

Making a RAID

If you ever need to create a RAID from scratch, either because too many disks have failed in an old set, or because you have a nice shiny new set of disks, budget a bit of time for it!

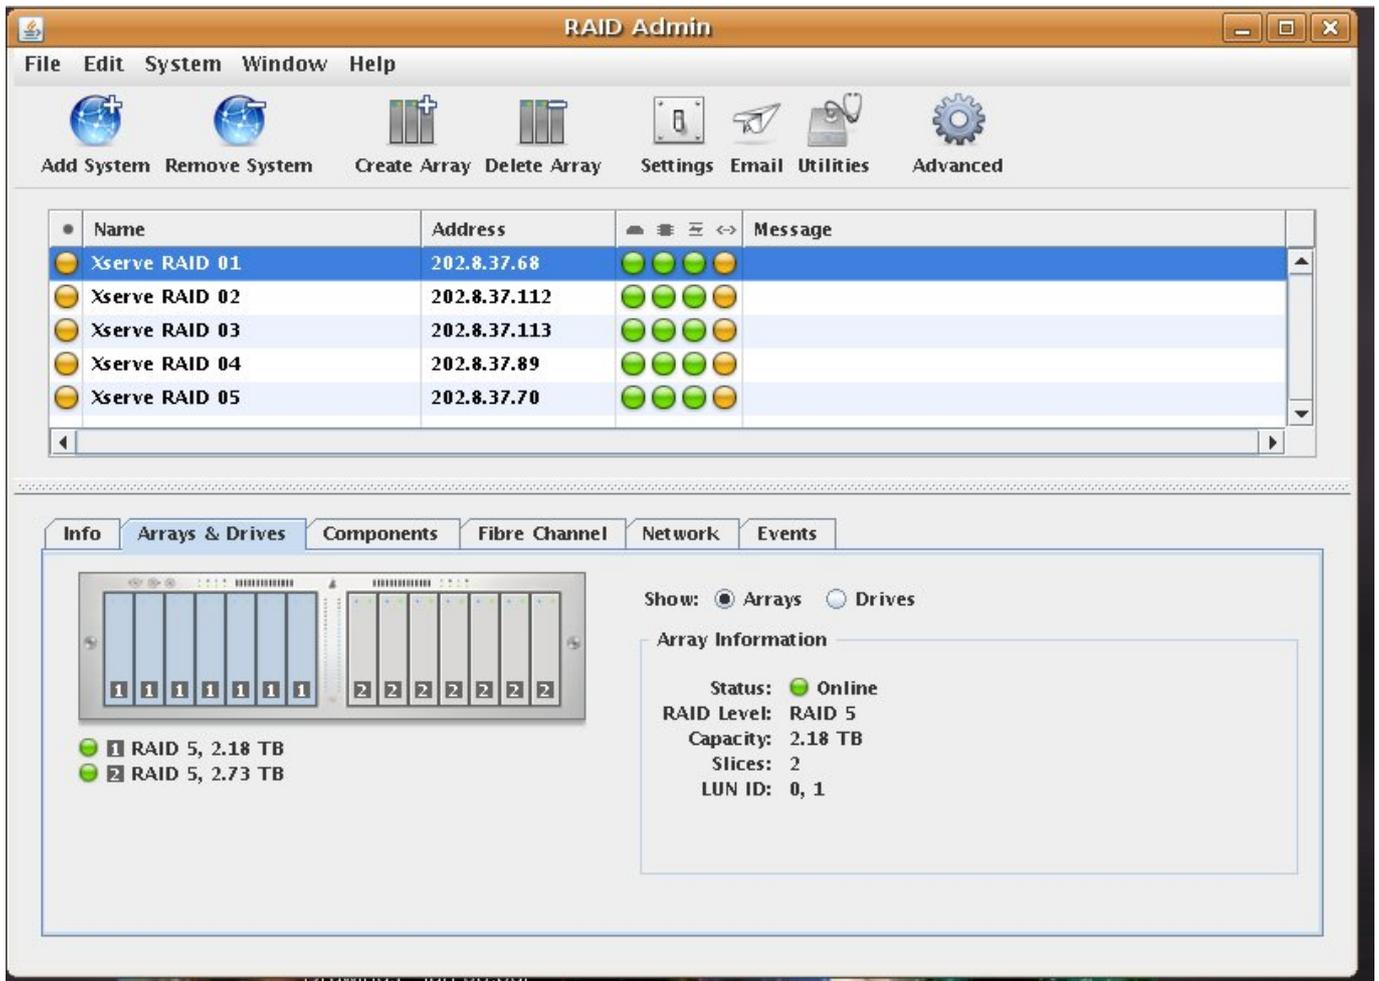
- Power the XRAID down (using appropriate unmounting and nfs-disabling procedures). [Check here](#).
- Insert the disks to be RAIDed into a chassis. It doesn't matter if the other half of the chassis still has a working RAID in it, just don't touch it!
- Power up the XRAID, making sure all disks are recognised.
- Log on to cuppa02 using a remote desktop for speediness
 - Log on to cuppa02, if this **hasn't** been done before, run `vncserver -geometry 1024x768` or just `vncserver` if the command has been run before but the node has since been restarted.
 - Run > Applications > Internet > Remote Desktop Viewer (in Ubuntu) on local workstation; type password
- Using the XRAID admin tools on cuppa02,

```
> java -jar RAID_Admin.jar
```

