

## FreeBSD GEOM mirror on Sparc64

As I promised earlier I will have a look at setting up a GEOM mirror on Sparc64 using FreeBSD 6.

Though it is quiet similar to the Intel world, there are some minor differences which I'm going to show you.

I'm referring myself to Ralf S. Engelschalls GEOM mirroring approach 1 using a whole disk.

His second method using disk slices won't work out on Sparc64 since the concept of disk slices does not exist on this platform.

Sparc64 allows only primary partitions, hence there is a maximum of eight instead of four as known from the intel world.

Considering this it should be possible however to do disk mirroring on single partitions according to his second concept.

Personally, I prefer to do mirroring using whole disks as I find it easier to setup and maintain and follows the concept of how ordinary hardware controllers implement it.

So here is my step by step tutorial on creating a GEOM mirror on Sparc64:

# 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 sparc :-)

# 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: in FreeBSD 6-current as of march 7 2006

# /dev/mirror/gm0 device has to be specified as just "mirror/gm0"!)

# (NOTICE: figure out what partitions you want with "sunlabel da0" before)

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

sunlabel -w -B mirror/gm0 auto # initialize

sunlabel -e 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

# **** ATTENTION ****
# You may try this step at your own risk. However there is a chance
# that it wont work with you either!
#
# As of my understanding, this step is only possible on the intel
# platform, where a special boot loader is stored inside the MBR.
# This boot loader does not exist on Sparc64, as Sun OpenBoot will
# call the boot loader directly from the primary partition at offset 0.
# I could only get this working by instructing OpenBoot manually to
# boot from the second disk.
#
# So after shutdown make sure to enter OpenBoot and start manually
# from the second disk.
# **** ATTENTION ****
#
# 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 ATA 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 Sun Netra T1 with an UltraSparc-III (440 MHz), 512 Megs of RAM and two IBM DDYST1835SUN18G 18 GB SCSI drives.
```

No computers have been harmed whatsoever during my experiments.

My next article will cover the creation of a GEOM disk mirror on the alpha architecture.