Red Hat Commercial Linux

Thank you, and congratulations! Red Hat Commercial Linux is the highest quality Linux system available! RHC Linux is more than just a collection of software compiled and put on a CD-ROM. RHC Linux is a powerful and flexible system, allowing you precise control over what is installed on your machine, at all times. RHC Linux also contains a number of graphical system administration tools, to help you manage file systems, printers, users, groups, and TCP/IP, all without config file editing or awkward command lines.

If you have Red Hat Commercial Linux on a CD-ROM, please read the `README` file and other files on the CD-ROM before continuing. Those files will contain information specific to the release you have.

Red Hat Software will strive to provide the best products and service in the Linux community. We welcome all bug reports and suggestions. Please see Appendix A for information about our on-line offerings and details on contacting us.

Contents

Preface – What is Red Hat Commercial Linux?

Red Hat Commercial Linux is a complete operating system based on the Linux kernel. It is easy to install – the installation script can walk you through the process in as little as 15 minutes. The installation system is very flexible – you can install and uninstall individual program "packages" with a couple clicks of your mouse button. It is easy to maintain – package installations can be verified and corrected. It is easy to administer – a set of graphical administration tools bundled into a "control-panel" reduce the hassle of basic every day administration. Complete source packages are provided for everything in the system.

This book documents the installation and system administration of RHC Linux and is divided into five chapters, Quick Start, Installation and Configuration, Installation Systems, Control Panel, and System Administration.

If you are itching to install RHC Linux right away, check out the Quick Start chapter. After reading that you should be able to do a basic installation without much trouble. For new Linux users, I would suggest reading through the entire manual first to get a good understanding of the various systems and installation procedures.

I would like to thank the following people for their generous help: Donnie Barnes, Erik Troan, Jim Fridlund, Greg Galloway, Andrew Majercik, Jon Mauney, Dustin Mollo, Damien Neil, and Tony Richardson. In addition I would like to thank all our beta testers for entrusting their systems to early versions of RHC Linux and for taking the time to submit bug reports from the front. I would also like to thanks Linus Torvalds and the hundreds of developers around the world for creating, truly, one of the wonders of distributed development.

Extra special super–duper thanks go to my wife, Lisa Lee, for putting up with me and all my Linux shenanigans. ♥

Red Hat Lore: Since everyone asks, I used to have a red hat that was my grandfather's old Cornell lacrosse hat from way way back. I lost it somewhere in Philadelphia a few years back. If you know where it is send it my way for a generous reward `:-)`.

Thanks again, and welcome to the wonderful world of Linux!

Marc Ewing
5 June 1995

Chapter 1

Quick Start

This chapter explains, *very briefly*, how to install RHC Linux. It is probably sufficient for anyone who is familiar with Linux.

1.1 Get Red Hat Commercial Linux

Either buy a CD-ROM or download it from an FTP site. If you are getting it from an FTP site be sure to read the README.FTP or other documentation on the FTP site. It will detail eveything you need to get and how to set things up for installation. Basically you need to have an existing Linux system, and a free ext2 partition to put everything. You'll need to download `rpps/*.ser` and `rpps/*.rpp` for all the series you need, `bootstrap/bootstrap.tar`, `config/*`, `bin/*`, `images/rootdisk.img` and a rootdisk from `images/*/*.img`.

1.2 Make a Boot Disk and a Root Disk

The `images` directory on the CD-ROM is the root of the boot disk images. Look at the file `version.idx` and choose a kernel version. Change to the proper kernel version directory. Read the `image.idx` file and choose a hardware configuration. Choose the boot image `bootXXXX.img` that corresponds to the image listed in the file. Use `dd` or `rawrite3.exe` (found in the `dos` directory) to write the proper boot image file to a formatted high density (1.44Mb) 3.5 inch floppy.

Do the same for the root disk, using the `images/rootdisk.img` image.

1.3 Install

Reboot with the boot floppy inserted, and the CD-ROM in your CD-ROM drive. Insert the root disk when asked, answer all the questions, and you are done!

Chapter 2

Installation

RHC Linux installation is a relatively painless process. If all your hardware is configured properly, is compatible with Linux, and Linux recognizes it at boot time, rest assured that the hard part is done!

2.1 Before You Begin

There are a few things you should do before you begin.

If you have not yet read the installation chapter (chapter 2) of Matt Welsh's book *Linux Installation and Getting Started*, you should probably do so now. It covers important issues like partitioning, file systems, swap space and LILO.

If Linux will coexist on your machine with OS/2, you will need to create your partitions with the OS/2 partitioning software. During the installation, do not create new partitions, but set the partition types properly for the Linux partitions.

If you will be connected to a network, be sure you know your IP address, network address, net mask, broadcast address, gateway IP address, and name server IP addresses.

You will need two formatted 3.5 inch high density (1.44Mb) floppy disks for the installation boot disk and root disk. You may want a third disk to create a boot disk for your final system.

Read all of the installation instructions before starting the install! This way you can be sure you will not run in to any surprises and you will be prepared for all the decisions you need to make.

2.2 Creating the Installation Boot and Root Disks

You must first create a boot disk and a root disk to use during the install. These must be 1.44 Mb 3.5 inch floppies. The disk images for these disks are located under the `images` directory on the CD.

If you are creating the boot disks from DOS the `dos/bootroot.exe` program will help you create the proper disks. Run it. Otherwise read on.

To write an image to a floppy, first format the floppy disk using either the DOS `format` command or `fdformat` under Linux. The use either `dos/rawrite.exe` (on the CD) for DOS, or `dd` for Linux, to write the image to the floppy disk.

All the images are located under the `images` directory on the CD. Your first decision is which kernel version to use. Read the `version.idx` file, and choose a kernel version. Change to the proper kernel version directory. For each version, there are a number of different kernel hardware configurations. Read the file `image.idx` and choose a hardware configuration. For example, see figure 2.1 for reference. The first entry on each line is the image number, which corresponds to image file `bootXXXX.img`. The next three extries (separated by semicolons) denote SCSI, Ethernet, and CD-ROM support.

Please note that all images have support for ATAPI CD-ROM drives, and that a SCSI CD-ROM requires only SCSI support. All of the images support all mice, SLIP, CSLIP, PPP, FPU emulation, all file systems except extfs and xiafs, console selection, ELF, SysV IPC, IP forwarding, firewalling and accounting, reverse ARP, and parallel printer support. None of the images have QIC tape support or sound card support. After the installation is complete you may want to rebuild a kernel that only includes the support you need.

Image Number	SCSI	Ethernet	CD-ROM (non-IDE, non-SCSI)
0000;	None;	None;	None (IDE/ATAPI only)
0001;	None;	None;	Sony, Mitsumi
...			
0006;	None;	WD, SMC, 3COM;	Matsushita/Panasonic, Sound Blaster Pro
...			
0020;	All;	Other;	None (IDE/ATAPI only)
...			
0023;	All;	Other;	Aztech/Orchid/Okano/Wearnes

Figure 2.1: Sample image.idx file

After choosing a kernel configuration, write the corresponding `bootXXX.img` file to the first floppy disk and label it. The root disk image is `images/rootdisk.img`. Write it to the second floppy disk and label it. You are now ready to proceed with the installation.

You may also want to write `images/rescue.img` to a third floppy. This can be used in conjuction with the boot disk to recover from a system that is sufficiently screwed up that it won't boot.

If you find that none of the hardware configurations is suitable, or that you wish to use a newer kernel with additional features, you can create your own custom boot disk. The instructions for doing so are in `bootfs/README` on the CD. You can also use any kernel or boot disk that mounts `/dev/fd0` as a ramdisk. In either case you must make sure that the kernel used supports your hardware (including CD-ROM and hard drives) and the iso9660 file system (for a CD-ROM install) or NFS file system (for a network install).

2.3 Installation and Configuration

2.3.1 CD-ROM Installs vs Network Installs

For a normal CD-ROM install, you boot the system with the RHC Linux CD-ROM inserted in the CD-ROM drive. If you wish to install over a network, you will need to mount the RHC Linux CD-ROM on a machine that supports ISO-9660 file systems with Rock Ridge extensions, and NFS. Export the CD-ROM file system via NFS.

You will use the same boot and root disks to do a network install, and you will need to know the NFS server's IP address and the path to the exported CD-ROM.

2.3.2 Booting the Installation System

Boot your machine off the boot floppy you just created. You will first see the `LILO boot:` prompt. The first time through, just hit enter. Watch the messages as the Linux kernel initializes your hardware. If the kernel boots and recognizes all your hardware properly you are set! If it misses something or gets something wrong, you may need to enter some parameters at the `LILO boot:` prompt, or use a different boot image. If you do need to specify some parameters, remember them! You will need them later.

