

## Re: [PATCH] sun disk label: fix signed int usage for sector count

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-08/msg09147.html>

---

- *From:* Jeff Mahoney <[jeffm@xxxxxxxx](mailto:jeffm@xxxxxxxx)>
  - *Date:* Sat, 26 Aug 2006 12:02:46 -0400
- 

-----BEGIN PGP SIGNED MESSAGE-----

Hash: SHA1

Jan Engelhardt wrote:

The current sun disklabel code uses a signed int for the sector count. When partitions larger than 1 TB are used, the cast to a sector\_t causes the partition sizes to be invalid:

Is not it that the sun disklabel does not even support [ptabs/partitions] more than 1 TB?

You're right. The Solaris side of things treats this as signed. I'm not sure how the user reporting this problem ended up with partitions larger than 1 TB. It looks like Solaris just treats the entire label as invalid and rejects the whole thing.

I'm not sure how many people out there are using Sun disk labels with sizes > 1TB. It seems like a pretty rare corner case, but there's no reason any data stored in those partitions wouldn't be invalid, and it will suddenly cut them off. Is this a rare enough occurrence that we don't care?

The following patch rejects the partitions, but it would be trivial to switch it to a warning. I'll post it separately if we agree this is the correct solution.

--Jeff

Solaris treats Sun disklabel partition sizes as a signed int, and rejects partitions that claim to be larger than 1 TB.

Linux likewise treats the sector count as a signed int, but doesn't do any checking, and passes it to put\_partition, which casts it causing it to be sign extended.

This patch performs checking to see if individual partitions are valid, and rejects the disk label entirely if they are not.

Re: [PATCH] sun disk label: fix signed int usage for sector count

Signed-off-by: Jeff Mahoney <jeffm@xxxxxxxx>

```
- --- linux-2.6.17/fs/partitions/sun.c 2006-01-02 22:21:10.000000000 -0500
+++ linux-2.6.17.fix/fs/partitions/sun.c 2006-08-26 11:54:54.000000000 -0400
@@ -74,10 +74,18 @@ int sun_partition(struct parsed_partitio
spc = be16_to_cpu(label->ntrks) * be16_to_cpu(label->nsect);
for (i = 0; i < 8; i++, p++) {
unsigned long st_sector;
-- int num_sectors;
+ unsigned int num_sectors;

st_sector = be32_to_cpu(p->start_cylinder) * spc;
num_sectors = be32_to_cpu(p->num_sectors);
+
+ if (num_sectors > INT_MAX) {
+ printk("Dev %s Sun disklable: partition %d has "
+ "invalid size > 1 TB\n", bdevname(bdev, b), i);
+ put_dev_sector(sec);
+ return 0;
+ }
+
if (num_sectors) {
put_partition(state, slot, st_sector, num_sectors);
if (label->infos[i].id == LINUX_RAID_PARTITION)
```

---

Jeff Mahoney

SUSE Labs

-----BEGIN PGP SIGNATURE-----

Version: GnuPG v1.4.2 (GNU/Linux)

Comment: Using GnuPG with SUSE - <http://enigmail.mozdev.org>

iD8DBQFE8HCmLPWxlyuTD7IRAIqZAJ0beMj4oMVsfv0T7RQ1cok751MfCQCfZuAW  
M5hesfVBbgyUmpmzsYBf06w=  
=o35O

-----END PGP SIGNATURE-----

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in  
the body of a message to majordomo@xxxxxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>

Re: [PATCH] sun disk label: fix signed int usage for sector count