

# Red Hat Linux 5.0

---

The Official Red Hat Linux Installation Guide

Red Hat Software, Inc.  
Research Triangle Park, North Carolina

Copyright © 1995, 1996, 1997 Red Hat Software, Inc.

Portions Copyright © Metro Link Incorporated 1990-1996

Red Hat, the Red Hat Shadow Man logo, RPM, the RPM logo, and Glint are trademarks of Red Hat Software, Inc.

Metro-X is a trademark of Metro Link, Incorporated.

Motif and UNIX are registered trademarks of The Open Group.

Alpha is a trademark of Digital Equipment Corporation.

SPARC is a registered trademark of SPARC International, Inc. Products bearing the SPARC trademarks are based on an architecture developed by Sun Microsystems, Inc.

Netscape is a registered trademark of Netscape Communications Corporation in the United States and other countries.

All other trademarks and copyrights referred to are the property of their respective owners.

ISBN: 1-888172-97-5

Revision 5.0, October 1997.

Red Hat Software, Inc.  
4201 Research Commons, Suite 100  
79 T. W. Alexander Drive  
P. O. Box 13588  
Research Triangle Park, NC 27709  
(919) 547-0012  
redhat@redhat.com  
<http://www.redhat.com>

While every precaution has been taken in the preparation of this book, the publisher assumes no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

The *Official Red Hat Linux 5.0 Installation Guide* may be reproduced and distributed in whole or in part, in any medium, physical or electronic, so long as this copyright notice remains intact and unchanged on all copies. Commercial redistribution is permitted and encouraged, but you may not redistribute it, in whole or in part, under terms more restrictive than those under which you received it.

# Contents

<b>1</b>	<b>New Features Of Red Hat Linux 5.0</b>	<b>1</b>
1.1	Installation-Related Enhancements . . . . .	1
1.2	New User-Mode Tools . . . . .	3
1.3	System Administration-Related Enhancements . . . . .	6
1.4	Miscellaneous New Features . . . . .	12
<b>2</b>	<b>Before You Begin</b>	<b>15</b>
2.1	Getting the Right Red Hat Linux Components . . . . .	16
2.2	Things You Should Know . . . . .	19
2.3	Installation Methods . . . . .	21
2.4	Disk Partitions . . . . .	25
2.5	Getting Documentation . . . . .	31
2.6	A Note About Kernel Drivers . . . . .	32
2.7	One Last Note... . . . .	33
<b>3</b>	<b>Starting the Installation</b>	<b>35</b>
3.1	The Installation Program User Interface . . . . .	35
3.2	Starting the Installation Program . . . . .	37

---

3.3	Beginning the Installation . . . . .	39
3.4	Selecting an Installation Method . . . . .	41
<b>4</b>	<b>Continuing the Installation</b>	<b>51</b>
4.1	Upgrading or Installing . . . . .	51
4.2	SCSI Support . . . . .	53
4.3	Creating Partitions for Red Hat Linux . . . . .	53
4.4	Initializing Swap Space . . . . .	65
4.5	Formatting Partitions . . . . .	66
4.6	Selecting and Installing Packages . . . . .	66
<b>5</b>	<b>Finishing the Installation</b>	<b>73</b>
5.1	Configuring a Mouse . . . . .	73
5.2	Configuring X Windows . . . . .	73
5.3	Configuring Networking . . . . .	75
5.4	Configuring the Clock . . . . .	76
5.5	Selecting Services for Start on Reboot . . . . .	76
5.6	Configuring a Printer . . . . .	77
5.7	Setting a Root Password . . . . .	81
5.8	Installing LILO . . . . .	82
5.9	Finishing Up... . . . .	84
<b>6</b>	<b>Post Installation Configuration</b>	<b>89</b>
6.1	Building a Custom Kernel . . . . .	89
6.2	X Window Configuration with XFree86 . . . . .	94
6.3	X Window Configuration with Metro-X . . . . .	96
6.4	Mouse configuration . . . . .	106

---

6.5	Sendmail . . . . .	107
6.6	Controlling Access to Services . . . . .	108
6.7	Anonymous FTP . . . . .	109
6.8	World Wide Web . . . . .	109
6.9	NFS Configuration . . . . .	110
6.10	Running Programs at Boot Time . . . . .	112
6.11	Color-ls Setup . . . . .	112
<b>7</b>	<b>Finding Documentation</b>	<b>113</b>
7.1	On Line Help . . . . .	113
7.2	Help from the Internet Community . . . . .	118
<b>8</b>	<b>Package Management with RPM</b>	<b>121</b>
8.1	RPM Design Goals . . . . .	122
8.2	Using RPM . . . . .	123
8.3	Impressing Your Friends with RPM . . . . .	128
8.4	Other RPM Resources . . . . .	130
<b>9</b>	<b>Control Panel</b>	<b>131</b>
9.1	User and Group Configuration . . . . .	132
9.2	File System Configuration . . . . .	136
9.3	Printer Configuration . . . . .	137
9.4	Network Configuration . . . . .	142
9.5	Time and Date . . . . .	147
9.6	Kernel Daemon Configuration . . . . .	148
<b>10</b>	<b>Glint</b>	<b>151</b>

---

10.1 Starting glint . . . . .	152
10.2 The Package Display . . . . .	152
10.3 Configuration . . . . .	155
10.4 Package Manipulation . . . . .	156
<b>11 System Administration</b>	<b>163</b>
11.1 Filesystem Structure . . . . .	163
11.2 Special Red Hat File Locations . . . . .	168
11.3 Users, Groups and User-Private Groups . . . . .	168
11.4 User Authentication with PAM . . . . .	173
11.5 The Boot Process, Init, and Shutdown . . . . .	176
11.6 Rescue Modes . . . . .	179
<b>A Getting Technical Support</b>	<b>183</b>
A.1 An Overview of Our Support Policy . . . . .	183
A.2 Getting Support . . . . .	186
A.3 Support FAQ (Frequently Asked Questions) . . . . .	189
<b>B Making Installation Diskettes</b>	<b>193</b>
B.1 Making a Diskette Under MS-DOS . . . . .	194
B.2 Making a Diskette Under a Linux-like O/S . . . . .	195
<b>C Information Specific to Red Hat Linux/SPARC</b>	<b>197</b>
C.1 Supported Hardware . . . . .	197
C.2 Installation Overview . . . . .	199
C.3 Console Commands . . . . .	199
C.4 Ramdisk-based Installation Criteria . . . . .	200

---

C.5	Choosing a Boot Method . . . . .	201
C.6	Choosing an Installation Method . . . . .	205
C.7	Installation Using a Serial Terminal . . . . .	207
C.8	Installation: Selecting System Components . . . . .	207
C.9	SILO Configuration . . . . .	208
C.10	Partitioning . . . . .	208
C.11	X Windows . . . . .	209
<b>D</b>	<b>Information specific to Red Hat Linux/Alpha</b>	<b>211</b>
D.1	Introduction . . . . .	211
D.2	Supported Hardware . . . . .	211
D.3	Installation Overview . . . . .	213
D.4	Choosing Diskette Images . . . . .	214
D.5	Installing . . . . .	215
<b>E</b>	<b>Alpha Miniloader Howto</b>	<b>221</b>
E.1	Introduction . . . . .	221
E.2	What is MILO? . . . . .	222
E.3	Pre-Built Standard MILO Images. . . . .	224
E.4	How To Build MILO . . . . .	225
E.5	How To Load MILO . . . . .	226
E.6	MILO's User Interface . . . . .	238
E.7	Running the Flash Management Utility . . . . .	242
E.8	Restrictions. . . . .	245
E.9	Problem Solving. . . . .	245
E.10	Acknowledgements. . . . .	247

---

<b>F</b>	<b>General Parameters</b>	<b>249</b>
F.1	CD-ROM parameters . . . . .	250
F.2	SCSI parameters . . . . .	251
F.3	Ethernet parameters . . . . .	251
<b>G</b>	<b>Metro X</b>	<b>253</b>
G.1	Special Features . . . . .	253
G.2	Metro X Configuration File . . . . .	255
<b>H</b>	<b>Red Hat Linux Frequently Asked Questions</b>	<b>265</b>
H.1	Introduction . . . . .	265
H.2	General Questions . . . . .	265
H.3	Installation Questions . . . . .	272
H.4	Using Red Hat Linux . . . . .	276
H.5	Copyright Notice . . . . .	276
<b>I</b>	<b>Red Hat and Freely Distributable Software.</b>	<b>277</b>
I.1	The Free Software Model . . . . .	278
I.2	The Evolution of “Free Software” . . . . .	278
I.3	The GNU Public License . . . . .	279
I.4	Will Linux Replace Commercial OSes? . . . . .	280
I.5	Red Hat Software, Inc.’s Role in Linux Development . . . . .	281
I.6	Further Reading . . . . .	281
<b>J</b>	<b>Licenses</b>	<b>283</b>
J.1	The BSD Copyright . . . . .	283
J.2	X Copyright . . . . .	284

J.3 Gnu Public License . . . . . 285



# Preface

Welcome! And thanks for your interest in Red Hat Linux. We have what we think is the best Linux distribution on the market today, and we work hard to keep it that way. Red Hat Linux 5.0 is the latest in a long line of software from Red Hat Software. We hope you like it, and that you enjoy using Red Hat Linux as much as we've enjoyed making it for you.

It's interesting to note that, while Linux is popular and well-known by a certain segment of the computer-using population, there are many people out there that are only now hearing about Linux. For this group of people, the following section should provide enough background to help you get acquainted with Linux and Red Hat Software. If you'd like to learn more about how Linux is developed and distributed, please see Appendix I on page 277 for more information.

## What is Linux?

Back in August of 1991, a student from Finland began a post to the comp.os.minix newsgroup with the words:

```
Hello everybody out there using minix -  
I'm doing a (free) operating system (just a hobby,  
won't be big and professional like gnu) for  
386(486) AT clones.
```

The student was Linus Torvalds, and the "hobby" he spoke of eventually became what we know today as Linux.

A full-featured POSIX-compliant operating system, Linux has been developed not just by Linus, but by hundreds of programmers around the world. The interesting thing about this is that this massive, world-wide development effort is largely uncoordinated. Sure, Linus calls the shots where the kernel is concerned, but Linux is more than just the kernel. There's no management infrastructure; A student in Russia gets a new motherboard, and writes a driver to support a neat feature the motherboard has. A system administrator in Maryland needs backup software, writes it, and gives it away to anyone that needs it. The right things just seem to happen at the right time.

Another interesting thing is that Linux can be obtained for absolutely no money. That's right, most of the software is available (at no charge) to anyone with the time and inclination to download it. But not everyone has that much time...

### What is Red Hat Linux?

Enter a group of programmers based in North Carolina. Their goal was to make it easier for people to give Linux a try. Like many other such groups, their approach was to bundle all the necessary bits and pieces into a cohesive *distribution*, relieving "newbies" from some of the more esoteric aspects of bootstrapping a new operating system on their PCs.

However, unlike other distributions, this one was fundamentally different. The difference? Instead of being a snapshot of a hard disk that had a working copy of Linux on it, or a set of diskettes from which different parts of the operating system could be dumped, this distribution was based on *packages*.

Each package provided a different piece of software, fully tested, configured, and ready to run. Want to try a new editor? Download the package and install it. In seconds, you can give it a try. Don't like it? Issue a single command, and the package is removed.

If that was all there was to it, this distribution would be pretty nifty. But being package-based meant there was one additional advantage:

This Linux distribution could be easily upgraded.

Software development in the Linux world is fast-paced, so new versions of old software come out continually. With other distributions, upgrading software was painful – a complete upgrade usually meant deleting everything on your hard drive and starting over.

By now you've probably guessed that the group of programmers in North Carolina is Red Hat Software, and the package-based distribution is Red Hat Linux.

Since Red Hat Linux's introduction in the summer of 1994, Linux and Red Hat Software have grown by leaps and bounds. Much has changed; support for more esoteric hardware, huge increases in reliability, and the growing use of Linux by companies around the world.

But much still remains the same. Linux is still developed by people world-wide; Linus is still involved. Red Hat Software is still located in North Carolina; still trying to make Linux easier for people to use.

And Red Hat Linux is still package-based; always has been, always will be.

Since the release of version 4.0, Red Hat Linux runs on three leading computing platforms: Intel compatibles PCs, Digital Alpha computers, and Sun SPARC equipment. Our unified source tree and the benefits of RPM (Red Hat Package Management) technology enable us to deploy Red Hat Linux for each platform with a minimum of effort. This in turn enables our users to manage and port software between these platforms as easily as possible.

We make Red Hat Linux available by unrestricted FTP from our site and many mirror sites on the Internet. Red Hat Linux is also available on CD-ROM. For current information on our product offerings and links to other Linux resources please check Red Hat Software's web site at <http://www.redhat.com>.

On most systems, Red Hat Linux is easy to install; the installation program can walk you through the process in as little as 15 minutes. The system itself is very flexible. With RPM, you can install and uninstall individual software packages with minimal effort. Because of RPM, Red Hat Linux is also easy to maintain – package installations can be verified and corrected, and packages can be installed and uninstalled simply and reliably. Furthermore, Red Hat Linux is easy to administer. Included are a rich set of administrative tools reduce the hassle of everyday system administration. Complete source code is provided for the freely distributable components of the system.

## **An Overview of This Manual**

This manual is organized to guide you through the process of installing Red Hat Linux quickly and easily. Toward that goal, let's take a quick look at each chapter to help you get acclimated:

**Chapter 1, *New Features Of Red Hat Linux 5.0*** contains information concerning new functionality that has been added to Red Hat Linux 5.0.

**Chapter 2, *Before You Begin*** contains information on tasks you should perform prior to starting the Red Hat Linux installation.

**Chapter 3, *Starting the Installation*** contains detailed instructions for starting the Red Hat Linux installation process.

**Chapter 4, *Continuing the Installation*** contains instructions on the main part of the installation process.

**Chapter 5, *Finishing the Installation*** contains instructions on the last steps required to complete the installation process.

**Chapters 6 – 11** explain how to configure various services once you have installed Red Hat Linux, how to find documentation on your system, and how to use the various system management and administration tools which accompany Red Hat Linux. They also include an explanation of what's special about your Red Hat Linux system, including where special files live and how to shut down Red Hat Linux.

**Appendixes** contain extra information about Red Hat Linux, including information for owners of Digital Alpha and Sun SPARC machines, frequently asked questions, etc.

## Quick Start Information

Those of you that have installed Red Hat Linux/Intel before and are in a hurry to get started need only boot from a boot diskette, and select the desired installation method. If you are installing from an FTP site, a hard disk, or you'll be using a PCMCIA card during the installation, you'll be prompted to insert a supplemental diskette. In either case, answer all questions as they are presented.

If you are attempting to install Red Hat Linux for either the Alpha or the SPARC, you really should read Chapter 2 on page 15. It will refer you to information specific to your non-Intel-based system.

---

## Upgrading from a Prior Version of Red Hat Linux

The installation process for Red Hat Linux 5.0 includes the ability to upgrade from prior versions of Red Hat Linux (2.0, 2.1, 3.0.3, 4.0, 4.1, and 4.2) which are based on RPM technology. Upgrading your system installs the modular 2.0.x kernel as well as updated versions of the packages that are installed on your machine. The upgrade process preserves existing configuration files using a `.rpmsave` extension (e.g., `sendmail.cf.rpmsave`) and leaves a log telling what actions it took in `/tmp/upgradelog`. As software evolves, configuration file formats can change, so you should carefully compare your original configuration files to the new files before integrating your changes.

## A Word From the Developers

We would like to thank all our beta testers for entrusting their systems to early versions of Red Hat Linux and for taking the time to submit bug reports from the front, especially those of you who have been with Red Hat since the “Halloween” release and earlier. We would also like to thank Linus Torvalds and the hundreds of developers around the world for creating, truly, one of the wonders of distributed development.

And, again, we’d like to thank *you* for your interest in Red Hat Linux!

*The Red Hat Development Team*

## Notes from the Editor

Starting with this manual, Red Hat Software has dedicated full-time resources to producing documentation. Those of you that have read previous versions of this manual will probably notice that some changes have been made; we hope you’ll like them. Our goal is to evolve this manual into a comprehensive document that can serve the needs of anyone wishing to install Red Hat Linux, no matter what their level of experience.

## **We Need Feedback!**

As we move towards this goal, there are sure to be mistakes. If you spot one, or if you've thought of a way to make this manual better, we'd love to hear from you! Please send mail to:

`docs@redhat.com`

Be sure to mention the manual's title, along with your suggestion. If you've found an error, please include the page number so we can find it easily. We may not be able to respond to every message sent to us, but you can be sure that we'll be reading them all!

## **I Couldn't Have Done it Without...**

Thanks go out to the past authors of this manual. A great deal of their work is still here. A great, big "Thank You" is also owed to the Red Hat Linux 5.0 development team for putting up with the many questions, comments, and pleading requests for reviews of this manual. Thanks, guys!

Last but far from least is the support group at Red Hat Software. They have given many insightful suggestions regarding this manual, based on extensive experience with thousands of Red Hat Linux customers. So if you find yourself breezing through the installation chapters, it's due in no small part to their input. I thank them.

*Edward C. Bailey*

# 1

## New Features Of Red Hat Linux 5.0

This chapter describes features that are new to Red Hat Linux 5.0.

### 1.1 Installation-Related Enhancements

There have been many changes made to make the Red Hat Linux installation process easier. Here's a list:

- Automatic probing of PCI devices
- Expert installation mode
- Kickstart installation mode
- New method of configuring partitions and mount points
- Daemons can be configured to run at boot or not
- Enhanced XFree86 configuration

### 1.1.1 Automatic Probing of PCI devices

The installation program can probe computers with PCI devices. If desired, auto-probing can be disabled by running the installation in expert mode (see Section 1.1.2).

### 1.1.2 Expert Installation Mode

By entering `expert` at the initial `boot :` prompt, you can run the installation program in expert mode. Expert mode causes the installation program to revert to an installation style very similar to that used with Red Hat Linux 4.2. This means that most of the new auto-probing code is replaced with dialog boxes that prompt the user to manually enter device information. These dialog boxes also include a field allowing the user to enter options to be passed to the driver module being loaded. Therefore, if a particular hardware configuration requires module options, an expert-mode installation should be used.

### 1.1.3 Kickstart Installation Mode

It's possible to automate Red Hat Linux installations by using kickstart mode. Instead of asking the user questions during the installation, kickstart mode reads a file that contains all the information necessary to install Red Hat Linux. This is handy when installing many similarly-configured systems. For more information on kickstart mode, please read `/doc/README.ks` on the Red Hat Linux CD-ROM.

### 1.1.4 Configuring Partitions and Mount Points

In the past, the installation program used `fdisk` to set up disk partitions, and a series of dialog boxes to define mount points. While that functionality is still available in Red Hat Linux 5.0, there is an alternative. Disk Druid provides both mount point definition and partition configuration in one concise screen. In addition, Disk Druid supports the concept of "growable partitions" – partitions that can grow beyond a minimum size to allocate all remaining space on a hard disk.

For more information on Disk Druid, please refer to Section 4.3.1 on page 54.

**SPARC**

Disk Druid is not available under the Red Hat Linux/SPARC installation program.

### 1.1.5 Configuring Services to Run at Boot Time

In prior versions of Red Hat Linux, if you installed a particular service (such `bind`, `samba`, etc.) while installing Red Hat Linux, that service configured to automatically start once the installation was complete and the system rebooted. Now, you will have the option, during the installation, to select which services should start and which should not. During the installation, you will be asked to select which services should be started whenever the system is booted. For more information, refer to Section 5.5 on page 76.

### 1.1.6 Enhanced XFree86 Configuration

The XFree86 configuration utility `xconfigurator` has been enhanced. It will now attempt to probe your system to automatically determine what type of video card is installed. In addition, `xconfigurator` will now determine the optimal display resolution and color depth for your hardware configuration.

## 1.2 New User-Mode Tools

One of the biggest stumbling blocks to a new Red Hat Linux user is that many tasks can only be performed by the root account, by knowing an arcane set of commands, or both. To make life easier, there are several new tools that can be used by any user to perform common tasks:

- User information tool
- User password tool
- Filesystem mount tool
- Network device tool

### 1.2.1 User Information Tool

The user information tool is used to set the “gecos” information in a user’s passwd record. Figure 1.1 shows it in action.



Figure 1.1: User Information Tool

As you can see, you can enter your:

- Full Name
- Office
- Office Phone
- Home Phone
- Shell

To use this tool, run `/usr/bin/userinfo`.

### 1.2.2 User Password Tool

As the name implies, the user password tool is used to change passwords. The tool displays a pop-up message identifying the user whose password is being changed, followed by prompts for the current and new password. To use the password tool, run `/usr/bin/userpasswd`.

### 1.2.3 Filesystem Mount Tool

This program makes it easy to mount and unmount filesystems. It will display **Mount** and **Format** buttons for each filesystem in `/etc/fstab` that has been configured with the `user` option. In Figure 1.2, the CD-ROM drive has been mounted (and can be unmounted with a click of the **Unmount** button). The **Format** button is active for all writeable, unmounted media, and will format the media when pressed.

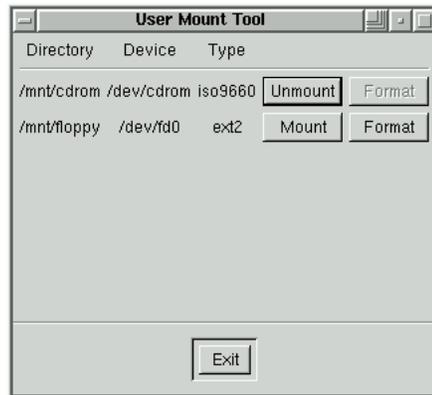


Figure 1.2: Filesystem Mount Tool

To use this tool, run `/usr/bin/usermount`.

### 1.2.4 Network Device Tool

It's easy to start and stop network interfaces when using this tool. Every network interface available to users is displayed, along with a colored button representing its status (see Figure 1.3 on the next page). Clicking on the interface's button toggles that interface (ie, turns it on when it is off, and vice versa). Clicking on the status info button forces an update on that interface's status (though the status will also update automatically). When done, click **Exit** to exit.

To use this tool, run `/usr/bin/usernet`.

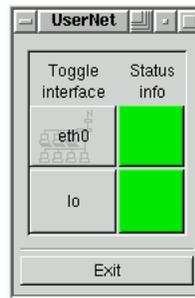


Figure 1.3: Network Device Tool

## 1.3 System Administration-Related Enhancements

Red Hat Linux 5.0 has these features for system administrators:

- glibc-based system
- Improved shadow utilities
- New initscript utility
- New backup script
- Improved network configuration utility
- New group for diskette access
- Anti-spam changes to sendmail

### 1.3.1 Glibc-Based System

Red Hat Linux 5.0 is now based on a newer library known as *glibc*. There are a number of advantages to *glibc*:

1. As an actively maintained library, *glibc* is more stable.
2. *Glibc* has been found to be faster by many people.

3. The same library can now be used on all Red Hat Linux platforms.
4. Glibc supports threads and internationalization.

### 1.3.2 Improved Shadow Utilities

Support for shadow passwords has been enhanced significantly for Red Hat Linux 5.0. The `shadow-utils` package contains a number of utilities that support:

- Conversion from normal to shadowed passwords and back (`pwconv`, `pwunconv`).
- Verification of the password, group, and associated shadow files (`pwck`, `grpck`).
- Industry-standard methods of adding, deleting and modifying user accounts (`useradd`, `usermod`, and `userdel`).
- Industry-standard methods of adding, deleting, and modifying user groups (`groupadd`, `groupmod`, and `groupdel`).
- Industry-standard method of administering the `/etc/group` file (`gpasswd`).

**Please Note:** There are a few additional points of interest concerning these utilities:

- The utilities will work properly whether shadowing is enabled or not.
- The utilities have been slightly modified to support Red Hat Software's user private group scheme. For a description of the modifications, please see the `useradd` man page. For more information on user private groups, please turn to Section 11.3.3 on page 171.
- The `adduser` script has been replaced with a symlink to `/usr/sbin/useradd`.

### 1.3.3 New Initscript Utility

The `chkconfig` utility provides a simple command-line tool for maintaining the `/etc/rc.d` directory hierarchy. It relieves system administrators from having to directly manipulate the numerous symlinks in `/etc/rc.d`.

Please see the `chkconfig` man page for more information.