When you see the "VFS: Insert root floppy ..." message, insert the root floppy and press Enter. If all goes well you will be presented with the welcome dialog box.

2.3.3 The Installation Program

Using the Dialog Boxes

Use of the dialog boxes is quite simple. Your arrow keys and the tab key should move you around intuitively. If you are presented with a menu of items from which you are supposed to select one, move the highlight bar over the item you wish to select, and press enter. If you are presented with a list of items and you are supposed to select any number of items (a *checklist*) move the hightlight bar and press the space bar to toggle selection of individual items. Press Enter when you are done selecting items.

Express Install

If you select an Express Install the installation program will walk you through each step of the installation in order. You can also let the Express Install select the series of packages you want to install. For the Mother's Day release, these series are: Base, Emacs, Net Utils, Devel, XFree86, X Utils, and iBCS.

During the Express Install, hitting escape will generally get you to the main RHS Install menu (which you will also get to if you do not choose the Express Install).

The RHS Install Menu

The main installation menu allows you to select any phase of the installation process at any time. You can also have it walk you through each part of the installation in order by selecting "Next" from the menu each time, or rejoin the Express Install. If you chose the steps individually, bear in mind that some steps are prerequisite to others. You should also be aware of the consequences of doing parts of the installation *twice*. In general, **if you make an error from which you are unsure if you can recover without problems, restart the installation.**

Using the Shell During Install

There is an item in the install menu to go to a shell. Selecting this starts up a shell from which you can tinker with the system. Please be aware that before the CD is mounted you will have very few commands at your disposal, and even after the CD is mounted, there are many commands that are unavailable. After packages are installed you will have access to more programs.

If the terminal doesn't look correct when the shell is started, try entering

```
stty sane
```

or

```
setterm -reset
```

To exit the shell, type Control-d.

If you don't know what you are doing you can easily corrupt the installation process while in the shell. *Be careful!*

Step 1: Mount RHS

For a regular install, select "CD". If your CD-ROM drive is not found, and you have some familiarity with Linux, you may be able to get it to mount using the "Device" or "MajMin" options. If you have downloaded RHC Linux from an FTP site you will also need to use the "Device" and "MajMin" options to mount the partition with RHC Linux files on it.

If you have an ATAPI based IDE CD-ROM drive you may need to specify a LILO boot option to force the probe to look harder. The parameter format is:

```
hdx=cdrom
```

where `hdx` can be any of hda, hdb, hdc, or hdd, depending on the interface and drive location of the CD-ROM drive.

If you are doing a network install choose "NFS". You will configure networking immediately so that the CD-ROM can be mounted (see section Configure network).

After configuring networking, you need to specify the server and path that contains the RHC Linux CD-ROM. You will find that things will start up faster if you use an IP address for your NFS server. In fact, if you have no name server, or if it happens to be down, you *must* use an IP address.

Step 2: Partitioning

Read the section on creating partitions in *Linux Installation and Getting Started*. If you have 8 Mb or less you must create and use a swap partition. For decent basic performance, you should have at least 24 Mb of combined RAM and swap space.

You will be given a choice of two partitioning programs `cfdisk` and `fdisk`. The `cfdisk` program is a full screen program and can be somewhat easier for the novice (or expert) to use.

You will then be presented with a list of drives found on your system. Select a drive to partition. After partitioning that drive you will be offered that chance to partition another drive.

After exiting the partitioning program, if you see a message like "Re-read table failed with error ..." you will need to reboot your machine so that Linux can reload the partition data. This generally happens if you changed the partition structure related to extended partitions. If you need to reboot, do so, and go through all the same steps you did to get to the partitioning step, but then skip the partitioning step (you've already done it).

Step 3: Select Swap

The swap partitions listed are those marked as Linux swap (82) during the partitioning phase. Each swap partition must be formatted with `mkswap` before it can be used. If you have not run `mkswap` on these partitions before, do so now. In fact, it is a good idea to just run `mkswap` every time. If you wish to check for bad blocks on your swap partitions, indicate so, and any blocks found to be bad will be ignored.

Step 4: File systems

You must format all file systems you wish to use in Linux. If you are reinstalling, or installing over an old Linux installation *you must still format all file systems*. This ensures that all filesystems are properly created and are empty.

Select file systems to format (with the space bar) and press enter. You will then have your choice of file system types. Unless you have a compelling reason not to, choose the `ext2` file system type. You will then have the option to check for bad blocks If you wish to check for bad blocks, indicate so, and any blocks found to be bad will be ignored.

When you have formatted all your partitions, select "Cancel".

Step 5: Fstab

At this stage you define the layout of your file system – that is, where the partitions you created earlier are mounted. The top part of the dialog box indicates the layout as it is currently set up (it will be empty at first). You add partitions to the file system by selecting "Add" and choosing the partition. For each partition you use you need to enter a mount point. Keep in mind that you must have a partition mounted at `/` (the *root* partition), and that no two partitions can have the same mount point.

As an example, say you've created a 32 Mb partition `/dev/hda2` for your `/` partition, and you've created a 300 Mb partition `/dev/hdb1` for `/usr`. Your layout might look like this:

```
/dev/hda2          /
/dev/hdb1          /usr
```

When you are done arranging your file system select "Done".

Step 6: Select Series

If you choose to let the Express Install select series for you you will not see this stage.

In this stage you select (with the space bar) the series you wish to install. The Base series is not optional, but everything else is. The X Utils series contains the Linux Installation manager and the `control-panel`, so if at all possible you will want to install it. It is a good idea to install the HOWTOs series as it contains lots of documentation. If you select the LaTeX series you must also select the TeX series. Also, if you select any of the X Windows based series, you must select the XFree86 series.

The iBCS series enables Intel Binary Compatibility, which will allow you to run iBCS compliant binaries for systems like SCO, UnixWare and other SVR4 systems. The WordPerfect demonstration series contains a demo version of SCO WordPerfect, and it requires the iBCS series.

If you do not install something now, it can be easily installed later using `rpp-install`, tt rpp-install-series, or the Linux Installation Manager. The Linux Installation Manager allows a much finer granularity of installation control than is available at this stage – individual packages from a series can be installed. However, *if you do not install the XFree86 and X Utils series, you will not be able to use the Linux Installation Manager or the control-panel tools.*

Step 7: Select X Server

If you are installing X Windows, you will need to select a server that works with your video card. Each server name indicates the video card chipset it can drive.

Step 8: Hostname

If you haven't already configured your hostname for a network install, you will be asked to enter a hostname. This is asked now because the next install phase (installing packages) makes use of the hostname. If you are connected to a network, use your fully qualified domain name (FQND) as your hostname (for example, `splat.redhat.com`). Otherwise, use any single token name you want (for example `splat`).

Step 9: Init Swap, Make FS, Mount FS, Bootstrap

This stage requires no input from you. Your swap partitions will be formatted and initialized, your filesystems will be created and mounted, and a bootstrap system will be installed.

Step 10: Install Series

This stage requires no input from you. As the packages are installed a progress meter is displayed and an estimated time to completion is given. Depending on the speed of your hardware and the number of series you selected, this step can take anywhere from 15 minutes to 2 hours.

Step 11: Install X Server, Install Kernel

This stage requires no input from you. The X Server you selected in stage 7 will be installed, and the kernel will be copied from your boot floppy (which you should have inserted) to your new system.

If you installed the development series, after the kernel is installed the kernel sources will also be installed. Kernel sources are usually needed if you are doing development as many programs reference kernel header files in their sources.

Step 12: Configure Mouse

If you have a mouse, indicate so, and then select the type of mouse you have. If it is a serial mouse you will be asked to select the serial port to which your mouse is connected.

All this stage does is create a symbolic link from `/dev/mouse` to your mouse device. If you select the wrong serial port, move your mouse or install a mouse of a different type, just change the link to point to the new device.

Step 13: X Configuration

Please read all notes in the dialog boxes for this stage. Choose auto probe unless you've done this before and auto probe doesn't work.

After the auto probe (or if you selected no auto probe) you will be asked to verify the parameters (listed below) of your video card. They will be filled with values gleaned from the probe. If they are wrong, or if you did not auto probe, you will need to enter them.

