Red Hat Linux Frequently Asked Questions

Donnie Barnes, djb@redhat.com

V3.0.4, March 28, 1996

Contents

1 Introduction **3**

2 General Questions **3**

2.1 Where can I get information on Red Hat Linux? . . . 3
2.2 What Media is Red Hat Linux Available on? . . . 3
2.3 The Red Hat FTP site is *slow*. Are there mirrors? . . . 4
2.4 How do I order from the Red Hat Software? . . . 4
2.5 How do I contact Red Hat Software? . . . 4
2.6 How do I get new updates to Red Hat as they happen? . . . 5
2.7 How do I get support if I bought the Caldera, PHT, etc CD-ROM? . . . 5
2.8 What is the relationship between Red Hat and Caldera, PHT, etc? . . . 5
2.9 What do I do if I have media (manual or CD-ROM) problems? . . . 5
2.10 What's the quickest and best way to get support for Red Hat? . . . 6
2.11 How do I report a problem? . . . 6
2.12 What versions of Red Hat Linux are there? . . . 7
2.13 What hardware does Red Hat Linux support? . . . 7
2.14 How do I get Linux to work with my Adaptec 2940? . . . 7
2.15 Why won't my Adaptec (274x, 284x) work? . . . 8
2.16 Why don't the 1.3.32 or 2940 images boot? . . . 8
2.17 What applications are there for Red Hat Linux? . . . 8
2.18 Does Red Hat Linux include source code? . . . 8
2.19 What version of X Windows is included with Red Hat? . . . 9
2.20 Are Accelerated-X and Metro-X easy to install? . . . 9
2.21 What packages do we ship in our latest release? . . . 9
2.22 Does Red Hat support UMSDOS installs? . . . 9

2.23 How do I use the rescue.img with RedHat 2.x? 9

3 Installation Questions **9**

3.1 Do I need swap space? 9

3.2 Why do I get an error enabling swap space in 2.1? 9

3.3 I can't get the mouse to work on my Laptop. Why? 10

3.4 Why doesn't the Linux Library work with Red Hat 2.x? 10

3.5 How do I get my IDE CD-ROM on my SB16 to work? 10

3.6 How do I get my Sony CDU31 or CDU33 to work? 11

3.7 How do I do a network install of Red Hat 3.0.3? 11

3.8 Why doesn't my Adaptec 154x work? 11

3.9 How do I run the install program again? 11

3.10 Why am I missing some programs? 12

3.11 Do I need to reinstall if I didn't get a LILO prompt? 12

3.12 How do I mount my CD-ROM? 13

3.13 Why does my Trident 9400 show up as an 8900? 14

3.14 Why does Linux only see part of my RAM? 14

3.15 How do I install to my DOS partition? 14

3.16 What does the boot disk message *really* mean? 14

3.17 How do I make my own *special* boot disk for 2.x? 16

3.18 How do I make my own *special* boot disk for 3.0.3? 17

3.19 How do I use the NYS/NIS system that comes with RedHat? 17

3.20 Why can't I get to Unixware now that I've installed Linux? 17

3.21 Why won't my Trantor T-130B work? 17

3.22 Why does my Diamond do odd things under X Windows? 18

3.23 Where do I go to find an "errata" sheet? 18

3.24 I installed Metro-X, now how do I use it? 18

4 Using Red Hat Linux **18**

4.1 Where are the printing commands (lpr, etc)? 18

4.2 Why does Red Hat 'adduser' put each user in his/her own group? 18

4.3 Why doesn't the xmh shipped with Red Hat work properly? 19

4.4 How do I get color ls to work? 19

4.5 I rebuilt my kernel, but I got vmlinux instead of vmlinuz. Why? 19
4.6 How do I make my own packages (RPMs)? 20
4.7 I compiled a program, but I can't run it. What gives? 20
4.8 Why does iBCS break when I recompile my kernel? 20
4.9 Why does Netscape display a ton of errors when I start it? 20
4.10 Why are include files missing when I try and compile? 21
4.11 Where is Openwin? 21
4.12 Where is package XYZ? 21
4.13 Why doesn't xtexsh work? 21
4.14 How do I add nameservers in the netcfg tool? 21

5 Upgrading **22**

5.1 How do I upgrade from 2.x or Caldera 1.0 to 3.0.3? 22
5.2 I upgraded 2.0 to 2.1 and now Glint doesn't work. Why? 22
5.3 I upgraded my kernel to 1.3.x and now PPP doesn't work. Why? 22

6 Copyright Notice **22**

1 Introduction

This is the second edition of the Official Red Hat Linux FAQ. It is maintained by support@redhat.com and all comments or suggestions for this FAQ should be sent to that address. To get updates to this FAQ, send mail to *faq@redhat.com* for an automated reply, or look in the `docs` directory of any Red Hat mirror.

2 General Questions

2.1 Where can I get information on Red Hat Linux?

Contact Red Hat Software for ordering information. The best way is via `http://www.redhat.com`. For other info, send mail to *info@redhat.com* for an automated reply (quite helpful...*really*). You can also poke around `ftp://ftp.redhat.com` and `http://www.redhat.com` for new announcements and product info.