### 1.3.4 New Backup Script

As the man page states, “`rhbackup` is a program useful for simple local and remote system backups. It ... uses `rsh` and `tar` as its basic network and archiving mechanism.”

For more information on `rhbackup`, please consult the `rhbackup` man page.

### 1.3.5 Improved network configuration utility

The network configuration utility `netcfg` has been enhanced. It now has the ability to create *clone* network interfaces. Consider the following example:

A laptop is connected to a corporate LAN most days. Sometimes, the laptop’s owner travels to a division of the company, and would like to use the laptop on the division’s LAN. In addition, at night, the laptop’s owner would like to connect the laptop to a LAN at home.

In the past, this situation would have required a fair bit of effort to accommodate. With clone interfaces, it’s only necessary to create two clones of the original Ethernet interface, and give them the appropriate information. This is done by selecting the interface to clone and pressing the **clone** button. After giving the clone device a name (such as “home”), you’ll see a window similar to the one in Figure 1.4 on the next page.

As you can see, you can specify an alternate IP address and netmask (the network and broadcast addresses are calculated automatically), or use `bootp` or `dhcp` to configure the interface. In addition, the clone device may be configured to start automatically at boot time, and can be set so that any user may bring the interface up or down.

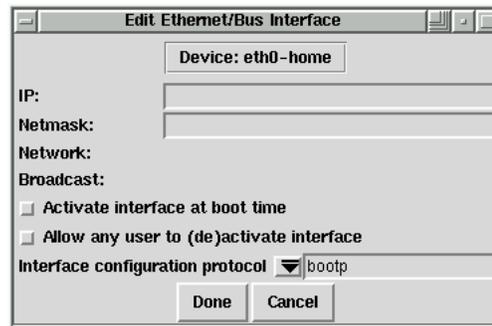


Figure 1.4: Clone Interface Setup

### 1.3.6 New Diskette Group

Red Hat Linux 5.0 also includes a new group that system administrators can use to give certain users access to the system's diskette drive. This can be done explicitly, or via the `pam_env` module.

### 1.3.7 Anti-spam changes to sendmail

Sendmail version 8.8.7 is included with Red Hat Linux 5.0. Also included are rulesets that discourage the distribution of unsolicited commercial e-mail (also known as "spam"). These rulesets were created by Claus Aßmann. Although some information is included below, please consult Claus' web page documenting his changes at:

<http://www.informatik.uni-kiel.de/%7Eca/email/check.html>

**Please Note:** In this release of Red Hat Linux, sendmail does not relay by default. Because of this, if your Red Hat Linux system must act as a relay, please review the following information.

#### Overview

This release has been configured to take advantage of the latest changes to sendmail to prevent spam by the addition of new sendmail rulesets.

The following rulesets have been added to the S98 (local additions) section of `/etc/sendmail.cf`:

- `junk`
- `check_mail`
- `check_rcpt`
- `remove_local`
- `junkIP`
- `check_relay`

The actions of these rulesets are controlled by the following files in the `/etc/mail` directory:

- `ip_allow`
- `name_allow`
- `relay_allow`
- `deny`

Let's take a look at each of these files.

`ip_allow`

This is a flat text file that contains the addresses or networks of hosts that are allowed to use this host as a relay.

The format for this file is:

```
hhh.hhh.hhh.hhh
ccc.ccc.ccc
bbb.bbb
aaa
```

These lines represents a host address, a class C network, a class B network and a class A network, respectively.

name\_allow

This is a flat text file that contains the hostnames or domain names of hosts that are allowed to use this host as a relay.

The format for this file is:

```
hostname.domain.com
domain.com
```

These lines represent a hostname and a domain name.

**Please Note:** Hosts not 'included' in either ip\_allow or name\_allow will be given an error message of "551: xxx, we do not relay" where xxx is the RCPT address that was given.

relay\_allow

This is a flat text file that contains the hostnames, domain names, host addresses or networks for which this machine will relay to:

The format for this file is:

```
hostname.domain.com
domain.com
hhh.hhh.hhh.hhh
ccc.ccc.ccc
bbb.bbb
aaa
```

The formats in this file are the same as for the ip\_allow and name\_allow files described earlier.

**Please Note:** Since ip\_allow, name\_allow and relay\_allow are plain text files, any changes to them will necessitate restarting sendmail. This can be accomplished by executing the followin commands:

```
/etc/rc.d/init.d/sendmail stop; /etc/rc.d/init.d/sendmail start.
```

deny

This is a flat text file that must be converted into a hash database file with `makemap(8)` before use. This file contains a set of email addresses, hostnames, domain names, host addresses, or network addresses separated by a tab from an error message.

The format for this file is:

```
user@domain.com<tab>error message 0
hostname.domain.com<tab>error message 1
domain.com<tab>error message 2
hhh.hhh.hhh.hhh<tab>error message 1
ccc.ccc.ccc<tab>error message 3
bbb.bbb<tab>error message 4
aaa<tab>error message 5
```

Each line represents the identification of a known spammer. When an email address is specified in this file, the error message provided will be returned to the spammer's MTA. For the other formats, the corresponding error message will be logged to `/var/log/maillog`, and "550 Access denied" is returned to the spammer's MTA.

The following command can be used to create the hash database version of this file:

```
makemap -v /etc/mail/deny < /etc/mail/deny
```

The output from this command is the database file `/etc/mail/deny.db`.

## 1.4 Miscellaneous New Features

These new features defy categorization:

- Modular Sound Drivers
- Sound Card Configuration Tool
- Window Manager Enhancements

### 1.4.1 Modular Sound Drivers

In Red Hat Linux 5.0, the standard OSS/Free sound drivers have been modified to be completely modular. This allows for such things as loading and unloading of the various sound drivers without recompiling the kernel tree or rebooting. The work was performed by Alan Cox and was sponsored by Red Hat Software, Inc.

For additional information, please consult the README files in the `sndconfig` documentation directory (`/usr/doc/sndconfig*`). The latest information can always be found at `ftp://ftp.redhat.com/pub/sound/`.

If you have any issues concerning the modular sound drivers, please send mail to `sound-bugs@redhat.com`. There is also a mailing list associated with the modular sound drivers (`sound-list@redhat.com`). To subscribe, send mail to `sound-list-request@redhat.com` with "subscribe" as the subject line.

#### Supported Sound Cards

At present, the following sound cards are supported:

- Sound Blaster 1.0
- Sound Blaster 2.0
- Sound Blaster Pro
- Sound Blaster 16
- Sound Blaster 16 PnP
- Sound Blaster AWE32/AWE64 (In SB-16 mode only)

### 1.4.2 Sound Card Configuration Tool

Also included in Red Hat Linux 5.0 is `sndconfig`, a screen-oriented utility that can properly configure modular sound card drivers.

There are a few things that you should know about `sndconfig`:

**Plug and Play Aware** – `sndconfig` is able to detect and configure Plug and Play sound cards such as the Sound Blaster 16 PnP. Please note, however, that the

present version of `sndconfig` creates a new version of the file `/etc/isapnp.conf` with only the sound card's settings uncommented. Therefore, if you have other Plug and Play devices on your system, you must do one of two things:

- You can manually add your sound card's PnP information to your existing `/etc/isapnp.conf`, and run `sndconfig` using the `--noprobe` option.
- You can run `sndconfig` without the `--noprobe` option, and uncomment the configuration information for your other PnP cards in the `/etc/isapnp.conf` file created by `sndconfig`.

In any case, note that `sndconfig` saves your original `/etc/isapnp.conf` file as `/etc/isapnp.conf.bak`.

**Modifies** `/etc/conf.modules` – `sndconfig` modifies the module configuration file `/etc/conf.modules` by adding information about the module options required for your sound card. Note that `sndconfig` saves your original `/etc/conf.modules` file as `/etc/conf.modules.bak`.

To set up your sound card, run `/usr/sbin/sndconfig`.

### 1.4.3 Window Manager Enhancements

Enhancements have been made to the supplied window managers and their menus.

# 2

## Before You Begin

While installing Red Hat Linux is a straightforward process, taking some time prior to starting the installation can make things go much more smoothly. In this chapter, we'll discuss the steps that should be performed before you start the installation.

**Please Note:** If you are currently running a version 2.0 (or greater) Red Hat Linux system, you can perform an upgrade. Skim this chapter to review basic issues relating to installation, and read the following chapters in order, following the directions as you go. The upgrade procedure starts out identically to the installation procedure; you will be directed to choose an installation or upgrade after booting the installation program and answering a few questions.

There are five things you should do prior to installing Red Hat Linux:

1. Make sure you have access to the Red Hat Linux components required for installation.
2. Make sure you know your computer's hardware configuration and networking information.
3. Decide, based on the first two tasks, what method you will use to install Red

Hat Linux.

4. Determine where on your hard drive(s) Red Hat Linux will reside.
5. Make sure you have sufficient documentation to effectively use your Red Hat Linux system after the installation.

Let's start by making sure you have everything you need to install Red Hat Linux.

## 2.1 Getting the Right Red Hat Linux Components

If you've purchased the Red Hat Linux boxed set, you're ready to go! However, mistakes occasionally happen, so now is a good time to double-check the contents of your boxed set. If you haven't purchased a Red Hat Linux boxed set, skip to Section 2.1.3 on page 18.

### 2.1.1 Contents of the Red Hat Linux Boxed Set

The Red Hat Linux boxed set contains the following items:

- The installation guide.
- Red Hat Linux CDs 1 and 2.
-  Boot and supplemental diskettes.
- License and Registration card.
- Sticker with registration number.

Let's take a quick look at each item:

#### Installation Guide

The Red Hat Linux Installation Guide is what you're currently reading. It contains the information necessary to install Red Hat Linux. In addition, it contains information about aspects of the operating system that are unique to Red Hat Linux.

## CDs 1 and 2

These two Compact Discs contain the entire Red Hat Linux distribution, including source code. CD 1 contains all the binary packages built for the type of computer (Intel, Alpha, or SPARC) that you have. CD 2 contains the source packages that were used to build the binary packages on CD 1.

## Boot and Supplemental Diskettes

### INTEL

These diskettes are used to start the installation process for Red Hat Linux/Intel. Depending on your computer's configuration and the type of installation you select, you may need only the boot diskette, the boot *and* supplemental diskettes, or none at all. When we discuss the different installation methods later in this chapter, we'll explain which diskettes are needed for each type of installation.

### ALPHA

Alpha owners should refer to Section D.4 on page 214 for information on which diskettes are required.

### SPARC

SPARC owners should turn to Section C.5.1 on page 202 for information on their diskette needs.

## License and Registration card

This card contains the license terms for Red Hat Linux, in addition to the license terms for any commercial software that may be included on the Red Hat Linux CD. This is where you'll also find any late-breaking information about Red Hat Linux 5.0.

## Registration Sticker

The alphanumeric registration string on this sticker is used to register you for installation support from Red Hat Software.

### 2.1.2 Missing Something?

If you've purchased the Official Red Hat Linux boxed set from Red Hat Software (or one of its distributors), and you're missing one or more of the items listed above, please let us know! Either call us at 1-888-733-4281 (+1-919-547-0012 outside the USA), or send mail to [sales@redhat.com](mailto:sales@redhat.com).

If your Red Hat Linux boxed set is complete, please skip ahead to section 2.1.4.

### 2.1.3 No Boxed Set? No Problem!

Of course, not everyone purchases a Red Hat Linux boxed set. It's entirely possible to install Red Hat Linux using a CD created by another company, or even via FTP. In these cases, you may need to create one or more diskettes to get started.

#### INTEL

For people installing Red Hat Linux/Intel, you'll need a boot diskette and, optionally, a supplemental diskette. It is also possible to start the installation directly from the CD, under certain conditions. We'll discuss this in more detail when we outline the various installation methods available.

#### ALPHA

People with Alpha-based systems should refer to Section D.4 on page 214 for additional information on the diskettes they may need.

#### SPARC

SPARC owners should refer to section C.5.1 on page 202 for information on which diskettes they'll need.

### 2.1.4 Checking for Updated Diskette Images

From time to time, we find that the installation may fail, and that a revised diskette image is required in order for the installation to work properly. In these cases, we make special images available via the Red Hat Linux Errata.

Since this is a relatively rare occurrence, you will in general save time if you try to use the standard diskette images first, and then review the Errata only if you experience any problems completing the installation.

There are two ways to review the Errata:

1. **World Wide Web** – By pointing your web browser at <http://www.redhat.com/errata>, you can read the Errata on-line, and download diskette images easily.
2. **Electronic Mail** – By sending an empty mail message to [errata@redhat.com](mailto:errata@redhat.com), you will receive a mail message containing the complete Errata. Also included are URLs to each updated package and diskette image in the Errata. By using these URLs, you can then download any necessary diskette images. Remember to use binary mode when transferring a diskette image!

For now, concentrate only on the Errata entries that include new diskette images (the filenames always end in `.img`). If you find an entry that seems to apply to your problem, get a copy of the diskette images, and create them using the instructions in Appendix B on page 193.

## 2.2 Things You Should Know

In order to prevent any surprises during the installation, you should collect some information before attempting to install Red Hat Linux. You can find most of this information in the documentation that came with your system, or from the system's vendor or manufacturer.

**Please Note:** The most recent list of hardware supported by Red Hat Linux can be found at Red Hat Software's World Wide Web site at <http://www.redhat.com/hardware>. It's a good idea to check your hardware against this list before proceeding.

### 2.2.1 Basic Hardware Configuration

You should have a basic understanding of the hardware installed in your computer, including:

- **hard drive(s)** – Specifically, the number, size, and type. If you have more than one, it's helpful to know which one is first, second, and so on. It is also good to know if your drives are IDE or SCSI.

- **memory** – The amount of RAM installed in your computer.
- **CD-ROM** – Most importantly, the unit's interface type (IDE, SCSI, or other interface) and, for non-IDE, non-SCSI CD-ROMs, the make and model number.
- **SCSI adapter (if one is present)** – The adapter's make and model number.
- **network card (if one is present)** – The card's make and model number.
- **mouse** – The mouse's type (serial, PS/2, or bus mouse), protocol (Microsoft, Logitech, MouseMan, etc.), and number of buttons; also, for serial mice, the com port it is connected to.

#### ALPHA

In addition to the latest hardware compatibility list on Red Hat Software's website, owners of Alpha-based systems should refer to Section D.2 on page 211 for more information on supported hardware configurations.

#### SPARC

In addition to the latest hardware compatibility list on Red Hat Software's website, SPARC owners should refer to Section C.1 on page 197 for a list of supported hardware.

## 2.2.2 Video Configuration

If you will be installing the X Window System, you should also be familiar with the following:

- **your video card** – The card's make and model number (or the video chipset it uses), and the amount of video RAM it has.
- **your monitor** – The unit's make and model number, along with allowable ranges for horizontal and vertical refresh rates.

## 2.2.3 Network-related Information

If you will be connected to a network, be sure you know your:

- **IP address** – Usually represented as a set of four numbers separated by dots, such as 10.0.2.15.
- **netmask** – Another set of four numbers separated by dots. An example netmask would be 255.255.248.0.
- **gateway IP address** – Yet another set of four dot-separated numbers. For instance, 10.0.2.254.
- **one or more name server IP addresses** – One or more sets of dot-separated numbers. 10.0.2.1 might be the address of a name server.
- **domain name** – The name given to your organization. For instance, Red Hat Software has a domain name of redhat.com.
- **hostname** – The name of your computer. A computer might be named pooh, for instance.

**Please Note:** The information given above is an example only! Do *not* use it when you install Red Hat Linux! If you don't know the proper values for your network, ask your network administrator.

## 2.3 Installation Methods

### ALPHA

Alpha owners should take a moment to read Appendix D on page 211. This appendix covers aspects of the installation process that differ from a typical Red Hat Linux/Intel installation.

### SPARC

People with SPARC systems should take a moment to read Appendix C on page 197. This appendix covers aspects of the installation process that differ from a typical Red Hat Linux/Intel installation.

You can install or upgrade Red Hat Linux via any of several different methods. Each method works best in different situations, and has different requirements. But before we discuss each installation method, let's take a look at an issue that may affect some of you.

### 2.3.1 PCMCIA Support During the Installation

**INTEL**

The PCMCIA section is specific to Intel-based computers only

Most Intel-based laptop computers support PCMCIA (also known as PC Card). Computers that support PCMCIA devices contain a controller having one or more slots in which a PCMCIA device can be installed. These devices may be modems, LAN adapters, SCSI adapters, and so on.

When installing Red Hat Linux/Intel on a PCMCIA-capable computer, it is important to note if a PCMCIA device will be used during installation. For example, if you want to install Red Hat Linux/Intel from a CD-ROM, and your CD-ROM drive is connected to a PCMCIA adapter, the installation program will require PCMCIA support. Likewise, if you are going to use one of the network-based installation methods, you will need PCMCIA support if your network adapter is PCMCIA-based.

PCMCIA support is dependent on two things:

1. The type of PCMCIA controller in your computer system.
2. The type of PCMCIA device that you wish you use during the installation.

While nearly every PCMCIA controller and most popular PCMCIA devices are supported, there are some exceptions. For more information, please consult the Red Hat Linux Hardware Compatibility List at <http://www.redhat.com/hardware>.

The main thing to keep in mind is that if you require PCMCIA support, you will need the supplemental diskette.

**ALPHA**

PCMCIA support is not available for the Alpha.

**SPARC**

PCMCIA support is not available for the SPARC.

### 2.3.2 Installing From a CD-ROM

If you have a Red Hat Linux CD-ROM, and your computer has a supported CD-ROM drive, you should consider this installation method. Installing directly from CD-ROM is the most straightforward approach. When installing from CD-ROM, the packages you select are read from the CD-ROM, and are installed on your hard drive.

#### How To Do It

As the name implies, you'll need a Red Hat Linux CD-ROM, a supported CD-ROM drive, and a means of starting the installation program.

<b>INTEL</b> Intel systems will need to use the boot diskette (and the supplemental diskette if PCMCIA support is required). There is an alternate method of installing from CD-ROM that uses no diskettes, but requires that the system be running DOS. We'll discuss this approach (known as <i>autoboot</i> in Chapter 3 on page 35. For now, note that PCMCIA support is not available when using <i>autoboot</i> .
--

If you've determined that this installation method is most applicable to your situation, please skip ahead to Section 2.4 on page 25.

### 2.3.3 Installing From an FTP Site

If you don't have a Red Hat Linux CD-ROM or a CD-ROM drive, but you do have network access, then an FTP installation may be for you. When installing via FTP, the Red Hat Linux packages you select are downloaded (using FTP) across the network to your computer, and are installed on your hard drive.

#### How To Do It

When doing an FTP install, you'll need LAN-based access to a network; a dialup connection via modem won't cut it. If your Local Area Network has Internet access, you can use one of the many FTP sites that mirror Red Hat Linux. You can find a list of mirror sites at `ftp://ftp.redhat.com/pub/MIRRORS`.

If your LAN doesn't have Internet access, all is not lost. If there is a computer on your LAN that can accept anonymous FTP requests, simply load Red Hat Linux on that system, and you're ready to go.

**INTEL**

For an FTP installation, you must use both the boot and supplemental diskettes. You will need to have a valid nameserver configured or you must specify the IP address of the FTP server you will be using. You will also need the path to the Red Hat Linux directory on the FTP server.

If you've determined that this installation method is most applicable to your situation, please skip ahead to Section 2.4 on the next page.

### 2.3.4 Installing From an NFS Server

If your system doesn't have a CD-ROM drive, but you do have network access, then an NFS installation may be for you. When installing via NFS, the Red Hat Linux packages you select are NFS-served to your computer from an NFS server system. The packages are then installed on your hard drive.

#### How To Do It

If you wish to perform an NFS installation, you will need to mount the Red Hat Linux CD-ROM on a machine that supports ISO-9660 file systems with Rock Ridge extensions. The machine must also support NFS. Export the CD-ROM file system via NFS. You will need to have a nameserver configured, or know the NFS server's IP address, as well as the path to the exported CD-ROM.

**INTEL**

For an NFS installation, you'll need a boot diskette only.

If you've determined that this installation method is most applicable to your situation, please skip ahead to Section 2.4 on the facing page.

### 2.3.5 Installing From a Hard Drive

If none of the other installation methods will work for you, but you have some means of getting the Red Hat Linux package files written to your system's hard drive, you can install from your hard drive. In this installation method, the Red Hat Linux packages you select are read from one partition on a hard drive, and are installed on another partition (or set of partitions).

#### How To Do It

The hard drive installation method requires a bit of up-front effort on your part, as you must copy all the necessary files to a partition before starting the Red Hat Linux installation program. You must first create a `RedHat` directory at the top level of your directory tree. Everything you will install should be placed in that directory. First copy the base subdirectory and its contents.

Next, copy the packages you want to install to another subdirectory called `RPMS`. You can use available space on an existing DOS partition or a Linux partition that is not required in the install procedure (for example, a partition that would be used for data storage on the installed system).

#### INTEL

If you are using a DOS filesystem, you will not be able to use the full Linux filenames for the RPM packages. The installation process does not care what the filenames look like, but it is a good idea that you keep track of them. You'll need a boot and supplemental diskette when installing from a hard drive.

## 2.4 Disk Partitions

In order to install Red Hat Linux, you must make disk space available for it. This disk space needs to be separate from the disk space used by other operating systems you may have installed on your computer, such as Windows, OS/2, or even a different version of Linux.

A disk can be divided into different *partitions*. Each partition can be accessed as if it was a separate disk. Furthermore, each partition has a *type* that is used to indicate how information is stored in the partition. For example, there are different partition types used by DOS, OS/2, and Linux.

Even if you will be installing Red Hat Linux on its own hard disk, or on a computer which contains no other operating system, you'll still need to create partitions for Red Hat Linux to use. In this case it's pretty easy, as there are no other partitions on the hard disk to worry about.

On the other hand, you may wish to install Red Hat Linux on a disk which already contains software or data from a different operating system. Things can get a little trickier in this situation, since a mistake can destroy your existing partitions, not to mention the data they contain!

During the installation process, you'll be given the chance to create partitions for Red Hat Linux. At this point, your main concern is making sure you have sufficient disk space available to create those partitions. Let's review the different ways to free up space for Red Hat Linux partitions.

### 2.4.1 Repartitioning Strategies

There are three possible scenarios you may face when attempting to repartition your hard disk:

- Unpartitioned free space is available.
- An unused partition is available.
- Free space in an actively used partition is available.

Let's look at each scenario in order.

#### Using Unpartitioned Free Space

In this situation, the partitions defined do not span the entire hard disk, leaving unallocated space that is not part of any defined partition. If you think about it, an unused hard disk also falls into this category; the only difference is that *all* the space is not part of any defined partition.

In this case, you can simply create the necessary partitions from the unused space.

### Using An Unused Partition

Last year you replaced that tiny 105MB hard drive on your Windows system with a 1.2GB monster. You partitioned it into two equal parts, figuring that you'd use the C: "drive" (really the drive's first partition) for Windows, and the D: "drive" (really the drive's second partition) for your collection of freeware programs downloaded from the Internet. Well, you'd been so used to using C: that you never put anything of substance on D:.

If you find yourself in this situation, you can use the space allocated to the unused partition. You'll need to delete the partition, and create the appropriate Linux partitions in its place.

### Using Free Space From An Active Partition

This is the most common situation. It is also, unfortunately, the hardest to deal with. The main problem is that you have enough free space, but it's presently allocated to a partition that is in use. If you purchased a computer with pre-installed software, the hard disk most likely has one massive partition holding the operating system and data.

Aside from adding a new hard drive to your system, you have two choices:

**Destructive Repartitioning** – Basically, you delete the single large partition, and create several smaller ones. As you might imagine, any data you had in that partition is destroyed. This means that making a complete backup is necessary. For your own sake, make two backups, use verification (if available in your backup software), and try to read data from your backup *before* you delete the partition. Note also that if there was an operating system of some type installed on that partition, it will need to be reinstalled as well.

After creating a smaller partition for your existing software, you can reinstall any software, restore your data, and continue with your Red Hat Linux installation.

**Non-Destructive Repartitioning** – Here, you run a program that does the seemingly impossible; it makes a big partition smaller without losing any of the files stored in that partition. Many people have found this method to be reliable and trouble-free; however, we recommend that you make backups of the partition being resized, just in case.

**INTEL**

What software should you use to perform this feat? You can use the `fips` utility to resize FAT (File Allocation Table) partitions. It's included on the Red Hat Linux/Intel CD-ROM in the `dosutils` directory. Before using it, however, we highly recommend reading the `fips` documentation, located in the `fipsdocs` subdirectory.