- Chipset
- Video RAM
- Clocks

You will then be presented with a list of monitors. Select one that matches you monitor. You may want to choose "Generic Multisync" or "Generic Monitor" if you don't find a match. You may also select "Custom", in which case you will be prompted for the following information:

- Monitor Bandwidth
- Monitor Horizontal Sync
- Monitor Vertical Refresh

You will then be presented with some video modes that *should* be within the parameters of your card and monitor. You may select as many as you wish. The VESA modes are recommended, and are most likely to work with most hardware. You will then review the details of each mode. Make sure eveything looks correct and within specs. The next dialog box asks you to select the primary video mode. This is the mode in which X will first start. You can cycle through all the modes after X starts by typing Control-Alt-KeyPad+. At this point if you wish to select different video modes, choose Cancel and you will return to the video mode selection dialog box.

The next dialog box presents a list of options that may or may not be useful for your hardware. If you do not know about any of these options, do not select any!

You will then be asked how many buttons are on your mouse. If you have a three button mouse, you will then be asked if yor mouse sends Left and Right signals when the middle mouse button is pressed. If you don't know the answer to this, say no. Later, if you find that your middle button is not working properly, uncomment the ChordMiddle line in `/etc/XF86Config`.

Usually, X configuration will work well enough to get X running, but it will probably not generate an optimal configuration. If the configuration is faulty (ie running `startx` doesn't work), you can try running `Xconfigurator` again, and reconfiguring.

If this configuration fails you will need to manually configure X by editing `/etc/XF86Config`. Read the `XF86Config` man page and the X Configuration section of Matt Melsh's `Linux Installation and Getting Started` for help.

Step 14: Configure Networking

You will already have completed this stage if you are doing a network install. In that case you will be told so and asked if you wish to use this configuration for your system. Normally you will select "yes", but if you wish to change something select "no".

The first question asks if you wish to configure networking and bring it up, or just configure it. There really is no difference between the two, so select "Config and Init".

The second question asks you if you have an internal or external (D-Link) ethernet device. At present D-Link configuration is not included, so choose internal. RHC Linux only supports the internal ethernet cards listed in the hardware compatibility list at present.

You will then be asked to fill in a number of networking parameters. Sometimes defaults will be provided. Table 2.2 lists the parameters that are requested and some sample values. *These are only examples!! Do not use these parameters!!*

After entering the parameters, you will have the oportunity to review them and change them before committing to them.

Parameter	Example
Hostname	acc
Domain Name	redhat.com
FQDN	acc.redhat.com
IP Address	199.183.24.5
Network Address	199.183.24.0
Network Mask	255.255.255.0
Broadcast Address	199.183.24.255
Default Gateway	199.192.24.1
Nameserver	192.100.81.101

Figure 2.2: Networking Parameters

Step 15: Configure Modem

If you have a modem you will be asked which serial port it is on.

All this stage does is create a symbolic link from `/dev/modem` to your modem device. If you select the wrong serial port or move your modem to a different serial port, just change the link to point to the new device.

Step 16: Configure Keyboard

Choose a keyboard map to use for your installed system. The first two letters are the country code. In the US, you should probably use "us.map".

Step 17: Configure Clock and Time Zone

You will be asked what time your CMOS clock keeps. If at all possible use GMT/UTC time. This will require setting your CMOS clock to GMT if it is not already on GMT, which may cause it to read the wrong time under DOS or other operating systems.

Next you need to select a time zone. Be sure to choose the most appropriate time zone.

Step 18: LILO Configuration

LILO is the "LInux LOader" and can be installed on the Master Boot Record of your bootable hard drive to allow Linux (and other operating systems) to boot directly off your hard drive. If you do not install LILO, you will need to use a boot disk to boot Linux.

Generally, you will install LILO on the MBR of the drive yor system tries to boot first, `/dev/hda` for IDE based systems, and `/dev/sda` for SCSI based systems. If you wish to use some other loader, such as the OS/2 Boot Manager, you should install LILO on your root partition, and configure the boot manager to boot it. If you do not want to install LILO anywhere on your hard drive, you can install it on a floppy disk (the `/dev/fd0` option). This will enable you to specify boot parameters while still leaving your hard drive's MBR untouched.

After selecting the LILO location you will have the oportunity to enter some hardware paramters. These are the same parameters used at the LILO boot prompt when you booted the boot disk. If you didn't need to specify anything there to have your hardware detected, you won't need to put anything here.

You will then have the oportunity to add an additional operating system to be booted by LILO (for example, a DOS system on another partition). You will have to select the type of system, and then the aprtition on which the system lives. Then you will enter a name for this system that you can use on the LILO command line to boot it.

You then are shown the `/etc/lilo.conf` file and asked if it looked OK. If you notice that it is wrong, you should cancel the LILO installation by choosing "no". At this point you can redo the LILO configuration step, or go to a shell, edit `/mnt2/etc/lilo.conf` (using any editor you installed – try `vi`), and run `lilo -r /mnt2`. If you give up on LILO all together you will have to create a boot disk.

Matt Welsh's *Linux Installation and Getting Started* has a lot of good information on LILO configuration.

Step 19: Add Users

In this stage you can create users. You have the option of *not* creating a home directory for them. You may not want to create a home directory for them if you will be mounting `/home` from another machine.

The users are created with their passwords locked. Before the users can log in you will need to set passwords for each user with the `passwd` command. Do this after finishing the installation and logging in as root.

Step 20: Set Root Password

At this stage you will input a password for the root user. You will be asked to re-type the password, to make sure you didn't make any typing mistakes the first time around.

Step 21: Finish

You must run this finishing step before you are done. It writes lots of configuration data to your new system. *If you skip it, your system will not boot!*

You Are Done!

Congratulations! You can now boot your new Linux system.

2.4 Additional Configuration

Your system is now fully functional but you may still want to do some additional configuration. This section describes some of the various system configuration tasks you may need to do.

2.4.1 Package Installation (or Uninstallation)

Remember that even if you selected every series available, there are still packages on the CD-ROM that are not installed on your system, for example, all the source packages. Use the rpp commands or the Linux Installation Manager to install additional packages, or uninstall packages you do not need.

2.4.2 Sendmail

A default `sendmail.cf` file will be installed in `/etc`. The default configuration should work for most SMTP-only sites. It will *not* work for UUCP sites; you will need to generate a new `sendmail.cf` if you need to use UUCP mail transfers.

To generate a `sendmail.cf`, install `m4` and the `sendmail` sources. Read the `README` in `/usr/src/mail/sendmail-8.6.` for more details on creating `sendmail` configuration files. Also, O'Reilly publishes a good sendmail reference called "sendmail" by Bryan Costales.

One common sendmail configuration is to have a single machine act as a mail gateway for all the machines on your network. For instance, at Red Hat Software we have a machine `redhat.com` that does all our mail. On that machine we simply need to add the names of machines for which `redhat.com` will handle mail to `/etc/sendmail.cw`. Here is an example:

```
# sendmail.cw - include all aliases for your machine here.
torgo.redhat.com
poodle.redhat.com
devel.redhat.com
```

Then on the other machines, `torgo`, `poodle`, and `devel`, we need to edit `/etc/sendmail.cf` to "masquerade" as `redhat.com` when sending mail, and to forward any local mail processing to `redhat.com`. Find the `DH` and `DM` lines in `/etc/sendmail.cf` and edit them thusly:

```
# who gets all local email traffic ($R has precedence for unqualified names)
DHredhat.com

# who I masquerade as (null for no masquerading)
DMredhat.com
```

With this type of configuration, all mail sent will appear as if it were sent from `redhat.com` and any mail sent to `torgo.redhat.com` or the other hosts, will be delivered to `redhat.com`.

2.4.3 /etc/hosts.allow

As a security measure, most network services are managed by a protective program called a *TCP wrapper*. The protected services are those listed in `/etc/inetd.conf` that specify `/usr/sbin/tcpd`. `tcpd` can allow or deny access to a service based on the origin of the request.

By default Red Hat Commercial Linux allows all service requests. To disable or limit services you can edit `/etc/hosts.allow`. Here is an example `/etc/hosts.allow` file:

```
ALL: redhat.com .redhat.com
in.talkd: ALL
in.ntalkd: ALL
in.fingerd: ALL
in.ftpd: ALL
```

This configuration allows all connections from `redhat.com` and `*.redhat.com` machines. It also allows talk, finger, and ftp requests from all machines.

`tcpd` allows much more sophisticated access control, using a combination of `/etc/hosts.allow` and `/etc/hosts.deny`. Read the tcpd(8) and hosts_access(5) man pages for complete details.

2.4.4 Anonymous FTP and World Wide Web

Both anonymous FTP and World Wide Web can be set up using the `control-panel`. See section 4 for details.

