



Klaus Knopper is the creator of Knoppix and co-founder of the LinuxTag expo. He currently works as a teacher, programmer, and consultant. If you have a configuration problem, or if you just want to learn more about how Linux works, send your questions to: klaus@linux-magazine.com

IBM Microdrive

 I'm running Knoppix 5.1.1 (CD version) on an IBM ThinkPad 600e laptop (192MB RAM; 400 MHz CPU). Everything boots up okay, but I have to use nodma cheat code. I'm trying to access an IBM Microdrive (512MB) through the PCMCIA slot using a PCMCIA adapter. While I can eventually access the drive and the files that it contains, the process takes a long time – so much so that it really isn't usable.

Once Knoppix has booted up, I insert the PCMCIA adapter containing the Microdrive; there are no beeps to indicate that it has been recognized and configured correctly. After about 5-10 minutes, a "PCMCIA disk [hde1]" icon appears on the desktop. Clicking on the icon to

ASK KLAUS!



mount it opens a window with a progress bar. The window's heading is `mounting /dev/hde1` (using source `/mnt/hde1`). However, the progress bar doesn't change and remains at 0% throughout the mounting process.

After another 5-10 minutes, the drive is mounted (indicated by the green arrow on the icon). Clicking the icon invokes Konqueror, but this takes another 15 minutes or so to display the contents of the Microdrive (hde1).

The stages described above take similar lengths of time when I issue the equivalent commands via a terminal window at the command line.

Any ideas what the problem is? Should I be entering any cheat codes at boot-up to aid PCMCIA access/configuration, or editing specific scripts? I've noted that Knoppix starts the `cardmgr` at start-up and also adds an entry for hde1 to the `/etc/fstab` file.



See Listing 1 for your `dmesg` output. While the flash disk adapter was detected (hde: IBM-DSCM-10512, CFA DISK drive, ide2 at 0x100-0x107,0x10e on irq 3, hde: 1052352 sectors (538 MB) w/60KiB Cache, CHS = 1044/16/63), it seems that the pcmcia subsystem has assigned a bad interrupt to that controller. IRQ 3 is usually used for serial devices, and if your notebook still has some of these, interrupt 3 is already allocated and now used twice. Although this would work for PCI devices, it just doesn't in the case of IDE controllers.

The `dmesg` output helped in identifying the problem – most people would have given up waiting after two minutes.

Now for the solution, we need to exclude interrupt 3 from the list of free interrupts available to pcmcia.

For 16-bit cards (such as your PCMCIA adapter), this is a `cardmgr` or `pcmciautils` config option.

Do a `cat /proc/interrupts` to find interrupts that are already in use and exclude them in `/etc/pcmcia/config.opts` or `/usr/lib/pcmciautils/config.opts` (depending on which pcmcia subsystem you use) like this:

```
exclude irq 3
```

Afterwards, you'll need to reload the pcmcia subsystem:

Listing 1: dmesg output

```
01 Probing IDE interface ide2...
02 hde: IBM-DSCM-10512,
   CFA DISK drive
03 ide2 at 0x100-0x107,0x10e
   on irq 3
04 hde: max request size: 128KiB
05 hde: 1052352 sectors (538 MB)
   w/60KiB Cache, CHS=1044/16/63
06 hde:<4>hde: lost interrupt
07 hde: lost interrupt
08 hde: lost interrupt
09 hde: lost interrupt
10 hde1
11 ide-cs: hde: Vpp = 0.0
12 hde: lost interrupt
13 hde: lost interrupt
14 hde: lost interrupt
15 hde: lost interrupt
16 hde: lost interrupt
```

```
/etc/init.d/pcmcia reload
```

or

```
/etc/init.d/pcmciautils reload
```

If there are really no interrupts left because other notebook components have already claimed them, a work-around could be to use a USB2 card reader instead of the pcmcia adapter.

Replacing Windows

? OK, I am a boring old “Git” who has reached his three score years and ten, and who started with DOS in the ‘80s and is still using Windows XP. What I want to know is how to set up and use Linux in one of its many forms (i.e., SUSE, Ubuntu, Mandriva, whatever) to augment or to replace Windows. I have tried on a number of occasions to get started by loading various Linux distros, but to date I cannot replicate a set-up that I can use to replace Windows.