2.2 What Media is Red Hat Linux Available on?

It is available on CD and via FTP at this point. To find out where you can FTP it from, see the question on Red Hat Mirrors.

2.3 The Red Hat FTP site is *slow*. Are there mirrors?

There are several mirrors of the Red Hat FTP site:

```
ftp.cms.uncwil.edu         /linux/redhat
ftp.wilmington.net         /linux/redhat
ftp.ibp.fr                 /pub/linux/distributions/redhat
ftp.gwdg.de                /pub/linux/install/redhat
ftp.msu.ru                 /pub/Linux/RedHat
ftp.cvut.cz                /pub/linux/redhat
sunsite.mff.cuni.cz        /OS/Linux/Distributions/Redhat
ftp.ton.tut.fi             /pub/Linux/RedHat
sunsite.doc.ic.ac.uk       /packages/linux/sun.../dist.../redhat
ftp.funet.fi               /pub/Linux/images/RedHat
gd.tuwien.ac.at            /pub/opsys/linux/redhat
ftp.dstc.edu.au            /pub/linux-redhat
gatekeeper.digital.com.au  /pub/redhat
ftp.lab.kdd.co.jp          /OS/Linux/packages/redhat
ftp.pht.com                /pub/linux/redhat
linux.ucs.indiana.edu      /pub/linux/redhat
ftp.uoknor.edu             /linux/redhat
ftp.cc.gatech.edu          /pub/linux/distributions/redhat
uiarchive.cso.uiuc.edu     /pub/systems/linux/dist.../redhat
ftp.caldera.com            /pub/mirrors/redhat
sunsite.unc.edu            /pub/Linux/distributions/redhat
gatekeeper.dec.com         /pub/linux/redhat
ftp.rge.com                /pub/systems/linux/redhat
```

The rest of this document has path references to files on these mirrors. Just concatenate that file name to any of the above mirror locations to find the file.

2.4 How do I order from the Red Hat Software?

```
Red Hat Software Sales can be reached at:
        phone: (203) 454-5500
               (800) 454-5502
        fax:   (203) 454-2582
        email: sales@redhat.com
        WWW:   http://www.redhat.com
```

2.5 How do I contact Red Hat Software?

```
Red Hat Software
25 Sylvan Rd South, Suite F
Westport, CT 06880
redhat@redhat.com
(203) 454-5500
```

```
(800) 546-7274
(203) 454-2582 (fax)
http://www.redhat.com
ftp://ftp.redhat.com
```

2.6 How do I get new updates to Red Hat as they happen?

Keep your eye on:

`current/i386/updates`

on any official mirror for new packages we've made since the current CD-ROM was pressed. Also check the non-free and contrib directories for packages that we can't distribute on CD-ROM or packages that users have sent to us. We will also make periodic announcements to redhat-announce-list mailing list with updates.

2.7 How do I get support if I bought the Caldera, PHT, etc CD-ROM?

Please do not ask for support unless you purchase a Red Hat Linux CD-ROM product produced by Red Hat Software. Yes, if you have Caldera then you have Red Hat Linux, but you bought it from Caldera. If you would like support for Caldera, please see their web site: `http://www.caldera.com` for their support options.

Also, Pacific Hi-Tech has sold CDs with "Official Red Hat Linux" on them. While it is a good CD, it was produced by them and is supported by them. Please don't contact us for help with PHT CDs.

2.8 What is the relationship between Red Hat and Caldera, PHT, etc?

All are separate companies. As with any maker of Linux CD-ROMs, Caldera, Pacific Hi-Tech, Yggdrasil, and LSL are all free to use Red Hat as the basis for their products. This does, however, make those vendors solely responsible for supporting their product if they want it supported. We can not support other vendors products directly, though users of Red Hat from other vendors are welcomed to participate in our mailing lists and make complete use of our WWW site and FTP site.

2.9 What do I do if I have media (manual or CD-ROM) problems?

If you are experiencing problems with the product you received (eg your manual is missing pages or your CD-ROM is scratched), please contact the Red Hat Software at `sales@redhat.com` or by phone using the above contact information. The problem will be rectified promptly. Please do not report these problems on the `redhat-list` or to `support` as they are likely isolated and your product will simply need to be replaced.

(No, this is not a common problem, but it does happen from time to time.)

2.10 What's the quickest and best way to get support for Red Hat?

You can search for answers to your questions first via `http://www.redhat.com/`. This has a search engine to search all the Red Hat mailing lists. If you don't find what you need, you can join the redhat-list and then send mail to *redhat-list@redhat.com* with your questions.

To subscribe to the list, send mail to *redhat-list-request@redhat.com* with the subject line reading:

subscribe

Due to relatively high traffic on that list, we also have the redhat-announce-list. It is used only for announcements of major concern to the Red Hat community and is low traffic.

To subscribe to the list, send mail to *redhat-announce-list-request@redhat.com* with the subject line reading:

subscribe

We also have a list for developers.