2.4.5 NFS Configuration

All the required NFS deamons are started at boot time by default, but the *portmapper* which controls access to NFS services, utilizes the `/etc/hosts.allow` and `/etc/hosts.deny` files for access control. The default installation enables all portmap service requests, but you may wish to disable or restrict such services. To do so edit `/etc/hosts.allow` and use `portmap: <pattern>` lines to control connections to `portmap`. In order to avoid deadlocks, the portmap program does not attempt to look up the remote host name or user name, nor will it try to match NIS netgroups. The upshot of all this is that only network number patterns will work for portmap access control. For example to allow all hosts in the `redhat.com` domain (whose network address is 199.183.24.0), you would add the following line to `/etc/hosts.allow`:

```
portmap: 199.183.24.0/255.255.255.0
```

See the hosts_access(5) and rpc.portmap(8) man pages for complete access control details.

To actually export a filesystem edit `/etc/exports`. For example,

```
/mnt/rhscd          *.redhat.com(ro)
/mnt/cdrom          *.redhat.com(ro)
```

allows hosts matching `*.redhat.com` to mount `/mnt/rhscd` and `/mnt/cdrom` read-only. After editing `/etc/exports` you will have to kill and restart the `rpc.nfsd` and `rpc.mountd` daemons so they will load the new configuration. This can be done by issuing:

```
/etc/rc.d/init.d/nfs stop
/etc/rc.d/init.d/nfs start
```

See the exports(5), rpc.nfsd(8) and rpc.mountd(8) man pages for complete NFS configuration details.

2.4.6 `rc.local`

The file `/etc/rc.d/rc.local` is executed at boot time, after all other initialization is complete. You can add additional initialization commands here. For instance, you may want to start up additional daemons, or initialize a printer or some serial ports.

Chapter 3

Installation Systems

This chapter is divided into two sections, the RPP system, and the Linux Installation Manager or lim. The RPP system is the foundation of RHC Linux. It allows you to install software packages (binary or source), track and verify the packages, and uninstall them. The lim is a graphical interface to the RPP system. In addition to being an easy point and click way of using the RPP system, the lim performs many safety and consistency checks on your requests before it executes them. You will learn to appreciate these features the first time you have to recover from an `rpp-uninstall bash`!

In case you were wondering, RPP stands for Red Hat Program Package (yeah, I know I left out the 'H' `:-)`. The words "RPP", "rpp", "package" and "Red Hat Program Package" are used interchangeably, and they all refer to the same thing.

NOTE: You may see mention of "relocatability", "prefix", "installation prefix", and "update packages". These features can and should be ignored as they will be removed in the next release of RHC Linux. All packages contain a default prefix, and there is no need to override it.

IMPORTANT: Do *not* execute two or more rpp programs at the same time! Do not execute any rpp programs while lim is running! Doing so can corrupt the rpp record files. A future version of the rpp system will incorporate locking to avoid this.

3.1 The RPP System

You may find that you never use the RPP system directly – the lim interface is more attractive and has a much lower learning curve – so it will only be discussed briefly. The RPP system is a collection of programs that operate on packages, with which you can install, uninstall, track and verify individual source and binary packages. The four most often used are `rpp-install`, `rpp-uninstall`, `rpp-query` and `rpp-verify`. `rpp-build` is used to actually create RPP packages.

Should you venture outside the lim into the underlying RPP system, please read the man pages for the rpp programs and rpp(5) for all the gory details.

3.1.1 What is a Red Hat Program Package?

Each rpp contains an *archive* and a bunch of other information. The archive is simply the collection of files the package contains. The other information includes: name, version, type (source or binary), size (and size to build for source packages), dependencies, date built and builder, and the description.

The file name convention for rpps is `name-version-{bin,src}.rpp`. This is only a useful convention, the file names can be whatever you like, but it does go a long way towards making the file names unique.

3.1.2 The RPP Record

When a package is installed, all information mentioned above is stored in the RPP *record file*. Additional information is also stored, including the install date, and the installer. The record file is stored in the RPP *record directory*, which is `/var/adm/rpp` by default and can be changed by setting the environment variable `RPP_RECORD_DIR`. If you do use some other directory, it must contain the following subdirectories: `verify`, `before-uninstall` and `after-install`. The rpp record directory contains other critical elements of the system, and you should be very careful messing around in it. Recovering from a corrupted record file is *very* difficult.

3.1.3 rpp-install

`rpp-install` installs an rpp. It is used thusly:

```
rpp-install [options] <rpp-file-name>
```

For example:

```
# cd /mnt/cdrom/rpps
# rpp-install findutils-4.1-bin.rpp
```

installs the package findutils-4.1-bin.rpp.

3.1.4 rpp-install-series

In the `rpps` directory of the CD-ROM there are a number of *series files* (named `*.ser`) which define series of packages. To install all the packages in a series use `rpp-install-series`:

```
# cd /mnt/cdrom/rpps
# rpp-install-series 60math.ser
```

`rpp-install-series` is a small `perl` script that simply parses the series file and executes `rpp-install` for each package listed.

3.1.5 rpp-uninstall

Yes, `rpp-uninstall` uninstalls installed packages. Use like so:

```
rpp-uninstall [options] name [version [type [prefix]]]
```

You must specify enough of name, vesion, type and prefix to uniquely specify a single installed package. `rpp-uninstall` will only uninstall a single package at a time. type is either `bin`, `src` or `-`, and any of name, version or prefix can be `-`. In all cases `-` acts as a wildcard. To uninstall the findutils package we just installed above, use:

```
# rpp-uninstall findutils
```

If there are other findutils packages installed (a source package, for instance), you will get a "Multiple matches" error, and nothing will be uninstalled. You will have to specify the package further:

```
# rpp-uninstall findutils - bin
```

3.1.6 rpp-query

`rpp-query` is used to display all kinds of information about packages (either installed or a package file), including the contents of the archive.

```
rpp-query [options] [name [version [type [prefix]]]]
```

where name, version, type and prefix are as in rpp-uninstall, except that they can match more than one package. Two important options are `-c`, which shows all information except the archive list, and `-L` which shows the archive list. If we hadn't uninstalled that findutils package we could get the list of its files with:

```
$ rpp-query -L findutils
/var/adm/locate
/usr/man/man5/locatedb.5
/usr/man/man1/xargs.1
/usr/man/man1/updatedb.1
/usr/man/man1/locate.1
/usr/man/man1/find.1
/usr/info/find.info-2
/usr/info/find.info-1
/usr/info/find.info
/usr/bin/xargs
/usr/bin/updatedb
/usr/bin/locate
/usr/bin/frcode
```

```
/usr/bin/find
/usr/bin/code
/usr/bin/bigram
/var/adm/rpp/verify/findutils-4.1-bin-0.vfy
```

That command would, in fact, list all files for all packages with the name “findutils”. Again, to get a single package, you may have to be more specific on the command line.

Here is the `-c` option in use:

```
$ rpp-query -c findutils
PACKAGE FORMAT VERSION 1
package name  : findutils
version       : 4.1
type          : binary
size          : 245375
prefix        : /
relocatability: not relocatable
dependencies  :
packaged by   : root@devel.redhat.com
packaged on   : Thu Jan 12 11:54:04 1995
installed by  : root@devel.redhat.com
installed by  : Thu Jan 12 11:54:48 1995
install prefix: /
description   :
This package includes find, xargs, locate, updatedb,
frcode, code, and bigram.  It includes man pages
and info files.

The locate database is stored in /var/adm/locate.
```

Another very important option is `-f`, which will tell you what package “owns” a particular file. The `-f` option can be combined with most other options. The usage is slightly different from above, but an example should suffice.

```
$ rpp-query -hf /usr/bin/bigram
PACKAGE FORMAT VERSION 1
package name  : findutils
version       : 4.1
type          : binary
size          : 245375
prefix        : /
relocatability: not relocatable
dependencies  :
packaged by   : root@devel.redhat.com
```

```
packaged on   : Thu Jan 12 11:54:04 1995
installed by  : root@devel.redhat.com
installed by  : Thu Jan 12 11:54:48 1995
install prefix: /
```

3.1.7 rpp-verify

With rpp-verify you can check the permissions, owner, group and other properties of one or more installed packages. It is used like rpp-query:

```
rpp-verify [options] [name [version [type [prefix]]]]
```

where name, version, type and prefix are as in rpp-query. For example to check all binary packages, use:

```
# rpp-verify - - bin
```

Any anomalies with respect to the info specified in the original package will be reported. If you specify the `-f` option the permissions, owner and group will be set to those originally specified in the package.