Issues like loading devices such as printer, etc., and getting online using a router rather than a modem are what I need most help with.

What I want is an “Idiots Guide” (i.e., you put the 1 there and the 2 there.) Someone must have produced such a guide by now. Where is it? Someone help me – I have only ten to twenty years left to do it.

💡 Sorry, but I’m probably not the best person to answer this. If you want to replace Windows, which you seem to know very well, by something that works like Windows and looks like Windows, please just keep using Windows. Linux is not a “replacement” – it is a different operating system with different applications and different ways of doing things, and it offers a lot of *choice* to do things in different ways instead of a single recipe that may work or fail.

I personally believe that GNU/Linux is way easier to install and use than Windows, but that’s just my own experience.

There is an online article called “Linux is NOT windows” [4], which may be a good starting point for under-

standing why GNU/Linux is not designed as a replacement for anything.

“Put in the installation DVD and follow the instructions,” would be the easy answer to your question, but I would not be happy with such an answer. I like to know what’s going on when working with software.

Video Card for Beryl

? I have a Fedora Core 6 system. The details as revealed with the `uname` command are as follows:

```
uname --all
Linux Linux 2.6.20-1.2962.fc6 ↗
#1 SMP Tue Jun 19 18:24:12 ↗
EDT 2007 i686
athlon i386 GNU/Linux
```

I originally installed the system with a 16Mb NVIDIA TNT video card as I intended to review Linux in place of my Windows XP system.

Now I have converted over to Fedora and already upgraded the memory, and now I’d like to upgrade my video card and use the Beryl interface. I have purchased a NVIDIA 7600 GS card with 256Mb on an AGP interface. I have installed the card and all seems good as far as the system is concerned; however, it looks like I need to install a later driver for the NVIDIA card. These are:

```
kmod-nvidia-100.14.11-1.↗
2.6.20_1.2962.fc6.i686.rpm
xorg-x11-drv-nvidia-100.14.11↗
-1.lvn6.2.i386.rpm
```

I got these off the Fedora site following a recommendation that I use the prepared rpm packages rather than the Linux download straight from the NVIDIA site.

I could have used `rpm -ivh` to install but opted to use a howto that suggested yum. I am now confused about what level NVIDIA I should install – i585 or i686 – and what steps I should take next.



In general, I usually recommend buying a video card that does not require proprietary drivers, especially since these can greatly influence your systems stability.

Beryl and its successor compiz-fusion work well with 3D acceleration on older ATI/Radeon cards and new intel on-board chipsets. (I also like graphics cards that do not require a noisy fan or heatsink.) But on request, I have also installed the NVIDIA driver for NVIDIA cards a couple of times.

Usually, I do not install vendor-specific RPMs, but use the original installation package from NVIDIA’s Linux driver website. The installation goes smoothly and automatically compiles and installs the correct kernel module, provided that the kernel sources of your running kernel are installed correctly. Also, it does the necessary changes to your `/etc/X11/xorg.conf`. The drawback is that you have to rerun the installation each time you update the kernel, but you would have to install or even rebuild the distribution-specific NVIDIA kernel module packages in that case as well, so there is not much difference.

If you install rpm packages instead of using NVIDIA’s installer, which of the packages to use depends on your kernel version. Yum or rpm should fail with an appropriate error message (instead of installing a non-working version) if you try to install the wrong module. Check your kernel version (`uname -a` may be sufficient). If your version is a kernel compiled for i585, then you will have to use the i585 NVIDIA kernel module; and if it is i686, then use the i686 type. In all other cases, running the original NVIDIA installer will most likely succeed, unless it is an unsupported card type.

Wireless Card

? I have an Intel Pro/Wireless 3945 in a LG M1 Express Dual notebook, and I cannot activate the card. Can you help me?



The IPW3945 driver is not yet part of the official kernel, but you can get it online [5]. Note that, apart from the kernel module, you also need a “regulatory userspace daemon” and the firmware, all of which are available with instructions on the same page. When the daemon some day be-



comes obsolete (hopefully), chances are that the driver will make it into the official kernel.

Remote Control

? I have been experimenting with the advice and knowledge from *Linux Magazine* and I am a bit confused. Following the advice about MythTV, I dared to configure it on my computer, an AMD Athlon 64 (model 3800, Socket AM2) with 1GB DDR2 and an XFX graphic card with an NVIDIA GeForce 8500 processor.

