



Special Note – New Meeting Location This Month

Our third meeting of this year will be at IBM Canada's offices in the TD Centre building at the corner of Portage and Main. We'll be meeting at the lobby on the main floor, and Steve Moffat will take us up to the meeting room just before the meeting starts.

This month's meeting is on November 12th at 7:30 PM. Please arrive a little early for the meeting, as it will take some time for Steve to get people up to the meeting room.

Parking is available either in the parkade behind the TD building (off Albert St.), or in the ground level lot just north of the TD building. Entrance to the lot is from Albert Street, behind the parkade. Either way, parking is \$1.25 flat rate for the evening.

As a reminder, the MUUG regular and Linux SIG meetings have been combined.

Contact Information

To contact the MUUG board for membership information or anything else, send e-mail to board@muug.mb.ca. We have a Web presence as well, at <http://www.muug.mb.ca/>, where you can find all kinds of information, including details of upcoming and past meetings and presentations and references related to them.

To contact the newsletter editor (and I know you want to shower him with dozens of well-written article submissions), e-mail editor@muug.mb.ca.

This Month's Meeting Topic

John Cormie of Plexxus Interactive (part of Musiplexx) will provide a live demonstration of the Caldera desktop for the November meeting. As usual, the main meeting topic will be preceded by a round-table discussion of any questions attendees might have on Unix-related issues they've encountered. Bring all your questions, easy or difficult!

LUG/nut 4 Installation Miscellaneous Notes from the Trenches

By Gilbert Detillieux

Shortly after getting my copy of the LUG/nut 4 CD, I used it to install RedHat 3.0.3 Linux. My PC for this install was a Pentium 133, with 32 MB of memory, a 698 MB EIDE disk (as `/dev/hdb`, in addition to a 1GB disk for Windows 95 as `/dev/hda`), an ATAPI 2x CD-ROM (as `/dev/hdd`), an ATI Graphics Expression (Mach64 based) video card with 2MB memory, an SB16 PnP sound card, and a D-Link DE-250CT Ethernet card.

Booting the Kernel

The installation itself went very smoothly, as per the instructions. I used a boot image from the `images/1213` directory (i.e. a 1.2.13 kernel). The image `boot0015.img` contained support for IDE/ATAPI CD-ROM and for various network cards, including the NE2000 (with which the DE-250 is compatible).

Note that the installation instruction sheet for LUG/nut 4 inaccurately states that there's an `images/2940` subdirectory on the CD, with support for the Adaptec 2940 PCI SCSI controller. Instead, there's a bare kernel file in the `images/custom` subdirectory, which can be copied onto the boot disk you've created beforehand.

After booting the kernel from the boot disk, and loading the two ramdisk floppies at the appropriate prompts, the installation program automatically started up. I selected the option to install from CD-ROM, and it automatically found the CD-ROM drive and mounted the CD. (Note that at boot prompt, I had entered "linux hdd=cdrom" so that it could find the drive, which was the slave device on the secondary IDE controller, despite the fact that there was no master device. On a single-disk system, this likely wouldn't be necessary, since the CD-ROM would be found as `/dev/hdb`.)

Installation

One very nice thing about RedHat's installation program is that it asks all the questions about packages to be installed up front, rather than as the installation proceeds. There are 386 packages to choose from, but they are fortunately arranged into 26 fairly logical sets. You can select or de-select entire sets, and for each selected set, you'll then have the option to select or de-select individual packages within. Once all the desired packages have been selected, the actual installation requires no attention, so you can take a coffee break (or a lunch break if you've got a slow system).

The instruction sheet recommends that you install the `a.out` compatibility package, which is a good idea. It's relatively small, and can be useful in case you run across some older, binary-only packages for Linux. Make sure you install any of the ELF-related support packages, since this is the way most Linux software is going. Also, if you're planning on doing the upgrade to the Linux 2.0 kernel afterwards, don't bother loading the kernel-source package (in the Development set) or the iBCS package (Other set).