The `-S` option verifies the BSD checksum of each package file. To find out what files (those "owned" by packages) on your system have been changed:

```
# rpp-verify -aS
...
sendmail 8.6.9 bin /
/var/log/sendmail.st: checksum different
/etc/sendmail.cf: checksum different
...
```

This shows that two configuration files in the `sendmail` package have been altered. This is probably OK, but if you found that `/bin/su` or some other security sensitive file were changed, you may want to investigate.

3.1.8 rpp-build

rpp-build is used to create source and binary packages. There are a fair number of details involved in preparing a package before running rpp-build. Please see rpp-build(8) and the BUILDING AN RPP section of rpp(5) for details.

3.2 The Linux Installation Manager

As already stated, the Linux Installation Manager is a graphical interface to the RPP system. Using the lim you can install and uninstall packages, one at a time or hundreds at a time. You can verify packages, list their files and view the package information. In addition, the lim will warn you about various bad moves that the rpp programs would let you do without warning.

3.2.1 Quick Start

lim is quite intuitive and easy to use. First, you will need to be root to do anything usefuly with lim. Then start X Windows using `startx` (if you haven't already) and type `lim` in an xterm. You can mess around all you want, and unless you click on Engage or select Verify & Correct from the Verify menu, nothing will be installed, uninstalled or otherwise changed. If you want, turn on the balloon help from the Help menu and you will get a little floating window of help as you move the pointer around the screen.

3.2.2 The Main Window

When you start up lim the main window appears (see Window 3.1) and it starts reading in installation information and searching for available packages. This may take a few seconds, especially on slower machines; you can be sure things are proceeding normally by watching the status bar at the bottom of the window.

The main window is divided into three panels, each of which contains a listbox and some buttons. The first panel is for *available packages*. It lists information for all the packages lim was able to find. The second panel is for *installed packages*. It lists information for all packages you have installed with lim and/or rpp-install. The third panel is for *action packages*. It lists all the actions that will happen when you click on the Engage button (more on that later).

Note that the title for each listbox shows the number of packages listed in the listbox, which is not necessarily the number of packages that the lim knows about. For instance, if you are only displaying source packages, the labels will tell you how many source packages were found. See section 3.2.7 for details on choosing which packages to display.

In all three listboxes you select items or ranges of items by clicking and dragging with the first mouse button. Additional bindings are listed in figure 3.1. Note that button 2 will pop up the package description (see section 3.2.3).

Mouse Action	Result
Shift-B1	extend nearest end of selection to pointer
Control-B1	select individual item retaining selection
Control-Shift-B1	select all items
B3	unselect item
Control-B3	toggle selection of item
Control-Shift-B3	deselect all items
B2	display description for item

Figure 3.1: Selecting Items in the Listboxes

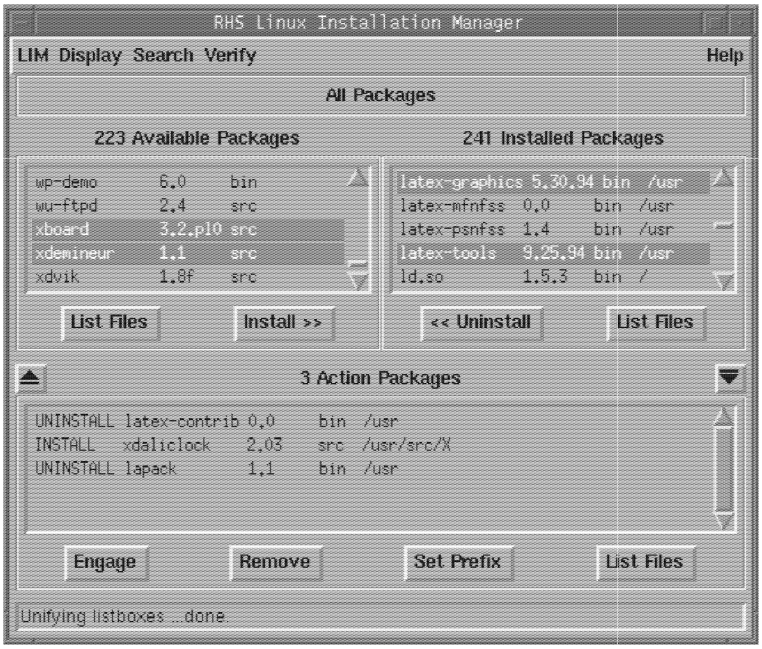

Window 3.1: lim User Interface

The two small arrow buttons move the "boundary" between the available and installed listboxes and the action listbox.

The Available and Installed Packages Panels

The available and installed panels function similarly. The List Files buttons will pop up a transcript window and list all the files in the selected packages. The Install and Uninstall buttons will move the selected packages to the action packages panel for installation and uninstallation, repectively. You can also move individual packages to the action list by double clicking on them.

The Action Packages Panel

The four buttons in the action packages panel are equally intuitive. The List Files button lists files in the selected packages. The Remove button removes the selected packages from the list. The Set Prefix button pops up a dialog box where you can specify the installation prefix for the selected packages. And finally, the Engage button causes the packages to actually be installed and/or uninstalled. Note that the Engage button is the only button that operates on all the packages in the list, not just those in the selection.

3.2.3 The Description

The description window (Window 3.2) displays all the information contained in a package. Clicking

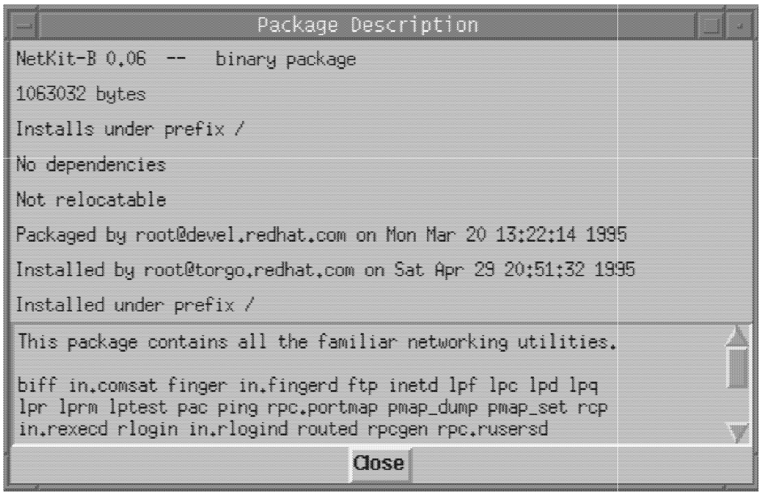

Window 3.2: lim Description Window

mouse button 2 over any list item will bring up the description window for that package. In addition, you can configure lim to always have the description window visible and to have the description track the pointer as it moves around the main window (see section 3.2.8 for details).

3.2.4 Warnings from Engage

When you click on Engage, lim performs a number of sanity checks before installing and/or uninstalling any packages. Most of these are warnings which you can choose to ignore, but you do have the option of cancelling the whole thing at any time.

Duplicate Actions lim checks to see if you are installing or uninstalling something twice.

Already Installed lim checks to see if you are installing a package that is already installed (possibly different version). This is usually OK for source packages. For binary packages this is probably a bad idea – the filenames of different versions of the same binaries are usually not unique, and you can imagine the possible implications of two packages "owning" the same file.

Critical Packages lim also checks to see if you are uninstalling any *critical packages.* These are packages that are so important that if you actually uninstall them, things can go horribly wrong. Some examples of critical packages are libc-linux-shared and bash. If any critical packages are being uninstalled, lim asks you if you want to do a *fake* uninstall, which removes the installation record, but not the actual files of the package. This is highly recommended. You will probably be uninstalling a critical package only to install a new version, and the new version should install right over the old files.

Disk Usage Finally, lim checks the the total change in disk usage, given the sizes of the packages you are installing and uninstalling and compares this with the free disk space as reported by df(1). You will be warned if it appears that you do not have enough disk space to accomodate the installation.

