo from gasp in `/usr/local/src/tempo`
 - installed subversion by RPM, `cfitsio-3.140` from source: `./configure --prefix=/usr/local`
 - `pgplot` from source `$PGPLOT_DIR`, g77 from RPM
 - `ln -s /usr/local/include/pgplot/libpgplot.so /usr/local/lib`
 - also compiled **gfortran** version (g77 and gfortran, for f90, not compatible) per [link](#)
 - in `/usr/local/include/pgplot-gfortran`
 - Built LAPACK + ATLAS from source ([HOWTO](#))
 - `../configure -Fa alg -fPIC --with-netlib-`

- lapack=/usr/local/src/lapack-3.2.1/lapack_LINUX.a
 - cd /usr/local/lib; ln -s /usr/local/src/ATLAS/ATLAS.x86_64/lib/lib* .
- Numpy, SciPy from svn
 - rm -Rf build ; python ./setup.py build --fcompiler=gnu95; python ./setup.py install --prefix=/usr/local/
 - + (Nose from [web](#))
- iPython, matplotlib (+tkinter), PyFFTW, ctypes, git from source
- presto from ~~svn~~ tar from github ; keep changes in old Makefile
 - (had to link libs2g.so to /usr/lib64), /usr/local/src/presto
 - ppgplot: ppgplot_libraries = ["cpgplot", "pgplot", "X11", "png", "m", "g2c"] + ppgplot_library_dirs = ["/usr/X11R6/lib"]
- Set UseDNS to NO in /etc/ssh/sshd_config for master+nodes, after very slow logins after IP changes to front node
 - Which turned out to be caused by outdated DNS server in named.conf and resolve.conf: rocks set var Kickstart PublicDNSServers 195.169.63.49 in python/setup.py
- NODE PACKAGES
 - cd /export/rocks/install/contrib/5.1/x86_64/RPMS
 - pgplot i386 & x86_64?
 - <http://rpm.pbone.net/index.php3?stat=3&search=pgplot&srodzaj=3>
 - check dependencies with yum; downloader from
 - <http://www.cyberciti.biz/faq/yum-downloadonly-plugin/>
 - look/google for EL5 or FC9/10, x86_64 (+ potentially i386)
 - check
 - <https://www.icts.uiowa.edu/confluence/display/ICTSit/ROCKS+5.1+Documentation>
 - to make your own
- ADMIN
 - User quota ([1](#), [2](#), chmod 644 quota file, quotaon -a)
 - groupadd vault; usermod -G vault leeuwen; #etc

DOING

- LIGHTPATH
 - installed eth2 on c4. edited /etc/sysconfig/network-scripts/ifcfg-eth2 to static IPADDR=192.87.39.129, NETMASK=255.255.255.248, MTU=9000
 - On command line, added /sbin/route add 145.100.26.152 gw 192.87.39.130 for Huygens

2013 BOOT Kernel 2.6.18-238.9.1.el5 got beyond RAID wait after 5 minutes.

.18. Bad IRQ, kernel panic (as was original problem)

TODO

- Read OAK topics and redo

From:

<https://www.astron.nl/astrowiki/> - **AstroWiki**

Permanent link:

https://www.astron.nl/astrowiki/doku.php?id=drop_setup&rev=1382447604

Last update: **2013/10/22 13:13**