**Please Note:** At the present time, the `fips` utility cannot resize the new FAT32-style partitions found on many newer computers. If your computer has a FAT32 partition, you can either use destructive repartitioning, or purchase commercial disk management software capable of shrinking FAT32 partitions.

## 2.4.2 Disk Partitions and Other Operating Systems

If your Red Hat Linux partitions will be sharing a hard disk with partitions used by other operating systems, most of the time you'll have no problems. However, there are certain combinations of Linux and other operating systems that require extra care. Information on creating disk partitions compatible with other operating systems is available in several HOWTOs and Mini-HOWTOs, available on the Red Hat Linux CD in the `doc/HOWTO` and `doc/HOWTO/mini` directories. In particular, the Mini-HOWTOs whose names start with `Linux+` are quite helpful.

**INTEL**

If Red Hat Linux/Intel will coexist on your machine with OS/2, you must create your disk partitions with the OS/2 partitioning software—otherwise, OS/2 may not recognize the disk partitions. During the installation, do not create any new partitions, but do set the proper partition types for your Linux partitions using the Linux `fdisk`.

## 2.4.3 One Last Wrinkle: Using LILO

**INTEL**

LILO (the LInux LOader) is the most commonly used method to boot Red Hat Linux on Intel-based systems. Being an operating system loader, LILO operates "outside" of any operating system, using only the Basic I/O System (or BIOS) built into the computer hardware itself.

**ALPHA**

Alpha owners may skip ahead to Section 2.4.4 on the next page.

**SPARC**

SPARC owners may skip ahead to Section 2.4.4 on the following page.

**BIOS-Related Limitations Impacting LILO**

LILO is subject to some limitations imposed by the BIOS in most Intel-based computers. Specifically, most BIOSes can't access more than two hard drives and they can't access any data stored beyond cylinder 1023 (the 1024th cylinder) of any drive. Note that some recent BIOSes do not have these limitations, but this is by no means universal.

All the data LILO needs to access at boot time (including the Linux kernel) are located in the `/boot` directory, which is normally part of the root part partition (known as `/`). Here are the guidelines you must follow if you are going to use LILO to boot your Red Hat Linux system:

**On First Two IDE Drives** – If you have 2 IDE (or EIDE) drives, `/boot` must be located on one of them. Note that this two-drive limit also includes any IDE CD-ROM drives on your primary IDE controller. So, if you have one IDE hard drive, and one IDE CD-ROM on your primary controller, `/boot` must be located on the first hard drive *only*, even if you have other hard drives on your secondary IDE controller.

**On First IDE Or First SCSI Drive** – If you have one IDE (or EIDE) drive and one or more SCSI drives, `/boot` must be located either on the IDE drive or the SCSI drive at ID 0. No other SCSI IDs will work.

**On First Two SCSI Drives** – If you have only SCSI hard drives, `/boot` must be located on a drive at ID 0 or ID 1. No other SCSI IDs will work.

**Partition *Completely* Below Cylinder 1023** – No matter which of the above configurations apply, the partition that holds `/boot` must be located entirely below cylinder 1023. If the partition holding `/boot` straddles cylinder 1023, you may face a situation where LILO will work initially (because all the necessary information is below cylinder 1023), but will fail if a new kernel is to be loaded, and that kernel resides above cylinder 1023.

As mentioned earlier, it is possible that some of the newer BIOSes may permit LILO to work with configurations that don't meet our guidelines. Likewise, some of LILO's more esoteric features may be used to get a Linux system started, even if the configuration doesn't meet our guidelines. However, due to the number of variables involved, Red Hat Software cannot support such extraordinary efforts.

#### 2.4.4 How Many Partitions?

Although you can install Red Hat Linux in a single large partition (subject to any of the partitioning considerations we've mentioned so far), it's a much better idea to split things up a bit. We recommend the following layout as a compromise between single-partition simplicity, and multi-partition flexibility:

**Please Note:** If you plan to install all the software packages accompanying Red Hat Linux, you may need to use even larger partitions.

**A swap partition** – Swap partitions are used to support virtual memory. If your computer has 16 MB of RAM or less, you *must* create a swap partition. Even if you have more memory, a swap partition is still recommended. The minimum size of your swap partition should be equal to your computer's RAM, or 16 MB (whichever is larger).

**A root partition** – The root partition is where / (the root directory) resides. It only needs to contain things necessary to boot your system, as well as system configuration files. A root partition of 50 MB to 80 MB works well for most systems.

INTEL

Don't forget the LILO constraints we mentioned in Section 2.4.3 on page 28!

**A /usr partition** – The /usr partition is where much of the software on a Red Hat Linux system resides. This partition should be between 200 MB and 500 MB, depending on how many packages you plan to install. If at all possible, try to be generous with the /usr partition. Any RPM-based packages you install later will (in general) use more space from /usr than from any other partition.

**A /home partition** – This is where users' home directories go; the size of /home depends on how many users you plan to have on your Red Hat Linux system and what they might store in their home directories.

**ALPHA**

The following partition is specific to Red Hat Linux/Alpha installations.

**A MILO partition** – Alpha owners that will be using MILO to boot their systems should create a 1.5 MB DOS partition where MILO can be copied after the installation is complete.

Additionally, you may wish to create any of the following:

**A /usr/local partition** – Traditionally, `/usr/local` has been used to hold things you wish to keep separate from the rest of your Red Hat Linux system, such as software that is not available as an RPM package. The size depends on the amount of software you anticipate putting on your system.

**A /usr/src partition** – There are two things that are stored in `/usr/src` on a Red Hat Linux system:

**Linux Kernel Sources** – The complete sources for the Linux kernel are stored here, and new kernels are built here.

**Sources For RPM-Based Packages** – If a source package file (aka SRPM) is installed, the files are stored here. Note that, unless specified otherwise, any packages built also use a build directory located here.

Again, the size of this partition would depend on the amount of software you anticipate building.

**A /tmp partition** – As the name implies, the `/tmp` partition is for temporary files. Creating a partition dedicated to `/tmp` is a good idea for larger, multiuser systems or network server machines. The reason is that many active users can fill the root partition (`/`), which is where `/tmp` is located. It's not necessary to dedicate a partition to `/tmp` on single-user workstations.

## 2.5 Getting Documentation

Red Hat Linux is a powerful, full-featured operating system. Unless you're a Linux wizard, you're going to need documentation to make the most of your Red Hat Linux system. Everyone should review Chapter 7 on page 113 for more information on available Linux documentation. While many people find the resources

described in chapter 7 to be very helpful, people who are just starting to use Linux will likely need additional information. The information that will be most helpful to you depends on your level of Linux expertise:

**New To Linux** – If this is your first time using Linux (or any Linux-like operating system, for that matter), you’ll need solid introductory information on basic Unix concepts. For example, O’Reilly and Associates (<http://www.ora.com/>) produce a wide variety of Linux and Unix-related books. Give their more general titles a try.

**Some Linux Experience** – If you’ve used other Linux distributions (or a Linux-like operating system), you’ll probably find what you’re looking for in some of the more in-depth reference material available. Red Hat Software’s *Doctor Linux* and *Linux Man* are great for overall documentation, while O’Reilly’s more specialized titles are valuable when you need a lot of information on a particular subject.

**Old Timer** – If you’re a long-time Red Hat Linux user, you probably don’t need us telling you what documentation to read. Thanks for reading this far!

## 2.6 A Note About Kernel Drivers

During installation of Red Hat Linux, there are some limits placed on the filesystems and other drivers supported by the kernel. However, after installation there is support for all file systems available under Linux. At install time the modularized kernel has support for (E)IDE devices, (including ATAPI CD-ROM drives), SCSI adapters, and network cards. Additionally, all mice, SLIP, CSLIP, PPP, PLIP, FPU emulation, console selection, ELF, SysV IPC, IP forwarding, firewalling and accounting, reverse ARP, QIC tape and parallel printers, are supported.

**Please Note:** Because Red Hat Linux supports installation on many different types of hardware, many drivers (including those for SCSI adapters, network cards, and many CD-ROMs) are not built into the Linux kernel used during installation; rather, they are available as *modules* and loaded as you need them during the installation process. If necessary, you will have the chance to specify options for these modules at the time they are loaded, and in fact these drivers will ignore any options you specify for them at the `boot :` prompt.

After the installation is complete you may want to rebuild a kernel that includes support for your specific hardware configuration. See Section 6.1 on page 89 for

information on building a customized kernel.

## **2.7 One Last Note...**

Please read all of the installation instructions *before* starting; this will prepare you for any decisions you need to make and should eliminate potential surprises.



# 3

## Starting the Installation

This chapter explains how to start the Red Hat Linux installation process. We'll cover the following areas in this chapter:

- Getting familiar with the installation program's user interface.
- Starting the installation program.
- Selecting an installation method.

By the end of this chapter, the installation program will be running on your system, and the appropriate installation method will have been selected.

### 3.1 The Installation Program User Interface

The Red Hat Linux installation program uses a screen-based interface that includes most of the on-screen “widgets” commonly found on graphical user interfaces. You'll see:

- Windows
- Buttons
- Text regions (with or without scrollbars)
- Check boxes
- Text input boxes

We'll point out these features as we encounter them in the installation. Note that, because the installation program must run on a wide variety of computers, the user interface is character-based. In addition, some computers may not have a mouse, so the installation program uses keystrokes for getting around.

### 3.1.1 Using the Keyboard to Navigate

You can navigate around the installation dialogs using a simple set of keystrokes. In most dialog boxes there is a *cursor* or *highlight* which you can move using the arrow keys (← → ↑ ↓). By using **Tab**, and **Alt-Tab**, you can cycle forward or backward through each widget on the screen. In addition, summaries of available function keys are presented at the bottom of each screen.

To “press” a button, position the highlight over the button (using **Tab**, for instance) and press **Space** (or **Enter**). To select an item from a list of items, move the highlight to the item you wish to select and press **Enter**. To select an item with a *check box*, move the highlight to the check box and press **Space** to select an item. To deselect, press **Space** a second time.

Pressing **F12** accepts the current values and proceeds to the next dialog; it is usually equivalent to pressing the **OK** button.

**Please Note:** Unless a dialog box is waiting for your input, do not press any keys during the installation process – it may result in unpredictable behavior.

### 3.1.2 A Note About Virtual Consoles

There is more to the Red Hat Linux installation program than the dialog boxes it presents as it guides you through the installation process. In fact, the installation program makes several different kinds of diagnostic messages available to

you, in addition to giving you a way to enter commands from a shell prompt. It presents this information on five *virtual consoles* which you can switch between using a single keystroke. These virtual consoles can be very helpful if you encounter a problem while installing Red Hat Linux. Messages displayed on the install or system consoles can help pinpoint the problem. Please see Figure 3.1 for a listing of the virtual consoles, the keystrokes to switch to them, and their contents.

Console	Keystroke	Contents
1	<b>Alt-F1</b>	installation dialog
2	<b>Alt-F2</b>	shell prompt
3	<b>Alt-F3</b>	install log (messages from install program)
4	<b>Alt-F4</b>	system log (messages from kernel, etc.)
5	<b>Alt-F5</b>	other messages

**Figure 3.1:** Virtual Console Information

In general, there should be no reason to leave virtual console #1 unless you are attempting to diagnose installation problems. But if you are the curious type, feel free to look around.

## 3.2 Starting the Installation Program

Now it's time to start installing Red Hat Linux. To start the installation, it is first necessary to boot the installation program. Before we start, please make sure you have all the resources you'll need for the installation. If you've already read through Chapter 2 on page 15, and followed the instructions, you should be ready.

### ALPHA

If you haven't created your diskettes yet, please read Section D.4 on page 214, and create them now. After you've created the necessary diskettes, please read Section D.5 on page 215 for information on starting the installation.

### SPARC

If you haven't prepared for the installation yet, please read Sections C.2 on page 199 through C.5. Determine how you will boot the installation program, and issue the boot command that will start the installation.

### 3.2.1 Booting the Installation Program

To start installing Red Hat Linux, insert your boot diskette into your computer's first diskette drive and reboot. After a short delay, a screen containing the `boot :` prompt should appear. The screen contains information on a variety of boot options. Each boot option also has one or more help screens associated with it. To access a given help screen, press the appropriate function key as listed in the line at the bottom of the screen. You should keep two things in mind:

- The initial screen will automatically start the installation program if you take no action within the first minute. To disable this feature, press one of the help screen function keys.
- If you press a help screen function key, there will be a slight delay as the help screen is read from diskette.

Normally, you'll only need to press `Enter` to boot. Watch the boot messages to see whether the Linux kernel detects your hardware. If it does not properly detect your hardware, you may need to restart the installation in "expert" mode. Expert mode disables most hardware probing, and gives you the option of entering options for the drivers loaded during the installation. Expert mode can be entered using the following boot command:

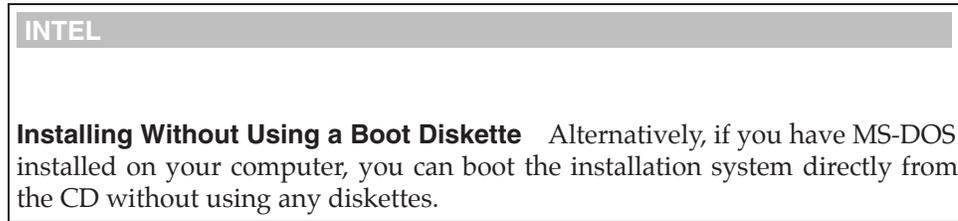
```
boot: expert
```

**Please Note:** The initial boot messages will not contain any references to SCSI or network cards. These devices are supported by modules that are loaded during the installation process.

Options can also be passed to the kernel. For example, to instruct the kernel to use all the RAM in a 128 MB system, enter:

```
boot: linux mem=128M
```

After entering any options, press `Enter` to boot using those options. If you do need to specify boot options to identify your hardware, please make note of them – they will be needed later.



To do this, use the following commands (assuming your CD is drive d:):

```
C:\> d:  
D:\> cd \dosutils  
D:\dosutils> autoboot.bat
```

Note that this method will not work if run in a DOS window – the `autoboot.bat` file must be executed with DOS as the only operating system. In other words, Windows cannot be running.

**Please Note:** The Red Hat Linux/Intel CD-ROM can also be booted by newer computers that support bootable CD-ROMs. Not all computers support this feature, so if yours can't boot from CD-ROM, you'll have to use a boot diskette (or autoboot from DOS) to get things started.

### 3.3 Beginning the Installation

After booting, the installation program begins by asking if you are using a color monitor (see Figure 3.2 on the next page). **Yes** is the default – you should see a cursor or highlight next to the **Yes** button. If you don't see a cursor or highlight press **Tab** once to select **No**, and then **Enter** to proceed.

Next, the installation program displays a welcome message. Press **Enter** to begin the installation. If you wish to abort the installation process at this time, simply eject the boot diskette now and reboot your machine.

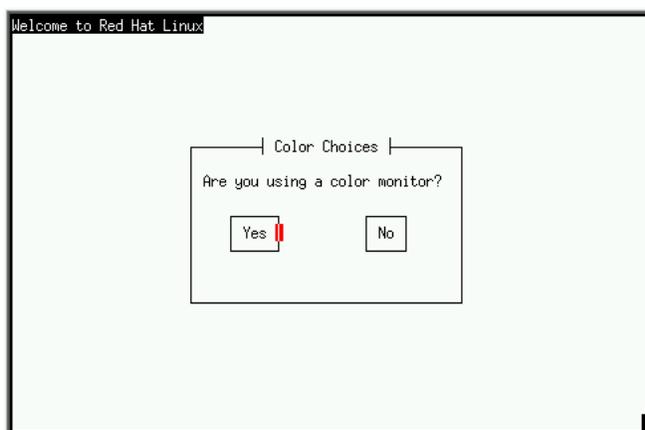


Figure 3.2: Beginning the Installation

### 3.3.1 Selecting a Keyboard Type

After the welcome dialog, the installation program gives you an opportunity to select a keyboard type (see Figure 3.3 on the facing page). Using the **↑** and **↓** keys, select the appropriate keyboard type for your Red Hat Linux system. Note the scroll bar to the right of the keyboard types – it indicates that there are more types than can be displayed at one time. You’ll be seeing scroll bars like this throughout the installation program.

After selecting the appropriate keyboard type, press **Enter**; the keyboard type you select will be loaded automatically both for the remainder of the installation process and each time you boot your Red Hat Linux system. If you wish to change your keyboard type after you have booted your Red Hat Linux system, you may use the `/usr/sbin/kbdconfig` command.

### 3.3.2 PCMCIA Support

Next, the installation program will probe your system to determine if your system requires PCMCIA (also known as PC Card) support. If a PCMCIA controller is found, you will be asked to insert the supplemental diskette – Select **OK** when you’ve done so.

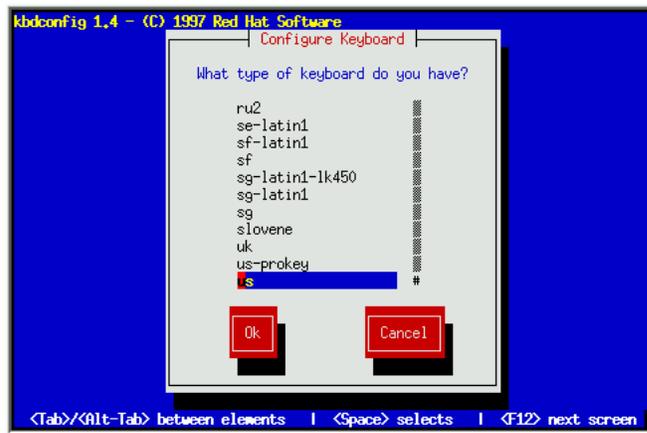


Figure 3.3: Selecting a Keyboard Type

The installation program will then display a progress bar as the supplemental diskette is loaded.

**Please Note:** If you are performing an installation in expert mode, you will be asked whether PCMCIA support is required. You must answer **Yes** if you are using a PCMCIA device for installation (for example, a PCMCIA ethernet card for installing via NFS, or a PCMCIA SCSI card for installing from a SCSI CD).

### 3.4 Selecting an Installation Method

Next, you are asked what type of installation method you wish to use (see Figure 3.4 on the next page). Highlight the appropriate choice and select **OK**, or press **Enter**. You can install Red Hat Linux via any of four basic methods (see Section 2.3 on page 21), some of which require the use of a supplemental diskette. To summarize, you can install Red Hat Linux from:

**CD-ROM** – If you have a CD-ROM drive and the Red Hat Linux CD-ROM. Does not require a supplemental diskette. Please refer to Section 3.4.1 on the next page to select the CD-ROM installation method.

**NFS** – If you are installing from an NFS server which is exporting the Red Hat

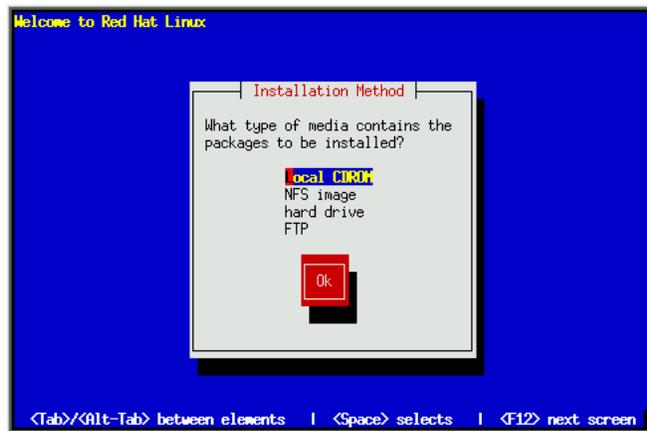


Figure 3.4: Selecting an Installation Method

Linux CD-ROM or a mirror of Red Hat Linux from the Red Hat FTP site. Does not require a supplemental diskette. Please refer to Section 3.4.2 on the facing page to select the NFS installation method.

**Hard Drive** – If you copied the Red Hat Linux files to a local hard drives. Requires a supplemental diskette. Please refer to Section 3.4.8 on page 47 to selection the hard drive installation method.

**FTP** – If you are installing directly from an FTP server. Requires a supplemental diskette. Please refer to Section 3.4.2 on the facing page to select the FTP installation method.

### 3.4.1 Installing From CD-ROM

If you are going to install Red Hat Linux from CD-ROM, select “CD-ROM”, and select **Ok**. The installation program will then probe your system, and attempt to identify your CD-ROM drive. If it cannot, you will be asked what type of CD-ROM you have. You can choose from the following types:

**IDE (ATAPI)** Most CD-ROM drives available for Intel-based PCs are IDE (also known as ATAPI). If this type is selected, the installation program will check all available IDE devices for a CD-ROM drive.

**SCSI** Select this if your CD-ROM is attached to a supported SCSI adapter; the installation program asks you to choose a SCSI driver. Choose the driver that most closely resembles your adapter. You may specify options for the driver if necessary; however, most drivers will detect your SCSI adapter automatically.

**Other** If your CD-ROM is neither an IDE nor a SCSI CD-ROM, it's an "other". Sound cards with proprietary CD-ROM interfaces are good examples of this CD-ROM type. The installation program presents a list of drivers for supported CD-ROMs – choose a driver and, if necessary, specify any driver options.

Once your CD-ROM drive has been identified, you will be asked to insert the Red Hat Linux CD-ROM into your CD-ROM drive. Select **Ok** when you have done so. After a short delay, the next dialog box will appear. Turn to Chapter 4 on page 51 to continue installing Red Hat Linux.

### 3.4.2 Installing Over the Network via NFS or FTP

NFS- and FTP-based installations are similar, in that they both require network-related information. Therefore, we'll discuss them together. Highlight the installation method you prefer, select **OK**, and press **[Space]**. If you're going to do an NFS installation, please skip ahead to 3.4.4 on page 45. If you're going to do an FTP installation, read on.

### 3.4.3 Installing via FTP

If you are going to install Red Hat Linux from an FTP site, highlight "FTP", and select **OK**. FTP installations require the supplemental diskette – when asked to insert it, do so, and select **OK**. The installation program will then display a progress bar as the supplemental diskette is loaded. Mark this place in the manual, because you'll be returning here later.

Next, turn to Chapter 4 on page 51, and follow the directions there, up to and including Section 4.4 on page 65. Then come back here.

Back so soon? At this point, you have performed the following steps:

- Selected an install or an upgrade.

- Loaded support for any SCSI devices on your computer.
- Set up disk partitions and mount points using either Disk Druid or `fdisk`.
- Identified available swap partitions, and initialized them.

Make sure you've bookmarked this spot, because you'll be coming back in a minute. Now skip ahead to Section 3.4.5 on the next page. This section will guide you through the necessary network configuration dialog boxes. See you in a minute...

OK, you're back. You should have entered all the necessary network information, and should now be looking at the "FTP Information" dialog box. Here's where you point the installation program at the FTP site of your choice (see Figure 3.5). Enter the name or IP address of the FTP site you are installing from, and the name of the directory there which contains the `RedHat` directory for your architecture. For example, if the FTP site contains the directory `/pub/mirrors/redhat/i386/RedHat`, enter `/pub/mirrors/redhat/i386`. If you are not using anonymous FTP, or if you need to use a proxy FTP server (if you're behind a firewall, for example), check the check box, and another dialog box will request the FTP account and proxy information.



Figure 3.5: Installing via FTP

If everything has been specified properly, you should see a message box indicating that `base/hdlist` is being retrieved. Now turn to Section 4.5 on page 66 to continue installing Red Hat Linux.

### 3.4.4 Installing via NFS

If you are going to install Red Hat Linux from an NFS-served filesystem, highlight “NFS image” and select **OK**.

### 3.4.5 Network Driver Configuration

Next, the installation program will probe your system and attempt to identify your network card. Most of the time, the driver can locate the card automatically. If it is not able to identify your network card, you’ll be asked to choose the driver that supports your network card and to specify any options necessary for the driver to locate and recognize it.

### 3.4.6 Configuring TCP/IP Networking

After the installation program has configured your network card, it presents several dialogs for configuring your system’s TCP/IP networking. Figure 3.6 contains example networking information.

Field	Example Value
IP Address	10.0.2.15
Netmask	255.255.255.0
Default Gateway	10.0.2.254
Primary Nameserver	10.0.2.1
Domain Name	redhat.com
Hostname	pooh.redhat.com

**Figure 3.6:** Sample Networking Information

Note that the information in figure 3.6 is a sample only! You should obtain the proper information from your network administrator.

