

Software RAID-1 with Mandrake9.0 problems – LONG

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.setup/2003-07/0435.html>

From: Theophanis Kontogiannis (*tkonto_at_aegean.gr*)

Date: 07/12/03

Date: Sat, 12 Jul 2003 12:04:17 +0300

Please Cc: tkonto@aegean.gr

Hello all

I have COMPAQ EVO and 2x40GB IBM EIDE, each installed on a sepperate bus of the motherboard.

I have configured RAID -1

```
[root@linux-lab root]# cat /etc/fstab
/dev/md0 / ext3 defaults 1 1
/dev/hda1 /boot ext3 defaults 1 2
/dev/hdd1 /boot_backup ext3 defaults 1 2
none /dev/pts devpts mode=0620 0 0
/dev/md1 /home ext3 defaults 1 2
/dev/hdc /mnt/cdrom auto
user,icharset=iso8859-1,codepage=850,noauto,ro,exec 0 0
/dev/fd0 /mnt/floppy auto
user,icharset=iso8859-1,sync,codepage=850,noauto,exec 0 0
none /proc proc defaults 0 0
none /tmp tmpfs defaults 0 0
/dev/hda5 swap swap defaults 0 0
/dev/hdd5 swap swap defaults 0 0
```

```
[root@linux-lab root]# cat /etc/raidtab
raiddev /dev/md0
raid-level 1
chunk-size 64k
persistent-superblock 1

nr-raid-disks 2
nr-spare-disks 0
    device /dev/hda4
    raid-disk 0
    device /dev/hdd4
```

raid-disk 1

raiddev /dev/md1
raid-level 1
chunk-size 64k
persistent-superblock 1

nr-raid-disks 2
nr-spare-disks 0
 device /dev/hda2
 raid-disk 0
 device /dev/hdd2
 raid-disk 1

The server works fine but at random times (usually after 4-6 days of normal operations) it blows. The errors I get on the console are:

```
ext3-fs error (device md(9,0)): read_inode_bitmap: Cannot read inode
bitmap_block_group=85, inode_bitmap=2785281
ext3-fs error (device md(9,0)) in ext3_new_inode: IO failure
raid1: hda4: unrecoverable I/O read error for block 2785281
```

and thousands replicas of these on the screen with changing the block number.

I reboot and I get:

```
Loading linux-secure
BIOS data check successful
Uncompressing Linux... OK, booting the kernel
Error only one processor found.
Unkown bridge resource 2: Assuming transparent
md: superblock update time inconsistency -- using the most recent one
md: md0: raid array is not clean -- starting background reconstruction
kmod: failed to exec /sbin/modprobe -s -k md-personality=3, errno=2
md: personality 3 is not loaded!
md: superblock update time inconsistency -- using the most recent one
md: md1: raid array is not clean -- starting background reconstruction
kmod: failed to exec /sbin/modprobe -s -k md-personality=3, errno=2
md: personality 3 is not loaded!
md: superblock update time inconsistency -- using the most recent one
md: md1: raid array is not clean -- starting background reconstruction
raid1: md1, not all disks are operational -- trying to recover array
md1: no spare disk to reconstruct array! -- continuing in degraded mode
md: superblock update time inconsistency -- using the most recent one
md: md0: raid array is not clean -- starting background reconstruction
raid1: md0, not all disks are operational -- trying to recover array
md1: no spare disk to reconstruct array! -- continuing in degraded mode
md0: no spare disk to reconstruct array! -- continuing in degraded mode
INIT version 2.83 booting
```

```
.....  
modprobe cannot locate module fb0  
starting up RAID devices: md0 md1  
Checking filesystems  
/dev/hda1 recovering journal  
/dev/hdd1 recovering journal  
/dev/hdd1 clean 11/126720 files, 8090/253015 blocks  
/dev/hda1 clean 50/126720 files, 9069/253015 blocks  
/dev/md1 recovering journal  
/dev/md1 clean 22831/10289152 files, 858639/20563168 blocks
```