To subscribe to the list, send mail to *redhat-devel-list-request@redhat.com* with the subject line reading:

subscribe

You can also peruse our bug tracking system via the WWW. See `http://www.redhat.com/bugs`. All open and closed bugs are listed there. We are working to improve the speed and layout of the bug tracking system. It will get better soon...

2.11 How do I report a problem?

Send mail to tech-sup@redhat.com or report the bug to the redhat-list. Again, if you didn't buy Red Hat from us or the ACC Bookstore, please see the previous question, "How do I get support...".

When reporting problems, please be as detailed as possible!!! For example, if you can't get Red Hat to see your SCSI drives, tell us:

- What version of Red Hat
- What kind of SCSI card
- what kernel
- what SCSI id's are the drives
- how many IDE drives do you have
- specific details of the problem

If you can't get your CD-ROM recognized, tell us:

- What version of Red Hat
- make and model of CD-ROM
- type of interface (IDE, SB, etc)

- bootdisk version
- hardware configuration

Again, be as detailed as possible about your problem!

2.12 What versions of Red Hat Linux are there?

There have been several now. Here's a rundown:

- Preview (or Beta) - Summer '94
- Halloween - Fall '94
- Mother's Day 1.0 - Summer '95
- Mother's Day 1.1 - Late Summer '95
- RedHat 2.0beta - Late Late Summer '95
- RedHat 2.0 - Early Fall '95
- RedHat 2.1 - Late Fall '95
- RedHat Linux/AXP 2.1-Beta - Late Fall '95
- RedHat Linux/AXP 2.1 - January '96
- RedHat Linux/Intel 3.0.3 - March '96

2.13 What hardware does Red Hat Linux support?

For a current list of supported hardware, please see:

`http://www.redhat.com/`

Look under the Red Hat support section.

For a better list of video cards, you should see:

`http://www.xfree86.org/`

2.14 How do I get Linux to work with my Adaptec 2940?

Red Hat 2.1 provides several boot images that should work with the 2940. We don't officially support it since there is no 2940 driver in the 1.2.13 kernel distribution, but we are trying to help. None of the images in the `1213` directory support the 2940. You can get images that do support it in the `2940` directory or in the `1332` directory. We suggest using the ones in the `2940` directory. If you can't get those to work, you can try 1.3.32 images. In both cases, please see the next question in the FAQ!

Some tips on getting the drivers to work:

- Disable sync negotiation on the card
- Disable disconnect/reconnect on the card
- Change the transfer rate on *all* devices to 5.0 M/s

All of the above changes need to be made in the 2940 BIOS.

2.15 Why won't my Adaptec (274x, 284x) work?

The Adaptec 274x and 284x are supported but may need to be reconfigured to work properly. It has been reported that the wide negotiation on these cards is not quite working, so it should be disabled.

2.16 Why don't the 1.3.32 or 2940 images boot?

There is a problem with the images in the `1332` and `2940` directory. Under DOS, the images are named `filename.img`. In reality, they are really named `filename.img.gz`. These images are compressed using gzip. If you know it, that's fine. All you need to do is copy the image to your hard drive. Then there is the other problem. In the `dosutils` directory there is a copy of `gzip.exe`. It is not the `gzip` program, it is actually a self extracting archive of the gzip program. You need to copy it to your hard drive as well and run it. It will extract itself and some other files.

Once you've gone this far, you can now uncompress your image and write it to floppy. Do the following:

```
move filename.img filename.gz
gzip -d filename.gz
rawrite
```

Then give rawrite `filename` as the file to write to floppy.

Also, you'll notice the IDE/SCSI CD-ROM images don't exist. We removed them due to space requirements. If you require one of them, you should just use the image numbered one higher (the Mitsumi image). It has drivers you don't need, but should still work).

2.17 What applications are there for Red Hat Linux?

There are many applications for Linux available. Please see `http://www.redhat.com/` and go to the "Linux Applications and Utilities" section. This section is maintained by a private individual that is not related to Red Hat in any way. You will find information on nearly every product available for Linux on his pages.

2.18 Does Red Hat Linux include source code?

Yes. We include the exact source that was used to build the distribution. From release 2.0 on, Red Hat Linux is built on a packaging system called RPM that ONLY uses pristine source (the same as what you'd find at the author's ftp site) and possibly a Red Hat patch. Complete information on RPM is available on `http://www.redhat.com`.

2.19 What version of X Windows is included with Red Hat?

Red Hat 3.0.3 ships with XFree86 3.1.2, and comes standard with a very nice configuration tool.

2.20 Are Accelerated-X and Metro-X easy to install?

Yes. Both will install just fine on a Red Hat system.

2.21 What packages do we ship in our latest release?

For the sake of brevity, we can't list them all here, but you can find a complete list on `http://www.redhat.com/` This will not only give you a list, but will allow you to query the packages to find out in more detail what each of those packages actually has in it. We include not only a description of the package, but a *complete* file list for each package as well.