The first dialog asks you for IP and other network addresses (see Figure 3.7 on the following page). Enter the **IP address** you are using during installation and press **(Enter)**. The installation program attempts to guess your **Netmask** based on your IP address; you may change the netmask if it is incorrect. Press **(Enter)**. The installation program guesses the **Default gateway** and **Primary nameserver**

addresses from your IP address and netmask; you may change them if they are incorrect.

If your local area network has a `bootp` server and you've added an entry for the computer you're installing Red Hat Linux on, you can check the **Use Bootp** check box.

In either case, choose **OK** to continue.

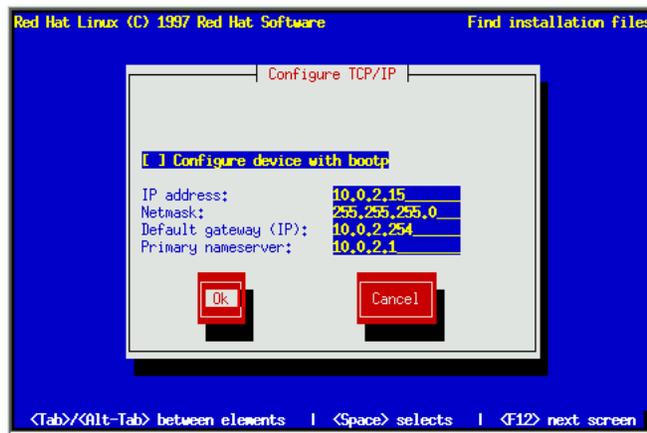


Figure 3.7: Configuring TCP/IP

After the first dialog box, you may see a second one. It will prompt you for a domain name, a hostname, and other networking information (see Figure 3.8 on the next page). Enter the **Domain name** for your system and press **Enter**; the installation program carries the domain name down to the **Host name** field. Enter the hostname you are using in front of the domain name to form a fully qualified domain name (FQDN). If your network has more than one nameserver, you may enter IP addresses for additional nameservers in the **Secondary nameserver** and **Tertiary nameserver** fields. Choose **OK** to continue.

If you're doing an FTP installation, head back Section 3.4.3 on page 43, and pick up where you left off. If you're doing an NFS installation, read on.



Figure 3.8: Configuring Networking

### 3.4.7 NFS Server Information

The next dialog requests information about the NFS server (see Figure 3.9 on the next page). Enter the name or IP address of your NFS server, and the name of the exported directory that contains the Red Hat Linux CD. For example, if the NFS server has the Red Hat Linux CD mounted on `/mnt/cdrom`, enter `/mnt/cdrom` in **Red Hat directory**. If the NFS server is exporting a mirror of the Red Hat Linux installation tree instead of a CD, enter the directory which contains the RedHat directory. For example, if your NFS server contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386`.

After a short delay, the next dialog box will appear. Turn to Chapter 4 on page 51 to continue installing Red Hat Linux.

### 3.4.8 Installing From a Hard Drive

If you are going to install Red Hat Linux from a locally-attached hard drive, highlight "hard drive" and select **OK**.

Before you started the installation program, you must first have copied all the necessary files to a partition on a locally-attached hard drive. If you haven't done this yet, please refer to Section 2.3.5 on page 25. Installing from a hard drive re-

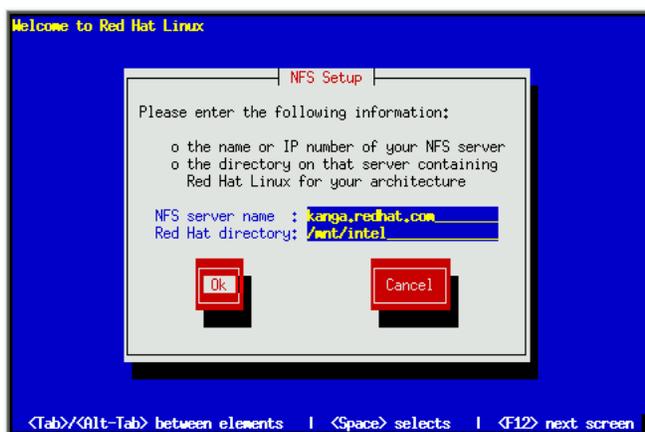


Figure 3.9: Installing via NFS

quires the supplemental diskette; when you are directed to, please insert it in your computer's diskette drive, and select **OK**. A progress bar will be displayed as the supplemental diskette is loaded.

Next, turn to Chapter 4 on page 51, and follow the directions there, up to and including Section 4.4 on page 65. Then come back here.

At this point, you have have performed the following steps:

- Selected an install or an upgrade.
- Loaded support for any SCSI devices on your computer.
- Set up disk partitions and mount points using either Disk Druid or `fdisk`.
- Identified available swap partitions, and initialized them.

At this point, a dialog box entitled "Select Partition" is displayed (see Figure 3.10 on the facing page). Enter the device name of the partition holding the RedHat directory tree. There is also a field labelled "Directory". If the RedHat directory is not in the root directory of that partition (for example, `/test/new/RedHat`), enter the path to the RedHat directory (in our example, `/test/new`).

If the installation program was unable to find the necessary files on the partition

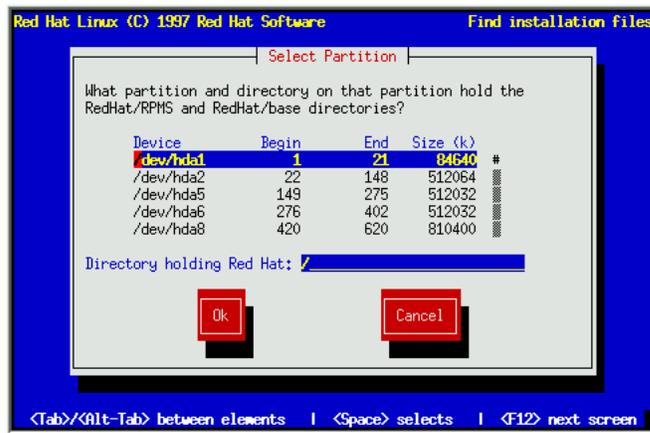


Figure 3.10: Selecting Partition for HD Install

and directory you've specified, you'll be returned to the "Select Partition" dialog box to make the necessary corrections.

If everything has been specified properly, you should see a message box indicating that the packages are being scanned. Now turn to Section 4.5 on page 66 to continue installing Red Hat Linux.



# 4

## Continuing the Installation

### 4.1 Upgrading or Installing

After you choose an installation method (and, if you're installing via the network, after you configure networking information), the installation program prompts you to either *install* or *upgrade* (see Figure 4.1 on the following page).

#### 4.1.1 Installing

You usually install Red Hat Linux on a clean disk partition or set of partitions, or over another installation of Linux.

**Please Note:** Installing Red Hat Linux over another installation of Linux (including Red Hat Linux) does *not* preserve any information from the prior installation. Make sure you save any important files!

If you wish to perform a full install, choose **Install**, and skip to section 4.2 on page 53.



Figure 4.1: Upgrading or Installing

### 4.1.2 Upgrading

The installation process for Red Hat Linux 5.0 includes the ability to upgrade from prior versions of Red Hat Linux (version 2.0 and later) which are based on RPM technology. Upgrading your system installs the modular 2.0.x kernel as well as updated versions of the packages which are currently installed on your machine. The upgrade process preserves existing configuration files by renaming them using a `.rpmsave` extension (e.g., `sendmail.cf.rpmsave`) and leaves a log telling what actions it took in `/tmp/upgradelog`. As software evolves, configuration file formats can change, so you should carefully compare your original configuration files to the new files before integrating your changes.

If you wish to upgrade your Red Hat Linux system, choose **Upgrade**.

**Please Note:** Some upgraded packages may require that other packages are also installed for proper operation. The upgrade procedure takes care of these *dependencies*, but it may need to install additional packages. You will be shown the names of the required packages, and you may then decide to install them or not.

## 4.2 SCSI Support

After you choose to perform an upgrade or a full install, the installation program will probe your system for SCSI adapters. In some cases, the installation program will ask you whether you have any SCSI adapters. If you choose **Yes**, the following dialog presents a list of SCSI drivers. Choose the driver that most closely resembles your SCSI adapter. The installation program then gives you an opportunity to specify options for the SCSI driver you selected; most SCSI drivers detect your hardware automatically, however.

## 4.3 Creating Partitions for Red Hat Linux

At this point, it's necessary to let the installation program know where it should install Red Hat Linux. This is done by defining *mount points* for one or more disk partitions in which Red Hat Linux will be installed. You may also need to create and/or delete partitions at this time.

The installation program then presents a dialog box that allows you to choose from two disk partitioning tools (see Figure 4.2 on the following page). The two choices you have are:

- **Disk Druid** – This is Red Hat Linux's new install-time disk management utility. It can create and delete disk partitions according to user-supplied requirements, in addition to managing mount points for each partition.



SPARC  
Disk Druid is not available on Red Hat Linux/SPARC.

- **fdisk** – This is the traditional Linux disk partitioning tool. While it is somewhat more flexible than Disk Druid, the downside is that `fdisk` assumes you have some experience with disk partitioning, and are comfortable with its somewhat terse user interface.

With the exception of certain esoteric situations, Disk Druid can handle the partitioning requirements for a typical Red Hat Linux installation.

Select the disk partitioning tool you'd like to use, and press **Enter**. If you choose Disk Druid, continue reading. If you'd rather use `fdisk`, please turn to Section 4.3.2 on page 61.

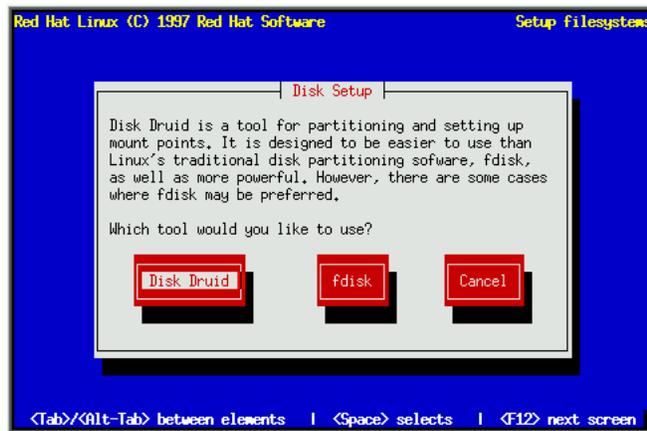


Figure 4.2: Selecting Disk Setup Method

### 4.3.1 Using Disk Druid

If you selected Disk Druid, you will be presented with a screen that looks like figure 4.3 on the facing page. While it may look overwhelming at first, it really isn't. Let's go over each of Disk Druid's three sections, and you'll see that there's nothing to be scared of.

#### The "Current Disk Partitions" Section

Each line in the "Current Disk Partitions" section represents a disk partition. You'll note that this section has a scroll bar to the right, which means that there might be more partitions than can be displayed at one time. If you use the  $\uparrow$  and  $\downarrow$  keys, you can see if there are any additional partitions there. Each line in this section has five different fields:

**Mount Point** – This field indicates where the partition will be mounted when Red Hat Linux is installed and running.

**Device** – This field displays the partitions device name.

**Requested** – The "Requested" field shows the minimum size requested when the partition was defined.

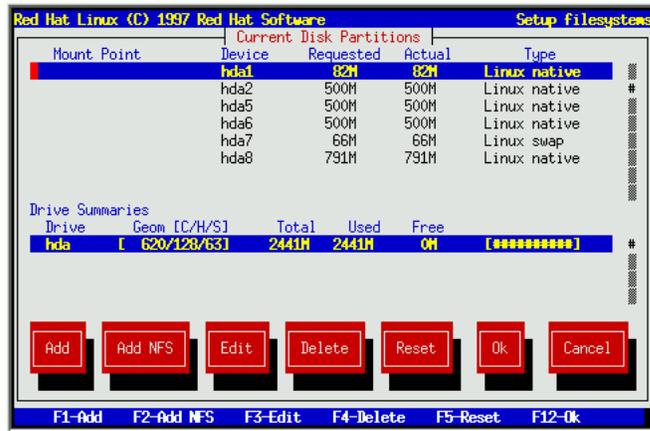


Figure 4.3: Disk Druid Main Screen

**Actual** – The “Actual” field shows the space currently allocated to the partition.

**Type** – This field shows the partition’s type.

**Another Type of Partition** As you scroll through the “Current Disk Partitions” section, you might see an “Unallocated Requested Partitions” title bar, followed by one or more partitions. As the title implies, these are partitions that have been requested but, for one reason or another, have not been allocated. A common reason for having an unallocated partition is a lack of sufficient free space for the partition. In any case, the reason the partition is unallocated will be displayed after the partition’s mount point.

### The “Drive Summaries” Section

Each line in the “Drive Summaries” section represents a hard disk on your system. Each line has the following fields:

**Drive** – This field shows the hard disk’s device name.

**Geom [C/H/S]** – This field shows the hard disk’s *geometry*. The geometry consists of three numbers representing the number of cylinders, heads, and sectors as reported by the hard disk.

**Total** – The “Total” field shows the total available space on the hard disk.

**Used** – This field shows how much of the hard disk’s space is currently allocated to partitions.

**Free** – The “Free” field shows how much of the hard disk’s space is still unallocated.

**Bar Graph** – This field presents a visual representation of the space currently used on the hard disk. The more pound signs there are between the square braces, the less free space there is. In Figure 4.3 on the page before, the bar graph shows no free space.

### Disk Druid’s Buttons

These buttons control Disk Druid’s actions. They are used to add and delete partitions, and to change partition attributes. In addition, there are buttons that are used to accept the changes you’ve made, discard those changes, or exit Disk Druid entirely. Let’s take a look at each button in order.

**Add** – The “Add” button is used to request a new partition. When selected, a dialog box will appear containing fields that must be filled in.

**Add NFS** – This button is used to add a read-only NFS-served filesystem to the set of mount points on your Red Hat Linux system. When selected, a dialog box will appear containing fields that must be filled in.

**Edit** – The “Edit” button is used to modify attributes of the partition currently highlighted in the “Current Disk Partitions” section. Selecting this button will cause a dialog box to appear. Some or all of the fields in the “Edit Partition” dialog box may be changed, depending on whether the partition information has already been written to disk or not.

**Delete** – The “Delete” button is used to delete the partition currently highlighted in the “Current Disk Partitions” section. Selecting this button will cause a dialog box to appear asking you to confirm the deletion.

**Reset** – This button is used to discard all changes you may have made while in Disk Druid, and return the list of partitions to those read from the partition table(s) on your hard disk(s). When selected, you’ll be asked to confirm whether you want the changes discarded or not. Note that any mount points you’ve specified will be lost, and will need to be reentered.

**Ok** – The “Ok” button causes any changes made to your systems partitions to be written to disk. You will be asked to confirm your changes before Disk Druid rewrites your hard disk partition table(s). In addition, any mount points you’ve defined are passed to the installation program, and will eventually be used by your Red Hat Linux system to define the filesystem layout.

**Cancel** – This button causes Disk Druid to abort without saving any changes you’ve made. When this button is selected, the installation program will present a dialog box that permits you to choose what action you’d like to take next.

**Please Note:** You will need to dedicate at least one partition to Red Hat Linux, and optionally more. This is discussed more completely in Section 2.4.4 on page 30.

Now let’s see how Disk Druid is used to set up partitions for your Red Hat Linux system.

### Adding a Partition

To Add a new partition, select the **Add** button, and press **Space** or **Enter**. A dialog box entitled “Edit New Partition” will appear (see Figure 4.4). It contains the following fields:

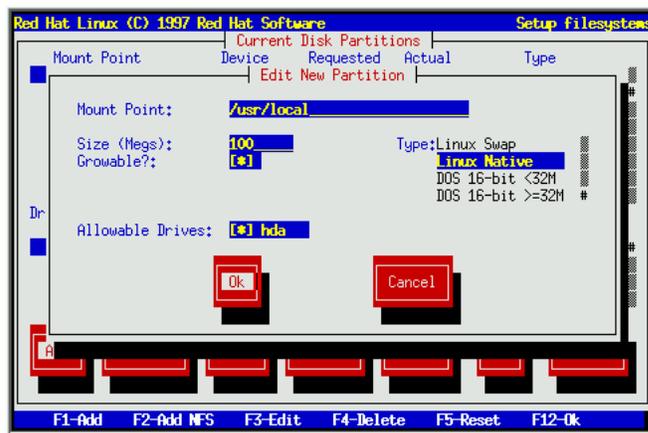


Figure 4.4: Creating a New Partition

- **Mount Point** – Highlight this field, and enter the partition’s mount point. For example, if this partition should be the root partition, enter `/`, enter `/usr` for the `usr` partition, and so on.
- **Size (Megs)** – In this field, enter the size (in megabytes) of the partition. Note that this field starts with a “1” in it, meaning that unless you change it, you’ll end up with a 1 MB partition. Delete it using the **[Backspace]** key, and enter the desired partition size.
- **Growable?** – This check box indicates whether the size you entered in the previous field is to be considered the partition’s exact size, or its minimum size. Press **[Space]** to check and uncheck the box. When checked, the partition will grow to fill all available space on the hard disk. In this case, the partition’s size will expand and contract as other partitions are modified.
- **Type** – This field contains a list of different partition types. Select the appropriate partition type by using the **[↑]** and **[↓]** keys.
- **Allowable Drives** – This field contains a list of the hard disks installed on your system, with a check box for each. If a hard disk’s box is *not* checked, then this partition will never be created on that hard disk. By using different check box settings, you can direct Disk Druid to place partitions as you see fit, or let Disk Druid decide where partitions should go.
- **Ok** – Select this button and press **[Space]** when you are satisfied with the partition’s settings, and wish to create it.
- **Cancel** – Select this button and press **[Space]** when you don’t want to create the partition.

**Problems When Adding a Partition** If your attempt to add a partition and Disk Druid can’t carry out your request, you’ll see a dialog box like the one in figure 4.5 on the facing page. In the box are listed any partitions that are currently unallocated, along with the reason they could not be allocated. Select the **Ok** button, and press **[Space]** to continue. Note that the unallocated partition(s) are also displayed on Disk Druid’s main screen (though you may have to scroll the “Current Disk Partitions” section to see them).

### Deleting a Partition

To delete a partition, highlight the partition in the “Current Disk Partitions” section, select the **Delete** button, and press **[Space]**. You will be asked to confirm the

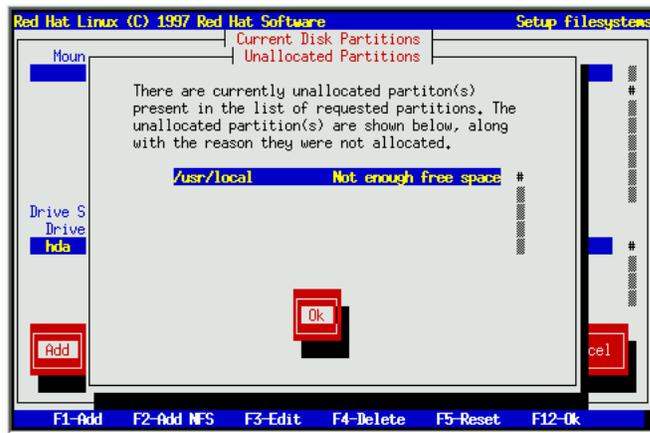


Figure 4.5: Unallocated Partition Warning

deletion.

### Editing a Partition

To change a partition's settings, highlight the partition in the "Current Disk Partitions" section, select the **Edit** button, and press **Space**. You will be presented with a dialog box very similar to the one shown in Figure 4.4 on page 57. Make the appropriate changes, select **Ok**, and press **Space**.

**Please Note:** If the partition already existed on your hard disk, you will only be able to change the partition's mount point. If you want to make any other changes, you will need to delete the partition and recreate it.

### Adding an NFS Mount

To add a read-only NFS-served filesystem, select the **Add NFS** button, and press **Space**. If you have not selected a network-related installation method, you will be presented with several dialog boxes concerning network configuration (Turn back to Section 3.4.5 on page 45 for more information). Fill them in appropriately. You will then see a dialog box entitled, "Edit Network Mount Point" (see Figure 4.6 on the next page. In this dialog box you will need to enter the NFS server name, the



press **[Space]**. Then turn to Section 4.4 on page 65.

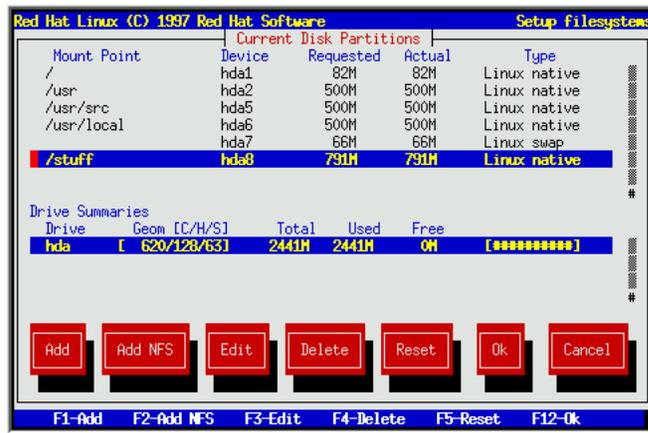


Figure 4.7: Partitions and Mount Points Defined

### 4.3.2 Using fdisk

If you'd rather use `fdisk` to manage partitions, this is the section for you. Once you've selected `fdisk`, you'll be presented with a dialog box entitled "Partition Disk" (see Figure 4.8 on the following page). In this box is a list of every disk on your computer. Move the highlight to the disk you'd like to partition, select **Edit**, and press **[Space]**. You will then enter `fdisk` and can partition the disk you selected. Repeat this process for each disk you want to partition. When you're done, select "Done".

#### An Overview of fdisk

`fdisk` includes online help which is terse but useful. Here are a few tips:

- The command for help is `m`.
- To list the current partition table, use the `p` command (see Figure 4.9 on page 63).



Figure 4.8: Selecting a Disk for Partitioning

- To add a new partition, use `n`.
- Linux `fdisk` creates partitions of type `Linux native` by default. When you create a swap partition, don't forget to change it to type `Linux swap` using the `t` command. The value for the `Linux swap` type is 82. For other partition types, use the `l` command to see a list of partition types and values.
- Linux allows up to four (4) partitions on one disk. If you wish to create more than that, one of the four may be an *extended* partition, which acts as a container for one or more *logical* partitions. Since it acts as a container, the extended partition must be at least as large as the total size of all the logical partitions it is to contain.
- It's a good idea to write down which partitions (e.g., `/dev/hda2`) are meant for which filesystems (e.g., `/usr`) as you create each one.
- **Please Note:** None of the changes you make take effect until you save them and exit `fdisk` using the `w` command. You may quit `fdisk` at any time without saving changes by using the `q` command.

```
This is the fdisk program for partitioning your drive. It is running
on /dev/hda.

Command (m for help): p

Disk /tmp/hda: 128 heads, 63 sectors, 620 cylinders
Units = cylinders of 8064 * 512 bytes

   Device Boot   Begin    Start    End  Blocks  Id System
/tmp/hda1             1         1     21   84640+  83 Linux native
/tmp/hda2             22        22    148  512064  83 Linux native
/tmp/hda3            149       149    620 1903104   5 Extended
/tmp/hda5            149       149    275  512032+  83 Linux native
/tmp/hda6            276       276    402  512032+  83 Linux native
/tmp/hda7            403       403    419  68512+  82 Linux swap
/tmp/hda8            420       420    620  810400+  83 Linux native

Command (m for help): █
```

Figure 4.9: Sample Output From `fdisk`

### Numbering Partitions

Linux refers to disk partitions using a combination of letters and numbers which may be confusing to inexperienced users. Here is a summary:

**First Two Letters** – The first two letters of the partition name indicate the type of device on which the partition resides. You'll normally see either `hd` (for IDE disks), or `sd` (for SCSI disks).

**The Next Letter** – This letter indicates which device the partition is on. For example, `/dev/hda` (the first IDE hard disk) or `/dev/sdb` (the second SCSI disk).

**The number** denotes the partition. The first four (primary or extended) partitions are numbered 1 through 4. Logical partitions start at 5. E.g., `/dev/hda3` is the third primary or extended partition on the first IDE hard disk; `/dev/sdb6` is the second logical partition on the second SCSI hard disk.

### Changing the Partition Table

When you are finished partitioning your disks, press **Done**; you may see a message indicating that the installation program needs to reboot. This is a normal occurrence after changing a disk's partition data; it usually happens if you created, changed, or deleted any extended partitions. After you press **OK**, your machine will reboot. Follow the same installation steps you did up until **Partitioning Disks**; then simply choose **Done**.