I have to say that I used the recommendations of the Mandrake installer (i.e. did not allow me to make RAID on the /boot) and the two disks are EXACTLY the same

after bootup dmesh gives me:

```
loop... 4771.02 BogoMIPS  
Memory: 248156k/253888k available (1346k kernel code, 5344k reserved, 504k  
data,  
 140k init, 0k highmem)  
Dentry cache hash table entries: 32768 (order: 6, 262144 bytes)  
Inode cache hash table entries: 16384 (order: 5, 131072 bytes)  
Mount-cache hash table entries: 4096 (order: 3, 32768 bytes)  
Buffer-cache hash table entries: 16384 (order: 4, 65536 bytes)  
Page-cache hash table entries: 65536 (order: 6, 262144 bytes)  
CPU: Before vendor init, caps: bfebfbff 00000000 00000000, vendor = 0  
CPU: L1 I cache: 12K, L1 D cache: 8K  
CPU: L2 cache: 512K  
CPU: Hyper-Threading is disabled  
CPU: After vendor init, caps: bfebfbff 00000000 00000000 00000000  
Intel machine check architecture supported.  
Intel machine check reporting enabled on CPU#0.  
CPU: After generic, caps: bfebfbff 00000000 00000000 00000000  
CPU: Common caps: bfebfbff 00000000 00000000 00000000  
Enabling fast FPU save and restore... done.  
Enabling unmasked SIMD FPU exception support... done.  
Checking 'hlt' instruction... OK.  
POSIX conformance testing by UNIFIX  
mtrr: v1.40 (20010327) Richard Gooch (rgooch@atnf.csiro.au)  
mtrr: detected mtrr type: Intel  
CPU: Before vendor init, caps: bfebfbff 00000000 00000000, vendor = 0  
CPU: L1 I cache: 12K, L1 D cache: 8K  
CPU: L2 cache: 512K  
CPU: Hyper-Threading is disabled  
CPU: After vendor init, caps: bfebfbff 00000000 00000000 00000000  
Intel machine check reporting enabled on CPU#0.  
CPU: After generic, caps: bfebfbff 00000000 00000000 00000000  
CPU: Common caps: bfebfbff 00000000 00000000 00000000  
CPU0: Intel(R) Pentium(R) 4 CPU 2.40GHz stepping 07  
per-CPU timeslice cutoff: 1462.88 usecs.
```

```
enabled ExtINT on CPU#0
ESR value before enabling vector: 00000000
ESR value after enabling vector: 00000000
Error: only one processor found.
ENABLING IO-APIC IRQs
Setting 8 in the phys_id_present_map
...changing IO-APIC physical APIC ID to 8 ... ok.
init IO-APIC IRQs
IO-APIC (apicid-pin) 8-0, 8-2, 8-5, 8-10, 8-11 not connected.
..TIMER: vector=0x31 pin1=-1 pin2=-1
...trying to set up timer (IRQ0) through the 8259A ... failed.
...trying to set up timer as Virtual Wire IRQ... works.
number of MP IRQ sources: 42.
number of IO-APIC #8 registers: 24.
testing the IO APIC.....
```

```
IO APIC #8.....
.... register #00: 08000000
..... : physical APIC id: 08
.... register #01: 00178020
..... : max redirection entries: 0017
..... : PRQ implemented: 1
..... : IO APIC version: 0020
.... register #02: 00000000
..... : arbitration: 00
.... IRQ redirection table:
NR Log Phy Mask Trig IRR Pol Stat Dest Deli Vect:
00 000 00 1 0 0 0 0 0 0 00
01 001 01 0 0 0 0 0 1 1 39
02 000 00 1 0 0 0 0 0 0 00
03 001 01 0 0 0 0 0 1 1 41
04 001 01 0 0 0 0 0 1 1 49
05 000 00 1 0 0 0 0 0 0 00
06 001 01 0 0 0 0 0 1 1 51
07 001 01 0 0 0 0 0 1 1 59
08 001 01 0 0 0 0 0 1 1 61
09 001 01 0 0 0 0 0 1 1 69
0a 000 00 1 0 0 0 0 0 0 00
0b 000 00 1 0 0 0 0 0 0 00
0c 001 01 0 0 0 0 0 1 1 71
0d 001 01 0 0 0 0 0 1 1 79
0e 001 01 0 0 0 0 0 1 1 81
0f 001 01 0 0 0 0 0 1 1 89
10 001 01 1 1 0 1 0 1 1 91
11 001 01 1 1 0 1 0 1 1 99
12 001 01 1 1 0 1 0 1 1 A1
13 001 01 1 1 0 1 0 1 1 A9
14 001 01 1 1 0 1 0 1 1 B1
15 001 01 1 1 0 1 0 1 1 B9
16 001 01 1 1 0 1 0 1 1 C1
17 001 01 1 1 0 1 0 1 1 C9
```