2.22 Does Red Hat support UMSDOS installs?

In a word, no. It is not a good performance idea. Aside from that, you don't "need" DOS on your machine anyway :-)

2.23 How do I use the rescue.img with RedHat 2.x?

Due to an oversight on our part, no documentation on the rescue disk was included with 2.0 or 2.1. The disk image is there and works, however. To use it:

- Boot your boot install floppy
- Insert your ramdisk1 install floppy when prompted
- Insert the rescue floppy when prompted for ramdisk2

3 Installation Questions

3.1 Do I need swap space?

Yes! If you have 8M of RAM, you *must* enable at least 8M of swap during the installation of Red Hat. If you have more RAM, you should still enable swap, though it isn't mandatory.

3.2 Why do I get an error enabling swap space in 2.1?

Red Hat 2.1 made it out of the office with a bug in the mkswap part of the install. It only occurs when you use swap partitions of about 40M or greater. The solution is to enable it by hand or use our updated image. The updated image is available from any Red Hat mirror in `redhat-2.1/updates/images/ramdisk2-beta.img` (Note: This image currently only does a text install...if you want to do the X install you must do the following "hand" fix with your original images). If you want to fix it by hand, you can do the following:

- Go through the mkswap step and let it fail. Note which device you are using for your swap space.
- Continue to the next step (Hostname).
- Switch to Virtual Console 2 (Control-Alt-F2).
- Type the following commands substituting your device for /dev/blah:

  ```
  mkswap /dev/blah
  swapon /dev/blah
  ```

- Switch back to VC1 (Control-Alt-F1) and continue the install.

3.3 I can't get the mouse to work on my Laptop. Why?

Some laptops will report that the "device is in use" when trying to use the mouse. This is most likely due to the fact that the IRQ is getting stolen by the PCMCIA card services package. To fix it, add the following line to /etc/sysconfig/pcmcia:

```
OPTS=irq_mask=0xefff
```

3.4 Why doesn't the Linux Library work with Red Hat 2.x?

The Linux Library CD-ROM is a couple months old and is a.out. You need to install the aout-libs package off your CD-ROM to get it to work. You can use `glint` or `rpm` to do the install.

It does report that you "do not have a Red Hat system". This is normal. It is really checking for the existence of our old packaging system, RPP. Since we don't use that, it reports that error. You can disregard it and continue the install.

3.5 How do I get my IDE CD-ROM on my SB16 to work?

The Creative Labs kit comes with the IDE Interface set to the third IDE port by default (IDE2, I think). This won't work under Linux as it only recognizes two IDE ports. What you need to do is move a jumper on the card to configure it for IDE1 (the secondary IDE port). I'm not sure if you have to make changes under DOS for this configuration to work, but your manuals should help with this (it *will* still work, but you may have to edit a file or two).

You may also have to specify command line parameters like:

```
boot:  linux hdN=cdrom
```

If you have one IDE drive, N is most likely "b" in this case (ie "linux hdb=cdrom"). If you have two, then it will be "c".

3.6 How do I get my Sony CDU31 or CDU33 to work?

Under 2.x, the LILO message when you boot the first floppy says:

```
Sony 31&33 CD-ROM REQUIRED       cdu31a=base_addr,[irq|0][,PAS]
```

"REQUIRED" means just that. There has been a driver change since whatever version you had that requires you to enter a parameter to get the drive recognized. The probe was breaking other things, so they took it out.

You need to do something like:

```
LILO boot:  linux cdu31a=0x340
```

If 340 is your base address...and yes, even if you have a 33 use 31a, not 33.

Also, at some point in the install it will ask if you need to enter additional hardware params for LILO. Answer yes, and then enter whatever it was you had to enter above to get the install to work.

3.7 How do I do a network install of Red Hat 3.0.3?

Get the RedHat-HOWTO from the `docs` directory of any Red Hat Mirror. It will tell you all you need to know.

3.8 Why doesn't my Adaptec 154x work?

It appears that some 1540's and 1542's are a bit hosed by the Buslogic probe. You will usually get an error referring to "DMA" in the Adaptec probe when this happens. To fix it, you need to enter command line parameters as follows:

```
boot:  linux aha154x=0x340 buslogic=0x334
```

This is assuming, of course, that the base address of your Adaptec is 0x340. If not, put your address in there instead. This will force the buslogic to only probe an address where your Adaptec is not.

3.9 How do I run the install program again?

Reboot your machine using the three disks that you made for the installation. You will most likely need to completely reinstall your system.

If you just want to reconfigure networking or install more packages, try using the control-panel tools.

3.10 Why am I missing some programs?

Try doing a 'df' at the shell prompt. If you have "/" or "/usr" showing up as 100% used, chances are your partition was too small for the entire install. Unfortunately, this does not result in any type of error messages unless you are watching closely (which most people don't when the thing tells you it is going to take 2 hours to install :-). You will most likely need to reinstall and do one of two things: a) make your root or "/usr" partition larger or b) install fewer packages.

Also, if none of the partitions you made for Linux is full, you are probably just missing some packages. You will need to use `glint` or `rpm` to install those.