The TV card is a Hauppauge WinTV NOVA T-500 PCI, with two TV tuners for Terrestrial Digital TV.

I installed Ubuntu from scratch, downloaded the driver for NVIDIA and, following the instructions on the MythTV page, installed the drivers for the Hauppauge card. Later, using synaptic, I installed all the MythTV packages and ... it worked! Everything is perfect, except for one thing: The remote control doesn't work! I've looked in a thousand places and I am unable to get it working. I was wondering if you could help me configure it.

The other small problem is that I have music and videos stored in a network hard disk and I have configured MythTV to access this disk, but Ubuntu doesn't load it automatically when booting. Ubuntu waits for me to access the disk first, for example, when clicking on the disk in Nautilus, and then mounts it without any problems. If I don't access the disk first, it doesn't mount. In *fstab*, I have it configured as "auto". What am I doing wrong?

I love Linux and *Linux Magazine*. Thank you for helping us see that Linux isn't that hard.



GNU/Linux itself is

surely not hard to use, but some applications are complex and have a pretty big learning curve. MythTV is supposed to be easy to use, but configuring it can be quite painful, especially since not all hardware-vendors are Linux-friendly.

The thing with DVB cards is that (sometimes) parts of the chipset are well-supported, and others aren't. Sometimes vendors sell the same card with the exact same version and product number, but with different chipsets inside – you only find out when looking into the output of *lspci -n* and *lspci -v*. A DVB card contains several components that require different kernel modules to support them (i.e. you may get a picture, but are unable to switch channels if only the controller chipset is supported, but the tuner isn't.)

The best description of your card that I could find is in the mythtv.org wiki [1]. Since the card already works except for the remote control, I believe that you are not a victim of the "diversit chipset" that is warned about on that page. There also is a section in the wiki about configuring the remote control under Ubuntu, which contains the right settings for all keys on the remote.

Regarding your network drive: This would be a */etc/fstab* line for mounting an NFS network drive (server IP 192.168.0.1) automatically at boot:

```
192.168.0.1:/data ↗
/media/nfs ↗
nfs bg,rsize=8192,↗
wsize=8192,hard,↗
nfsvers=3 0 0
```

/media/nfs must exist, and at least *portmap* must have started locally prior to mount accessing this nfs drive.

We are using nfs version 3 here, with some read/write optimization and safe data transfer (option *hard* always retries until a transaction is complete, while *soft* would time out if the network is down for some time), *bg* is used to keep up trying to mount in background even if the server is unavailable or not reach-

able at boot time. Note that you don't need the "auto" keyword in the options list, as it is the default.

If you add "user" in the options list, you don't have to be root to mount */media/nfs*.

If the nfs mount does not work on boot, the *bg* option will make sure it is retried later.

Please also check that there is an init script that actually *does* network mounts. It may be the case for your setup that there is */etc/init.d/mountall.sh*, but nfs is explicitly excluded in there. Then, depending on the distribution, there should be a *../init.d/mountnfs.sh* script that is started later after the network has been set up correctly.

For Samba, the */etc/fstab* entry would look like

```
//192.168.0.1/sharename ↗
/media/sharename smbfs ↗
credentials=/etc/↗
sharename.credentials 0 0
```

where *sharename* is the share defined in */etc/samba/smb.conf* on the server running at 192.168.0.1, and */etc/sharename.credentials* contains these two lines:

```
username=username_on↗
_samba_server
password=username_password↗
_on_samba_server
```

of course, */etc/sharename.credentials* should not be readable for anyone but root on the client. Instead of *smbfs*, you may want to use *cifs*, which is a much-extended version of the Samba client filesystem [2]. ■

INFO

- [1] Hauppauge_WinTV_Nova-T_500_PCI card: http://www.mythtv.org/wiki/index.php/Hauppauge_WinTV_Nova-T_500_PCI
- [2] Server Message Block: http://en.wikipedia.org/wiki/Server_Message_Block
- [3] Dual-booting howto: <http://www.go2linux.org/dual-boot-two-linux-distros-debian-and-mandriva>
- [4] "Linux is NOT Windows": <http://linux.oneandoneis2.org/LNW.htm>
- [5] IPW3945 driver: <http://ipw3945.sourceforge.net/>

Send your Linux questions to klaus@linux-magazine.com.