IRQ to pin mappings:

IRQ1 -> 0:1
IRQ3 -> 0:3
IRQ4 -> 0:4
IRQ6 -> 0:6
IRQ7 -> 0:7
IRQ8 -> 0:8
IRQ9 -> 0:9
IRQ12 -> 0:12
IRQ13 -> 0:13
IRQ14 -> 0:14
IRQ15 -> 0:15
IRQ16 -> 0:16
IRQ17 -> 0:17
IRQ18 -> 0:18
IRQ19 -> 0:19
IRQ20 -> 0:20
IRQ21 -> 0:21
IRQ22 -> 0:22
IRQ23 -> 0:23

..... done.

Using local APIC timer interrupts.

calibrating APIC timer ...

..... CPU clock speed is 2391.5075 MHz.

..... host bus clock speed is 132.8612 MHz.

cpu: 0, clocks: 1328612, slice: 664306

CPU0<T0:1328608,T1:664288,D:14,S:664306,C:1328612>

Waiting on wait_init_idle (map = 0x0)

All processors have done init_idle

PCI: PCI BIOS revision 2.20 entry at 0xeca9e, last bus=5

PCI: Using configuration type 1

PCI: Probing PCI hardware

Unknown bridge resource 2: assuming transparent

PCI: Using IRQ router PIIX [8086/24c0] at 00:1f.0

PCI->APIC IRQ transform: (B0,I2,P0) -> 16

PCI->APIC IRQ transform: (B0,I29,P0) -> 16

PCI->APIC IRQ transform: (B0,I29,P1) -> 19

PCI->APIC IRQ transform: (B0,I29,P3) -> 23

PCI->APIC IRQ transform: (B0,I31,P0) -> 18

PCI->APIC IRQ transform: (B0,I31,P1) -> 17

PCI->APIC IRQ transform: (B5,I8,P0) -> 20

PCI->APIC IRQ transform: (B5,I9,P0) -> 18

PCI->APIC IRQ transform: (B5,I10,P0) -> 21

isapnp: Scanning for PnP cards...

isapnp: No Plug & Play device found

Linux NET4.0 for Linux 2.4

Based upon Swansea University Computer Society NET3.039

Initializing RT netlink socket

apm: BIOS not found.

