

XRaid Drives for VLBI Recorders

Please note this document is not complete and errors may be present. Send comments and corrections to cwest@astro.swin.edu.au and I will correct them for the next revision.

Author: Craig West

Revision: 0.1

Removing the drives

Removing the drives should be a simple exercise, but care must be taken. Firstly you need to power down the host machine (or machines). These are the computers that are connected to the XRaids via fibre channel – see the descriptions below.

Once all the affected machines are physically powered down, you can proceed to power down the XRaids. Note: You should see the two green lights at the bottom of the central column on the front of the Xraid are no longer lit.

To **power down** the xraid manually you can press and hold in the power button on the rear of the XRaid chassis. You will need to hold it in for a few seconds, lights will flash on the rear of the chassis and when the lights go out you can release the button.

Before removing any drives, always make sure you have powered down both the machine and the XRaid. Don't remove the fibre channel cables and think that will solve the problem, it could cause the data on the drive to be corrupted.

Removing the drives can now take place. No lights on the drive is a good sign. Press the front of the drive and a handle will pop out. Put your hand in the handle and pull. You should not have a drive in your hand. The XRaid has two halves, ensure you remove all drives from a given half before power is reapplied.

Note: Failure to remove a whole raid set, or removing part of the raid set while it is powered will result in the entire raid set being compromised. It can take 40+ hours to rebuild a raid5 setup. Please be careful.

Installing the drives

This is one of the parts of the whole system where you need to be very careful. By not following the directions carefully you could destroy the raid set and waste valuable time in rebuilding it. The drives need to be pushed in firmly, it will be mentioned many times, but it can't be stressed enough. No hammers are needed, but a push that is firm will not destroy the drives.

Firstly, ensure the XRAid is powered down and the drives are removed from the section you are going to install the new drives into <see above>.

Now you should always install the drives in the correct order. Depending on where the drives are from they are labelled different. The drives always come in groups of 7. Each set currently starts with some alpha-numeric string to define the set, and then a number at the very end that defines each drive in the set. So anything numbered 1-7 or 8-14 should be installed together. That is two separate groups of numbers 1-7 is a raid set and 8-14 is a different raid set. It is hoped that we will use a more common and logical drive set in the future. Number then from left to right in increasing order.

In order to load the drives you simply push the drives into the chassis, they will only go in one way, if they don't go all the way in first time, rotate them around and try again. They side with the disk label (not the one on the handle, but the real drive label) should face to the left. Once you push the drive in you should hear a soft click. Push all the drives in so they are flush with the front of the case when their handles are closed. Now you need to re-push the drive again to ensure it is all the way in. To do this I recommend pushing the handle quite firmly and then releasing (the handle should now be out). Now push the handle firmly again, and it should stay in. This is a double firm push on EACH drive one at a time. It might seem like over kill but I can't stress enough how important it is to get them in correctly.

Once all the drives are inserted, you may only need to do one side of the XRAid, or perhaps both, you can power the unit up. Don't power it up until you have finished installing all the drives.

To power up the unit press and hold the power button on the rear of the XRAid chassis for about 1 second, you will see the lights turn on and it will power up. Once all the lights have turned on for the drives, it takes a short while for this to happen, you can then turn the computer back on.

Mounting the drives

Mounting the drives sounds simple in theory, you should just need to type `mount /data/xraid0` and `mount /data/xraid1`. However depending on drive sizes and the like it will not be as simple as this. At the time this document was written there were only 2 mount points available for the xraids – listed above. This will be corrected in the future.

If you need to manually mount a different drive, please identify that drive correctly first (see below) and then mount it as root using a command like:

```
# mount /dev/sdc1 /data/xraid0
```

It should be noted that both `/dev/sdc1` and `/data/xraid1` should not be already mounted when you do this. You can see if they are mounted by checking the output from “`mount`”. If they are already mounted, they should be firstly unmounted using the “`umount`” command.

```
# mount
... other drives ...
/dev/sdc1 on /data/xraid0 type ext3 (rw,noexec,nosuid,nodev)
```

This shows that `/dev/sdc1` is mounted as `/data/xraid0`

The format at the chosen site may be slightly different to this, but we will work to ensure they are the same and examples updated in this document.

Unmounting the drives

Unmounting the drives is easy. You can type `umount /data/xraid0` (or associated drive). You will not be able to unmount the drives if there is a program still running in that directory.

As this document doesn't detail how to remove the drives without power down a machine (future versions will mention this). It is ok for you to simply power down the machine without unmounting the drive. You may want to ensure that no-one is using the machine first.

Identifying the partitions

One of the hardest part of knowing how to mount the drives is knowing what they are labelled, and what the system can tell you. Because we are using Linux 2.4 kernels we are unable to access any drives that are bigger than 2TB. This causes a problem when using 7 drives that are 400 or 500GB in size (with raid5). As such we modify the drives with the Apple XRaid to “slice” the drive in 2. This then presents two drives to the system that are smaller than 2TB each... problem solved.