3.11 Do I need to reinstall if I didn't get a LILO prompt?

Not necessarily. You should be able to use the boot disk that you made for the install to boot to your root partition on the hard drive. Insert the disk, reboot the machine, and at the boot prompt, type:

`linux ramdisk=0 root=/dev/xxxx` where xxxx is the hard drive and partition number. If you installed to the first partition of your first IDE drive, this would be hda1. If you installed to the third partition on your second drive, this would be hdb3. The second partition of your second SCSI drive would be sdb2.

If you don't remember, try what you *think* might work. The worst that will happen is a kernel panic or the machine will hang. In either case, reboot and try another partition.

Once you get it booted, you need to edit /etc/lilo.conf and then run 'lilo'. For more info on how to do that properly, you can then read the LILO-HOWTO in "/usr/doc".

An example lilo.conf to boot linux and DOS is below:

```
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50

image=/vmlinuz
                label=linux
                root=/dev/sdb2
                read-only
other=/dev/sda1
                label = dos
```

You may also need a line in the "other" section with "table = /dev/hda" or whatever drive your DOS partition is on. The only lines from above that should change if you boot linux and DOS are "boot", "root", and "other". Some of what is there is optional, like the append and what YOU actually label the dos partition (some people prefer msdos).

One caveat: Make sure the "boot" line points to your first IDE drive, OR your first SCSI drive IF AND ONLY IF you only have NO IDE drives. That is the only way the machine has to know where to find it. You need to also be sure that the partition in the "boot" line is marked "Active" by either the DOS fdisk, or the linux fdisk (or cfdisk). Here are the "rules" of where "/" must live:

1. If you have 2 IDE drives, your root partition MUST live on one of them, NO MATTER WHAT. This also includes any IDE CD-ROM drives on your PRIMARY controller. If you have one IDE hard drive, and one IDE CD-ROM on the PRIMARY controller, you must install your root partition on the IDE hard drive.
2. If you have 1 IDE drive and SCSI drives, your root partition MUST live on the IDE drive or the SCSI drive THAT IS AT ID 0. No other ID's will work.
3. If you have SCSI only, your root partition MUST live on a drive at ID 0 or ID 1. No other ID's will work.

These rules are non negotiable.

BUT I WANT DOS TO BOOT BY DEFAULT!!!

Easy. Just change the above /etc/lilo.conf to:

```
boot=/dev/sda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50

other=/dev/sda1
        label = dos
image=/vmlinuz
        label=linux
        root=/dev/sdb2
        read-only
```

and then run 'lilo' after saving the file.

3.12 How do I mount my CD-ROM?

If you installed from CD-ROM, you should be able to do:

`mount /mnt/cdrom`

If you get an error that says the mount point does not exist, do:

`mkdir /mnt/cdrom`

If you did an install via NFS: You need to know what type of CD-ROM you have. The best way is to watch the boot messages as you reboot, or look at /var/adm/messages (`more /var/adm/messages`). You most likely have an IDE CD-ROM that will show up as /dev/hdc or /dev/hdd, or you have a Sound Blaster type that is /dev/sbpcd, or you have a SCSI that is /dev/scd0.

In any case, you can do the following:

`mount -t iso9660 /dev/xxxx /mnt/cdrom`

and substitute xxxx for YOUR device. You can then: `cd /mnt/cdrom` and then: `ls` and you will be looking at the contents of the CD-ROM.

3.13 Why does my Trident 9400 show up as an 8900?

XFree86 3.1.2 currently doesn't support the advanced features of the 9400. The timing for the release of the 9400 was just bad for the release dates of 3.1.1, so they weren't able to add the support in. Now that the card is pretty popular and a new release is expected soon of XFree86, we expect to see direct support in the next version of XFree86. For now, you can run resolutions all the way to 1024x768 at 60 hz. If you have trouble using Xconfigurator to set it up, you will probably find that xf86config will do the job (it just doesn't look as nice doing it :-)

3.14 Why does Linux only see part of my RAM?

There are a couple things that could be wrong. On *some* 386's, you need to compile your kernel with "Limit memory to 16M?" enabled. Some AMI BIOS motherboards have a memory option to split memory into two segments. I'm not sure what the option is called, but it has something to do with "linear". You'll find it in your BIOS setup. You may need to specify the amount of memory to the kernel at boot time from lilo. To try this, type: `linux mem=32M` at the lilo prompt. If that works, you can add that line to the append in your /etc/lilo.conf and rerun lilo. The example from above would be:

```
append = "mem=32M"
```

Don't forget to run 'lilo' after editing the file.

3.15 How do I install to my DOS partition?

You can't. None of the current Red Hat releases support UMSDOS installs. If you have one drive that has one big DOS partition with free space you want to use, you have two options. One is to backup all your DOS data somehow and then repartition the drive into smaller partitions and then restore the DOS data back to a DOS partition. You then have a free partition for linux. The second choice is to use fips. fips is a program that will let you change your partition size of a DOS partition without harming the data on it. It is *STRONGLY* suggested that you backup your data before using fips. Red Hat Software makes no warranties or claims as to the suitability of using fips for this purpose nor can we guarantee that data will remain unharmed. I haven't heard of anyone having a bad experience with fips, but I won't guarantee it either.