Starting kswapd

VFS: Diskquotas version dquot_6.5.0 initialized

```
pty: 256 Unix98 ptys configured
Serial driver version 5.05c (2001-07-08) with HUB-6 MANY_PORTS MULTIPORT
SHARE_I
RQ SERIAL_PCI ISAPNP enabled
ttyS00 at 0x03f8 (irq = 4) is a 16550A
Uniform Multi-Platform E-IDE driver Revision: 7.00alpha2
ide: Assuming 33MHz system bus speed for PIO modes; override with idebus=xx
ICH4: IDE controller on PCI bus 00 dev f9
PCI: Enabling device 00:1f.1 (0005 -> 0007)
ICH4: chipset revision 1
ICH4: not 100% native mode: will probe irqs later
  ide0: BM-DMA at 0x24a0-0x24a7, BIOS settings: hda:DMA, hdb:pio
  ide1: BM-DMA at 0x24a8-0x24af, BIOS settings: hdc:DMA, hdd:DMA
PDC20269: IDE controller on PCI bus 05 dev 48
PDC20269: chipset revision 2
PDC20269: not 100% native mode: will probe irqs later
PDC20269: (U)DMA Burst Bit DISABLED Primary PCI Mode Secondary PCI Mode.
  ide2: BM-DMA at 0x1440-0x1447, BIOS settings: hde:pio, hdf:pio
  ide3: BM-DMA at 0x1448-0x144f, BIOS settings: hdg:pio, hdh:pio
hda: WDC WD1200JB-00DUA1, ATA DISK drive
hdc: LTN486S, ATAPI CD/DVD-ROM drive
hdd: WDC WD1200JB-00DUA1, ATA DISK drive
ide0 at 0x1f0-0x1f7,0x3f6 on irq 14
ide1 at 0x170-0x177,0x376 on irq 15
hda: 234375000 sectors (120000 MB) w/8192KiB Cache, CHS=14589/255/63,
UDMA(100)
hdd: 234375000 sectors (120000 MB) w/8192KiB Cache, CHS=14589/255/63,
UDMA(100)
hdc: ATAPI 48X CD-ROM drive, 120kB Cache, DMA
Uniform CD-ROM driver Revision: 3.12
ide-floppy driver 0.99b
Partition check:
hda: hda1 hda2 hda3 < hda5 > hda4
hdd: hdd1 hdd2 hdd3 < hdd5 > hdd4
RAMDISK driver initialized: 16 RAM disks of 32000K size 1024 blocksize
ide-floppy driver 0.99b
md: md driver 0.90.0 MAX_MD_DEVS=256, MD_SB_DISKS=27
md: Autodetecting RAID arrays.
[events: 00000073]
[events: 00000073]
[events: 00000001]
[events: 00000001]
md: autorun ...
md: considering hdd4 ...
md: adding hdd4 ...
md: adding hda4 ...
md: created md0
md: bind<hda4,1>
md: bind<hdd4,2>
md: running: <hdd4><hda4>
md: hdd4's event counter: 00000001
```

