

# How to load XRAID disk sets

- Ensure no one has left open files on the disks: Run `lsof` in the directories where the data is located.
- You may also want to see whether anyone is using the machine. Use the “w” command.
- Disable NFS on the computer to which the destination xraid is attached (may not be necessary, but always check). The safest way to do this, especially on cuppa02 (as it also exports /home and /nfs/apps) is to edit the nfs exports table `/etc/exports` and comment out the xraid entries (by adding a '#' at the start of the line), then reload the nfs server (hopefully this will not cause any interruption to its operation)

```
>sudo /etc/init.d/nfs-kernel-server reload
```

- We don't typically use NFS to export disk data, but we can by simply editing `/etc/exports`.
- If one of your partitions IS NFS mounted at the time, this may just produce errors. In this case you may have to manually figure out which node(s) the partition is mounted on. I don't know if there's an easy way to do this. You'll need to visit each node, check if the partition is mounted (eg `ls /nfs/xraid0?/?_?`), and if so, unmount it (`sudo umount /nfs/xraid0?/?_?`).
- Stop the disks and unmount the file systems.

```
> sudo umount /exports/xraid0X/*
```

- If any of the xraid file systems cannot be unmounted, you will have to figure out which process is hanging on before proceeding.
- Power down the xraid (this could be done in software from the comfort of your office through “XRAID Admin Tools”, which aren't yet installed).
- Remove disks and replace in appropriate cases in the correct order!
- Insert new disks, #1 on the left, #7 on the right. Check contacts on disks (and in XRAID if possible) before inserting into XRAID chassis. Disks slide in and then require a final push, you should hear them lock (thump!) into place, at which stage with the handle depressed they will be flush with the chassis.
- Power on the XRAID. This should be done in the cluster room so you can watch for alarms/red lights. The power button is on the BACK of the chassis. Hold it in for a couple of seconds.
- If all disks come up green, you can go back to your office! Otherwise, open XRAID Admin Tools and identify the source of the problem. If a disk has failed, it might be worth trying to power-cycle and re-insert. Sometimes the disk itself is fine, but not properly seated in the chassis. If a spare disk is required, you should be able to insert it in place of the failed one (even while the chassis is running) and the array will automatically rebuild using the new disk. If using an OLD disk or trying to re-insert an improperly seated disk, you may have to explicitly make it available to the array, because it will already contain RAID information. See [Rebuilding XRAIDs](#).
- Once the XRAID is running, you can reload the SCSI devices and mount their file systems. *There should be no reason to reboot the host node.*

```
> sudo /nfs/apps/vlbi/refresh_xraid
```

- The program will print out a list of detected APPLE devices and any error messages if the refresh was not successful.
- Create the mount points for the xraid (if the SCSI device has changed we'll need new mount points accordingly)

```
> sudo /nfs/apps/vlbi/udevrules.pl
```

- Mount the new file systems using

```
> sudo mount /exports/xraid/?_?
```

where ?? is l\_1, l\_2, l\_3, r\_1, r\_2, r\_3 as required (left and right referring to disk banks in a chassis).

- We *believe* that this will mount the first left device in **l\_1** and so forth. However you should check this by ensuring that the data you expect to be on the device is in fact there before you try doing anything with it! ABSOLUTELY NO WARRANTY and all that.
- If you are concerned, try mounting /dev/sd?1 to the exports directory you want it to be, and make sure the left/right set of lights come on when you run the command. I believe devices always come up in the same order within a disk set, but disk sets are not always recognised in the same order (left-right, right-left, alternating) in a chassis.
- Now restart NFS if desired, by editing /etc/exports and removing the # character from the start of each xraid line, then run

```
> sudo /etc/init.d/nfs-kernel-server reload
```

and on each node where you need to use these data, mount the device over NFS (also see below, though it may be out of date)

```
> mount /nfs/xraid0?/?_?
```

(Note that sudo access is not required for this).

## NFS mounting

**Note** This is not recommended now. The correlator will start datastream processes on the correct local host node provided this is specified in the corresponding line of the machines file (this was apparently not always the case in the past).

- It is now possible to mount the Xraids over nfs as any user, with e.g. `mount /nfs/xraid01/l_1`. To make life easier, there are shell scripts to do it in /home/corr/LBA/scripts/ `mountxraids.sh` and `unmountxraids.sh`
- On cuppa01, type `cssh cuppa` to get a multi-window command prompt, so the scripts can be run simultaneously on all cuppas.
- To change disks, after unmounting from NFS, it seems to be necessary to restart the nfs server on the local host node:

```
sudo /etc/init.d/nfs-kernel-server restart
```

in order to

```
umount /exports/xraid01/l_1
```

etc. without getting “device busy” errors. Be a bit careful doing this... try to check that no one is

running stuff first. If other people are logged on may get stale file handle problems.

[Back to XRAID menu](#)

From:

<https://www.atnf.csiro.au/vlbi/dokuwiki/> - **ATNF VLBI Wiki**

Permanent link:

<https://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/correlator/loaddisks?rev=1229413215>

Last update: **2008/12/16 18:40**