The hard part is knowing what these drives are called. You can find this out easily enough, and if you know what size drives you loaded things should be easier again. What will happen when the machine boots (provided the cables are connected correctly, see below) is that the left/upper drive set will be detected first, then the right/lower set. As each set is detected, it will detect if the drive has 1 or 2 slices. This means that if you put a drive set of 180GB drives in the left side, and a set of 500GB drives in the right, that the system will located the single slice 1TB partition first, then the 2 sliced 1.5TB drives that are associated with the right hand side.

To confirm this you can view the `/proc/scsi/scsi` file. Just cat that file, the example presented here is the indicative of the above description and should be similar to what is seen on the recorder machines.

```
> cat /proc/scsi/scsi
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: ATA      Model: HDS724040KLSA80  Rev: KFAO
  Type:   Direct-Access          ANSI SCSI revision: 05
Host: scsi1 Channel: 00 Id: 00 Lun: 00
  Vendor: ATA      Model: HDS724040KLSA80  Rev: KFAO
  Type:   Direct-Access          ANSI SCSI revision: 05
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: APPLE    Model: Xserve RAID      Rev: 1.50
  Type:   Direct-Access          ANSI SCSI revision: 05
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: APPLE    Model: Xserve RAID      Rev: 1.50
  Type:   Direct-Access          ANSI SCSI revision: 05
Host: scsi3 Channel: 00 Id: 00 Lun: 01
  Vendor: APPLE    Model: Xserve RAID      Rev: 1.50
  Type:   Direct-Access          ANSI SCSI revision: 05
```

You can ignore the scsi0 and scsi1 hosts, we are only interested in the scsi2 and scsi3. As you will see they have APPLE and Xserve RAID written there... just what we are looking for. The scsi2 host shows 1 drive connected to it, on channel 0, id 0 and lun 0. This is the 180GB drive set. The scsi3 shows 2 drives, channel 0, id 0 and two different luns 0 and 1. It is the Lun that is modified when you slice the drives. Hence scsi3 has a sliced drive on it.

The drives will be labelled as `/dev/sdc`, `/dev/sdd` and `/dev/sde`

We are able to confirm this... but for now this document won't.

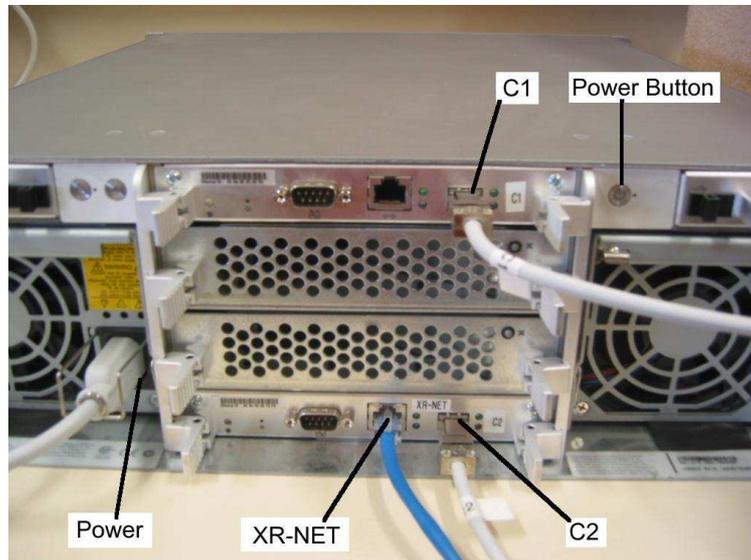
There should only be a single partition on each drive. As such you will be able to mount the drives as `/dev/sdc1`, `/dev/sdd1` and `/dev/sde1`.

Physical location of parts

There are two main parts to the XRaid system that need to be setup. This setup gets a little more complicated for the ATNF telescopes as they have 2 recorders, see below for further information.

The photo below shows the back of the XRaid. You may or may not have the XR-NET connected, it is not needed to follow the instructions here and is only useful for remotely turning the XRaid on and off. This document uses the labels shown in this image. The C1 and C2 cables are the Fibre Channel cables. The top one in this image is also considered the left or upper side of the XRaid when you view it from the front. The C2 is the lower or right side.

Please also take note of the Power button location.



< LEFT OR UPPER DRIVE SET > < RIGHT OR LOWER DRIVE SET >



<COLUMN>

Solving problems:

The biggest problem you are likely to come across is not having permission to write to the XRaid drive. That is fixed fairly easily. Login to the computer as root, and change the permissions to the affected drive.

```
chmod a+wx /data/xraid?   Where ? is the xraid drive number
```

The other problem you may have is finding there is no free space on the drive, perhaps someone forgot to format the drive first. Here is the command line we use to format the drives. This assumed that you know which drive you are formatting and that you have correctly identified it <see above section>. You also need to ensure the drive is unmounted before you can format it.

```
mke2fs -j -m -0 -T largefile /dev/sd?1   ? is the drive letter
```

Are you unsure which drive you are writing to? I have a new method where by when ever I format the drives I put a file in the root directory called label. This file tells you the name of the drive and if it is a sliced drive which slice. E.g. “Disk Set D1-7 (slice a)”