3.16 What does the boot disk message *really* mean?

Okay, there are a lot of questions on this, so I'll go into it in better detail.

```
Line 1:  Our name and the version you have.
Line 3-4: short description of what to do.

The rest of the message is for lilo parameters.  These
should only be used if you have a Sony CDU31 or 33 -OR-
if you are having problems with the auto detection of other
hardware.
```

```
Be sure that for every parameter you enter that has commas that
you do NOT put a space after the comma.  The only spaces
on the command line should be after 'linux' and between
multiple parameters (ie: mcd=0x340,11 bmouse=5).

Line 6:  This line is almost NEVER necessary and should be used
only as a last resort.  If your hard drives are not being
recognized or giving odd errors, you can try these parameters.
If your first hard drive acting up, you would enter:
        linux hda=xxx,xxx,xxx
where the 'xxx' is your parameters from BIOS.
Line 7:  Use this if your Mitsumi CD-ROM does not get recognized
properly.  An example:  linux mcd=0x340,11
Line 8:  This line is required for the sony cdu31a or 33a.
Use "cdu31a" for BOTH the 31 and the 33 (ie. do NOT use
"cdu33a" as a parameter on the command line).
An example:  linux cdu31a=0x340,0
Line 9:  If you have an Aztech CD268, try:
        linux aztcd=0x340
(or whatever your port is).
Line 10:  If your ethernet card is not recognized, you can
try to use these parameters.  For an ethernet card on
base address 0x340 and irq 10, enter:
        linux ether=10,0x340
Line 11:  If your bus mouse isn't recognized, try this parameter.
Line 12:  For an old Seagate ST0x controller, use this parameter.
Line 13:  For a Future Domain TMC-8xx controller, use this.
Line 14:  For a trantor T128, use this line.
Line 15:  For an NCR-5380, use this line.  For a card at 0x330
and IRQ 11 and DMA 5, you would use:
        linux ncr5380=0x330,11,5
Line 16:  Use this line for an Adaptec 1505, 1510, 1515, 1520,
and 1522 (and Sound Blaster SCSI).  The reconnect and parity
are optional.  For a CD-ROM hooked to ID 1 an a controller
with base address of 0x330 and IRQ 11 you would use:
        linux aha152x=0x330,11,1,reconnect,parity
(you can try it without the last two)
Line 17:  Use this if you are having problems with an Adaptec
1542 (rare).
Line 18:  Use this if you are having trouble with a Buslogic
controller (also rare).
Line 19:  Use this for Pro Audio Spectrums and Pro Audio Studio
cards that aren't recognized.
Line 20:  Use this if you have the "normal" Sound Blaster CD
interface and it doesn't work.  This is the regular ole run
of the mill Sound Blaster type that controls the Panasonic
56x drives, Matsushta Kotobuki drives, TEAC CD55A, and a
```

```
few others.  If you have a Sound Blaster with something
like the Panasonic connected using base address 0x240, you
would use:
        linux sbpcd=0x240,SoundBlaster
Line 21:  If you have an IDE (ATAPI compliant) CD-ROM that is
not getting recognized by the auto probe, you can use this
parameter to tell it where to look.  Normally, if you have
IDE hard drives, they take up the first available drive
letters.  If you had 2 IDE hard drives, your CD-ROM would be
hdc.  If you had one IDE hard drive, it would be hdb.  It's
usually okay to just try them all.  You would use:
        linux hdb=cdrom
if you had one IDE hard drive.
Line 22: A reminder that you must put a "0x" in front of
all Hex values (such as base addresses).
Line 23: Instructions on what to enter.
```

3.17 How do I make my own *special* boot disk for 2.x?

There are people who have linux installed by some means, but need their own *special* boot disk to install their system. This is possible in some cases. Do not try this with kernels greater than version 1.3.48. It will not work. See below for instructions for later kernels.

First of all, we recommend you only try this with the same kernel source version that we ship (ie 1.2.13 for 2.0 and 2.1). You need to build your kernel on your machine and make sure and say 'y' to the following options during the `make config`:

- CONFIG_MODVERSIONS
- CONFIG_BINFMT_ELF

Then build the kernel with `make dep`, `make clean`, and finally `make zImage`. Once done, your new kernel will be `/usr/src/linux/arch/i386/boot/zImage`. You now need to put it on an install floppy. Under linux, the steps to do this would be something like:

```
dd of=/dev/fd0 if=/path/to/boot/images/boot0000.img
mkdir /mnt/floppy
mount -t ext2 /dev/fd0 /mnt/floppy
cp /usr/src/linux/arch/i386/boot/zImage /mnt/floppy/vmlinuz
lilo -r /mnt/floppy
umount /mnt/floppy
```

You should then be able to reboot with the new floppy and the new kernel.

If you must use a kernel version greater than 1.3.48, you can follow the following steps:

- Make sure and say 'y' to the ramdisk question in `make config`.