(run from `/nfs/apps/vlbi/RAID_Admin_151`) check that all disks are available (click on each disk, and **Make Available** using **Disk Utilities** if necessary).

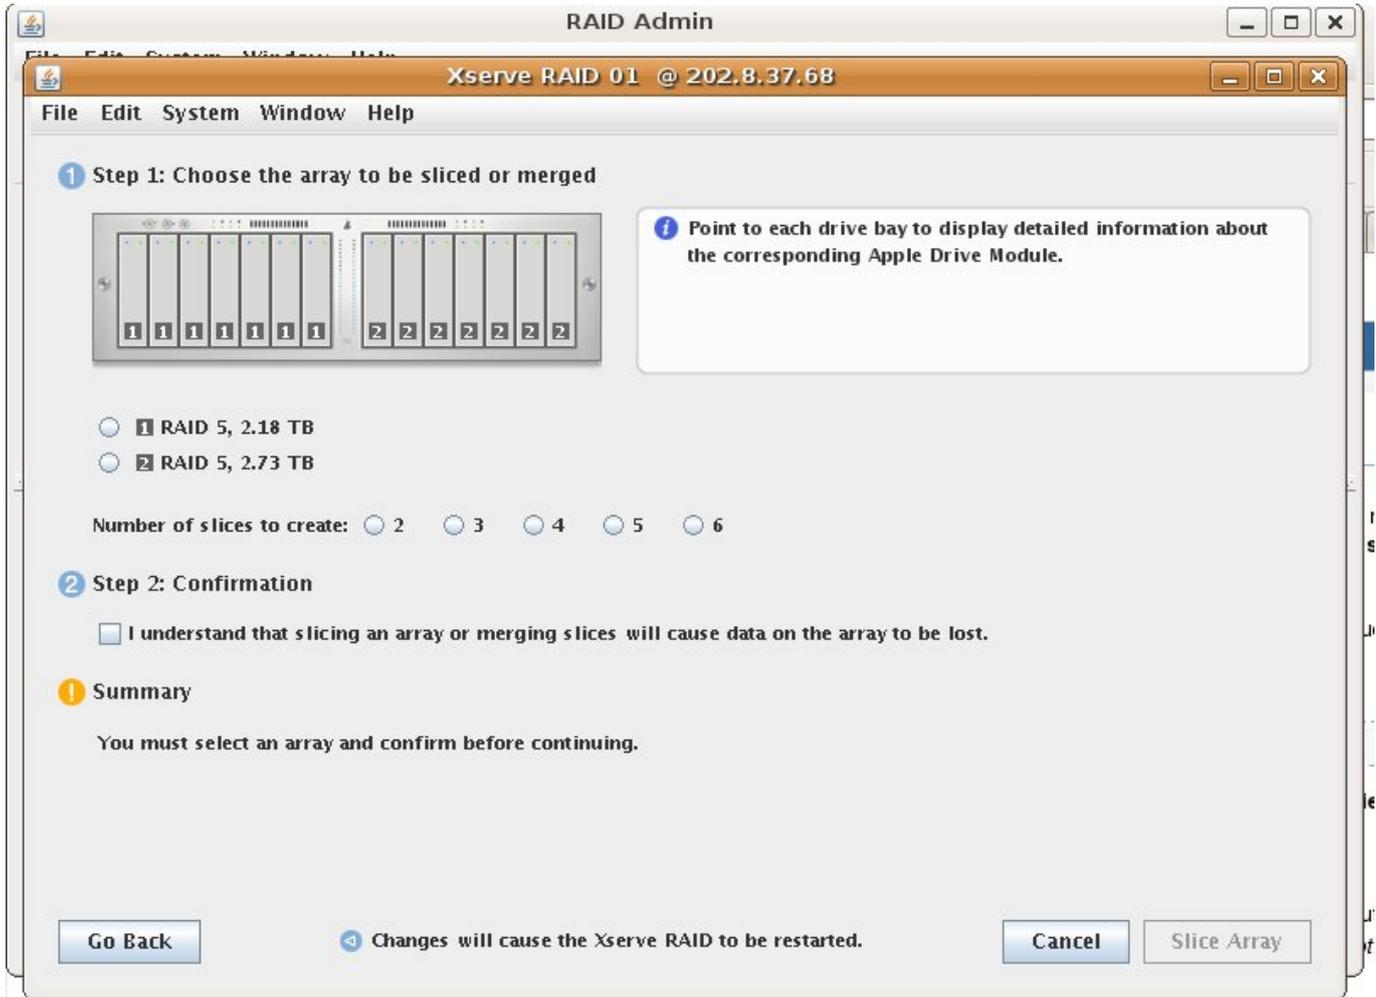
You'll need a password at some point. Ask Aidan or Claire if you don't know it.