comp.os.linux.setup: Software RAID-1 with Mandrake9.0 problems – LONG

```
md: hda4's event counter: 00000073
md: superbblock update time inconsistency -- using the most recent one
md: freshest: hda4
md: kicking non-fresh hdd4 from array!
md: unbind<hdd4,1>
md: export_rdev(hdd4)
md: md0: raid array is not clean -- starting background reconstruction
md: RAID level 1 does not need chunksize! Continuing anyway.
kmod: failed to exec /sbin/modprobe -s -k md-personality-3, errno = 2
md: personality 3 is not loaded!
md :do_md_run() returned -22
md: md0 stopped.
md: unbind<hda4,0>
md: export_rdev(hda4)
md: considering hdd2 ...
md: adding hdd2 ...
md: adding hda2 ...
md: created md1
md: bind<hda2,1>
md: bind<hdd2,2>
md: running: <hdd2><hda2>
md: hdd2's event counter: 00000001
md: hda2's event counter: 00000073
md: superbblock update time inconsistency -- using the most recent one
md: freshest: hda2
md: kicking non-fresh hdd2 from array!
md: unbind<hdd2,1>
md: export_rdev(hdd2)
md: md1: raid array is not clean -- starting background reconstruction
md: RAID level 1 does not need chunksize! Continuing anyway.
kmod: failed to exec /sbin/modprobe -s -k md-personality-3, errno = 2
md: personality 3 is not loaded!
md :do_md_run() returned -22
md: md1 stopped.
md: unbind<hda2,0>
md: export_rdev(hda2)
md: ... autorun DONE.
NET4: Linux TCP/IP 1.0 for NET4.0
IP Protocols: ICMP, UDP, TCP, IGMP
IP: routing cache hash table of 2048 buckets, 16Kbytes
TCP: Hash tables configured (established 16384 bind 16384)
Linux IP multicast router 0.06 plus PIM-SM
NET4: Unix domain sockets 1.0/SMP for Linux NET4.0.
RAMDISK: Compressed image found at block 0
Freeing initrd memory: 134k freed
VFS: Mounted root (ext2 filesystem).
md: raid1 personality registered as nr 3
Journalled Block Device driver loaded
md: Autodetecting RAID arrays.
[events: 00000001]
[events: 00000073]
```

```
[events: 00000001]
[events: 00000073]
md: autorun ...
md: considering hda2 ...
md: adding hda2 ...
md: adding hdd2 ...
md: created md1
md: bind<hdd2,1>
md: bind<hda2,2>
md: running: <hda2><hdd2>
md: hda2's event counter: 00000073
md: hdd2's event counter: 00000001
md: superblock update time inconsistency -- using the most recent one
md: freshest: hda2
md: kicking non-fresh hdd2 from array!
md: unbind<hdd2,1>
md: export_rdev(hdd2)
md: md1: raid array is not clean -- starting background reconstruction
md: RAID level 1 does not need chunksize! Continuing anyway.
md1: max total readahead window set to 124k
md1: 1 data-disks, max readahead per data-disk: 124k
raid1: device hda2 operational as mirror 0
raid1: md1, not all disks are operational -- trying to recover array
raid1: raid set md1 active with 1 out of 2 mirrors
md: updating md1 RAID superblock on device
md: hda2 [events: 00000074]<6>(write) hda2's sb offset: 82252672
md1: no spare disk to reconstruct array! -- continuing in degraded mode
md: considering hda4 ...
md: adding hda4 ...
md: adding hdd4 ...
md: created md0
md: bind<hdd4,1>
md: bind<hda4,2>
md: running: <hda4><hdd4>
md: hda4's event counter: 00000073
md: hdd4's event counter: 00000001
md: superblock update time inconsistency -- using the most recent one
md: freshest: hda4
md: kicking non-fresh hdd4 from array!
md: unbind<hdd4,1>
md: export_rdev(hdd4)
md: md0: raid array is not clean -- starting background reconstruction
md: RAID level 1 does not need chunksize! Continuing anyway.
md0: max total readahead window set to 124k
md0: 1 data-disks, max readahead per data-disk: 124k
raid1: device hda4 operational as mirror 0
raid1: md0, not all disks are operational -- trying to recover array
raid1: raid set md0 active with 1 out of 2 mirrors
md: updating md0 RAID superblock on device
md: hda4 [events: 00000074]<6>(write) hda4's sb offset: 30177984
md0: no spare disk to reconstruct array! -- continuing in degraded mode
```

```
md1: no spare disk to reconstruct array! — continuing in degraded mode