### 4.3.3 Filesystem Configuration

Since Red Hat Linux may be distributed across several disk partitions, you need to choose which disk partitions go with which *filesystem*.

#### Selecting the Root Partition

If you are performing a full install, the installation program presents a list of the Linux partitions it can find and prompts you to choose a *root partition* (see Figure 4.10 on the facing page. The root partition (or *root filesystem*) is the one which holds everything necessary to boot your Red Hat Linux system; it is mounted on `/` (the *root directory*) when your system boots. Select the root partition you wish to use and press **Enter**.

If you are performing an upgrade, the installation program tries to find your root partition automatically; if it finds one, it goes on to the next step.

#### Other Partitions

The next dialog box contains a list of all disk partitions with filesystems readable by Red Hat Linux, including partitions for MS-DOS or Windows. This gives you the opportunity to assign these partitions to different parts of your Red Hat Linux filesystem. The partitions you assign will be automatically mounted when your Red Hat Linux system boots. Select the partition you wish to assign and press **Enter** (or choose **Edit**); then enter the *mount point* for that partition, e.g., `/usr` (see Figure 4.11 on page 66).

Additionally, Red Hat Linux allows you to mount read-only NFS volumes when your system boots; this allows directory trees to be shared across a network. Choose

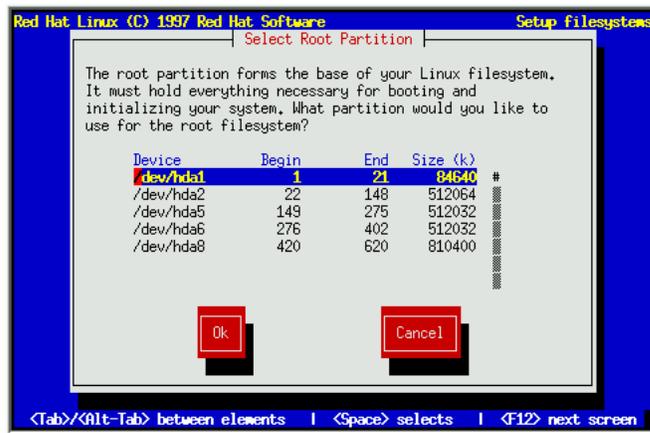


Figure 4.10: Selecting Root Partition

**Add NFS.** If you have not selected a network-related installation method, you will be presented with several dialog boxes concerning network configuration (Turn back to Section 3.4.6 on page 45 for more information). Fill them in appropriately. You will then see a dialog box entitled “Edit Network Mount Point”. Enter the NFS server’s hostname, the path to the NFS volume, and the local mount point for that volume (see Figure 4.12 on page 67).

## 4.4 Initializing Swap Space

After you’ve created partitions for Red Hat Linux, the installation program looks for swap partitions (see Figure 4.13 on page 68). If it finds any, it asks whether you want to initialize them. Select the partition(s) you wish to initialize as swap space using **[Space]**; if you wish to check the partitions for bad blocks, make sure the **Check for bad blocks during format** box is checked. Choose **OK**, and press **[Space]**.

If the installation program can’t find a swap partition and you’re sure one exists, make sure you have set the partition type to **Linux swap**; see Section 4.3 on page 53 for information on how this is done with Disk Druid or `fdisk`.

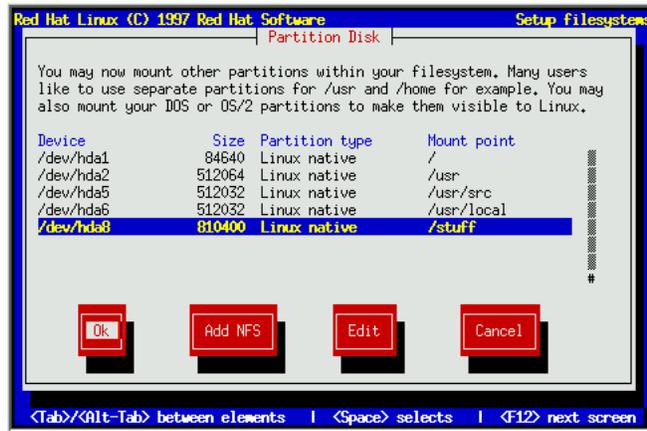


Figure 4.11: Filesystem Configuration

## 4.5 Formatting Partitions

The next dialog box presents a list of partitions to format (see Figure 4.14 on page 69). All newly created partitions should be formatted. In addition, any already-existing partitions that contain old data you no longer need should be formatted. However, partitions such as `/home` or `/usr/local` must not be formatted if they contain data you wish to keep. Select each partition to format and press **Space**. If you wish to check for bad blocks while formatting each filesystem, select **Check for bad blocks during format**. Select **OK**, and press **Space**.

## 4.6 Selecting and Installing Packages

After your partitions have been configured and formatted, you are ready to select packages for installation. You can select *components*, which group packages together according to function, individual packages, or a combination of the two.

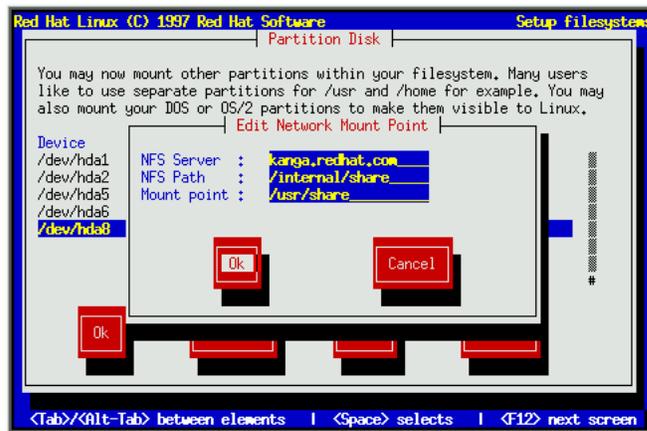


Figure 4.12: Adding an NFS Mount

### 4.6.1 Selecting Components

Components group packages together according to the functionality they provide. For example, **C Development**, **Networked Workstation**, or **Web Server**. Select each component you wish to install and press **Space**. Selecting **Everything** (which can be found at the end of the component list) installs all packages included with Red Hat Linux (see Figure 4.15 on page 70).

If you wish to select or deselect individual packages, choose **Select individual packages**.

### 4.6.2 Selecting Individual Packages

After selecting the components you wish to install, you may select or deselect individual packages. The installation program presents a list of the package groups available; select a group to examine and press **Enter**. The installation program presents a list of the packages in that group, which you may select or deselect using **Space** (see Figure 4.16 on page 71). You may view a detailed description of a package by pressing **F1**. When you are finished selecting individual packages, press **OK** in the **Select Group** dialog box.

**Please Note:** Some packages (such as the kernel and certain libraries) are required



Figure 4.13: Initializing Swap Space

for every Red Hat Linux system and are not available to select or deselect.

### 4.6.3 Package Dependencies

Many software packages, in order to work correctly, depend on other software packages or libraries that must be installed on your system. For example, many of the graphical Red Hat system administration tools require the `python` and `pythonlib` packages. To make sure your system has all the packages it needs in order to be fully functional, Red Hat Linux checks package *dependencies* each time you install or remove software packages.

After you have finished selecting packages to install, the installation program checks the list of selected packages for dependencies. If any package requires another package which you have not selected to install, the program presents a list of these *unresolved dependencies* and gives you the opportunity to resolve them (see Figure 4.17 on page 71). If you simply press **OK**, the program will resolve them automatically by adding all required packages to the list of selected packages.



Figure 4.14: Formatting Partitions

#### 4.6.4 Package Installation

After all package dependencies are resolved, the installation program presents a dialog box telling you that a log file containing a list of all packages installed will be written to `/tmp/install.log` on your Red Hat Linux system. Select **Ok** and press **Space** to continue.

At this point, the installation program will format every partition you selected for formatting. This can take several minutes (and will take even longer if you directed the installation program to check for bad blocks).

Once all partitions have been formatted, the installation program starts to install packages. A window entitled "Install Status" is displayed with the following information:

**Package** – The name of the package currently being installed.

**Size** – The size of the package (in kilobytes).

**Summary** – A short description of the package.

**Package Installation Progress Bar** – A bar showing how complete the current package installation is.

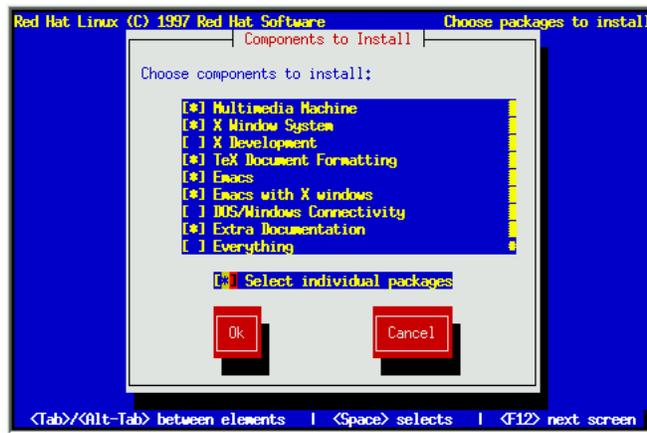


Figure 4.15: Selecting System Components

**Statistics Section** – This section has three rows labeled “Total”, “Completed”, and “Remaining”. As you might guess, these rows contain statistics on the total number of packages that will be installed, statistics on the number of packages that have been completely installed, and statistics on the packages that have not yet been installed. The information tracked on these three rows includes:

**Packages** – The number of packages.

**Bytes** – The size.

**Time** – The amount of time.

**Overall Progress Bar** – This bar changes color showing how close to completion the entire installation is.

If you’re doing an FTP installation, a message box will pop up as each package is retrieved from the FTP site.

At this point there’s nothing left for you to do until all the packages have been installed. How quickly this happens depends on the number of packages you’ve selected, and your computer’s speed. Once all the packages have been installed, please turn to the next chapter to finish your installation of Red Hat Linux.



Figure 4.16: Selecting Packages



Figure 4.17: Unresolved Dependencies



# 5

## Finishing the Installation

### 5.1 Configuring a Mouse

Next, the installation program will probe your system and try to find a mouse. If one is found, a dialog box is displayed showing the port on which the mouse was found. Press **Space** to continue.

If necessary you'll be asked to give additional information about the mouse, such as the mouse's protocol and number of buttons.

If you wish to change your mouse configuration after you have booted your Red Hat Linux system, you may use the `/usr/sbin/mouseconfig` command.

### 5.2 Configuring X Windows

After setting up your mouse, if you installed the X Windows packages, you will have the opportunity to configure your X server. If you did not choose to install the X Window System, you may skip to Section 5.3 on page 75.

### 5.2.1 Configuring the Metro-X Server

If you are installing from an Official Red Hat Linux/Intel CD, you may install one copy of MetroLink's Metro-X server, a high-quality commercial X server for Linux. Please see the MetroLink license (paper copy) accompanying your Red Hat Linux CD for license information. If you choose to use Metro-X, the install program launches the Metro-X configuration utility `configX`; please see Section 6.3 on page 96 for more information on how to use `configX`.

### 5.2.2 Configuring an XFree86 Server

If you wish to use XFree86, the installation program launches the `Xconfigurator` utility, which is described in detail in Section 6.2 on page 94.

`Xconfigurator` first probes your system in an attempt to determine what type of video card you have. Failing that, `Xconfigurator` will present a list of video cards. Select your video card from the list and press **Enter**. If your video card does not appear on the list, XFree86 may not support it. However, if you have technical knowledge about your card, you may choose **Unlisted Card** and attempt to configure it anyway.

Once you have selected your video card, the installation program installs the appropriate XFree86 server, and `Xconfigurator` presents a list of monitors. If your monitor appears on the list, select it and press **Enter**. Otherwise, select **Custom**. If you do select **Custom**, `Xconfigurator` prompts you to select the horizontal sync range and vertical sync range of your monitor (these values are generally available in the documentation which accompanies your monitor, or from your monitor's vendor or manufacturer).

**Please Note:** It is not recommended to select a monitor "similar" to your monitor unless you are certain that the monitor you are selecting does not exceed the capabilities of your monitor. If you do so, it is possible you may overclock your monitor and damage or destroy it.

Next, `Xconfigurator` prompts you for the amount of video memory installed on your video card. If you are not sure, please consult the documentation accompanying your video card. It will not damage your video card by choosing more memory than is available, but the XFree86 server may not start correctly if you do.

If the video card you selected might have a video clockchip, `Xconfigurator` presents a list of clockchips. The recommended choice is **No Clockchip Setting**,

since XFree86 can automatically detect the proper clockchip in most cases.

Finally, `xconfigurator` prompts you to select the video modes you wish to use; select one or more modes by pressing `[Space]`. `xconfigurator` then writes a configuration file containing all of your choices to `/etc/X11/XF86Config`.

## 5.3 Configuring Networking

Next, the installation program gives you an opportunity to configure (or reconfigure) networking. If you are installing from CD-ROM or from a local hard disk, the installation program asks if you want to configure networking. If you choose **No**, your Red Hat Linux system will be a standalone workstation. If you choose **Yes**, you may configure networking as described below.

If you are installing Red Hat Linux via NFS or FTP, you have already entered temporary networking information. The install program offers you three choices (see Figure 5.1):



Figure 5.1: Network Configuration Options

- **Keep this setup** – Keeps the network configuration you used during the installation. All the networking information you entered previously becomes part of your system's permanent configuration.

- **Reconfigure network now** – The installation program presents the network configuration dialogs in Section 3.4.6 on page 45. The values you used during installation will be filled in as defaults. Choose this if your system will be installed on a network other than the one you used to install Red Hat Linux.
- **Don't setup networking** – Don't set up networking at all. Your system will not have networking configured. Choose this if you installed your system over a network, but it will be used as a standalone workstation.

### 5.3.1 Network Configuration Dialogs

If you elected to configure networking at this time, you will be presented with a series of dialog boxes. Please turn to Section 3.4.6 on page 45 for more information.

## 5.4 Configuring the Clock

Next, the installation program presents a dialog to help you configure your Red Hat Linux system's timezone (see Figure 5.2 on the facing page).

If you wish to set the hardware (CMOS) clock to GMT (Greenwich Mean Time, also known as UTC, or Coordinated Universal Time), select **Hardware clock set to GMT**. Setting your hardware clock to GMT means your Red Hat Linux system will properly handle daylight savings time, if your timezone uses it. Most networks use GMT.

Select the timezone your system will be operating in from the list, and press **Enter**.

If you wish to change your timezone configuration after you have booted your Red Hat Linux system, you may use the `/usr/sbin/timeconfig` command.

## 5.5 Selecting Services for Start on Reboot

Next you'll see a dialog box entitled "Services" (see Figure 5.3 on page 78). Displayed in this box is a list of services with a check box by each. Scroll through this list, and check every service that you would like automatically started every time your Red Hat Linux system boots. If you're not sure what a particular service is,

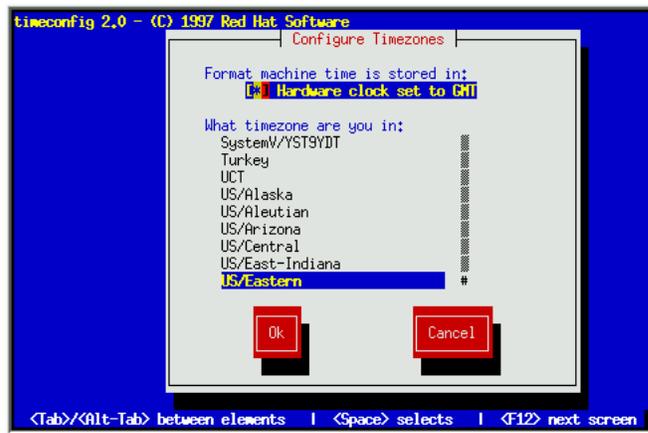


Figure 5.2: Configuring Timezones

move the highlight to it and press **F1**. You'll then get a brief description of the service.

Note that you can run `/usr/sbin/ntsysv` or `sbin/chkconfig` after the installation to change which services automatically start on reboot.

## 5.6 Configuring a Printer

After you have set up networking, the installation program asks whether you would like to configure a printer. If you choose **Yes**, a dialog box will ask you to indicate how the printer is connected to your computer (see Figure 5.4 on page 79).

Here is a brief description of the three types of printer connections available:

**Local** – This printer is directly connected to your computer.

**Remote lpd** – This printer is connected to your local area network (either through another computer, or directly), and is capable of communicating via `lpr` / `lpd`.

**LAN Manager** – This printer is connected to another computer which shares the printer via LAN Manager (or SMB) networking.



Figure 5.3: Selecting Services

After selecting a printer type, you'll be presented with a dialog box entitled "Standard Printer Options" (see Figure 5.5 on page 80). Enter the name of the queue and the spool directory you'd like to use, or accept the default information.

The dialog box you'll see next depends on the printer connection type you selected. Turn to the section that corresponds to your printer connection type:

**Local** – Section 5.6.1.

**Remote lpd** – Section 5.6.2 on the next page.

**LAN Manager** – Section 5.6.3 on the facing page.

### 5.6.1 Locally Attached Printers

If you selected "Local" as your printer's connection type, you'll see a dialog box similar to the one in Figure 5.6 on page 81.

Enter the printer device name in the field provided. As a convenience, the installation program attempts to determine which printer ports are available on your computer. Select **Next**, and press **Space**. Now turn to Section 5.6.4 on page 80 to continue.

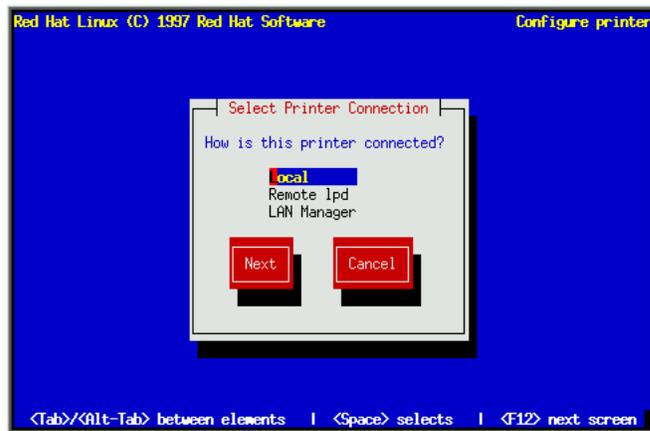


Figure 5.4: Selecting Printer Type

### 5.6.2 Remote lpd Printers

If you selected “Remote lpd” as your printer’s connection type, you’ll see a dialog box similar to the one in Figure 5.7 on page 82.

Enter the name of the computer to which the printer is directly connected in the “Remote hostname” field. The name of the queue *on the remote computer* that is associated with the remote printer goes in the “remote queue” field. Select **Next**, and press **Space**. Now turn to Section 5.6.4 on the following page to continue.

### 5.6.3 LAN Manager Printers

If you selected “LAN Manager” as your printer’s connection type, you’ll see a dialog box similar to the one in Figure 5.8 on page 83.

Enter the necessary information in the fields provided. Select **Next**, and press **Space**. Now turn to Section 5.6.4 on the following page to continue.



Figure 5.5: Standard Printer Options

#### 5.6.4 Finalizing Printer Setup

Next, you'll see a dialog box entitled "Configure Printer" (see Figure 5.9 on page 84). Select the printer type that most closely matches your printer. Select **Next**, and press **Space** to continue.

After selecting the printer type, you will see a dialog box similar to the one in Figure 5.10 on page 85. Set the paper size and resolution appropriately. The **Fix stair-stepping of text** check box should be checked if your printer does not automatically perform a carriage return after each line.

Finally, you'll see a dialog box that contains all the information pertaining to your printer (see Figure 5.11 on page 86). Verify that the information is correct. If everything looks OK, select **Done**. If you need to make changes, select **Edit**. You can also select **Cancel** if you'd rather not configure a printer at this time.

If you select **Done**, you will be given the option to configure another printer, or you may continue with the installation.



Figure 5.6: Local Printer Device

## 5.7 Setting a Root Password

The installation program will next prompt you to set a *root password* for your system (see Figure 5.12 on page 86). You'll use the root password to log into your Red Hat Linux system for the first time.

The root password must be at least six characters long; the password you type is not echoed to the screen. You must enter the password twice; if the two passwords do not match, the installation program will ask you to enter them again.

You ought to make the root password something you can remember, but not something that is easy for someone else to guess. Your name, your phone number, `qwerty`, `password`, `root`, `123456`, and `anteater` are all examples of poor passwords. Good passwords mix numerals with upper and lower case letters and do not contain dictionary words: `Aard387vark` or `420BMttNT`, for example. Remember that the password is case-sensitive. Write down this password and keep it in a secure place.

**Please Note:** The *root* user (also known as the *superuser*) has complete access to the entire system; for this reason, logging in as the root user is best done only to perform system maintenance or administration, so that critical system files are not inadvertently changed, moved, or deleted. Please see Section 9.1 on page 132 for instructions on how to add a user account for yourself after you reboot your

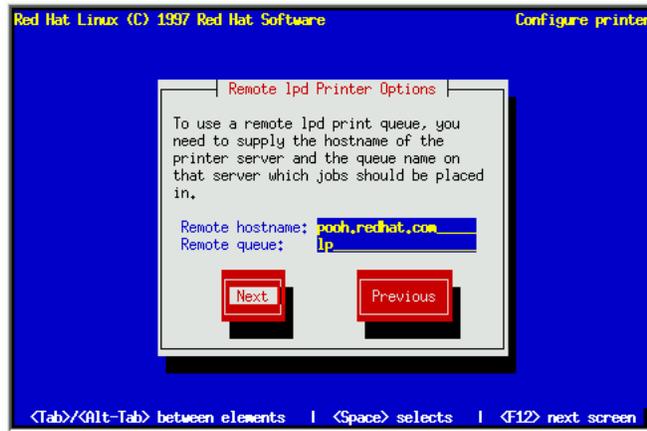


Figure 5.7: Remote lpd Printer Options

system.

## 5.8 Installing LILO

In order to be able to boot your Red Hat Linux system, you usually need to install LILO (the Linux LOader). You may install LILO in one of two places:

**The Master Boot Record (MBR)** is the recommended place to install LILO, unless another operating system loader (e.g., System Commander or OS/2's Boot Manager) is already installed there. When your machine boots, LILO starts and presents the `boot :` prompt; you can boot Red Hat Linux or any other operating system you configure LILO to boot (see below).

**The first sector of your root partition** is recommended if you are already using another boot loader on your system (such as OS/2's Boot Manager); then you can setup that boot loader to start LILO and boot Red Hat Linux.

A dialog box will appear that will let you select the type of LILO installation you desire (see Figure 5.13 on page 87). Select the location you wish to install LILO and press **OK**. If you do not wish to install LILO, press **Skip**.

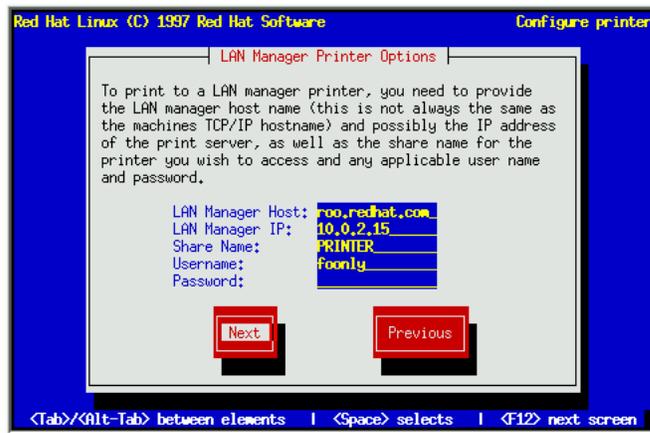


Figure 5.8: LAN Manager Printer Options

**Please Note:** If you choose **Skip**, you will not be able to boot your Red Hat Linux system directly, and will need to use another boot method (such as a boot diskette). Use this option at your own risk!

### 5.8.1 Adding Options to the LILO Boot Command Line

Finally, the installation program will ask if you wish to add default options to the LILO boot command (see Figure 5.14 on page 87). Any options you enter will be passed to the Linux kernel every time it boots. If you have an LBA drive, check **Use linear mode**. Select **OK** and press **Space** when finished.