3.2.5 The Transcript Window

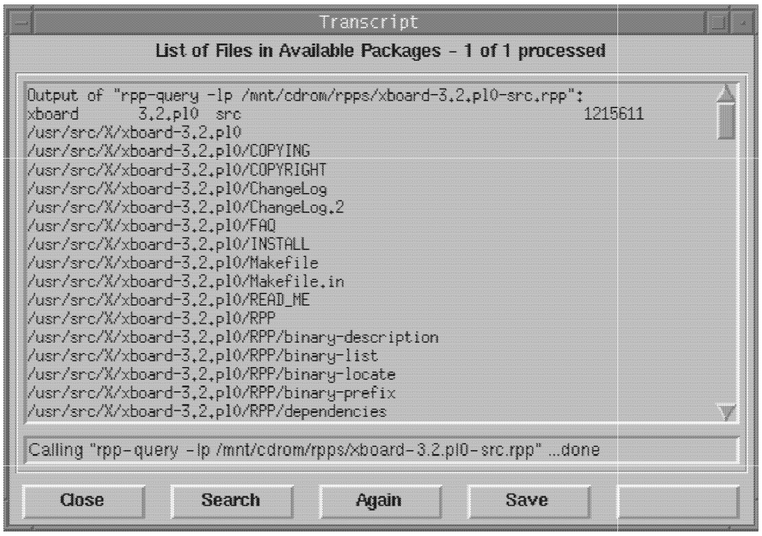

Window 3.3: lim Transcript Window

The transcript window (Window 3.3) is used whenever you click on Engage or List Files, or select something from the Verify menu. It consists of a label, a scrolling transcript area, a status bar and four buttons: Close, Search, Save and Cancel. While the program(s) generating the transcript are running, only the Cancel button is active, and clicking on it informs lim that you want it to stop (sometimes lim must complete what it is currently doing or do some backing out in order to keep things clean – therefore the Cancel button might not have an immediate effect).

Once the programs are done or canceled the other buttons become active. Close will close the transcript window, Search will let you search for a regular expression in the transcript, and Save allows you to save the transcript to a file.

3.2.6 Menus

lim The Set Paths item pops up a dialog box to set various paths (see section 3.2.9). The Options item pops up a dialog box to set various options that control lim operation (see section 3.2.8).

Display The display menu allows you to choose which packages or series of packages you wish to view (see section 3.2.7).

Search The four items in the Search menu let you search through all the information loaded about packages for a certain regular expression (if you don't know what a regular expression is, you can just type plain text). The current selections in the appropriate listboxes will be lost, and each package that matches the regular expression will be selected.

Verify You can verify (and correct) installed packages by selecting one or more packages from the install package listbox and choosing Verify Only or Verify & Correct from the Verify menu. A transcript window will appear and the packages will be verified. See section 3.1.7 for details on what verification and correction entails.

Help The Balloon Help item toggles the operation of balloon help. When balloon help is active, a small window with some help text will float near the pointer as it enters different buttons, listboxes and other widgets. Note that balloon help does not interfere with the normal operation of lim. The other help item brings up a window showing what the mouse buttons and keys do in the listbox windows.

3.2.7 Package Display

You can choose which packages or series of packages you wish to display with the items in the Display menu.

Window 3.4: The Display Menu

The first three items in the Display menu (Window 3.4) operate on all packages the lim knows about. If you select "Binary Packages", for instance, the listboxes will display all binary packages.

The second three items and all the following items are used to Display a *series* of packages. These are simply convenient groups of related packages. For instance the development series may contain the C compiler and debugger, and other useful development tools. You can choose to display the series packages, or the sources to the series packages, or both.

Note that there are two options (see section 3.2.8) that govern which packages are displayed in the Available listbox (in addition to your specifications from the Display menu). You can choose to list only the highest available version of a package, and you can choose to exclude packages which are already installed (this is the default).

3.2.8 Options

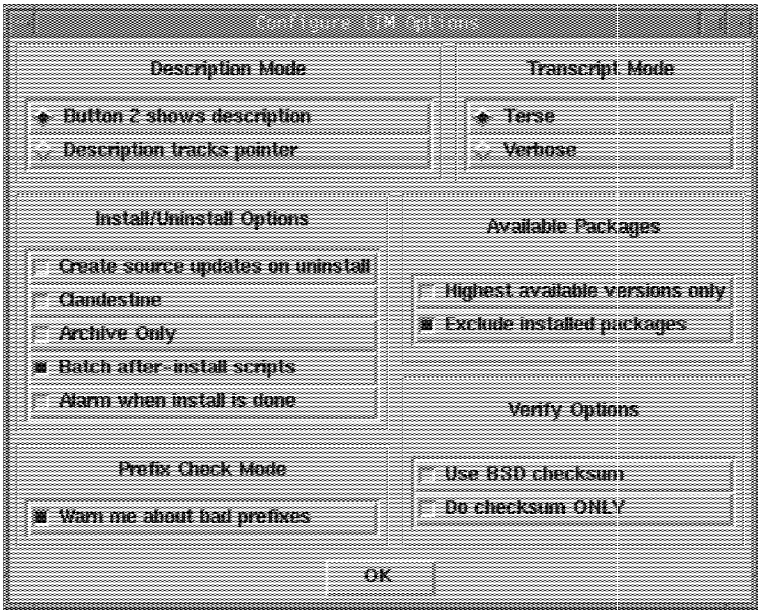

Window 3.5: lim Options Dialog Box

Selecting Options from the lim menu brings up a dialog box (Window 3.5) for configuring various lim options.

Description Mode controls whether the description changes when you press button 2 or the description automatically tracks the pointer.

Transcript Mode is either verbose or terse. Verbose causes lim to invoke the rpp programs with verbose options.

Available Packages Options There are two options for available packages. The first is "highest available versions only". When this option is selected the available packages listbox will only display the highest version available for any given package (with the same name and type). The second is "exclude installed packages". When this option is selected (which it is by default), packages that are already installed, ie are listed in the installed listbox, will not be displayed in the available packages listbox.

Verify There are two options that are used when you verify packages. You can verify the BSD checksum in addition to the ownerships and permissions, or you can check *only* the BSD checksum. The BSD checksum is useful for finding files that have been altered since they were installed.

Prefixes This option controls whether or not lim warns you when you try to use prefixes that don't conform to conventions, or install a non-relocatable package under a prefix other than the default prefix. As the prefix and relocatability features will be removed in the next release of RHC Linux, you shouldn't really bother with this.

Uninstallation You can have lim attempt to generate updates automatically when you uninstall a source package. When this option is enabled, lim will automatically generate update packages for source packages you have modified. As the update package feature will be removed in the next release of RHC Linux, you shouldn't really bother with this either.

Installation You can have installations be clandestine (no record of the installation) or archive only (no control files, just the package archive), which are both usually a bad idea. See rpp-install(1) for details on these two options. You can also batch any after-install scripts. Some packages have scripts that run after installation, and some of those scripts may want to interact with you. If you are installing a lot of packages at once, batching the after-install scripts will cause them to run after all packages have been installed. This way, you don't have to sit and watch all the packages install themselves, responding every now and then to an after-install script.

3.2.9 Paths

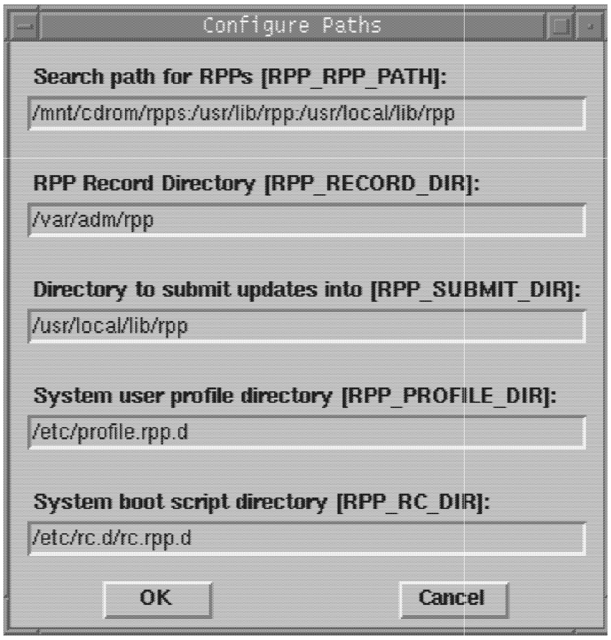

Window 3.6: lim Paths Dialog Box

Selecting Set Paths from the lim menu brings up a dialog box (Window 3.6) to configure the various paths lim uses to find things.

The RPP search path uses the same syntax as your familiar PATH environment variable. lim uses the search path to look for packages and series files. You may wish to set things up such that the search path includes the `rpps` directory on the RHC Linux CD-ROM (a standard CD-ROM installation will already be set up this way).

The RPP record directory is where all the installation records are kept. There is usually no reason to change this.

The update directory is realted to update packages, which will not be supported in the future. Do not use it. The same is true for the user profile directory and boot script directory.

3.2.10 Requirements

The interpreter for lim is rhswish, which is a standard Tk wish program with lots of additions. You need to have rhswish installed. Since rhswish is based on Tcl/Tk, you need to have the Tcl and Tk shared libraries installed, as well as the Xpm shared libraries. And since lim calls rpp programs and verify-install, it requires that those be installed too. If you installed the Base and X Utils series, eveything will already be installed.