md: ... autorun DONE.
EXT3-fs: INFO: recovery required on readonly filesystem.
EXT3-fs: write access will be enabled during recovery.
kjournald starting. Commit interval 5 seconds
EXT3-fs: recovery complete.
EXT3-fs: mounted filesystem with ordered data mode.
Freeing unused kernel memory: 140k freed
Real Time Clock Driver v1.10e
usb.c: registered new driver usbdevfs
usb.c: registered new driver hub
usb-uhci.c: $Revision: 1.275 $ time 17:28:41 Sep 20 2002
usb-uhci.c: High bandwidth mode enabled
PCI: Setting latency timer of device 00:1d.0 to 64
usb-uhci.c: USB UHCI at I/O 0x2440, IRQ 16
usb-uhci.c: Detected 2 ports
usb.c: new USB bus registered, assigned bus number 1
hub.c: USB hub found
hub.c: 2 ports detected
PCI: Setting latency timer of device 00:1d.1 to 64
usb-uhci.c: USB UHCI at I/O 0x2460, IRQ 19
usb-uhci.c: Detected 2 ports
usb.c: new USB bus registered, assigned bus number 2
hub.c: USB hub found
hub.c: 2 ports detected
usb-uhci.c: v1.275:USB Universal Host Controller Interface driver
PCI: Setting latency timer of device 00:1d.7 to 64
hcd.c: ehci-hcd @ 00:1d.7, Intel Corp. 82801DB USB EHCI Controller
hcd.c: irq 23, pci mem d0044000
usb.c: new USB bus registered, assigned bus number 3
ehci-hcd.c: restricting 64bit DMA mappings to segment 0 ...
ehci-hcd.c: USB 2.0 support enabled, EHCI rev 1. 0
hub.c: USB hub found
hub.c: 6 ports detected
usbdevfs: remount parameter error
EXT3 FS 2.4-0.9.17, 10 Jan 2002 on md(9,0), internal journal
Adding Swap: 2097136k swap-space (priority -1)
Adding Swap: 2097136k swap-space (priority -2)
kjournald starting. Commit interval 5 seconds
EXT3 FS 2.4-0.9.17, 10 Jan 2002 on ide0(3,1), internal journal
EXT3-fs: mounted filesystem with ordered data mode.
kjournald starting. Commit interval 5 seconds
EXT3 FS 2.4-0.9.17, 10 Jan 2002 on ide1(22,65), internal journal
EXT3-fs: mounted filesystem with ordered data mode.
kjournald starting. Commit interval 5 seconds
EXT3 FS 2.4-0.9.17, 10 Jan 2002 on md(9,1), internal journal
EXT3-fs: mounted filesystem with ordered data mode.
eepro100.c:v1.09j-t 9/29/99 Donald Becker
http://www.scyld.com/network/eepro100.html
eepro100.c: $Revision: 1.36 $ 2000/11/17 Modified by Andrey V. Savochkin
```

```
<saw@sa
w.sw.com.sg> and others
eth0: OEM i82557/i82558 10/100 Ethernet, 00:0B:CD:20:99:42, IRQ 20.
Board assembly 287579-001, Physical connectors present: RJ45
Primary interface chip i82555 PHY #1.
General self-test: passed.
Serial sub-system self-test: passed.
Internal registers self-test: passed.
ROM checksum self-test: passed (0x04f4518b).
8139too Fast Ethernet driver 0.9.25
eth1: SMC1211TX EZCard 10/100 (RealTek RTL8139) at 0xd0257000,
00:10:b5:d6:21:37
, IRQ 21
eth1: Identified 8139 chip type 'RTL-8139B'
IPVS: Connection hash table configured (size=4096, memory=32Kbytes)
IPVS: Each connection entry needs 120 bytes at least
IPVS: ipvs loaded.
eth1: Setting 100mbps full-duplex based on auto-negotiated partner ability
41e1.
Intel 810 + AC97 Audio, version 0.21, 17:25:35 Sep 20 2002
PCI: Setting latency timer of device 00:1f.5 to 64
i810: Intel ICH4 found at IO 0x2400 and 0x2000, IRQ 17
i810_audio: Audio Controller supports 6 channels.
ac97_codec: AC97 Audio codec, id: 0x4144:0x5372 (Analog Devices AD1981)
i810_audio: AC'97 codec 0 supports AMAP, total channels = 2
Linux agpgart interface v0.99 (c) Jeff Hartmann
agpgart: Maximum main memory to use for agp memory: 196M
agpgart: Detected an Intel 845G Chipset.
agpgart: detected 8060K stolen memory.
agpgart: AGP aperture is 128M @ 0xf0000000
memory : ce738f60
device eth0 entered promiscuous mode
Installing knfsd (copyright (C) 1996 okir@monad.swb.de).
.....
```

and the system boots up correctly and runs!
Any ideas on what is going on?

Once more I apologise for the long logs but I wanted to give as much detail as possible.

Sincerely,
T. Kontogiannis
University of the Aegean
Karlovasi – Samos
GREECE