Updates

The instruction sheet briefly mentions how to upgrade to a Linux 2.0 kernel, and refers to an online document that gives further instructions. However, neither document mentions the other update packages, which should be installed first, for the sake of completeness.

continued

The CD has an updates/RPMS directory, which contains 19 update packages. After you've booted your installed RedHat Linux system, you can very easily install these using the following two rpm commands:

```
rpm -Uvh /mnt/cdrom/updates/RPMS/*.rpm
rpm -Uvh /mnt/cdrom/updates/RPMS/NetKit*.rpm
```

(Of course, make sure that you've mounted the CD beforehand.) The second rpm command is required since the first one will give a warning and won't install the NetKit-B update, since it requires the portmap update first. After the first command installs the other packages, portmap will have been updated, and the second command should now successfully install NetKit-B.

Of course, if you're not installing a complete system, with network support, you may want to be more selective, and only install the update packages you need.

Upgrading to Linux 2.0

After doing the above, the upgrade to Linux 2.0 should be easy, and can be boiled down to the one following command:

```
rpm -Uvh
/mnt/cdrom/updates/2.0-kernel/RPMS/*.rpm
```

[Ed. note: enter this all as one line]

Note that there are 13 update packages here, not 11, as indicated in the doc/HOW-TO-UPGRADE-TO-2.0.txt file.

After installing the updates, you can go to the /usr/src/linux directory, and build your new kernel. (Check the README file there, as well as the Kernel-HOWTO in the /usr/doc/HOWTO directory, if you're not familiar with the steps involved in building a new kernel.)

A word of caution about the 2.0 kernel: Depending on your hardware, you could encounter serious problems with this kernel, particularly on certain systems with buggy IDE controllers. I

happened to hit the jackpot here, and wound up with so much file system damage to my root partition that I decided to reinstall from scratch. This time, I loaded the 2.0.11 kernel source, which I got on the Net, and built my kernel with that. This proved to be more successful.

Note that if you're going to use a 2.0 kernel, you should probably pick up a more recent source on the Net in any case. In addition to many bug fixes, many changes were made since the 2.0.0 release to improve POSIX compliance. I highly recommend going with the 2.0.11 release or later.

That Darn D-Link!

One of the most frustrating problems I came across (second only to the IDE bug, above) had nothing to do with Linux, per se, but only seemed to show up with Linux. My D-Link DE-250, which had been working fine under Windows 95, refused to work with Linux's NE2000 driver. As soon as I'd configure the interface, and try to send a packet out, the whole system would hang completely, requiring a hard reset.

I finally discovered the problem, when I ran D-Link's SETUP250.EXE program, to change the card's configuration. There are a couple settings related to timing. One setting, IO 16 Control, was set to Late rather than Normal. Although the "Late" setting was OK for the Windows 95 driver, it didn't work for Linux, nor did it work for D-Link's own diagnostic in the SETUP250.EXE program (which also hung). After changing the setting to "Normal," both the diagnostic routine and Linux were happy (and so was I)!

So, if you have a D-Link card, make sure you check its setup configuration before you spend a lot of time and effort trying to get it to work with Linux.

Flying at Mach64

The XFree86 3.1.2 server for the

Mach64, which comes with RedHat 3.0.3, supports many of the older Mach64 based video cards, but not the ATI Graphics Expression. When I tried using it to probe the card, it did find the card and the Mach64 chipset, but failed to identify both the clock and the (built-in) RAMDAC for this particular card. I picked up the latest beta (XFree86 3.1.2G, available at <http://www.XFree86.org/>) server for the Mach64, and it worked fine. It knew how to deal with the Expression's clock and built-in RAMDAC.

After that, using xf86config to configure the server for this card and my multi-sync monitor was a breeze. The configuration file that was produced worked almost flawlessly on the first try. The only problem was, ironically, with the 640x480 mode - the server picked the best available scan rate, but it must have been just far enough from the monitor's acceptable tolerances that it couldn't lock onto a stable image. After commenting out some of the "Modeline" commands, both the server and the monitor were able to agree on a stable image at a 60 Hz vertical scan rate. Also, the configuration included an 800x600 mode at 32 bits/pixel, which the card didn't support, so the server simply rejected it. (Removing that mode from the configuration file was trivial, and made the server happy again.)