### 5.8.2 Alternatives to LILO

If you do not wish to use LILO to boot your Red Hat Linux system, there are a few alternatives:

**LOADLIN** can load Linux from MS-DOS; unfortunately, it requires a copy of the Linux kernel (and an initial ram disk, if you have a SCSI adapter) to be available on an MS-DOS partition. The only way to accomplish this is to boot your Red Hat Linux system using some other method (e.g., from LILO on a



Figure 5.9: Configure Printer

diskette) and then copy the kernel an MS-DOS partition. LOADLIN is available from

`ftp://sunsite.unc.edu/pub/Linux/system/boot/dualboot/` and sunsite's various mirror sites.

**SYSLINUX** is an MS-DOS program very similar to LOADLIN; it is also available from `ftp://sunsite.unc.edu/pub/Linux/system/boot/dualboot/` and sunsite's various mirror sites.

**Some commercial bootloaders**, such as System Commander, are able to boot Linux (but may still require LILO to be installed in your Linux root partition).

## 5.9 Finishing Up...

After you have completed LILO installation, the installation program will reboot your system. Don't forget to remove any diskette that might be in the diskette drive.

At the boot: prompt, press **Enter** and watch your Red Hat Linux system boot. When it presents the login: prompt, enter `root`. At the Password: prompt, type the root password you set for your machine in Section 5.7 on page 81.

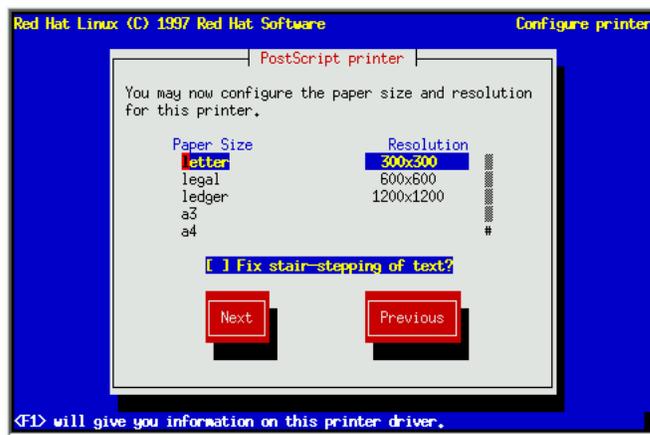


Figure 5.10: Printer Settings

Congratulations! Your Red Hat Linux installation is complete!

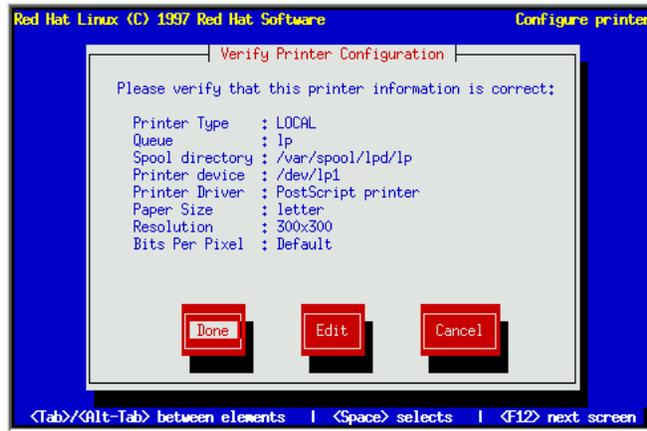


Figure 5.11: Verifying Printer Information



Figure 5.12: Root Password



Figure 5.13: Installing LILO



Figure 5.14: LILO options



# 6

## Post Installation Configuration

After installation your system is fully functional, but you may still want to do some additional configuration. This chapter describes some of the various system configuration tasks you may need to do.

### 6.1 Building a Custom Kernel

With the introduction of modularization in the Linux 2.0.x kernel there have been some significant changes in building customized kernels. In the past you were required to compile support into your kernel if you wanted to access a particular hardware or filesystem component. For some hardware configurations the size of the kernel could quickly reach a critical level. To require ready support for items that were only occasionally used was an inefficient use of system resources. With the capabilities of the 2.0.x kernel, if there are certain hardware components or filesystems that are used infrequently, driver modules for them can be loaded on demand. For information on handling kernel modules see Chapter 9, Section 9.6.

### 6.1.1 Building a modularized kernel

Only Red Hat Linux/Intel and Red Hat Linux/SPARC support modular kernels; Red Hat Linux/Alpha users must build a monolithic kernel (see Section 6.1.3 on page 93).

These instructions provide you with the knowledge required to take advantage of the power and flexibility available through kernel modularization. If you do not wish to take advantage of modularization, please see Section 6.1.3 for an explanation of the different aspects of building and installing a monolithic kernel. It is assumed that you have already installed the `kernel-headers` and `kernel-source` packages and that you issue all commands from the `/usr/src/linux/` directory.

It is important to begin a kernel build with the source tree in a known condition. Therefore, it is recommended that you begin with the command `make mrproper`. This will remove any configuration files along with the remains of any previous builds that may be scattered around the source tree. Now you must create a configuration file that will determine which components to include in your new kernel. Depending upon your hardware and personal preferences there are three methods available to configure the kernel.

- `make config` An interactive text program. Components are presented and you answer with **Y** (yes), **N** (no), or **M** (module).
- `make menuconfig` A graphic, menu driven program. Components are presented in a menu of categories, you select the desired components in the same manner used in the Red Hat Linux installation program. Toggle the tag corresponding to the item you want included; **Y** (yes), **N** (no), or **M** (module).
- `make xconfig` An X Windows program. Components are listed in different levels of menus, components are selected using a mouse. Again, select **Y** (yes), **N** (no), or **M** (module).

**Please Note:** In order to use `kernel.d` (see Section 9.6 for details) and kernel modules you must answer **Yes** to **kernel.d support** and **module version (CONFIG.MODVERSIONS) support** in the configuration.

**Please Note:** If you are building a Linux/Intel kernel on (or for) a machine that uses a “clone” processor (for example, one made by Cyrix or AMD), it is recommended to choose a **Processor type** of **386**.

If you wish to build a kernel with a configuration file (`/usr/src/linux/.config`) that you have already created with one of the above methods, you can omit the

`make mrproper` and `make config` commands and use the command `make dep` followed by `make clean` to prepare the source tree for the build.

The next step consists of the actual compilation of the source code components into a working program that your machine can use to boot. The method described here is the easiest to recover from in the event of a mishap. If you are interested in other possibilities details can be found in the Kernel-HOWTO or in the `Makefile` in `/usr/src/linux` on your Linux system.

- Build the kernel with `make boot`.
- Build any modules you configured with `make modules`.
- Move the old set of modules out of the way with:

```
rm -rf /lib/modules/2.0.29-old
mv /lib/modules/2.0.29 /lib/modules/2.0.29-old
```

Of course, if you have upgraded your kernel, replace `2.0.29` with the version you are using.

- Install the new modules (even if you didn't build any) with `make modules_install`.

If you have a SCSI adapter and made your SCSI driver modular, build a new `initrd` image (see Section 6.1.2; note that there are few practical reasons to make the SCSI driver modular in a custom kernel).

In order to provide a redundant boot source to protect from a possible error in a new kernel you should keep the original kernel available. Adding a kernel to the LILO menu is as simple as renaming the original kernel in `/boot`, copying the new kernel to `/boot`, adding a few lines in `/etc/lilo.conf` and running `/sbin/lilo`. Here is an example of the default `/etc/lilo.conf` file shipped with Red Hat Linux:

```
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
prompt
timeout=100
image=/boot/vmlinuz
        label=linux
```

```
root=/dev/hda1
read-only
```

Now you must update `/etc/lilo.conf`. If you built a new `initrd` image you must tell LILO to use it. In this example of `/etc/lilo.conf` we have added four lines at the bottom of the file to indicate another kernel to boot from. We have renamed `/boot/vmlinuz` to `/boot/vmlinuz.old` and changed its label to `old`. We have also added an `initrd` line for the new kernel:

```
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
prompt
timeout=100
image=/boot/vmlinuz
    label=linux
    initrd=/boot/initrd
    root=/dev/hda1
    read-only
image=/boot/vmlinuz.old
    label=old
    root=/dev/hda1
    read-only
```

Now when the system boots and you press **(Tab)** at the LILO boot : prompt two choices will be shown;

```
LILO boot:
linux    old
```

To boot the new kernel (`linux`) simply press **(Enter)**, or wait for LILO to time out. If you want to boot the old kernel (`old`), simply enter `old` and press **(Enter)**.

Here is a summary of the steps;

- `mv /boot/vmlinuz /boot/vmlinuz.old`
- `cp /usr/src/linux/arch/i386/boot/zImage /boot/vmlinuz`
- `edit /etc/lilo.conf`
- `run /sbin/lilo`

You can begin testing your new kernel by rebooting your computer and watching the messages to ensure your hardware is detected properly.

### 6.1.2 Making an initrd image

An `initrd` image is needed for loading your SCSI module at boot time. The shell script `/sbin/mkinitrd` can build a proper `initrd` image for your machine if the following conditions are met:

- The loopback block device is available.
- The `/etc/conf.modules` file has a line for your SCSI adapter; for example:

```
alias scsi_hostadapter BusLogic
```

To build the new `initrd` image, run `/sbin/mkinitrd` with parameters such as this:

```
/sbin/mkinitrd /boot/newinitrd-image 2.0.12
```

where `/boot/newinitrd-image` is the file to use for your new image, and `2.0.12` is the kernel whose modules (from `/lib/modules`) should be used in the `initrd` image (not necessarily the same as the version number of the currently running kernel).

### 6.1.3 Building a monolithic kernel

To build a monolithic kernel you follow the same steps as building a modularized kernel with a few exceptions.

- When configuring the kernel only answer **Yes** and **No** to the questions (don't make anything modular).
- Omit the steps:

```
make modules
make modules_install
```

- Edit the file `/etc/rc.d/rc.sysinit` and comment out the line `depmod -a` by inserting a `"#"` at the beginning of the line.

## 6.2 X Window Configuration with XFree86

There are four methods for configuring XFree86 on your machine:

- Xconfigurator
- XF86Setup
- xf86config
- by hand

Xconfigurator, XF86Setup, and xf86config are functional equivalents and should work equally well. If you are unsure of anything in this process, a good source of additional documentation is:

<http://www.xfree86.org>

Xconfigurator is a full-screen menu driven program that walks you through setting up your X server. XF86Setup is a graphical setup program distributed with XFree86 which requires the VGA 16-color server in order to run. xf86config is a line oriented program distributed with XFree86. It isn't as easy to use as Xconfigurator, but it is included for completeness. If these utilities fail to provide a working XF86config file, you may have an unsupported card or you may need to write the config file by hand. Usually the former is the case, so check and make sure your card is supported before attempting to write the config file yourself. If your card is not supported by XFree86 you may wish to consider using a commercial X server, such as the Metro-X server which accompanies Official Red Hat Linux/Intel. If you have questions about whether or not your video card is supported you can check out <http://www.xfree86.org> for information on XFree86, and <http://www.metrolink.com> for information on Metro-X.

### 6.2.1 The X Server

Provided you selected the proper video card at install time, you should have the proper X server installed. When later running Xconfigurator or xf86config, you need to make sure you select the same video card or the autoprobe will fail.

If you think you installed the wrong X server for your video card, you will have to install the correct one before it can be configured. For instance, if the CD is

mounted on `/mnt/cdrom`, and you need to install the S3 server, enter the following commands:

```
cd /mnt/cdrom/RedHat/RPMS
rpm -ivh XFree86-S3-3.1.2-1.i386.rpm
ln -sf ../../usr/X11R6/bin/XF86_S3 /etc/X11/X
```

This will install the S3 server and make the proper symbolic link.

## 6.2.2 Xconfigurator

To configure X Windows you must first select your video card. Scroll down the list of supported cards until you locate the card in your machine. Section 6.2.2 may help you determine the video server that matches your hardware. If your card is not listed it may not be supported by XFree86. In this case you can try the last card entry on the list (Unlisted Card) or a commercial X Windows server, such as Metro X.

The next step is to select your monitor. If your monitor is not listed you can select one of the generic monitor entries or “Custom” and enter your own parameters. Custom monitor configuration is recommended only for those who have a sound understanding of the inner workings of CRT displays. The average user should probably use one of the generic selections from the list. After selecting a monitor you need to tell Xconfigurator how much video memory you have. Move the highlight to the appropriate list entry and then press **Enter** or **F12** to continue. For the next step it is recommended that you select the default (No Clockchip Setting) entry, but experienced users may want to select a specific clockchip.

### Selecting your Server

If you are unsure what chip-set you have, the best way to find out is usually to look at the card. Figure 6.1 lists which chip-sets and boards require which servers. Pick the one that best matches your hardware.

### Finishing Up

If later you want to increase your refresh rate for your monitor, you can edit the config file by hand or you can run Xconfigurator again and pick a monitor from

Server	Chipset
VGA16	All VGA boards with 256K (16 color only)
SVGA	Trident 8900 & 9400, Cirrus Logic, C & T, ET4000, Others
Mach8	ATI boards using the Mach8 chip-set
Mach32	ATI boards using the Mach32 chip-set
Mach64	ATI boards using the Mach64 chip-set
8514	IBM 8514/A Boards and true clones
S3	#9 Boards, most Diamonds, some Orchids, Others
AGX	All XGA graphics boards
P9000	Diamond Viper (but not the 9100) and Others
W32	All ET4000/W32 cards, but not standard ET4000's

Figure 6.1: X Servers

our list that more closely matches the specs of your monitor.

The final configuration step consists of selecting the video modes that you want to include in your XF86config file. Use the arrow keys to move the cursor up and down the list under each color depth (8, 16 and 24 bit). Use the **Spacebar** to select individual resolutions and the **Tab** key to move between color depth fields. When you have selected the video modes you want to use move the cursor to the "OK" button and press **Enter**, or use the **F12** shortcut. An information screen will give you the most current information on selecting video modes, starting and stopping the X server.

### 6.3 X Window Configuration with Metro-X

By special arrangement, a license to install a single copy of the Metro-X enhanced X server is included with the Official Red Hat Linux/Intel CD. Please note that, unlike the XFree86 server (also included), Metro-X is commercially licensed software. You should review and accept its license terms before installing it. Metro-X supports many video cards that are not supported by XFree86. The configuration utility can be run in graphical mode (`configX`) or text mode (`configX.curses`).

### 6.3.1 System Requirements

8 megabytes of RAM minimum; 16 megabytes are recommended. The base server package requires 12 megabytes of disk space.

### 6.3.2 configX.curses

`configX.curses` is a simple graphical configuration tool for configuring Metro-X. The `configX.curses` utility does not provide support for configuring multiple screens or choosing the number of buttons on your mouse. If you need these facilities you must use `configX`. To run `configX.curses` type the command:

```
configX.curses
```

The opening screen has information on Metro Link, Inc. and brief instructions on what you will need to configure Metro-X. Press the Enter key to begin the configuration. The screen will display the current configuration information at the top and the configuration menu at the bottom. Use the arrow keys to move the highlight up and down the menu and the **Enter** key to select the highlighted item.

#### Graphics Card

When you select **Graphics Card** from the configuration menu, a list of supported cards will be displayed. Select your video card as above. After selecting your video card you will be asked to select a video mode from a list of defaults. The resolution used is a matter of personal preference. The smaller the numbers are the larger the items on the display will be. Some combinations of video cards and video modes support virtual screens larger than the physical display. If your combination is one of these, you will be asked to select the virtual screen size.

#### Mouse

Configuring your mouse with Metro-X requires that you know what type of mouse you have and how it is connected to your computer. When you select a mouse type, the default device will be provided in the configuration window. If you are in doubt, accept the default and make adjustments later if needed.

## Exiting

Selecting **Save and Exit** from the menu will save any changes you have made to the Metro-X configuration and exit the `configX.curses` program.

Selecting **Quit without Saving** from the menu will discard any changes you have made to the Metro-X configuration and exit the `configX.curses` program.

## 6.3.3 configX

Metro-X can also be configured for your system through a graphical utility called *ConfigX*, which is used to specify the mouse, keyboard, monitor, graphics card, color depth, and screen resolution(s) you are using.

### Start-up

1. Login as root and type the following command: `/usr/bin/X11/configX`
2. The explanation of special keys appears. (This message only appears prior to the initial configuration of Metro-X, when no Xconfig file exists.)
3. Press **Enter** to continue.

### Special Keys

When Metro-X is first installed, it doesn't yet recognize your mouse. Special keys must be used to move around the screens and make your selections during the initial configuration process.

### Mouse configuration

If the default mouse is not correct for your system:

1. Press **↓** to display the list of mice.
2. Press **↑** or **↓** to scroll to your mouse entry, then press **Enter** to select it.
3. Press **Tab** to move to the next field.

Key	Function
<b>Tab</b>	Moves forward from field to field.
<b>Shift-Tab</b>	Moves backward from field to field.
<b>↑ ↓ ← →</b>	Changes the value in the current field.
<b>Enter</b>	Sets/ unsets a toggle button.
<b>Page Down</b>	Switches to the next configuration page.
<b>Page Up</b>	Switches to the previous configuration page.
<b>Esc</b>	Cancels out of the current screen or pop-up.
<b>Alt-F</b>	Jumps to the File pull-down menu.
<b>Alt-S</b>	Jumps to the Screen pull-down menu.
<b>Alt-C</b>	Jumps to the Configuration pull-down menu.

**Figure 6.2:** Function of Special Keys

4. Press **↓** to display the list of mouse buttons. (The default entry is based on your mouse selection, and is usually the appropriate choice.)
5. Press **↑** or **↓** to scroll to your button entry, then press **Enter** to select it.
6. Press **Tab** to move to the next field.
7. If you need to change the path for your mouse device, edit it here.
8. Press **Tab** to move to the next field.

### Keyboard configuration

Metro-X includes international keyboard support, which allows you to select the keyboard appropriate for your language and country.

1. The initial keyboard setting is *Default: Default Console Mapping*, which will work with most keyboards.
2. Press **↓** to display the list of keyboards.
3. Press **↑** or **↓** to scroll to your keyboard entry, then press **Enter** to select it.
4. Press **Tab** to move to the next field.

### Monitor configuration

The default monitor setting is *Generic 640x480 @ 60Hz 13 inch*.

1. Press **↓** to display the list of monitors.
2. Press **↑** or **↓** to scroll to your monitor entry, then press **Enter** to select it.
3. Press **Tab** to move to the next field.
4. If you need to change the screen width (mm), edit it here.
5. Press **Tab** to move to the next field.
6. If you need to change the screen height (mm), edit it here.
7. Press **Tab** to move to the next field.

### Screen Saver

Some graphics cards and monitors support the Display Power Management System (DPMS), which minimizes the power used by the monitor when user input ceases. If your hardware has this feature, the phases of DPMS can be controlled in this section. The typical DPMS setup has a cascade effect, with the various phases being activated in succession with your choice of delay times between each.

- X Screen Saver - Preliminary phase to DPMS, which blanks the screen but provides no power reduction.
  - Standby Mode - First phase of DPMS, which provides moderate power reduction by disabling the horizontal sync signal to the monitor.
  - Suspend Mode - Second phase of DPMS, which provides significant power reduction by disabling the vertical sync signal to the monitor.
  - Monitor Off - Final phase of DPMS, which provides maximum power reduction by disabling both the horizontal and vertical sync signals to the monitor.
1. In the *X Screen Saver Delay* field, press **Enter** to toggle this option on or off.
  2. If toggled on, press **Tab** to move to the *Minutes* field, then use **←** or **→** to set the value.

3. Press **Tab** to move to the next field.
4. In the *Standby Mode Delay* field, press **Enter** to toggle this option on or off.
5. If toggled on, press **Tab** to move to the *Minutes* field, then use **←** or **→** to set the value.
6. Press **Tab** to move to the next field.
7. In the *Suspend Mode Delay* field, press **Enter** to toggle this option on or off.
8. If toggled on, press **Tab** to move to the *Minutes* field, then use **←** or **→** to set the value.
9. Press **Tab** to move to the next field.
10. In the *Monitor Off Delay* field, press **Enter** to toggle this option on or off.
11. If toggled on, press **Tab** to move to the *Minutes* field, then use **←** or **→** to set the value.

### Graphics Card Configuration

Press **Page Down** to move to the second configuration screen, which controls the graphics card, the number of colors, the virtual resolution, and the physical resolution.

Each graphics card configuration consists of five elements:

- The graphics card (make and model).
- The number of colors supported on screen simultaneously (color depth).
- The virtual resolution of the screen (width x height).
- The physical resolution of the screen (width x height).
- The screen refresh rate (Hz).

Non-interlaced modes give a better screen appearance than refresh rates which are labeled "Interlaced". The highest refresh rate that your monitor supports will give the best image. At higher screen resolutions a higher refresh rate is particularly important to prevent the screen from visibly flickering. The highest refresh rate supported by the selected monitor is automatically displayed. **Graphics Card**

1. The default graphics card setting is *IBM VGA*.
2. Press **↓** to display the list of graphics cards.
3. Press **↑** or **↓** to scroll to your graphics card entry, then press **Enter** to select it.
4. Press **Tab** to move to the next field.

### Colors

Not all graphics cards can support all color depths (number of colors). Only the options for the color depths supported by your card will be sensitive in this field. The other options will be grayed out.

- Only *one* color option can be selected (due to design limitations of X).
- In some cases the color options available will be limited by the physical resolution(s) you choose, based on the capabilities of your graphics card. Not all color depths are available at high resolutions.
  1. The default setting is at the lowest value for your card.
  2. Decide which available color setting you want to use.
  3. Press **Tab** to move to that color setting, then press **Enter** to select it. (The previous setting will be canceled.)
  4. Press **Tab** to move to the next field.

### Virtual Resolution

The *Virtual Resolution* setting allows you to have an on-screen work area that is actually larger than your physical screen. Windows can extend beyond the screen's physical boundaries and are viewed by scrolling into the extended areas. Only the options for the virtual resolutions supported by your card will be sensitive in this field. The other options will be grayed out. If you don't want to use this feature, select the *Disabled* option.

- Only *one* virtual resolution can be selected.

- In some cases the virtual-resolution options available will be limited by the physical resolution(s) you choose, since the virtual resolution must always be greater than the highest physical resolution.
  1. The default setting for virtual resolution is *Disabled*.
  2. Decide which available virtual resolution you want to use.
  3. Press **Tab** to move to that setting, then press **Enter** to select it. (The previous setting will be canceled.)
  4. Press **Tab** to move to the next field.

### Physical Resolution

Not all graphics cards and monitors can support all physical resolutions. Only the options for the resolutions supported by your hardware will be sensitive in this field. The other options will be grayed out.

- *Multiple* physical resolutions can be selected. You can switch between these resolutions while running Metro-X by pressing a hot-key combination:
  - Press **Ctrl-Alt-+** to change to the next resolution.
  - Press **Ctrl-Alt-** to change to the previous resolution.
- If you want only *one* resolution, make sure only *one Physical Resolution* toggle button is selected.
- In some cases the physical resolution(s) you choose may limit the selections you can make in *Colors* and *Virtual Resolution*, depending on the capabilities of your graphics card.
  1. The default setting is at the lowest value for your card.
  2. Decide which available resolution(s) you want to use.
  3. Press **Tab** to move to the setting, then press **Enter** to select it.
  4. Press **Tab** to move to the corresponding field to set the screen refresh rate (Hz) for that resolution, then use **←** or **→** to select a value from the frequency options available. The default setting is the highest frequency your monitor supports, and is usually the appropriate choice.
  5. Repeat this procedure for each physical resolution you want to set.

## Multiple-Screen Configuration

Metro-X supports up to four screens which can all be controlled simultaneously with a single keyboard and mouse. This allows you to run many applications without overlapping windows, and is particularly useful for image-processing work. For each screen used, a separate configuration must be created for that screen's monitor and graphics card. See section G.1.4 for graphics card requirements.

### Adding a Screen

Metro-X defaults to a single-screen setup. The first screen configured is labeled *Screen 0*. When you add a second screen, all of its configuration default settings will duplicate the current settings of Screen 0. So if your hardware is the same on all screens, you can save time by configuring Screen 0 before adding the other screens. Similarly, as more screens are added, the settings of whichever one is currently active when you add another screen will determine the default settings of the new screen.

1. Press **[Alt-S]** to display the Screen pull-down menu.
2. Type *A* for *Add*.
3. The *Monitor* and *Graphics Card* icons for Screen 0 become outlines, and two full-color icons appear for Screen 1.
4. Refer to the procedures described earlier in this chapter if any changes need to be made to the new screen's configuration.
5. Use the **[Page Down]** and **[Page Up]** keys to move to the configuration pages of each screen.