Also, these images are huge. If anyone knows how to rescale images, please do!! :)

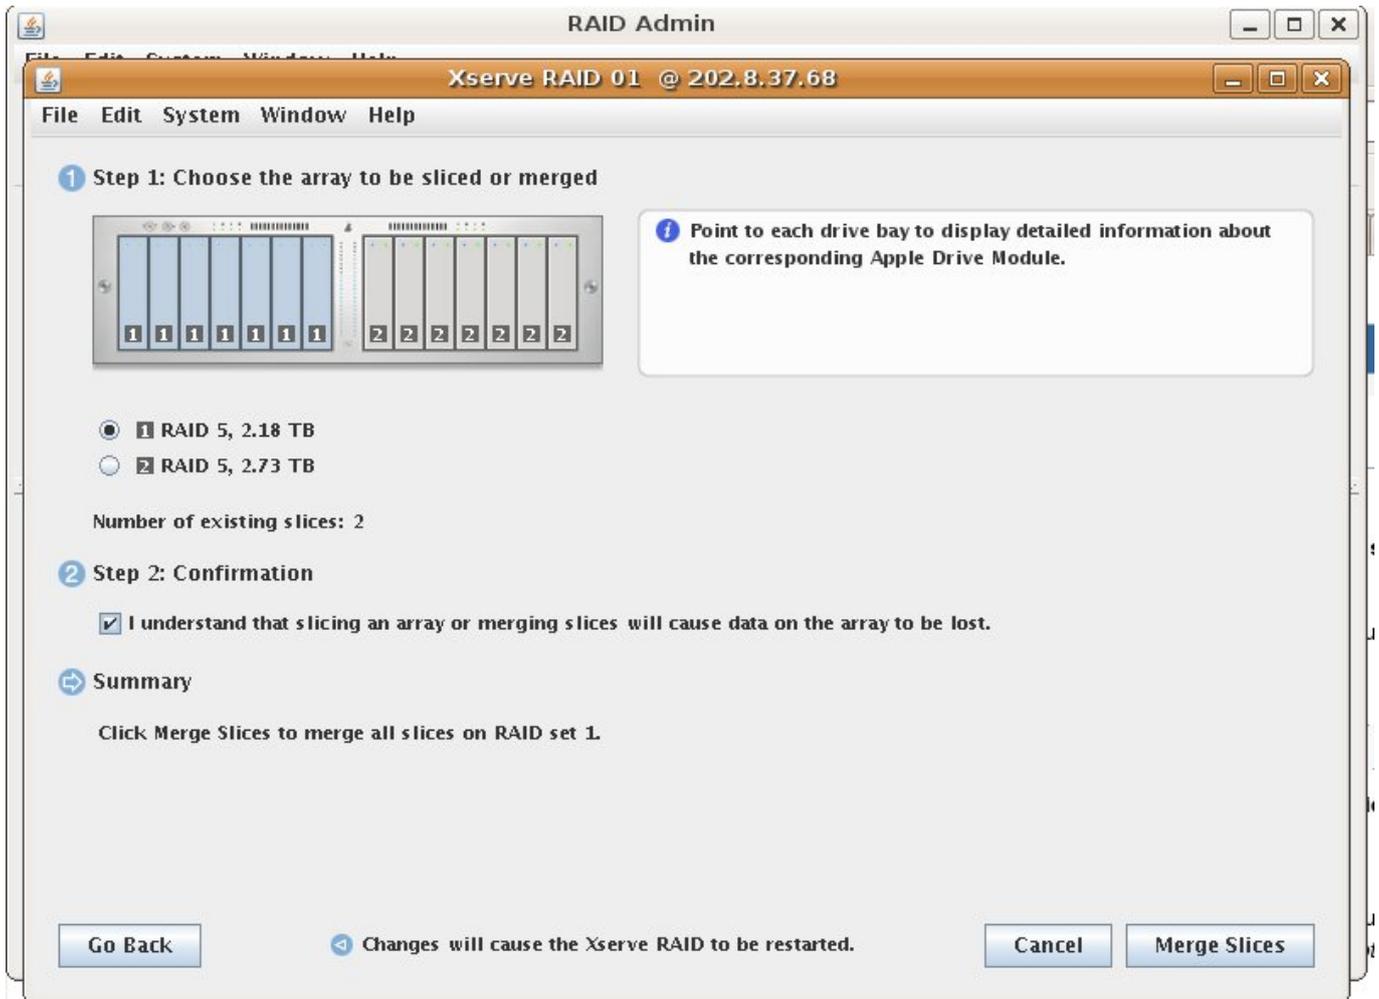


- Click the **Create Array** button in the main dialogue.
- Select the disks to be RAIDed (usually all 7 on either side), and the type of RAID. You can look up what each RAID type means, but for VLBI we use RAID5, which makes one disk effectively redundant. *The build process takes around 36 hours for 500GB disks, well over 2 days for 750GB disks. If you use the **background initialisation** option (default) you can continue to use the chassis while that disk set is being built.*

While building, you can Slice the disks. A “slice” is like a hardware partition. In VLBI we *used to use* 2 slices on 500GB disk sets, 3 slices on 750GB disk sets, creating slices of approximately 1.4TB (the XRAID admin tools will give you a size estimate of the whole array, I assume you can divide by 2 or 3!). However with newer kernels we're now able to work with devices larger than 2TB, so we no longer slice our disks (giving us 3TB for 500GB sets, and 4.5TB for 750GB sets).



- To slice the disks (no longer standard procedure), click the **Advanced** button (or menu), and select **Slicing**. Choose the number of slices you need (see above). This can be done concurrently with the building process, but only takes a few minutes. If you want to Unslice a set of disks to make a bigger array, or to change the number of slices you need to (first) in this menu press the button **Merge Slices**.



Once the RAID is built, you still need to partition and format the slices before they can be used. Each slice is seen on the computer as a separate device (ie 1.5TB hard disk, effectively). If you are unslicing a previously used VLBI disk, this still needs to be done as the computer will otherwise only see 1 1.5TB partition, and it won't have a valid partition table.

- ssh to the computer to which the XRAID is attached.
