

Ext3 failing on RAID5 array, with 4-port SATA_SIL (SIL3114)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-04/msg06765.html>

- *From:* "Simon Matthews" <simon.d.matthews@xxxxxxxxx>
 - *Date:* Sun, 30 Apr 2006 14:42:57 -0700
-

I keep getting failures on an ext3 filesystem on a RAID5 array. The journal aborts and the fs goes into ro mode.. The first error message appears to be:

```
EXT3-fs error (device md1): ext3_new_block: Allocating block in system zone - block = 10125314
```

This was working -- with an array that is composed of 2 SATA disks and one IDE disk. The SATA disks were on a 2 port SIIG card that uses the SiI 3112.

From LSPCI, the working card looks like:

```
SiI 3112 [SATALink/SATARaid] Serial ATA Controller (rev 02)
```

However, when I rebuilt the array with a 4-port SIIG card (using the SiI 3114) and an additional SATA disk, I keep getting failures -- with the journal aborting.

I tried all combinations: using only 3 SATA disks, using the original 2 SATA disks and the IDE disk (ie. the same disks as the earlier, working array). Still the same errors.

Does anyone have any suggestions? Is this a problem with the SIL 3114?

I am running a RedHat kernel:

```
Linux outlaw 2.6.9-34.EL #1 Fri Feb 24 16:44:51 EST 2006 i686 i686  
i386 GNU/Linux
```

From dmesg (with the 4-port SIIG SATA card installed) (see the end of

this for the ext3 journal failure). I have tried to pick out only the relevant lines:

```
sata_sil version 0.54  
ACPI: PCI interrupt 0000:00:0b.0[A] -> GSI 10 (level, low) -> IRQ 10  
ata1: SATA max UDMA/100 cmd 0xE08C0080 ctl 0xE08C008A bmdma 0xE08C0000 irq 10
```

Ext3 failing on RAID5 array, with 4-port SATA_SIL (SIL3114)

```
ata2: SATA max UDMA/100 cmd 0xE08C00C0 ctl 0xE08C00CA bmdma 0xE08C0008 irq 10
ata3: SATA max UDMA/100 cmd 0xE08C0280 ctl 0xE08C028A bmdma 0xE08C0200 irq 10
ata4: SATA max UDMA/100 cmd 0xE08C02C0 ctl 0xE08C02CA bmdma 0xE08C0208 irq 10
ata1: dev 0 cfg 49:2f00 82:74eb 83:7fea 84:4023 85:74e8 86:3c02 87:4023 88:203f
ata1: dev 0 ATA, max UDMA/100, 488397168 sectors: lba48
ata1: dev 0 configured for UDMA/100
scsi1 : sata_sil
ata2: dev 0 cfg 49:2f00 82:74eb 83:7fea 84:4023 85:74e8 86:3c02 87:4023 88:203f
ata2: dev 0 ATA, max UDMA/100, 488397168 sectors: lba48
ata2: dev 0 configured for UDMA/100
scsi2 : sata_sil
ata3: dev 0 cfg 49:2f00 82:346b 83:7d01 84:4023 85:3469 86:3c01 87:4023 88:207f
ata3: dev 0 ATA, max UDMA/133, 586072368 sectors: lba48
ata3: dev 0 configured for UDMA/100
scsi3 : sata_sil
ata4: no device found (phy stat 00000000)
scsi4 : sata_sil
Vendor: ATA Model: HDS722525VLSA80 Rev: V36O
Type: Direct-Access ANSI SCSI revision: 05
Vendor: ATA Model: HDS722525VLSA80 Rev: V36O
Type: Direct-Access ANSI SCSI revision: 05
Vendor: ATA Model: ST3300631AS Rev: 3.04
Type: Direct-Access ANSI SCSI revision: 05
SCSI device sda: 488397168 512-byte hdwr sectors (250059 MB)
SCSI device sda: drive cache: write back
sda: sda1
Attached scsi disk sda at scsi1, channel 0, id 0, lun 0
SCSI device sdb: 488397168 512-byte hdwr sectors (250059 MB)
SCSI device sdb: drive cache: write back
sdb:<6>parport0: PC-style at 0x378 (0x778) [PCSPP,TRISTATE]
sdb1
Attached scsi disk sdb at scsi2, channel 0, id 0, lun 0
SCSI device sdc: 586072368 512-byte hdwr sectors (300069 MB)
SCSI device sdc: drive cache: write back
sdc:<6>parport_pc: Via 686A parallel port: io=0x378
USB Universal Host Controller Interface driver v2.2
ACPI: PCI interrupt 0000:00:04.2[D] -> GSI 5 (level, low) -> IRQ 5
uhci_hcd 0000:00:04.2: UHCI Host Controller
sdc1 sdc2
Attached scsi disk sdc at scsi3, channel 0, id 0, lun 0
```

<some lines deleted>

```
hub 2-0:1.0: 2 ports detected
md: Autodetecting RAID arrays.
md: autorun ...
md: considering sdc1 ...
md: adding sdc1 ...
md: adding sdb1 ...
md: adding sda1 ...
md: adding hdc1 ...
```

Ext3 failing on RAID5 array, with 4-port SATA_SIL (SIL3114)

```
md: created md1
md: bind<hdc1>
md: bind<sda1>
md: bind<sdb1>
md: bind<sdcl>
md: running: <sdcl><sdb1><sda1><hdc1>
md: md1: raid array is not clean --- starting background reconstruction
raid5: automatically using best checksumming function: pIII_sse
pIII_sse : 1964.000 MB/sec
raid5: using function: pIII_sse (1964.000 MB/sec)
md: raid5 personality registered as nr 4
raid5: device sdcl operational as raid disk 2
raid5: device sdb1 operational as raid disk 1
raid5: device sda1 operational as raid disk 0
raid5: device hdc1 operational as raid disk 3
raid5: allocated 4207kB for md1
raid5: raid level 5 set md1 active with 4 out of 4 devices, algorithm 2
RAID5 conf printout:
--- rd:4 wd:4 fd:0
disk 0, o:1, dev:sda1
disk 1, o:1, dev:sdb1
disk 2, o:1, dev:sdcl
disk 3, o:1, dev:hdc1
md: ... autorun DONE.
md: syncing RAID array md1
md: minimum _guaranteed_ reconstruction speed: 1000 KB/sec/disc.
md: using maximum available idle IO bandwidth (but not more than 200000
KB/sec) for reconstruction.
md: using 128k window, over a total of 244195904 blocks.
md: resuming recovery of md1 from checkpoint.
NET: Registered protocol family 10
Disabled Privacy Extensions on device c0366c20(lo)
IPv6 over IPv4 tunneling driver
divert: not allocating divert_blk for non-ethernet device sit0
ACPI: Power Button (FF) [PWRF]
EXT3 FS on dm-0, internal journal
kjournald starting. Commit interval 5 seconds
EXT3 FS on hda1, internal journal
EXT3-fs: mounted filesystem with ordered data mode.
Adding 1048568k swap on /dev/VolGroup00/LogVol01. Priority:-1 extents:1
st0: Block limits 1 - 16777215 bytes.
ip_tables: (C) 2000-2002 Netfilter core team
ip_tables: (C) 2000-2002 Netfilter core team
eth0: link up, 100Mbps, full-duplex, lpa 0xC5E1
i2c /dev entries driver
eth0: no IPv6 routers present
lp0: using parport0 (polling).
lp0: console ready
Installing knfsd (copyright (C) 1996 okir@xxxxxxxxxxxx).
kjournald starting. Commit interval 5 seconds
EXT3 FS on md1, internal journal
```

