

FreeBSD GEOM mirror on Alpha

Actually my experience in GEOM disk mirroring on Alpha dates back to FreeBSD 5.3, though it seems to work out perfectly after upgrading to FreeBSD 6.0.

Again, the concept is similar to setting up a mirror on intel boxes, though - again - there are some minor differences which may cause some headaches!

I'm referring once more to Ralf S. Engelschalls GEOM mirroring approach using whole disks.

As with the Sun platform there is no concept of disk slices on Alpha. Alpha is also limited to primary partitions also and lacks the absence of the "PC MBR".

So his second method using disk slices may be adapted to primary partitions, though I did never try this on my own. For the sake of simplicity I prefer doing mirroring on whole disks pretty much the same way as physical controllers implement it.

So here is my step by step tutorial on creating a GEOM mirror on Alpha:

```
# assuming the following:
```

```
# first (ie. non-raid) scsi disk   : /dev/da0
```

```
# second (ie. future raid) scsi disk: /dev/da1
```

```
# geom mirror device (whole disk) : /dev/mirror/gm0
```

```
# make sure the second disk is treated as a really fresh one
```

```
# (not really necessary, but makes procedure more deterministically ;-)
```

```
dd if=/dev/zero of=/dev/da1 bs=512 count=79
```

```
# forget the whole PC MBR fdisk stuff, it simply doesn't exist on alpha :-)
```

```
# place a GEOM mirror label onto the second disk
```

```
gmirror label -v -n -b round-robin gm0 /dev/da1
```

```
# activate GEOM mirror kernel layer
```

```
# (makes the /dev/mirror/gm0 device available)
```

```
gmirror load
```

```
# place a BSD disklabel onto /dev/mirror/gm0
```

```
# (ATTENTION: consider the need to provide the full device path, eg. /dev/da1, etc
```

```
# on some releases of FreeBSD -- there seems to be an inconsistency on how different
```

```
# releases on different platforms handle this!)
```

```
# (NOTICE: figure out what partitions you want with "bsdlabel /dev/da0" before)
```

```
# (NOTICE: *no* partition offset is required at all for "a" partition!)
```

```
bsdlabel -w -B /dev/mirror/gm0      # initialize
```

```
bsdlabel -e /dev/mirror/gm0        # create custom partitions
```

```
# manually copy filesystem data from first to to second disk
```

```
# (same procedure for partitions "g", etc)
```

```
# (NOTICE: in my sample a have a single partition containing the / fs,
```

```
# just extend as needed if you have/need more partitions for /usr, etc)
```

```
newfs -U /dev/mirror/gm0a
```

```
mount /dev/mirror/gm0a /mnt
```

```
dump -L -0 -f- / | (cd /mnt; restore -r -v -f-)
```

```
# adjust new system configuration for GEOM mirror based setup
cp -p /mnt/etc/fstab /mnt/etc/fstab.orig
sed -e 's/dev/da0/dev/mirror/gm0/g' /mnt/etc/fstab
echo 'swapoff="YES"' >>/mnt/etc/rc.conf # for 5.3-RELEASE only
echo 'geom_mirror_load="YES"' >>/mnt/boot/loader.conf

# instruct boot stage 2 loader on first disk to boot
# with the boot stage 3 loader from the second disk
# (mainly because BIOS might not allow easy booting from second disk
# or at least requires manual intervention on the console)
echo "1:da(1,a)/boot/loader" >/boot.config

# reboot system
# (for running system with GEOM mirror on second disk)
shutdown -r now

# make sure the first disk is treated as a really fresh one
# (also not really necessary, but makes procedure more deterministically ;-))
dd if=/dev/zero of=/dev/da0 bs=512 count=79

# forget the whole PC MBR fdisk stuff, it simply doesn't exist on sparc :-))

# switch GEOM mirror to auto-synchronization and add first disk
# (first disk is now immediately synchronized with the second disk content)
gmirror configure -a gm0
gmirror insert gm0 /dev/da0

# wait for the GEOM mirror synchronization to complete
sh -c 'while [ "`gmirror list | grep SYNCHRONIZING`" != . ]; do sleep 1; done'

# reboot into the final two-disk GEOM mirror setup
# (now actually boots with the MBR and boot stages on first disk
# as it was synchronized from second disk)
shutdown -r now
```

The equipment used in this article is a AlphaStation 500/266 with an Alpha 266 Mhz CPU, 128 Megs of RAM and two IBM DDYS-T36950N S96H 36 GB SCSI drives.

No computers have been harmed whatsoever during my experiments.