General Impressions

The installation of RedHat, in general, was quite easy, and the system seems quite complete and easy to use. For someone more familiar with Slackware, like me, there are some subtle differences in the setup, and it takes a bit of getting used to. Particularly, the configuration of the rc.* files used by init is more complicated than under Slackware, albeit closer to the standards followed by most System V UNIXes. The default LILO configuration was annoying compared to the more sensible defaults in

Slackware.

Basic network configuration was pretty easy, but doing some of the extra stuff, such as configuring the NIS client support, was not obvious and rather poorly documented. (I still haven't configured amd, the auto-mount daemon, which is even more poorly documented than the NIS support.)

Finally, there were a few signs of sloppiness in the setup of the system. Some of the symlinks in /usr, to corresponding directories in /var, were simply missing, which could cause some programs to fail if they use the old path.

The /etc/csh.cshrc file, used by tcsh, contained an error in the setting of the HOSTNAME environment variable, and an even more annoying error in the checking of the TERM variable. The latter resulted in an error message in the output of any rsh command to that host, which caused the rcp command to fail due to a protocol error.

The default /etc/crontab file would periodically truncate several log files once they exceed a certain size, mailing the old content to "root" first. However, if sendmail is not yet correctly configured, you can lose this content. They should either set up a better procedure to "cycle" the log files, or not bother doing anything with them by default.

Overall, the quality of the software was very good, and the installation procedures are improving. Linux is well on its way to becoming the UNIX system for the masses... and you can't argue with the price!

Drive Letters Explained

The following e-mail was sent to Jerry Pournelle, author and Byte columnist, and explains in detail DOS's drive lettering algorithm.

From: Claude Felizardo
<claude@ultrapan.west.enterworks.com>

Date: Tue, 24 Sep 1996 11:31:28 -0700 (PDT)

Subj: drive letter assignment on PCs (posted with author's permission)

Mr. Pournelle,

Just read your Byte column about trying to add a ZIP drive to your system and I think I can explain the problems you are experiencing regarding musical drive letters.

When a PC boots, drive letters are allocated in a very complex but deterministic order. The basic rules are as follows:

1. All primary partitions before any logical partitions
2. IDE before SCSI
3. Loadable drivers last as specified in startup files (ie, config.sys and autoexec.bat)

When the PC boots, the BIOS tries to identify all of the hard drives present to determine which drive to boot from. IDE drives are always scanned first (primary master, primary slave, secondary master, secondary slave) followed by SCSI drives. Because the startup files haven't been read yet, those drives are added later.

If you have more than one SCSI host adapter in your system, the order is determined by the bios on the card(s) usually the card with the lowest address first. Embedded SCSI controllers on the motherboard are generally searched last, depending on motherboard BIOS. SCSI drives on the same string are searched in ascending order according to the device id (and presumably logical unit). Not sure what happens with devices above id=7 since the host adapter is usually at 7.

In the case of removable drives, as I'm sure you know, it depends if there is a cartridge physically present at boot time. As it turns out, as you'll see below, it also depends on how the cartridge was partitioned.

Keep in mind that a physical drive can have up to 4 primary partitions, only one of which may be active. Any other primary partitions will not be accessible until the system is rebooted. The extended partition can have any number of logical drives defined.

Once all of the physical drives have been identified the boot record on the first drive is loaded and the code executed. This is where OS/2's Boot Manager or Linux's LILO can be installed to allow the user to decide which OS to boot. DOS (including Win95) can only be booted from a primary partition. OS/2 and linux and presumably NT can be booted from any kind of partition, primary or logical.

Once the OS is loaded, it scans the various logical partitions and assigns drive letters depending on what file systems it supports. DOS only understands FAT so it skips any HPFS or EXT2 (linux) partitions. OS/2 will read FAT and HPFS so it will mount what it can see. Linux can read FAT, EXT2 and HPFS but then it's unix which doesn't use drive letters so it has no problems. Also, DOS can only be booted off of the first primary partition.

Once the set of readable partitions has been created, all active primary partitions are assigned a drive letter starting with C: for each physical drive present in the order described above. Logical partitions in the extended (or secondary) partition are skipped during this pass.

Next, all readable logical partitions are assigned drive letters. This is what causes drive letters to seem to skip from drive to drive. At this point, the OS can start to read its startup files (config.sys and autoexec.bat) to load any additional drivers such as CD drives (be it scsi, ide, or otherwise), parallel port drives, network drives, etc.