- Make sure and say 'y' to the UMSDOS and MSDOS questions in `make config`.
- Run "rdev -r <new kernel¿ 49152" on your new kernel.
- Mount the ramdisk 1 floppy image.
- Remove dev/ram2 on the floppy.
- Make a new ram2 with `mknod dev/ram2 b 1 1`
- Unmount the floppy

There will be some weirdness and an error or two displayed as the second ramdisk loads, but it should work.

3.18 How do I make my own *special* boot disk for 3.0.3?

For 3.0.3, the boot disks are now UMSDOS filesystem. This means you can use mcopy or the copy command from DOS to put a new kernel on your disk. If you are using a kernel version less than 1.3.48, build according to the directions for 2.x (above). Then just copy it to any boot floppy as `linux` (not "vmlinuz" as your normally would). You can `rawrite` or `dd` any boot image from the 1213 directory and just overwrite the kernel.

For versions later than 1.3.48, make sure and say 'y' to the ramdisk question in `make config`. Then copy your kernel to a boot image (you can write out any boot image) as above. Then, get the `ramdisk.img` from the 1357 directory and write it to disk. You can then boot using the boot disk and the parameters:

```
boot: linux ramdisk=256 prompt_ramdisk=1 load_ramdisk=1
```

Note that this will not work for PCMCIA installs unless you also build the PCMCIA Card Services package using 1.3.x kernel sources and also install the modules and binaries on your ramdisk. This is *not* trivial.

3.19 How do I use the NYS/NIS system that comes with RedHat?

Please see the NYS-Tips. You can find it at `http://www.redhat.com/`.

3.20 Why can't I get to Unixware now that I've installed Linux?

You need to boot Unixware off of a boot floppy and run 'fdisk /mbr'. That should allow you to get back into Unixware, but will not let you access Linux.

3.21 Why won't my Trantor T-130B work?

The Trantor T-130B is really an NCR-5380 in disguise. But, it will most likely not autodetect. You will need to use the NCR boot image along with a boot parameter. The registers are offset by 8 from a normal card. So, for a T-130B at base address 0x350 and IRQ 5, you would use:

```
ncr5380=0x358,5
```

3.22 Why does my Diamond do odd things under X Windows?

If you experience problems with XFree86 and your Diamond video card, you may try putting:

```
"Option Diamond
```

In your /etc/X11/XF86Config (or in /etc/XF86Config if you use that).

3.23 Where do I go to find an "errata" sheet?

With every release there seem to be some things that break or are overlooked. This things are getting fewer in number, but as of 3.0.3 some problems do still exist. To find solutions to known problems, please see:

`http://www.redhat.com/`

and look for the product page for your particular version. These sheets only exist for Red Hat Linux 3.0.3 and later.

3.24 I installed Metro-X, now how do I use it?

To get the `startx` command to use Metro-X instead of XFree86 you must make `/etc/X11/X` a symbolic link to `/usr/X11R6/bin/Xmetro` This is summarized by:

```
rm /etc/X11/X<br>
ln -s /usr/X11R6/bin/Xmetro /etc/X11/X
```

4 Using Red Hat Linux

4.1 Where are the printing commands (lpr, etc)?

Under 2.0, the printing commands are in with the networking packages. If you didn't install networking, you don't have them. This is easily remedied by installing the NetKit-B RPM off your CD-ROM or by FTP'ing the newer NetKit-B-lpr package from a 2.1 mirror near you (see above for a list).

Under 2.1 you are most likely just missing the NetKit-B-lpr package from your installation. Install that and everything will be fine.

4.2 Why does Red Hat 'adduser' put each user in his/her own group?

It's called the user group scheme, and it's described in the manual. What it basically does is allow you to use groups to allow multiple people to use a common set of files and edit them, transparently. It goes like this:

- You'd like to have a group of people work on a set of files in say, the "/usr/lib/emacs/site-lisp" directory. You trust a few people to mess around in there, but certainly not everyone.

- So you do "chown -R root.emacs /usr/lib/emacs/site-lisp". And you add the proper users to the group.
- To allow the users to actually create files in the directory you do "chmod 775 /usr/lib/emacs/site-lisp".
- This is fine, but when a user creates a new file it gets the group of the users default group (usually users).
- To prevent this you do 'chmod 2775 /usr/lib/emacs/site-lisp' which causes everything in the directory to get created with the "emacs" group.
- This is fine, but the new file needs to be mode 664 for another user in the emacs group. To do this you make the default umask 002.
- Well, this all works good, except that if your default group is "users", every file you create in your home directory will be writable by everybody in "users" (usually everyone).
- To fix this, you make each user have a "private group" as their default group.

At this point, by making the default umask 002 and giving everyone a private default group, you can easily set up groups which users can take advantage of without doing any magic. Just create the group, add the users, and do the above chown and chmod on the group's directories.