Chapter 4

The Red Hat Control Panel

The control-panel is a launching pad for a number of different system administration tools which make your life a lot easier by letting you configure things without remembering config file formats and awkward command line options.

To start the control-panel, start X Windows with `startx` and type `control-panel` in an xterm. You will need to be root to run the control-panel tools successfully.

Double clicking on the icon for each tool starts it up. Please note that you are not prevented from starting two instances of any tool, but doing so is a *very bad idea*. If you do accidentally start a second copy of a tool, you should quit it immediately. Also, do not manually edit any files managed by the control-panel tools while the tools are running. Similarly, do not run any other programs that may change those files while the tools are running.

4.1 User and Group Configuration

The user and group tool (Window 4.1) is used to manage the users and groups on your system. In a sense, it is a graphical editor for `/etc/passwd` and `/etc/group`. With this tool you can add and remove users and groups, set login shells, full names, home directories, uid, gids, etc.

4.1.1 Adding a User

To add a user, click on the Add button in the main window. A dialog box will pop up (Window 4.2) with a template for a user.

First enter the user name. This is *not* the user's First and Last name, it is the id they will use to log on to the system. Do not include any spaces or colons, and do not enter more than 8 characters. Press the enter key when you are finished, and some of the other fields will be filed in with default values. You may change them if you wish, but there is no need to. Next, enter the user's Full Name, and select a login shell using the drop-down menu. And finally you need to do something with the password. There are three options here, Edit, None, or Lock. None is a very bad option – with no

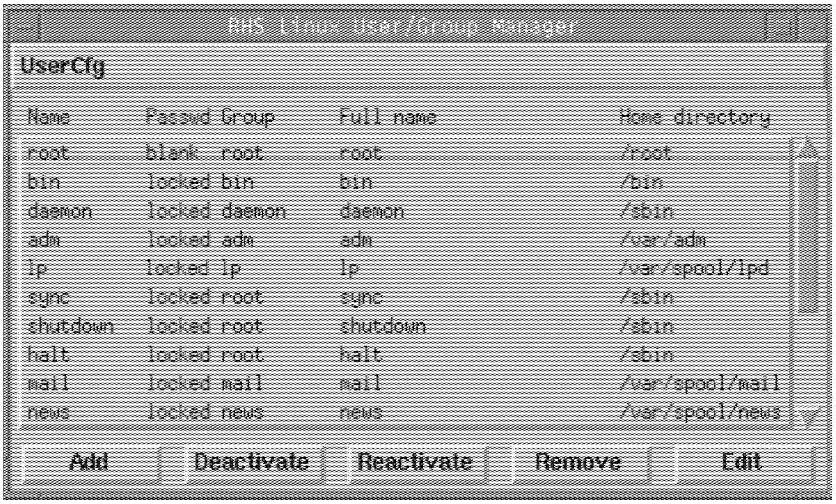

Window 4.1: User and Group Configuration Panel

Window 4.2: Add User Dialog Box

password anyone can login to your system using this id. Choosing Lock will prevent anyone from logging in with this id. Usually you will choose Edit. A small dialog box will appear where you will need to enter a password, and the confirm it by typing it a second time. Click on OK after this, and then click on Add in the user entry dialog box to add the user.

If this user does not already have a home directory (and it won't unless you have already created one), you will be asked if you want to create one. Unless you will be mounting home directories from another machine, you should click on Yes.

4.1.2 Removing a User

To remove a user, select a user in the main window by clicking on it, and click on Remove. After confirming the removal, you will be asked if you wish to remove the user's home directory and mail spool. You will normally answer Yes. After these files are deleted, you will be asked if you want to search the entire file system for files that have been "orphaned" by the removal of this user. Orphaned files are those that are not owned by any user in `/etc/passwd`. Since you have just removed a user from `/etc/passwd` any file that user owned will now be orphaned (except the home directory and mail spool, if you deleted them). If you answer yes, the search will proceed in the background, and all files found will be `chown`-ed to `nobody`. To find these files, execute the following command:

```
find / \( -group nobody -o -user nobody \) -print
```

Deactivating a User

Sometimes you may just want to temporarily "remove" a user, with the intention of reinstating the user later. The Deactivate and Reactivate functions handle this. When you deactivate a user, the user's password is locked by prepending a "`*`" to it, preventing the user from logging in. You also have the option of "collapsing" the user's home directory with `tar` and `gzip`, which can save some space. When you reactivate a user, the password locking is removed, and if the home directory was collapsed, it is expanded back to normal.

4.1.3 Creating a New Group

To create a new group, first switch to group editing mode by selecting Edit groups from the menu. Clicking on Add brings up a dialog box to specify the group details. First enter a name for the new group. Just like a user name, the group name should have no spaces or colons, and should be no longer than 8 characters. To add members to the group click on the Add button on the left. A list of user names will appear from which can select the group members. When you have selected all the group members, click on Add, then click on the lower Add button.

4.2 File System Configuration

The file system configuration tool (Window 4.3) allows you to easily examine and manipulate file system mount points, types, options, etc. It is very useful for manipulating a large number of file systems in a very dynamic way.

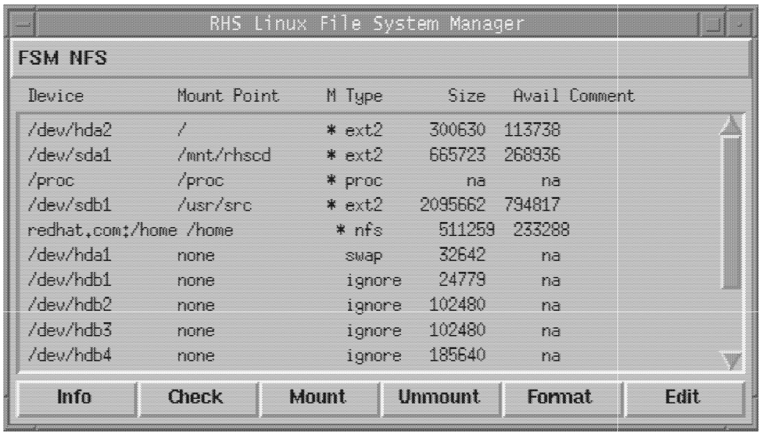

Window 4.3: File System Configuration Panel

4.3 Printer Configuration

The printer configuration tool maintains `/etc/printcap`, print spool directories, and print filters. It is designed to use GhostScript (`gs`), to convert PostScript files to a format understandable by your printer. Thus, you must print PostScript documents. You can use `nenscript` to convert plain text to PostScript.

Keep in mind that `lpr` prints to the `lp` printer (by default) and nenscript prints to the `PostScript` printer (by default). You will probably want `lp` to be one of your printer "names".

Needless to say, ghostscript and nenscript should be installed.

4.4 Network Configuration

The network confiuration panel (Window 4.4) is designed to allow easy manipulation of parameters such as IP address, gateway address, network address, etc, as well as name servers, and `/etc/hosts`. Network devices can be added, removed, configured, brought up, brought down, etc. Ethernet, loopback, SLIP and PPP devices are supported, although SLIP and PPP support should be considered experimental.

If you will be using SLIP or PPP configuration, you should edit `/etc/dip-script` to "chat" properly with your SLIP/PPP server. Do not edit the `@@@xxx@@@` entries – those are configured automatically.

The things to pay attention to are the `wait` lines, which must recognize stuff your SLIP/PPP server says.

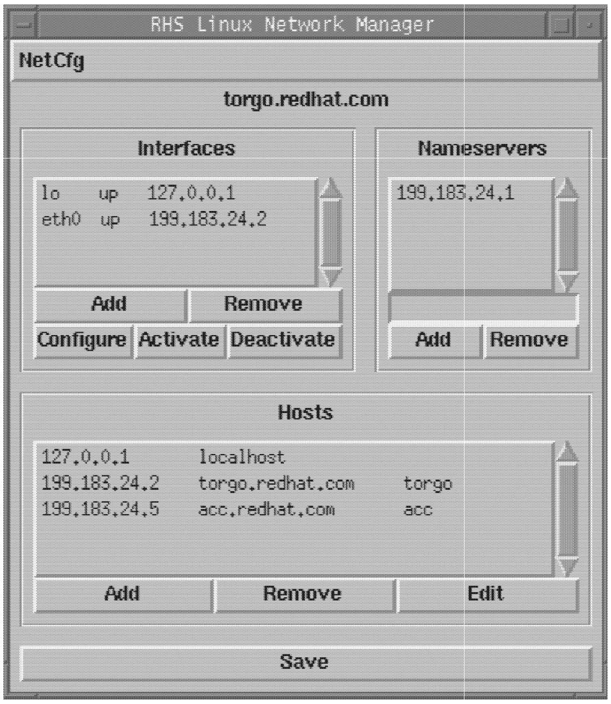

Window 4.4: Network Configuration Panel

4.5 Time and Date

The time machine (Window 4.5) allows you to change the time and date by clicking on the apropriate part of the time and date display and clicking on the arrows to change the value.

