

# How to replace a failed mirrored disk under Linux

Greg Baker (gregb@ifost.org.au)

February 11, 2005

\$Id: diskreplace.tex,v 1.4 2005/02/10 13:03:58 gregb Exp \$

## 1 IDE Disks

Usually a server based around IDE disks will be relatively small, and have only two IDE disks. If it has been set up sensibly, the disks will be on different controllers – the disks will be `/dev/hda` and `/dev/hdc`.

### 1.1 Primary boot disk failure

If `/dev/hda` has failed, then the system might not start up correctly.

1. Check the BIOS and see if there is an option to boot off the other disk. If you can, then proceed as if it were a secondary failure (see section ??) – swap the device file names around of course!
2. Try swapping the disks around. The two disks should be identical. If this works, follow section ?. Unfortunately some of the standard installer programs (such as Red Hat Enterprise 3) don't actually get this right, and so the end result might be an unbootable system (“replace system disk” message).
3. Try the install CD in some kind of rescue mode. For Red Hat, this is usually “linux rescue”. If it can find the existing installation, then `chroot /mnt/sysimage /bin/bash` (and you can follow the instructions in section ?).
4. If it can't find the existing installation, then you could poke around with `lsraid` and see what's gone wrong.
5. Wimp out. Do a non-destructive re-install.
6. Really, really wimp out. Ask the sysadmin for a recovery CD made with `mkcdrec`.

## 1.2 Secondary boot disk failure

The IDE protocol can't inform the computer that a disk has been replaced. So if you've just replaced a disk in a removable bay, you need to reboot.

1. `cat /proc/mdstat` – are there failed devices? Does each device only have one U (up) with all the others down \_ (probably looks like [U\_]).
2. Recreate the partitions on `/dev/hdc`. (It might be possible to use `dd if=/dev/hda of=/dev/hdc`, but it's rather slow). Alternatively ...
  - Run `fdisk -l /dev/hda`
  - Run `fdisk /dev/hdc`, and use the `n` command to create partitions. Use the `t` command to change the type of each partition (should be `fd` (Linux raid auto)) and the `a` command to set the bootable flag (they appear as “starred” partitions in `fdisk -l`).
3. Look at `/etc/raidtab`. Hopefully each multi-disk device is made up of partitions in order (`/dev/md0` is a mirror of `/dev/hda1` and `/dev/hdc1`, and so on<sup>1</sup>).
4. Work through each multi-disk device, and `raidhotadd` each partition as need be.
  - `raidhotadd /dev/md0 /dev/hdc1`
  - `raidhotadd /dev/md1 /dev/hdc2`
  - `raidhotadd /dev/md2 /dev/hdc3`
  - `watch cat /proc/mdstat` (if you want to watch it resync)
5. Try `grub-install`. If it comes up with errors about non-existent BIOS disks, you'll have to do it manually. Run `grub`, and then ...
  - `root (hd0,0)`
  - `setup (hd0)`
  - `root (hd2,0)`
  - `setup (hd2)`

(Depending on the version of `grub`, you might need to replace the `hd2` with `hd1` in the above. It doesn't hurt to have the `grub` loader installed on too many disks.)

---

<sup>1</sup>It is possible that `/etc/raidtab` doesn't match the way the array actually is.

## 2 SCSI Disks

To be written ...

## 3 Log

\$Log: diskreplace.tex,v \$

Revision 1.4 2005/02/10 13:03:58 gregb

Wrong tag

Revision 1.3 2005/02/10 13:01:04 gregb

Working on it.

Revision 1.2 2005/02/10 12:52:14 gregb

Added other sections and log comments.