

System upgrade and how to prevent downtime

In today's world there is a great need for 24h/7 systems. When the server is down, the company the system is working for may already start counting its losses. In that case, the loss is equal to real money.

Before you start system upgrade it's highly recommended to take a backup of your data. System mirror does not help much when the system is not bootable because of a software error. If something goes wrong after upgrade both halves of the mirror have the same data and you end up having a problem.

Restoring from a backup takes time and needs some experience. What if backup tape is broken?

In this article I present how to do a system upgrade on a FreeBSD server with gmirror(8) enabled.

System state and configuration

There are two ATA disks ad0 and ad1 mirrored with gmirror(8).

The system before upgrade is 6.0-Release. Look at the command listings.

```
mufasa# atacontrol list
```

```
ATA channel 0:
```

```
  Master: ad0 <ST3160815A/3.AAD> ATA/ATAPI revision 7
```

```
  Slave:  ad1 <ST3160815A/3.AAD> ATA/ATAPI revision 7
```

```
ATA channel 1:
```

```
  Master: acd0 <LG CD-ROM CRD-8521B/1.01> ATA/ATAPI revision 0
```

```
  Slave:   no device present
```

```
ATA channel 2:
```

```
  Master:   no device present
```

```
  Slave:   no device present
```

```
ATA channel 3:
```

```
  Master:   no device present
```

```
  Slave:   no device present
```

```
mufasa# gmirror list
```

```
Geom name: gm0
```

```
State: COMPLETE
```

```
Components: 2
```

```
Balance: round-robin
```

```
Slice: 4096
```

```
Flags: NONE
```

```
GenID: 0
```

```
SyncID: 1
```

```
ID: 2378501857
```

```
Providers:
```

```
1. Name: mirror/gm0
```

```
  Mediasize: 160041885184 (149G)
```

```
  Sectorsize: 512
```

```
  Mode: r6w6e7
```

```
Consumers:
```

```
1. Name: ad0
```

```
  Mediasize: 160041885696 (149G)
```

```
  Sectorsize: 512
```

```
  Mode: r1w1e1
```

```
  State: ACTIVE
```

```
  Priority: 1
```

```
  Flags: NONE
```

```
  GenID: 0
```

```
  SyncID: 1
```

```
  ID: 227080134
```

```
2. Name: ad1
```

```
  Mediasize: 160041885696 (149G)
```

```
  Sectorsize: 512
```

```
Mode: rlwle1
State: ACTIVE
Priority: 0
Flags: NONE
GenID: 0
SyncID: 1
ID: 2997696838
```

```
mufasa# gmirror status
  Name  Status Components
mirror/gm0 COMPLETE ad0
      ad1
```

System upgrade

The system upgrade is well documented in the FreeBSD's handbook so I won't go into details here.

1. cvsup proper sources
2. cd /usr/src/ && make buildworld
3. make bildkernel KERNCONF=KERNELNAMEHERE
4. Stop here!

Before make installworld and installkernel let's break the mirror so you have 2 identical system copies on two separate drives.

Breaking the mirror

```
mufasa# gmirror deactivate gm0 ad1
mufasa# gmirror status
  Name  Status Components
mirror/gm0 DEGRADED ad0
mufasa# gmirror list
Geom name: gm0
State: DEGRADED
Components: 2
Balance: round-robin
Slice: 4096
Flags: NONE
GenID: 0
SyncID: 2
ID: 2378501857
Providers:
1. Name: mirror/gm0
  Mediasize: 160041885184 (149G)
  Sectorsize: 512
  Mode: r6w6e7
Consumers:
1. Name: ad0
  Mediasize: 160041885696 (149G)
  Sectorsize: 512
  Mode: rlwle1
  State: ACTIVE
  Priority: 1
  Flags: NONE
  GenID: 0
  SyncID: 2
  ID: 227080134
```

In the output above you can see that ad1 is deactivated from the configuration. All write operation will be performed on gm0 device which has only one consumer ad0.

Countinuing with the system upgrade.

5. make installkernel KERNCONF=KERNELNAMEHERE

6. make install world
7. mergemaster
8. reboot

```
mufasa# uname -a
FreeBSD mufasa.x.y 7.2-RELEASE FreeBSD 7.2-RELEASE #5: Wed May 13 16:46:25 CEST 2009
root@mufasa.x.y:/usr/obj/usr/src/sys/MUFASA i386
```

Check if the system booted without any problems and if all services are up and running.

When the system boots successfully then the only thing left to do is activating the second half of the mirror -

```
mufasa# gmirror activate gm0 ad1
```

and watch as the synchronization progresses...

```
mufasa# gmirror status
Name  Status Components
mirror/gm0  DEGRADED  ad0
          ad1 (1%)
```

After synchronization you should see gm0 provider, consumers ad0 and ad1.

What is something goes wrong?

If the just upgraded OS does not boot after restart because (let's say) kernel problems, you can boot the system from the second drive, ad1 - the second half of the mirror.

The FreeBSD booting process is well described in the handbook so please treat it as the reference.

Choosing the boot drive

As soon as you turn on the system watch carefully for the BIOS messages displayed on the screen. As soon as you see the line

```
FreeBSD/i386 boot
Default: x:xx(x,x)/boot/loader
```

hit any key i.e. "?" - question mark key. If you are successful you should be in "**boot:**" prompt. This is the FreeBSD boot2 boot manager (see boot man page man boot).

Type in:

```
boot: 1:ad(1,a)/boot/loader -s
```

In the line above you have specified to boot from ad1 device into single user mode. The boot process should hang just before mounting the / file system.

Manual root filesystem specification:

```
<fstype>:<device> Mount <device> using filesystem <fstype>
e.g. ufs:/dev/da0s1a
?          List valid disk boot devices
<empty line> Abort manual input
```

```
mountroot>
```

At the prompt type

```
mountroot> ufs:ad1s1a
```

System boot should proceed after mounting / file system from ad1 drive and stop before mounting other filesystems like /usr; /var; /tmp but you should be already in single user mode having /bin/sh as a shell.

At the prompt remount / file system with read and write privileges, then mount /usr.

```
# mount -u -o rw /  
# mount /dev/ad1s1f /usr
```

Now you are able to edit /etc/fstab file, mount the other file systems and boot FreeBSD into multiuser mode.

```
# vi /etc/fstab
```

Hit ESC key and type in

```
:%s/mirror/gm0/ad1/g
```

This simple command will change a phrase mirror/gm0 into ad1 -
/dev/mirror/gm0s1a changes into /dev/ad1s1a

Exit vi typing in:

```
:wq!
```

Exit shell and the system should start booting into multiuser mode.

These few steps presented above should be enough to bring the system up after unsuccessful (unlucky?) system upgrade. When time is a cost you don't want to spend it recovering system from a tape. Booting from a second half of a mirror is safer as you have all your data consistent proceeding system upgrade.