

Sun OpenBoot: Fail-Over Booting A FreeBSD GEOM Mirror

After my recent article about [FreeBSD GEOM mirroring on Sparc64](#) I was asked how to boot from a degraded array if the primary hard disk had failed.

The point was how to prevent the admin to interact with OpenBoot in that case.

I'm sure this will only apply to people not very familiar with Sun's OpenBoot, as most other people will surely know about this.

While I'm pointing you at the solution, please have also a look at the [OpenBoot Command Reference Manual](#).

#1 Get GEOM mirror device names

First make sure you know which devices you used for building the mirror. The command 'gmirror list' will show it to you. Get the consumer names for your GEOM provider. The following output has been shortened a bit.

```
Geom name: gm0
State: COMPLETE
Components: 2
Balance: round-robin
Consumers:
1. Name: da0
2. Name: da1
```

#2 Check where devices are connected to

Now you should grep through message to find where your GEOM consumers are connected to. I have shortened the output again for the relevant parts.

```
# dmesg | grep "da[0|1]"
da0 at sym0 bus 0 target 0 lun 0
da0: Fixed Direct Access SCSI-2 device
da1 at sym0 bus 0 target 1 lun 0
da1: Fixed Direct Access SCSI-2 device
```

#3 Enter OpenBoot firmware console

Shutdown your operating system to enter OpenBoot firmware console. You can even 'STOP-A' the machine, though I wouldn't recommend that. If you just turned on your box and it is about to do the usual memory checks at all, it's safe to issue 'STOP-A'.

#4 Probe your scsi devices

If you know how the controllers and devices are attached and how they are enumerated in OpenBoot, you can skip to step 5. Otherwise issue a 'probe-scsi-all' command, which should give something like that:

```
/pci@1f,4000/scsi@2
Target 5
Unit 0 Removable Tape HP C1537A L706
Target 6
Unit 0 Removable Read Only device TOSHIBA XM5701TASUN12XCD2395

/pci@1f,4000/scsi@3
```

Target 0

Unit 0 Disk FUJITSU MAG3182L SUN18G 1111

Target 1

Unit 0 Disk FUJITSU MAG3182L SUN18G 1111

Look at the devices which match your output from step 2. In my example it's the devices on scsi bus 3, target 0 and target 1.

#5 Enumerate your device aliases

Run the command 'devalias' to list your device aliases. Usually there should be matching device aliases for your hard disks. In my case this was (output shortened):

```
disk0          /pci@1f,4000/scsi@3/disk@0,0
disk1          /pci@1f,4000/scsi@3/disk@1,0
```

You may or may not have additional device aliases matching scsi bus 3. You may even have a device alias call simply 'disk' which reads the same as 'disk1'. This depends on your configuration setup. Refer to the OpenBoot manual if you want to learn more about it.

If the device aliases wouldn't exist for some obscure reason, you should create them:

```
devalias disk0 /pci@1f,4000/scsi@3/disk@0,0
devalias disk1 /pci@1f,4000/scsi@3/disk@1,0
```

Make sure you don't accidentally overwrite other device alias names, especially if you are unsure about whether you need them or not.

You could easily use other names beside 'disk0' or 'disk1'. They should only be simple enough to be used in the next step.

#6 Set your boot device priority

Now you should instruct OpenBoot about your boot device priority. Check it out first using 'printenv boot-device', which should return something like:

```
boot-device = disk net
```

Set it to your GEOM consumers which match to the device aliases you found in the previous steps, eg:

```
{ 1 } ok setenv boot-device disk0 disk1 net
boot-device = disk0 disk1 net
```

Now whenever booting up your box, it should boot from the 'disk0' GEOM consumer (which is /dev/da0 in FreeBSD) and fail-over to 'disk1' (which is /dev/da1) if 'disk0' failed. It would even try to boot from the network when the 'net' device is also given as in my example.

Consider that you can set boot-device using real device names (eg. /pci@1f,4000/scsi@3/disk@0,0) though I wouldn't recommend it at all. Use device aliases whenever possible.