### Deleting a Screen

When a screen is deleted, all screens with higher screen numbers immediately drop down one slot. For instance: You currently have screens 0, 1, and 2 configured and you delete screen 1. That screen and all of its configuration information disappears. Screen 2 then becomes screen 1, but retains its own configuration. You are left with screens 0 and 1.

1. Use the **Page Down** or **Page Up** key to activate the set of icons for the screen you want to delete.
2. Press **Alt-S** to display the *Screen* pull-down menu.
3. Type *D* for *Delete*.
4. One set of *Monitor* and *Graphics Card* icons disappears. Higher-numbered screens (if any) immediately fill the gap of the deleted screen.

### Screen Layout

The *Screen Layout* window lets you specify how the multiple screens are arranged in your workspace so mouse movement from screen to screen will be appropriate.

1. Press **Alt-S** to display the *Screen* pull-down menu.
2. Type *L* for *Layout*.
3. The *Screen Layout* window appears.
4. Press **←** and **→** to view the various layout combinations for your number of screens.
5. When the layout you want is displayed, press **Enter** to select it.

### Font Path

If needed, you can edit your font path from within the Metro-X Configuration Utility. This will be necessary if you are using a font server instead of or in addition to your existing fonts.

1. Press **Alt-C** to display the *Configuration* pull-down menu, then press **Enter**.
2. The *Font Path* window appears.
3. Edit the file as necessary for your system. Metro-X looks for fonts in the order in which the directories are listed in this file. If you are using a font server, enter this as the first line:

```
tcp/<fontserver machine>:7000
```

If your port number is not the default 7000, change it to the appropriate port number for your system.

4. Press **Tab** to move to and amongst the bottom buttons, then press **Enter** to select one. If the information is correct, select the *OK* button. If you want to disregard the changes you've made, select the *Cancel* button. If you've edited your font path and want to go back to the original settings, select the *Load Defaults* button.

### Save and Exit

1. Press **Alt-F** to display the File pull-down menu.
2. Type *S* for *Save and Exit*.
3. Your changes are saved, the Metro-X Configuration Utility closes, and your system is now ready to run Metro-X!

### Quit Without Saving

1. Press **Alt-F** to display the File pull-down menu.
2. Type *Q* for *Quit*.
3. A precautionary dialog box appears: *Save current configuration before exiting?*
4. Press **Tab** to move to the *No* button, then press **Enter** to confirm that you want to exit the Metro-X Configuration Utility without saving your changes.

For more details on features of Metro X and the format of the `xconfig` file see Appendix G.

## 6.4 Mouse configuration

To configure your mouse (or reconfigure your mouse after installation) enter the command `/usr/sbin/mouseconfig`. Scroll down the list with the arrow keys until your mouse type is highlighted. If you have *any* mouse connected to a PS/2 style port you should select **PS/2** as your mouse type. If you have a mouse connected to a mouse port on an ATI video card you should select **ATI Bus Mouse**. If

you have a 2 button mouse and want to emulate the third (middle) mouse button with a simultaneous click of the left and right mouse buttons use **Tab** to move the cursor to the **Emulate 3 Buttons** checkbox and press **Spacebar** to toggle the selection.

## 6.5 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 new `sendmail.cf`, you will need to install `m4` and the `sendmail` source package. Read the `README` file in the `sendmail` sources for more details on creating `sendmail` configuration files. Also, O'Reilly & Associates publishes a good `sendmail` reference entitled *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 `mail.redhat.com` that does all our mail. On that machine we simply need to add the names of machines for which `mail.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 `mail.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 I send unqualified names to
# (null means deliver locally)
DRmail.redhat.com

# who gets all local email traffic
DHmail.redhat.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 `mail.redhat.com`.

Please be aware that if you configure your system to masquerade as another any email sent from your system to your system will be sent to the machine you are masquerading as. For example, in the above illustration, log files that are periodically sent to `root@poodle.redhat.com` by the cron daemon would be sent to `root@mail.redhat.com`.

## 6.6 Controlling Access to Services

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, and the configuration in `/etc/hosts.allow` and `/etc/hosts.deny`. By default Red Hat 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.

## 6.7 Anonymous FTP

Setting up anonymous FTP is simple. All you need to do is install the `anon-ftp` rpm package (which you may have already done at install time). Once it is installed, anonymous FTP will be up and running.

There are a few files you might wish to edit to configure your FTP server.

`/etc/ftpaccess` This file defines most of the access control for your ftp server. Some of the things that you can do are: set up logical “groups” to control access from different sites, limit the number of simultaneous FTP connections, configure transfer logging, and much more. Read the `ftpaccess` man page for complete details.

`/etc/ftphosts` The `ftphosts` file is used to allow or deny access to certain accounts from various hosts. Read the `ftphosts` man page for details.

`/etc/ftpusers` This file lists all the users that are *not* allowed to ftp into your machine. For example, `root` is listed in `/etc/ftpusers` by default. That means that you can not ftp to your machine and log in as `root`. This is a good security measure, but some administrators prefer to remove `root` from this file.

## 6.8 World Wide Web

The world wide web is one of the hottest aspects of the Internet today. Red Hat Linux lets you get in on the action in two ways – as a web browser, and as a web server. Let’s look at both.

### 6.8.1 World Wide Web Browsers

A variety of web browsers are available for Linux, including freely distributable browsers such as `arena`, `lynx`, and `grail`. The most popular commercial browsers are those from Netscape Communications Corporation. It is available for evaluation and non-commercial use from their ftp site:

`http://www.netscape.com/download/client_download.html`

It is also available in two fully licensed versions directly from Netscape.

Netscape Communicator 4.03 (part # 100-04790-00)  
Netscape Communicator 4.03 Professional (part # 100-04811-00)

For more information, please see:

<http://merchant.netscape.com/netstore/clients/index.html>

Netscape can be reached at:

Corporate Sales: (650) 937-2555  
Personal Sales: (650) 937-3777  
Government Sales: (650) 937-3678

**Please Note:** Please see the appropriate documents on Netscape's web site for the latest in licensing and right-to-use policies on these products.

## 6.8.2 World Wide Web Server

If you installed the Apache web server (from the `apache` package), then your Web service is already up and running! Just point your web browser at `http://localhost`.

The default page shown is `/home/httpd/html/index.html`. You can edit this file (or completely replace it) to your liking. All the CGI programs, icons, and html pages are stored in `/home/httpd`, but this can be changed in the apache configuration files, all of which are stored in `/etc/httpd/conf/`. Logs of all httpd activity are kept in `/var/log/httpd/`. Setting up your web site is as easy as adding your own HTML pages and links to the `/home/httpd/html/` directory. For more information on customizing your web server we recommend a reference such as *HTML: The Definitive Guide* by Chuck Musciano & Bill Kennedy, published by O'Reilly & Associates.

## 6.9 NFS Configuration

NFS stands for Network File System, and is a way to share files between machines as if they were on your local hard drive. Linux can be both an NFS server and an NFS client, which means that it can *export* filesystems to other systems, and *mount* filesystems exported from other machines.

### 6.9.1 Mounting NFS Filesystems

Use the `mount` command to mount a NFS filesystem from another machine:

```
mount porky:/mnt/export /mnt/local
```

In this command, `porky` is the hostname of the NFS fileserver, `/mnt/export` is the filesystem that `porky` is exporting, and `/mnt/local` is a directory on my local machine where we want to mount the filesystem. After the `mount` command runs (and if we have the proper permissions from `porky`) we can enter `ls /mnt/local` and get a listing of the files in `/mnt/export` on `porky`.

### 6.9.2 Exporting NFS Filesystems

The file that controls what filesystems you wish to export is `/etc/exports`. Its format is:

```
directory      hostname(options)
```

the `(options)` are optional. For example:

```
/mnt/export    speedy.redhat.com
```

would allow `speedy.redhat.com` to mount `/mnt/export`, but:

```
/mnt/export    speedy.redhat.com(ro)
```

would just allow `speedy` to mount `/mnt/export` read-only.

Each time you change `/etc/exports`, you need to tell the NFS daemons to examine it for new information. One simple way to accomplish this is to just stop and start the daemons:

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

The following will also work:

```
killall -HUP rpc.nfsd rpc.mountd
```

See the following man pages for more details: `nfsd(8)`, `mountd(8)`, and `exports(5)`. Another good reference is *Managing NFS and NIS Services* by Hal Stern, published by O'Reilly & Associates.

## 6.10 Running Programs at Boot Time

The file `/etc/rc.d/rc.local` is executed at boot time, after all other initialization is complete, and whenever you change runlevels. You can add additional initialization commands here. For instance, you may want to start up additional daemons, or initialize a printer. In addition, if you require serial port setup, you can edit `/etc/rc.d/rc.serial`, and it will be executed automatically at boot time.

The default `/etc/rc.d/rc.local` simply creates a nice login banner with your kernel version and machine type.

## 6.11 Color-ls Setup

`color-ls` is now part of the standard `ls` command. To get colorized directory listings, simply issue the command (in bash):

```
alias ls='ls --color=tty'
```

# 7

## Finding Documentation

Red Hat Linux includes thousands of pages of online documentation to help you learn how to use the system. The man pages, info documents, and plain text files included provide information on almost every aspect of Linux. For more help, Red Hat also includes all of the documentation produced by the Linux Documentation Project electronically.

### 7.1 On Line Help

When you are looking for general help on commands and error messages, the best place to start is right on your system.

#### 7.1.1 Man Pages

Most every command on your system has an associated “man” page. This is documentation that you can get to instantly should you have questions or problems. If you were having trouble with the command `ls`, you could enter `man ls`. This

will bring up the man page for `ls`. The man page is viewed through the `less` program, so all of the options to `less` will work while in a man page. Some important key strokes are:

- `q` to quit
- `Enter` to page down line by line
- `Space` to page down page by page
- `b` to page back up by one page
- `/` followed by a string and `Enter` to search for a string
- `n` to find the next occurrence of the previous search

Sometimes viewing man pages isn't too friendly on line. Providing you have a working printer, you can print man pages as well. If you don't have postscript printing capability and just want to print ASCII, you can print man pages with:

```
man COMMAND | lpr
```

If you do have a postscript printer, you will probably want to print with:

```
man -t COMMAND | lpr
```

In both of those commands substitute "COMMAND" for the command you are trying to get help for.

Also, sometimes things have more than one man page. Here is a table of what is located where:

Section	Contents
1	user commands
2	system calls
3	library calls
4	devices
5	file formats
6	games
7	miscellaneous
8	system commands
9	kernel internals

So, let's say that you want to see the man page for `swapon`. You do `man swapon`. You will actually get the man page for the system call `swapon(2)`, which is the function you use in a C program to turn swap on. Unless you are writing your own program to do it, this probably isn't what you want. So, using the chart above, you can see that what you want is probably a "system command" and is located in section 8. You can then do `man 8 swapon`. All of this is because `man` searches the man directories in order, and then returns when it finds the first match.

You can also search the man pages for strings. You do this using `man -k string_to_search_for`. This won't work, however, unless the `make-whatis` database has been created. Under Red Hat Linux, this is done by a cron job overnight. If you don't leave your system running overnight the database won't get created. If that is the case, run the following command as the `root` user:

```
/etc/cron.weekly/makewhatis.cron
```

Once you've done that, you could enter `man -k swapon`. That command would return:

```
# man -k swapon
swapon, swapoff (2) - start/stop swapping to file/device
swapon, swapoff (8) - enable/disable devices and files for \
    paging and swapping
```

So you can see that there are pages in section 2 and 8 both referring to `swapon` (and `swapoff` in this case).

## 7.1.2 Package Documentation

Many packages of software have `README` files and other documentation as part of the source package. Red Hat Linux uses a standard place to install those documents so that you don't have to install the sources to look at the documents. All of those documents are stored in subdirectories of `/usr/doc`. The subdirectory depends on the package. Each package that has extra documentation will create a directory called `packagename-version-releasename`. For example, the `tin` package might be version 1.22 and release number 2. The path to its documentation would be `/usr/doc/tin-1.22-2`.

For the most part, the documents in this directory are ASCII. You can usually view them with `more filename` or `less filename`. This is nice, but what if you

want to see if there is documentation for a specific command or file and you don't know the package it came from? It doesn't matter! You can simply enter:

```
rpm -qdf /etc/sendmail.cf
```

This will report all the documentation from the package containing the file `/etc/sendmail.cf`. Commands like this are covered more in depth in the RPM-HOWTO, available from `www.redhat.com`.

Also, what if it's a command you need help with and the man page isn't good? You could do something like:

```
rpm -qdf `which COMMAND`
```

Again, "COMMAND" is the actual command you need help with. This will work only when the command is on your path.

### 7.1.3 HOWTOs and FAQs

Most of the contents of the Linux Documentation Project (LDP) are available in `/usr/doc` on your system.

`/usr/doc/HOWTO` contains the ASCII versions of all the available HOWTOs at the time we pressed the CD-ROM. They are gzipped, so you have to use `gunzip` to unzip them or use a command like:

```
zcat HAM-HOWTO.gz | more
```

The latter will work, but is a bit less flexible than unzipping and then using `more`. However, unzipping also requires more disk space unless you re-gzip the document when done.

`/usr/doc/HOWTO/mini` contains the ASCII versions of all the available mini-HOWTOs. They are not compressed and can be viewed with `more` or `less`. `/usr/doc/HTML` contains the HTML versions of all the HOWTOs and the *Linux Installation and Getting Started* guide. To view things here, just use a WWW browser like `redbaron` from Red Hat Software, or `arena`. You would do something like:

```
cd /usr/doc/HTML  
arena index.html
```

`/usr/doc/FAQ` contains ASCII version (and some HTML versions) of some popular FAQs, including the RedHat-FAQ. All of them can be viewed using more or less.

### 7.1.4 The “locate” Command

When you don't know the full name of a command or file, but need to find it, you can usually find it with `locate`. `locate` uses a database to find all files on your system. Normally, this database gets built from a cron job every night. This won't happen, however, if your machine isn't booted into Linux all the time. So, if that is the case, you may occasionally want to run the following command:

```
/etc/cron.daily/updatedb.cron
```

You will need to be root on your system when doing that. That will allow `locate` to work properly.

So, if you know you need to find all the “finger” files, you could run:

```
locate finger
```

It should return something like:

```
/usr/bin/finger  
/usr/lib/irc/script/finger  
/usr/man/man1/finger.1  
/usr/man/man8/in.fingerd.8  
/usr/sbin/in.fingerd
```

One thing to note, however, is that `locate` not only returns hits based on file name, but also on path name. So if you have a `/home/djb/finger/` directory on your system, it would get returned along with all files in the directory.

### 7.1.5 “info” Pages

While `man` is the most ubiquitous documentation format, `info` is much more powerful. It provides hypertext links to make reading large documents much easier

and many features for the documentation writer. There are some very complete `info` documents on various aspects of Red Hat (especially the portions from the GNU project).

To read `info` documentation, use the `info` program without any arguments. It will present you with a list of available documentation. If it can't find something, it's probably because you don't have the package installed that includes that documentation. Install it with RPM and try again.

If you're comfortable using emacs, it has a built in browser for `info` documentation. Use `Ctrl-h` `Ctrl-i` to see it.

The `info` system is a hypertext based system. Any highlighted text that appears is a link leading to more information. Use `Tab` to move the cursor to the link, and press `Enter` to follow the link. Pressing `p` returns you to the previous page, `n` moves you to the next page, and `u` goes up one level of documentation. To exit `info`, press `Ctrl-x` `Ctrl-c` (control-x followed by control-c).

The best way to learn how to use `info` is to read the `info` documentation on it. If you read the first screen that `info` presents you'll be able to get started.

## 7.2 Help from the Internet Community

### 7.2.1 Red Hat Mailing Lists

If you can't find help for your problem on line and you have WWW access, you should see <http://www.redhat.com/support/mailing-lists/>. Here you can search the archives of the redhat-list. Many questions have already been answered there.

The subscription address for each of our lists is:

```
redhat-list-request@redhat.com
redhat-digest-request@redhat.com
redhat-announce-list-request@redhat.com
redhat-devel-list-request@redhat.com
rpm-list-request@redhat.com
xpp-list-request@redhat.com
sparc-list-request@redhat.com
```

To subscribe, send mail to the address of the list you want to subscribe to with `subscribe` in the `Subject`: line.

To unsubscribe, send mail to the address of the list you want to unsubscribe from with `unsubscribe` in the `Subject`: line.

Then to send mail to the list, you just send it to the address above without the `-request` in the name. The only exception to this is the `redhat-digest@redhat.com` which is only a digest version the `redhat-list@redhat.com`. In that case the address is `redhat-list@redhat.com`.

### **7.2.2 USENET Newsgroups**

Another good source of help is the `comp.os.linux` hierarchy on USENET. If you are familiar with news, you should check it out.



# 8

## Package Management with RPM

Red Hat Package Manager (RPM), is an open packaging system available for anyone to use, and works on both Red Hat Linux as well as other Linux and UNIX systems. Red Hat Software encourages other vendors to take the time to look at RPM and use it for their own products. RPM is distributable under the terms of the GPL.

For the end user, RPM provides many features that make maintaining a system far easier than it has ever been. Installing, uninstalling, and upgrading RPM packages are all one line commands, and all the messy details are taken care of for you. RPM maintains a database of installed packages and their files, which allows you to perform powerful queries and verification of your system. During upgrades RPM handles configuration files specially, so that you never lose your customizations – a feature that is impossible with straight `.tar.gz` files.

For the developer, RPM allows you to take source code for software and package it into source and binary packages for end users. This process is quite simple and is driven from a single file and optional patches that you create. This clear delineation of “pristine” sources and your patches and build instructions eases the

maintenance of the package as new versions of the software are released.

## 8.1 RPM Design Goals

Before trying to understand how to use RPM, it helps to have an idea of what the design goals are.

**Upgradability** With RPM you can upgrade individual components of your system without completely reinstalling. When you get a new release of an operating system based on RPM (such as Red Hat Linux), you don't need to reinstall your machine (as you do with operating systems based on other packaging systems). RPM allows intelligent, fully-automated, in-place upgrades of your system. Configuration files in packages are preserved across upgrades, so you won't lose your customizations.

**Powerful Querying** RPM is also designed to have powerful querying options. You can do searches through your entire database for packages or just certain files. You can also easily find out what package a file belongs to and where it came from. The files an RPM package contains are in a compressed archive, with a custom binary header containing useful information about the package and its contents, allowing you to query individual packages quickly and easily.

**System Verification** Another powerful feature is the ability to verify packages. If you are worried that you deleted an important file for some package, simply verify the package. You will be notified of any anomalies. At that point, you can reinstall the package if necessary. Any configuration files that you modified are preserved during reinstallation.

**Pristine Sources** A crucial design goal was to allow the use of "pristine" software sources, as distributed by the original authors of the software. With RPM, you have the pristine sources along with any patches that were used, plus complete build instructions. This is a big advantage for several reasons. For instance, if a new version of a program comes out, you don't necessarily have to start from scratch to get it to compile. You can look at the patch to see what you *might* need to do. All the compiled-in defaults, and all of the changes that were made to get the software to build properly are easily visible this way.

This goal may only seem important for developers, but it results in higher quality software for end users too. We would like to thank the folks from the BOGUS distribution for originating the pristine source concept.

## 8.2 Using RPM

RPM has five basic modes of operation (not counting package building): installing, uninstalling, upgrading, querying, and verifying. This section contains an overview of each mode. For complete details and options try `rpm --help`, and see the `rpm` man page and the (previously mentioned) RPM book.

### 8.2.1 Installing

RPM packages typically have file names like `foo-1.0-1.i386.rpm`, which includes the package name (`foo`), version (`1.0`), release (`1`), and architecture (`i386`). Installing a package is as simple as:

```
$ rpm -ivh foo-1.0-1.i386.rpm
foo #####
```

As you can see, RPM prints out the name of the package (which is not necessarily the same as the file name, which could have been `1.rpm`), and then prints a succession of hash marks as the package is installed, as a sort of progress meter.

Installing packages is designed to be simple, but you can get a few errors:

#### Package Already Installed

If the package is already installed, you will see:

```
$ rpm -ivh foo-1.0-1.i386.rpm
foo package foo-1.0-1 is already installed
error: foo-1.0-1.i386.rpm cannot be installed
```

If you really want to install the package anyway, you can use `--replacepks` on the command line, which tells RPM to ignore the error.

### Conflicting Files

If you attempt to install a package that contains a file that has already been installed by another packages, you'll see:

```
# rpm -ivh foo-1.0-1.i386.rpm
foo          /usr/bin/foo conflicts with file from bar-1.0-1
error: foo-1.0-1.i386.rpm cannot be installed
```

To cause RPM to ignore that error, use `--replacefiles` on the command line.

### Unresolved Dependency

RPM packages can “depend” on other packages, which means that they require other packages to be installed in order to run properly. If you try to install a package for which there is such an unresolved dependency, you'll see:

```
$ rpm -ivh bar-1.0-1.i386.rpm
failed dependencies:
    foo is needed by bar-1.0-1
```

To handle this error you should install the requested packages. If you want to force the installation anyway (a bad idea since the package probably will not run correctly), use `--nodeps` on the command line.

## 8.2.2 Uninstalling

Uninstalling a package is just as simple as installing:

```
$ rpm -e foo
```

Notice that we used the package *name* “foo”, not the name of the original package *file* “foo-1.0-1.i386.rpm”.

You can encounter a dependency error when uninstalling a package if some other installed package depends on the one you are trying to remove. For example:

```
$ rpm -e foo
removing these packages would break dependencies:
    foo is needed by bar-1.0-1
```

To cause RPM to ignore that error and uninstall the package anyway (which is a bad idea since the package that depend on it will probably fail to work properly), use `--nodeps` on the command line.

### 8.2.3 Upgrading

Upgrading a package is almost just like installing.

```
$ rpm -Uvh foo-2.0-1.i386.rpm
foo #####
```

What you don't see above is the fact that RPM automatically uninstalled any old versions of the `foo` package. In fact you may want to always use `-U` to install packages, since it works fine even when there are no previous versions of the package installed.

Since RPM performs intelligent upgrading of packages with configuration files, you may see a message like:

```
saving /etc/foo.conf as /etc/foo.conf.rpmsave
```

This means that your changes to the configuration file may not be “forward compatible” with the new configuration file in the package, so RPM saved your original file, and installed a new one. You should investigate and resolve the differences between the two files as soon as possible to ensure that your system continues to function properly.

Since upgrading is really a combination of uninstalling and installing, you can encounter any errors from those modes, plus one more: If RPM thinks you are trying to upgrade to a package with an *older* version number, you will see:

```
$ rpm -Uvh foo-1.0-1.i386.rpm
foo    package foo-2.0-1 (which is newer) is already installed
error: foo-1.0-1.i386.rpm cannot be installed
```

To cause RPM to “upgrade” anyway, use `--oldpackage` on the command line.

## 8.2.4 Querying

Querying the database of installed packages is accomplished with `rpm -q`. A simple use is `rpm -q foo` which will print the package name, version, and release number of the installed package `foo`:

```
$ rpm -q foo
rpm-2.0-1
```

Instead of specifying the package name, you can use the following options with `-q` to specify what package(s) you want to query. These are called *Package Specification Options*.

- `-a` queries all currently installed packages.
- `-f <file>` will query the package owning `<file>`.
- `-F` is the same as `-f` except it takes filenames via stdin (e.g. `find /usr/bin | rpm -qF`).
- `-p <packagefile>` queries the package `<packagefile>`.
- `-P` is like `-p` except it takes package filenames from stdin (e.g. `find /mnt/cdrom/RedHat/RPMS | rpm -qP`).

There are a number of ways to specify what information to display about queried packages. The following options are used to select the information you are interested in. These are called *Information Selection Options*.

- `-i` displays package information such as name, description, release, size, build date, install date, vendor, and other miscellaneous information.
- `-l` displays the list of files that the package “owns”.
- `-s` displays the state of all the files in the package. There are only two possible states, normal and missing.
- `-d` displays a list of files marked as documentation (man pages, info pages, README’s, etc).
- `-c` displays a list of files marked as configuration files. These are the files you change after installation to adapt the package to your system (sendmail.cf, passwd, inittab, etc).

For those options that display file lists, you can add `-v` to your command line to get the lists in a familiar `ls -l` format.