- 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. **NOTE** refresh_xraid will simply hang if only one side of the chassis has disks in, in which case it may be necessary to run `kill -9` on the process ID (find it using `pgrep refresh_xraid`).
- 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
```

- To check that all expected devices are present, run

```
sudo fdisk -l
```

This will give you a list of available devices. XRAID devices will likely

show up as sdc, sdd, through sdh. You can tell which is which by looking at the links in `/dev/disk/by-label`.

- If the device does not show up as expected (ie more than 1.5TB) follow [these instructions](#). Note that the "load disks" script will mount the disks, so you'll need to unmount them again prior to partitioning.

Failing all else, do

```
sudo shutdown -r now
```

(or `-r 1m` etc). This solves most problems except where someone else is using the computer at the time! Always check first!

- **Partition the devices** using Chris's script (note: uses parted because fdisk can't make things bigger than 2TB). New procedure is

```
sudo /nfs/apps/vlbi/partitiondisk.pl /dev/sdX
```

and check that it did what you wanted with

```
sudo parted /dev/sdX
```

and print or `sudo parted and print all`.

- Run `sudo fdisk -l` again to make sure you now have devices `sdcl`, `sdd1` etc which are the appropriate sizes (will give silly numbers if >2TB).
- **Format the partitions**
 - We are now using **xfs** as it appears to give better write speed performance than the alternatives, and is very quick to remove data as well. To format to xfs,

```
mkfs.xfs -f -l size=64m -L "ATNF V006B" /dev/sdcl
```

as appropriate, and for optimised performance, mount using

```
mount -o noatime,nodiratime,logbufs=8 /dev/sdcl /exports/xraid01/l_1
```

Can modify `fstab` to remember these options if required.

- Try to mount your brand new RAID! (see [Loading XRAID disks](#) for mounting procedures).

[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/createnew>

Last update: **2011/11/18 13:56**