The system clock is not changed until you click on the Set System Clock button.

Click on Reset Time to set the time machine time back to that of the system.

IMPORTANT: Changing the time can seriously stress your system, and could possibly cause problems. Try to quit as many applications and processes as possible before changing the time or date.

4.6 Searching Help Systems

The help tool allows you to search man pages, info pages, and plain text doc files (or any subset of these) for a regular expression. Search hits are listed and you can double click on them to view them. This tool is very basic, and should see drastic improvement over time.

The `HOWTO` rpp contains many, many useful documents. It installs help files under `/usr/doc`. If you

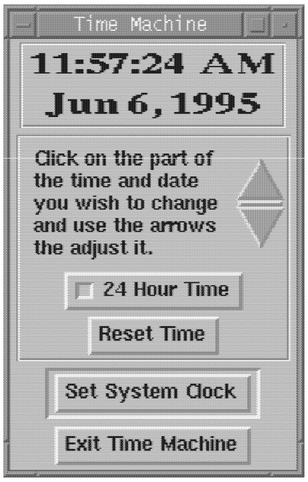

Window 4.5: Time and date Panel

have a question, it is probably answered in one of those documents. Additional help and information is availble using the `man` and `info` programs.

Chapter 5

System Administration

This chapter is an overview of the RHC Linux system. I'll illustrate things that you may not know about the system and things that are somewhat different from other systems. Note that most system administration task are performed with the control-panel, which is covered in chapter 4.

Please see chapter 4 of *Linux Installation and Getting Started* more details on basic Linux system administration.

5.1 Filesystem Structure

Red Hat Software is commited to the Linux File System Standard, and at the moment RHC Linux is "partially compliant". As of this writing the most recent FSSTND document was interim release 1.1, and a new version is expected very soon. We will continue to track the standard and bring RHC Linux towards full compliance.

While compliance with the standard means many things, the two most important are compatibility with other compliant systems, and the ability to mount the /usr partition read-only. The latter fact means the /usr partition can be mounted from a CD-ROM or from another machine via read-only NFS. Check the latest FSSTND document for the most up to date and complete definition of the standard.

RHC Linux strays from the FSSTND in the following areas:

`sed` is in `/usr/bin`, `/etc/magic` is moved to `/usr/lib/magic`. There is a symbolic link from `/etc/utmp` to `/var/adm/utmp`, for compatibility. RHC Linux has no `/usr/share` or `/usr/local/share`. The FSSTND `/usr/man` hierarchy is not implemented. TeX data is placed in `/var/tex` instead of `/var/lib/texmf`.

5.2 Users, Groups and User-Private Groups

Managing user and groups has traditionally been tedious. RHC Linux has a few tools and conventions that make user and groups easier to manage, and more useful.

5.2.1 Standard Users

Table 5.2.1 lists the stadard users set up by the installation process (this is essentially the /etc/passwd file). The group id (GID) in this table is the *primary group* for the user. See section 5.2.3 for details on how groups are used.

User	UID	GID	GECOS	HOME Directory	Shell
root	0	0	root	/root	/bin/bash
bin	1	1	bin	/bin	
daemon	2	2	daemon	/sbin	
adm	3	4	adm	/var/adm	
lp	4	7	lp	/var/spool/lpd	
sync	5	0	sync	/sbin	/bin/sync
shutdown	6	0	shutdown	/sbin	/sbin/shutdown
halt	7	0	halt	/sbin	/sbin/halt
mail	8	12	mail	/var/spool/mail	
news	9	13	news	/var/spool/news	
uucp	10	14	uucp	/var/spool/uucp	
operator	11	0	operator	/root	/bin/bash
games	12	100	games	/usr/local/games	
gopher	13	30	gopher	/usr/lib/gopher-data	
ftp	14	50	FTP User	/usr/rhs/ftp	
nobody	99	99	nobody	/root	

Figure 5.1: Standard Users

5.2.2 Standard Groups

Table 5.2.2 lists the standard groups as set up by the installation process (this is essentially the /etc/group file). See section 5.2.3 for details on how groups are used.

5.2.3 User Private Groups

RHC Linux uses a user private group (UPG) scheme, which makes UNIX groups much easier to use. The UPG scheme can be summarized with three points:

User Private Group Each user has their own primary group, of which only they are a member.

Group	Group ID	Members
root	0	root
bin	1	root,bin,daemon
daemon	2	root,bin,daemon
sys	3	root,bin,adm
adm	4	root,adm,daemon
tty	5	
disk	6	root
lp	7	daemon,lp
mem	8	
kmem	9	
wheel	10	root
mail	12	mail
news	13	news
uucp	14	uucp
man	15	
games	20	
gopher	30	
dip	40	
ftp	50	
nobody	99	
users	100	

Figure 5.2: Standard Groups

umask = 002 The traditional UNIX umask is 022, which prevents other users *and other members of a user's primary group* from modifying a user's files. Since every user has their own private group in the UPG scheme, this "group protection" is not needed. A umask of 002 will prevent users from modifying other users' private files.

SGID bit on Directories By setting the SGID bit on a directory (with `chmod g+s` *directory*), files created in that directory will have their group set to the directory's group.

The UPG scheme makes managing group projects very simple. Let's say you have a big project called FOO, with many people editing the FOO files in a FOO directory. Make a group called `foo`, `chgrp -R` the FOO directory to `foo`, and add the all the FOO users to the `foo` group. Now, all the FOO users will be able to edit the FOO files and create new files in the FOO directory, and these files will always retain their `foo` group, and thus they will always be editable by other FOO users.

If you have multiple projects like FOO, and users who are working on multiple projects, these users will never have to change their umask or group when they move from project to project. The SGID bit on each project's main directory "selects" the proper group.

Since each user's HOME directory is owned by the user and their private group, it is safe to set the SGID bit on the HOME directory. However, by default, files are created with the primary group of the user, so there is no real effect.

The easiest way to manage users and groups is through the Users and Groups module of the control-panel (see section 4 for details on the control-panel and section 4.1 for details on the Users and Groups module).

You can also use `adduser` to create a new user from the command line.

5.3 The Boot Process, Init, and Shutdown

RHC Linux uses SysV style init scripts (scripts which run at boot time), which live in /etc/rc.d. These scripts take care of initializing networking, nfs, cron, syslog, sendmail, lpd, etc. You should *not* mess around with these files unless you know what you are doing – you could easily prevent your system from booting.

The following runlevels are used in RHC Linux:

- **0** Halt.
- **1** Single user mode.
- **2** Multiuser mode, without NFS.
- **3** Full multiuser mode.
- **6** Reboot.

If your machine gets into a state where it will not boot due to a bad /etc/inittab, or will not let you login because you have screwed up /etc/passwd, or simply forgotten your password, boot into single user mode by typing `linux 1` at the LILO boot prompt. The system will come up and you will be given a shell from which you can fix things.

To shutdown RHC Linux, issue the `shutdown` command. If you want to turn off your computer, you should shutdown and *halt*:

```
shutdown -h now
```

If you want to shutdown and reboot use:

```
shutdown -r now
```

Do not run `reboot` *or* `halt` *directly!*. Those commands will not bring the system down gracefully, and you may cause file system corruption by using them.

5.4 RHC Linux Configuration Files

In addition to the files pertaining to the RPP system that reside in /var/adm/rpp, there are two other special locations that are reserved for RHC Linux configuration and operation.

The control-panel and related tools put lots of stuff in /usr/lib/rhs. There is probably nothing here that you would want to edit, it is mostly small scripts, bitmaps and text files.

The other location, /etc/sysconfig, is used to store configuration information. The major users of the files in this directory are the scripts that run at boot time. It is possible to edit these by hand, but it would be better to use the proper control-panel tool.

Appendix A

Red Hat Software Contact Information

Here are various contact points at Red Hat Software:

1. **General info** – `info@redhat.com` (finger or mail (auto-reply))
2. **Technical Support** – `support@redhat.com`
3. **Other** – `redhat@redhat.com`
4. **Anonymous FTP** – `ftp.redhat.com`
5. **World Wide Web** – `www.redhat.com`
6. **Mailing List** – See `info@redhat.com` for information

Red Hat Software
PO Box 4325
Chapel Hill, NC 27515