The order is usually dependent on the order they appear in the startup files but there are usually optional parameters to override this. For example, the "/

L:n" option to specify the drive letter for a CD drive.

Adding a physical drive (or removing or reassigning scsi IDs) can play havoc with the drive letters. If your first drive has a logical partition as well as a primary, and you add a new disk with a primary partition, the new primary will grab your D:

Original system:
drive 1
c: primary
d: logical

Adding 2nd drive with a primary
drive 1 drive 2
c: primary d: primary
e: logical f: logical

If you had more than one logical on that first drive then all of your logical partitions will get bumped appropriately. However, if your new drive only has logical partitions:

Adding 2nd drive with no primary
drive 1 drive 2
c: primary e: logical
d: logical f: logical

The only problem will be some wasted space at the front of the drive where FDISK seems to want to reserve space for an invisible primary partition that is not used. Just make sure that is marked inactive!

This explains why the parallel port version of the ZIP drive is always added at the end but the SCSI version gets inserted at seemingly random places. Most likely the cartridge has been partitioned with a primary partition meaning it will be allocated a drive letter BEFORE any logical partitions on any of your FIXED drives. It depends on whether or not a cartridge was left in the drive at power up. One way you may get around this is to repartition your cartridges.

In the case of your other system where adding the scsi zip didn't seem to affect drive letters, your first hard drive probably has no logical partitions and your 2nd drive has only a single primary.

What really screws up the works is when you add (E)IDE to a system which was previously all SCSI. Because IDE drives are allocated first, the entire SCSI chain gets pushed back resulting in none of previously installed OS's working. DOS won't boot on anything but the 1st primary.

Both OS/2 and windows have hard coded drive letters in their startup files so everything gets hosed. You pretty much have to reinstall everything from scratch.

What seems to work best is to only have one primary partition, everything else should contain just logical partitions. If using non FAT partitions, place these at the end of your FAT partitions so that drive letters between DOS and OS/2 are consistent (I boot OS/2 at E:). For CD drives and network drives, reserve space for the removable hard drives so that they can be added anytime without scrambling things.

Anyways, hope that explains it. sorry this got so long. I've had a ton of fun reinstalling DOS; Win3.0, 3.1; OS/2 2.0, 2.1, 3.0; Linux; and Win95 on my system which has grown from a single SCSI to 4 scsi drives and most recently the addition of an EIDE drive (aahhhh!). Adding a SCSI CD drive was a snap, trying to add a IDE cd drive via a sound card from different vendors was pure hell.

While I have not added a removable hard drive to my system yet, I have talked to friends with similar configurations to take a guess at what I'll have to do.

If I've gotten anything wrong, please let me know. Your column is one of the first that I read in each issue of Byte which I've read since high school

in the late 70's. Some of my all time favorite books have been those that you have written in collaboration with Mr. Niven. Looking forward to when Legacy of the Herot hits paperback.

Sincerely, Claude Felizardo

Claude A. Felizardo
enterWorks.com, a Telos Company
claude@west.enterWorks.com
(818) 351-2341 x303
fax: (818) 351-0751
claude@kelvin.jpl.nasa.gov
claude@pacbell.net

Red Hat Linux

More at <http://www.li.org/News/>

Durham, NC -- Red Hat Software, Inc. is pleased to announce the availability of Red Hat Linux release 4.0 for Intel, Alpha and Sparc computers. This release is the first synchronized release of *any* operating system for these three platforms, and represents a huge leap in reducing your administrative and support costs for your heterogeneous network.

This release offers many substantial improvements over previous releases, including additional hardware support, simplified installation, rewritten network configuration tools, dramatic performance improvements, and many more.

New features in Red Hat Linux 4.0

Modular 2.0.18 kernel! One kernel for all hardware, support for much more hardware. Kernel is distributed as an RPM package. This means that to upgrade the kernel you only need to upgrade the kernel RPM package, which is just a single command.

Only one installation floppy! Two for PCMCIA and FTP installs. Both floppies are included in the Official Red Hat Linux/Intel boxed set product.