4.3 Why doesn't the xmh shipped with Red Hat work properly?

We only ship xmh because it is part of the standard X11 distribution. xmh is basicaly shipped as-is from the X11 distribution, and is most likely broken. We've found that there are several better alternatives. If you want to use mh mail, install the mh RPM and the exmh RPM. exmh is a much better GUI mail reader than xmh.

4.4 How do I get color ls to work?

Some of you may be accustomed to other distributions and the fact that color ls is installed by default. Red Hat does not do that, but it is possible. The best way to find info on it is to do:

`man color-ls`

`man dircolors` In short, you can probably put something like `eval 'dircolors'` in your `.bash_profile` file. You also need `alias ls="color-ls --color=yes"` in there if you want it enabled by default..

4.5 I rebuilt my kernel, but I got vmlinux instead of vmlinuz. Why?

You need to do `make zImage` or `make zlilo` instead of just `make`. `make zlilo` is the best choice for most beginners as it will not only build the kernel, but also install it as the default for you if the make goes well. Be careful doing this, however, as you will lose your old kernel. `make zImage` will build a compressed kernel and place it in `/usr/src/linux/arch/i386/boot`. You must copy it to your root dir and run lilo yourself to use this one.

If you used make, you got a file called vmlinux. This is an *uncompressed* kernel image, and will NOT work! LILO will only boot a compressed kernel.

4.6 How do I make my own packages (RPMs)?

Simple. Read the RPM-HOWTO, available in the `docs` directory on any Red Hat Mirror.

4.7 I compiled a program, but I can't run it. What gives?

Lets say you do the following as root:

```
gcc -o hello hello.c
hello
```

You'll get "hello: command not found"

Why?

By default, '.' is not in your path. You must run:

```
./hello
```

(dot, slash, h, e, l, l, o with no spaces)

For security reasons '.' should not be in your path.

4.8 Why does iBCS break when I recompile my kernel?

You need to answer 'y' to the line that asks:

```
CONFIG_MODVERSIONS [n]
```

If you didn't do that, you'll need to rebuild your kernel and enable it.

4.9 Why does Netscape display a ton of errors when I start it?

You didn't read the Netscape installation instructions!

When you untar the distribution of Netscape, you will get an nls directory. Depending on where you untar the distribution, you will need to set some environment variables. You should add lines similar to the following to your `.bash_profile` file:

```
export XKEYSYMDB=${XKEYSYMDB-/home/foobar/lib/X11/XKeysymDB}
export XNLSPATH=${XNLSPATH-/home/foobar/lib/X11/nls}
```

Change the above paths accordingly to reflect the actual location where you installed everything.

4.10 Why are include files missing when I try and compile?

You most likely didn't install the kernel sources. Find the sources for the kernel on your system and install them. Then make sure `/usr/src/linux` exists. You may need to make a symlink from `/usr/src/linux` to the directory with your kernel version.

Then you need to `cd /usr/src/linux` and do: `make config` `make` and then you can hit control-C after a few seconds (it makes some symlinks that you might need right at the beginning).

4.11 Where is Openwin?

There is currently no Openwindows package for RedHat 2.x. It is available in Redhat 3.0.3 in xview-3.2p1-2.i386.rpm

4.12 Where is package XYZ?

Before asking where some particular piece of software is (like `diff` for example), please see any RedHat mirror and look in `current/i386/RedHat/rpm-contents.gz`. This file has a query of all available RPMs and their contents. `grep` through it for what you want and it will probably lead you to an RPM that you don't have installed. Use `glint` or `rpm -i` to install the proper RPM.

Also check `current/i386/{updates,contrib,non-free}`.

4.13 Why doesn't xtexsh work?

xtexsh is based on older versions of Tcl/Tk than what we ship. Unfortunately, it has not been ported to the newer versions as of yet and it does not work at all.

4.14 How do I add nameservers in the netcfg tool?

netcfg is the victim of poor GUI design in the nameserver department. To add a nameserver:

- Click the 'Add' Button
- Click on the '0.0.0.0' entry in the *list* (not in the entry box)
- Edit the '0.0.0.0' in the entry box to your nameservers IP address
- Hit Enter
- That's it!

5 Upgrading

5.1 How do I upgrade from 2.x or Caldera 1.0 to 3.0.3?

If you have 2.0 installed now and purchased the 2.1 CD-ROM, you can run the “upgrade” script on the CD-ROM. There are a few caveats, however. First of all, you should change to runlevel 1 to do this. This can be done with `telinit 1` as root. This will take you off the network (if any) and kill any unnecessary processes.

You will be presented with a list of config files that have been changed at the end of the upgrade. You may want to move some of them back to their original locations if the formats of the files haven’t changed. A quick look at the new one and your old one should reveal whether or not this has happened.

5.2 I upgraded 2.0 to 2.1 and now Glint doesn’t work. Why?

You most likely didn’t get the new `pythonlib` RPM installed. Install it and it should work again.

5.3 I upgraded my kernel to 1.3.x and now PPP doesn’t work. Why?

If you upgrade your kernel to a 1.3 kernel, you must use ppp-2.2.0. You can most likely find an RPM for this version on your CD or on ftp.redhat.com.

6 Copyright Notice