

## Re: Support HDIO\_GETGEO on device-mapper volumes

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-02/msg07449.html>

---

- *From:* "Molle Bestefich" <[molle.bestefich@xxxxxxxxxx](mailto:molle.bestefich@xxxxxxxxxx)>
  - *Date:* Tue, 21 Feb 2006 14:30:14 +0100
- 

Phillip Susi wrote:

Right. Unfortunately, it seems this was specifically removed from the kernel because it was problematic,

Ok. Got an archive link or other references?

IIRC, grub just uses the geometry it finds in the MBR, or you can override with the geometry command.

Last I checked, grub used HDIO\_GETGEO when running under Linux.

Which makes it pretty useless if what you say above is true. Except of course in the particular cases where the bogus values that the Linux kernel makes up happens to be correct.

Horribly unsafe way for a boot loader to function, if you ask me. (Not really grub's fault, more the Linux kernel it seems.)

In the case where fdisk must initialize a brand new MBR, then it should assume sensible defaults and allow you to override.

There are no sensible defaults.

There's the translation that the BIOS (and thus DOS, PartitionMagic, yadda) uses, and then there's everything else, which is wrong.

It won't do you any good to fill in "sensible defaults", since if it doesn't comply with what the BIOS decides (or you decide in the BIOS setup as is the case on my Compaq), those old programs just won't work.

And Linux doesn't care about CHS values anyway, so no point in filling

## Re: Support HDIO\_GETGEO on device-mapper volumes

in hocus pocus "sensible" stuff on that account.

Most of the time though, people will already have a drive formatted by windows which will have the correct geometry written in the MBR, so the linux tools will just use that.

Yes, well, but it WOULD be nice if Linux was able to partition the disk in a compatible way.

It's a stupid feature that one HAS to first boot a Windows CD, partition the harddrive and install Windows and then install Linux afterwards. Should be doable the other way around.

Moreover, the MBR values are not guaranteed to be correct if you want to partition your disk in a way compatible with the BIOS in your system. The disk might have been moved from another PC (using another translation), in which case the MBR values will be wrong. Another example is dd'ing an entire drive to a larger drive – the BIOS might use different translation for the two drives, and the MBR values will be incorrect.

Ok, quite theoretical perhaps, but the main idea is that if you want to partition a disk to be compatible with DOS, Windows etc. on the PC the disk is sitting in, you have to use BIOS, not MBR values.

(MBR values can be good for other things, for example preparing a disk for a different system.)

Ok, the better approach might be to destroy HDIO\_GETGEO entirely and tell userspace apps to use `/sys/firmware/edd/int13_dev` instead. It provides for a cleaner implementation.

But then again, that breaks compatibility with current tools.

They are already broken since the HDIO\_GETGEO isn't returning the values from the bios.

Ok.

What is the status of edd? Is that officially supported or is it deprecated?

## Re: Support HDIO\_GETGEO on device-mapper volumes

Can't imagine why it would be deprecated.

If it is ok to use, then it would probably be a good place for fdisk to get that info from, though it's going to have to figure out which int13\_dev ( if any ) corresponds to the device you are partitioning.

Comparing the first and last sector should do.  
If not, we can error out without mapping certain devices.  
If that's a problem to any users, they can use a tool to write unique sigs to the start ~2kB or end ~2kB of disks..

Sensible defaults for well known devices like /dev/hda should cover most cases, but not all.

Nah, I don't think that's a good idea.  
Figure out the right mappings or error out, don't guess.

For instance, I recently configured a server's bios to disable the primary ide channel because that disk was bad and I needed to boot from the second mirrored drive, so /dev/hdb actually corresponded to bios device 80, not /dev/hda.

And there are far more weird and obscure cases than this.

A good way to get the ball rolling could be if we could all agree that there are a number of tools that need the BIOS CHS numbers. Then we can decide where the universal logic that attempts to guess which devices are which should go. Then a proper implementation could be discussed and conjured, and dmraid, <whatever>fdisk, perhaps LILO/GRUB etc. could be made dependant on it. And THEN the HDIO\_GETGEO could be ripped...

Aye, well put. So let's see here. As far as I can see, the only tools that need this information are ones that can create MBRs, so that's fdisk, cfdisk, parted, and friends. Everything else can/should just use the values listed in the MBR.

Agree?

Re: Support HDIO\_GETGEO on device-mapper volumes

GRUB and LILO might want to see CHS values too, in case you're preparing a boot loader on a PC without int13 LBA extensions. Ok, that's pretty far out, but could happen :-). If nothing else, then because the user accidentally chose CHS mode.

dmraid might want to see the CHS values, in case there's an ATARAID adapter which might or might not be capping drives to their last GB, and dmraid wants to know whether the capping feature is enabled. Not sure if this is a sensible way to gather that information, but it might be the only one for some adapters (who knows).

Can't think of anything else on top of my head (not saying that there isn't any :-)).

-

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

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

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