### 8.2.5 Verifying

Verifying a package compares information about files installed from a package with the same information from the original package. Among other things, verifying compares the size, MD5 sum, permissions, type, owner and group of each file.

`rpm -V` verifies a package. You can use any of the *Package Selection Options* listed for querying to specify the packages you wish to verify. A simple use is `rpm -V foo` which verifies that all the files in the `foo` package are as they were when they were originally installed. For example:

- To verify a package containing particular file:

```
rpm -Vf /bin/vi
```

- To verify ALL installed packages:

```
rpm -Va
```

- To verify an installed package against an RPM package file:

```
rpm -Vp foo-1.0-1.i386.rpm
```

This can be useful if you suspect that your RPM databases are corrupt.

If everything verified properly there will be no output. If there are any discrepancies they will be displayed. The format of the output is a string of 8 characters, a possible "c" denoting a configuration file, and then the file name. Each of the 8 characters denotes the result of a comparison of one attribute of the file to the value of that attribute recorded in the RPM database. A single "." (period) means the test passed. The following characters denote failure of certain tests:

**5** MD5 checksum

**S** File size

**L** Symbolic link

- T** File modification time
- D** Device
- U** User
- G** Group
- M** Mode (includes permissions and file type)

If you see any output, use your best judgment to determine if you should remove or reinstall the package, or somehow fix the problem.

### 8.3 Impressing Your Friends with RPM

RPM is a very useful tool for both managing your system and diagnosing and fixing problems. The best way to make sense of all the options is to look at some examples.

- Let's say you delete some files by accident, but you aren't sure what you deleted. If you want to verify your entire system and see what might be missing, you would enter:

```
rpm -Va
```

If some files are missing, or appear to have been corrupted, you should probably either re-install the package or uninstall, then re-install the package.

- Let's say you run across a file that you don't recognize. To find out which package owns it, you would enter:

```
rpm -qf /usr/X11R6/bin/xjewel
```

The output would look like:

```
xjewel-1.6-1
```

- We can combine the above two examples in the following scenario. Say you are having problems with `/usr/bin/paste`. You would like to verify the package that owns that program but you don't know which package that is. Simply enter:

```
rpm -Vf /usr/bin/paste
```

and the appropriate package will be verified.

- If you are using a program and want to find out more information about it, you can enter the following to find out what documentation came with the package that “owns” that program (in this case `ispell`):

```
rpm -qdf /usr/bin/ispell
```

The output would be:

```
/usr/man/man4/ispell.4
/usr/man/man4/english.4
/usr/man/man1/unsq.1
/usr/man/man1/tryaffix.1
/usr/man/man1/sq.1
/usr/man/man1/munchlist.1
/usr/man/man1/ispell.1
/usr/man/man1/findaffix.1
/usr/man/man1/buildhash.1
/usr/info/ispell.info.gz
/usr/doc/ispell-3.1.18-1/README
```

- You find a new koules RPM, but you don’t know what it is. To find out some information on it, enter:

```
rpm -qip koules-1.2-2.i386.rpm
```

The output would be:

```
Name       : koules Distribution: Red Hat Linux Colgate
Version    : 1.2           Vendor: Red Hat Software
Release    : 2             Build Date: Mon Sep 02 11:59:12 1996
Install date: (none)     Build Host: porky.redhat.com
Group      : Games        Source RPM: koules-1.2-2.src.rpm
Size       : 614939
Summary    : SVGAlib action game; multiplayer, network
Description:
This arcade-style game is novel in conception and
excellent in execution. No shooting, no blood, no guts,
no gore. The play is simple, but you still must develop
skill to play. This version uses SVGAlib to run on a
graphics console.
```

- Now you want to see what files the koules RPM installs. You would enter:

```
rpm -qlp koules-1.2-2.i386.rpm
```

The output is:

```
/usr/man/man6/koules.6  
/usr/lib/games/kouleslib/start.raw  
/usr/lib/games/kouleslib/end.raw  
/usr/lib/games/kouleslib/destroy2.raw  
/usr/lib/games/kouleslib/destroy1.raw  
/usr/lib/games/kouleslib/creator2.raw  
/usr/lib/games/kouleslib/creator1.raw  
/usr/lib/games/kouleslib/colize.raw  
/usr/lib/games/kouleslib  
/usr/games/koules
```

These are just several examples. As you use the system you will find many more uses for rpm.

## 8.4 Other RPM Resources

For more information on RPM, check out the man page, the help screen (`rpm --help`), and the RPM documents available at

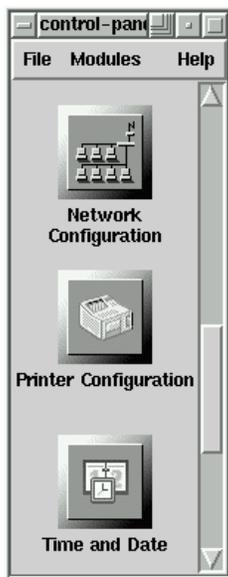
```
http://www.rpm.org/
```

There is also an RPM book available. It's called *Maximum RPM*, and is available from Red Hat Software and your local bookstore. It contains a wealth of information about RPM for both the end-user and the package builder. An on-line version of the book is available at <http://www.rpm.org/>.

There is also a mailing list for discussion of RPM related issues, called `rpm-list@redhat.com`. The list is archived on <http://www.redhat.com/support/mailling-lists/>. To subscribe, send mail to `rpm-list-request@redhat.com` with the word `subscribe` in the subject line.

# 9

## Control Panel



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

To start the control-panel, start the X Window System as root with `startx` and type `control-panel` in an xterm. You will need to be root to run the control-panel tools successfully. You can do this as well if you already have X running as a normal user. Just type `su -c control-panel` and then type the root password when prompted. If you plan to do other tasks as root, you could type `su -` followed by the root password when prompted. You will then be given a root shell. Here, you will need to type `DISPLAY=:0 control-panel &` to get the control panel to work. The "DISPLAY" part tells the control panel to use your display. This syntax works only for a Bourne-compatible shell like bash. If you run tcsh, you would enter `setenv DISPLAY :0; control-panel &`.

Double clicking on an icon starts up a tool. Please note that

you are not prevented from starting two instances of any tool, but doing so is a very bad idea because you may try to edit the same files in two places and end up overwriting your own changes. 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.

## 9.1 User and Group Configuration

The tool shown in Figure 9.1 manages the users and groups on your system. In a sense, it is a graphical editor for `/etc/passwd` and `/etc/group`, and for `/etc/shadow` if you have enabled shadow passwords. With this tool you can add and remove users and groups, set login shells, full names, home directories, uids, gids, etc.

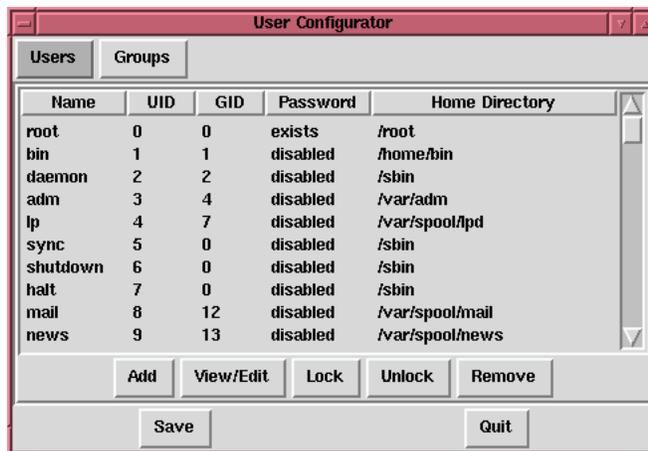
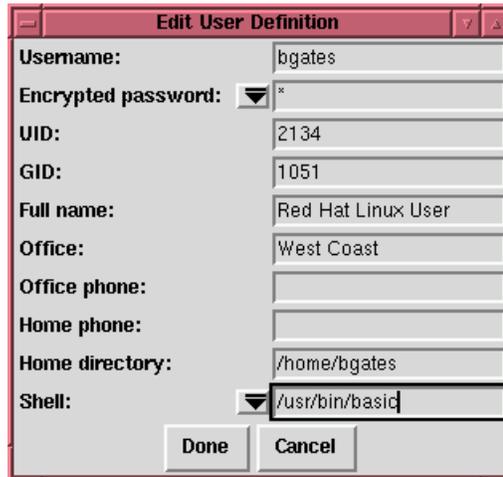


Figure 9.1: User and Group Configuration Panel

### 9.1.1 Adding or Editing a User

The same interface is used to add and edit users. To add a user, click on the **Add** button in the main window. A dialog box will pop up like the one in figure 9.2 will

appear, with some defaults already filled in. To edit an existing user, double-click on the user in the main window, or click on the user and then on the **View/Edit** button.



Username:	bgates
Encrypted password:	*
UID:	2134
GID:	1051
Full name:	Red Hat Linux User
Office:	West Coast
Office phone:	
Home phone:	
Home directory:	/home/bgates
Shell:	/usr/bin/basic

Figure 9.2: Add User Dialog Box

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 filled in with default values. You may change them if you wish, but it is not necessary. You can, at your option, enter the user's full name, office, and phone numbers. The login shell can be changed from the default (normally `/bin/bash`) either by choosing a selection from the drop-down menu or by simply typing in the shell like the other fields.

You need to do something about the password. There are five options on the **Encrypted Password** menu: Original, Change, No Password, Lock, and Unlock. No Password is a very bad option — with no password anyone can login to your system using this id. Choosing Lock will prevent anyone from logging in with this id but keep any password already assigned so that it can be unlocked later; Unlock will unlock a previously locked password. Original is blank (No Password) when you are adding a new user — the same dialog box is used for editing existing users, when Original is the user's original password. Usually you will choose Change. A small dialog box will appear where you will need to enter a password,

and then confirm it by typing it a second time. Click on **Done** in the password entry window to confirm that you have type the password correctly.

The numbers that the tool provides for **UID** and **GID** will work. If you change them, you are on your own; the tool is designed to assume that if you change them, you know what you are doing and want to change them. If you don't know what UID or GID stand for, leave them alone.

If you have configured shadow passwords on your system you will have a Shadow Management button that will allow you to configure password and account expiry. By default (fields shown as blank have the default value), no expiry will be done. If you are unfamiliar with shadow password account management, you will probably wish to accept the default. Click on **Done** when you are done editing the expiry information.

When you are done editing the user, click on **Done** in the Edit User Definition dialog box to add the user.

If this user does not already have a home directory (and they won't unless you have already created the home directory), one will be created.

### 9.1.2 Removing a User

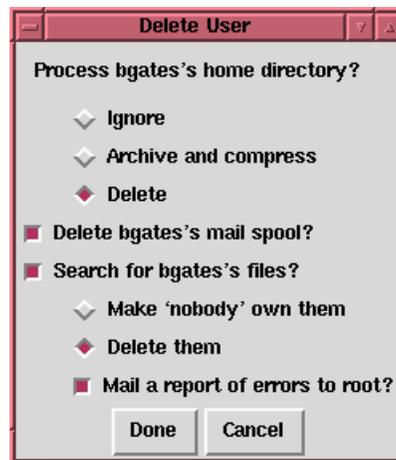


Figure 9.3: Delete User Dialog Box

To remove a user, select a user in the main window by clicking on it, and then click on **Remove**. You will be presented with the dialog box presented in Figure 9.3. There are three types of action you may take. You may ignore the user's home directory, leaving it alone, you can archive and compress it into a `.tar.gz` file, or you can delete it completely. You can delete the user's mail spool if you like. These first two actions will take place immediately after you confirm the deletion, before the tool accepts more input. You can search for files owned by the user on parts of the filesystem other than the user's home directory and either give ownership to the user `nobody` or delete them. If you choose to search for the user's files, you have the option of mailing a report of the errors to root. This is important because the searching will be done in the background as you continue to work; it may take a long time. To find files owned by `nobody`, 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 **Lock** and **Unlock** 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. When you lock a user, you are presented with the same options as when you delete a user, but you are unlikely to want to delete any files.

### 9.1.3 Creating a New Group

To create a new group, first switch to group editing mode by selecting the **Groups** tab at the top of the main window. 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. You can set a group password; if you do, users who are not members of the group may assume group membership by using the `newgrp` program and providing it with the password. To add members to the group, enter their names in the User list field. When you have selected all the group members, click **Done**

### 9.1.4 Editing an Existing Group

To edit an existing group, double-click on the group in the main window, or click on the group and then on **View/Edit**. The same interface is used to edit an existing group as to create a new group.

## 9.2 File System Configuration

The file system configuration tool shown in Figure 9.4 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. You probably don't want to do this unless you are an experienced system administrator. If you do know what you are doing, it should be fairly straightforward as to how to mount, unmount, and add devices.

The **Reload** entry in the **FSM** menu causes the file system configuration tool to re-load `/etc/fstab` from your hard drive. If you edit `/etc/fstab` by hand while the file system configuration tool is running (which you shouldn't do), you probably want to reload.

To use the buttons along the bottom of the window, select a filesystem in the main window, and click on one of the buttons. The buttons perform the following functions:

- **Info:** Displays information on the filesystem, including the device, partition type, filesystem type, mount point and options, comment, size, percent used, etc.
- **Check:** Performs a filesystem check (`fsck`) on the partition. You can do this only on *unmounted* partitions — if the partition is mounted, you will get an error and you will have to unmount it first.
- **Mount:** Mounts the selected filesystem.
- **Unmount:** Unmounts the selected filesystem.
- **Format:** Creates a new filesystem on the selected partition. This will erase all data on the selected partition! You can do this only on *unmounted* partitions — if the partition is mounted, you will get an error and you will have to unmount it first.

- **Edit:** Brings up a dialog box where you can edit the mount point, mount options, comment, etc.

### 9.2.1 Adding NFS Mounts

To mount a filesystem via NFS, select **Add Mount** from the **NFS** menu. A dialog box will appear and you will have to fill in the following values:

- **Device:** Enter the host name and path, separated by a colon. For example, `foo.bar.com:/usr/exported` indicates the `/usr/exported` directory on `foo.bar.com`.
- **Mount Point:** Enter the directory on your machine where you want to mount the NFS filesystem. For example, `/mnt/foo`.
- **Options:** Enter the mount options for this filesystem. The default is `soft, intr, rw`. The `rw` means the filesystem is read-write, and `soft, intr` are options that make your system a little more resilient when the remote server goes down. See the `mount` man page for a complete list of available options.
- **Comment:** This optional field can be used to store a small comment.

After filling everything out properly, click on **OK**. At this point the entry is made in your `/etc/fstab`, but the filesystem is not actually mounted. To mount it, select it in the main window and click on **Mount**.

## 9.3 Printer Configuration

The printer configuration tool (`printtool`) maintains `/etc/printcap`, print spool directories, and print filters. The filters allow you to print many different types of files, including:

- plain text (ASCII) files
- PostScript files
- TeX `.dvi` files

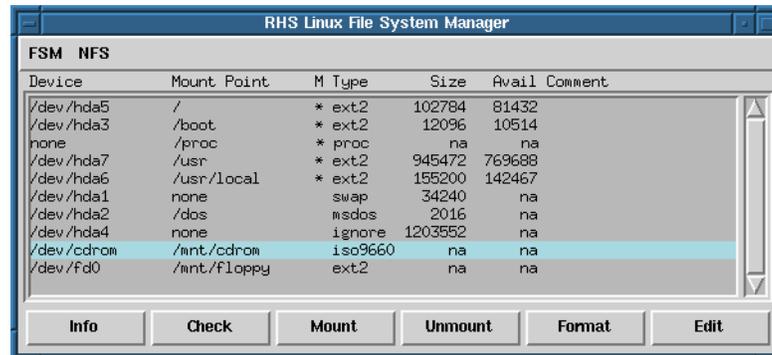


Figure 9.4: File System Configuration Panel

- GIF, JPEG, TIFF, and other graphics formats
- RPMs

Simply printing a GIF or RPM file using `lpr` does “the right thing”.

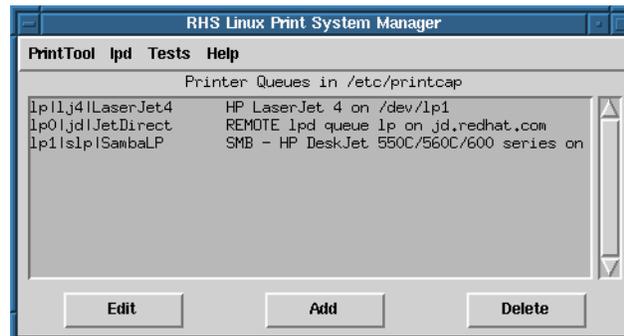


Figure 9.5: Print Tool

In order to create a new *print queue*, choose **Add**. Then, select what type of printer is being added. There are 3 types of print queues which can be configured with `printtool`:

- **Local**: print queues are for printers attached to a printer or serial port on your Red Hat Linux system.

- **Remote:** print queues are attached to a different system which you can access over a TCP/IP network.
- **SMB:** print queues are attached to a different system which uses LAN-Manager-type (SMB) networking.

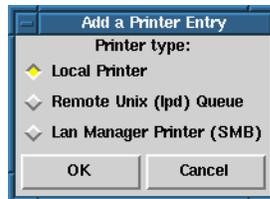


Figure 9.6: Selecting a Printer Type

After you choosing the printer type, a dialog requests further information about the print queue. All types of print queues require the following information:

- **Queue Name:** What the queue will be called. Multiple names can be specified with the | (pipe) character separating entries.
- **Spool Directory:** This is the directory on the local machine where files are stored before printing occurs. Be careful to not have more than one printer queue use a given spool directory.
- **File Limit:** Maximum size print job accepted, in kilobytes (1 kb = 1024 bytes). A size of 0 indicates no limit should be imposed.
- **Input Filter:** Filters convert printed files into a format the printer can handle. Press **Select** to choose the filter which best matches your printer.

In addition to configuring print queue able to print graphical and PostScript output you can configure a *text-only* printer, which will only print plain ASCII text. Most printer drivers are also able to print ASCII text without converting it to PostScript first; simply choose **Fast text printing** when you configure the filter. **Please Note:** This only works for non-PostScript printers.

- **Suppress Headers:** Check this if you don't want a header page to be printed at the beginning of each print job.

For *local* printers, the following are also required:

- **Printer Device:** Usually `/dev/lp1`, the name of the port which the printer is attached to. Serial printers are usually on `/dev/ttyS?` ports. You will need to manually configure serial parameters.

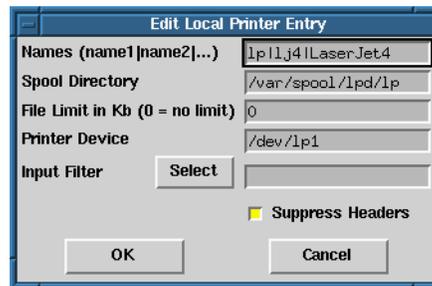


Figure 9.7: Adding a Local Printer

For *remote* printers, fill in the following information:

- **Remote Host:** Hostname of the remote machine hosting the printer.
- **Remote Queue:** Name of the queue to print to on the remote machine.

The remote machine must be configured to allow the local machine to print on the desired queue. Typically `/etc/hosts.lpd` controls this.

For *SMB* printers, fill in the following information:

- **Hostname of Printer Server:** Name of the machine to which the printer you want to use is attached.
- **IP number of Server:** The IP address of the machine to which the printer you want to use is attached; this is optional.
- **Printer Name:** Name of the printer on which you want to print.
- **User:** Name of user you must login as to access the printer (typically `guest` for Windows servers, or `nobody` for samba servers).
- **Password:** Password (if required) to use the printer (typically blank). Someone should be able to tell you this if you do not already know it.

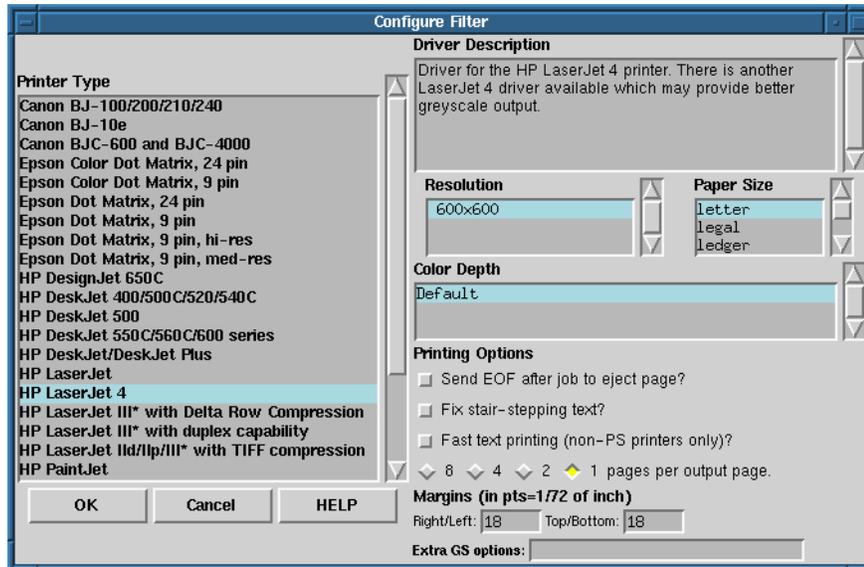


Figure 9.8: Configuring a Print Filter

**Please Note:** If you need to use a username and password for an SMB (LAN Manager) print queue, they are stored unencrypted in a local script and must be passed on the command line to the smbclient program. Thus, it is possible for another person to learn the username and password. It is therefore recommended that the username and password for use of the printer not to be the same as that for a user account on the local Red Hat Linux system, so that the only possible security compromise would be unauthorized use of the printer. If there are file shares from the SMB server, it is recommended that they also use a different password than the one for the print queue.

After you have added your print queue, you may need to restart the printer daemon (`lpd`). To do so, choose **Restart lpd** from the `lpd` menu.

You may print a *test page* for any print queue you have configured. Select the type of test page you would like to print from the **Tests** menu.

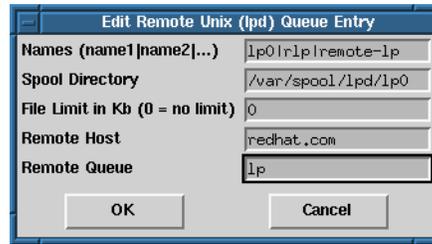


Figure 9.9: Adding a Remote Printer

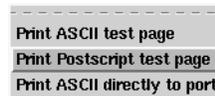


Figure 9.10: Printing a Test Page

## 9.4 Network Configuration

The network configuration tool (`netcfg`) shown in Figure 9.11 is designed to allow easy manipulation of parameters such as IP address, gateway address, and network address, as well as name servers and `/etc/hosts`.

Network devices can be added, removed, configured, activated, deactivated and aliased. Ethernet, arcnet, token ring, pocket (ATP), PPP, SLIP, PLIP and loopback devices are supported. PPP/SLIP/PLIP support works well on most hardware, but some hardware setups may exhibit unpredictable behavior. When using the Network Configuration Tool click **Save** to write your changes to disk, to quit without making any changes select **Quit**.

### 9.4.1 Managing Names

The **Names** panel of the Network Configuration tool serves two primary purposes: setting the hostname and domain of the computer, and determining which name server will be used to look up other hosts on the network. The Network tool is not capable of configuring a machine as a nameserver. To edit a field or add information to a field simply click on the field with the left mouse button and type the new information.

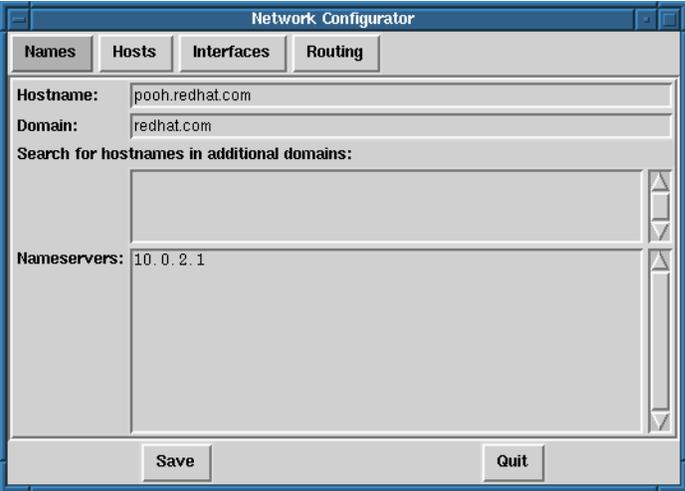


Figure 9.11: Network Configuration Panel

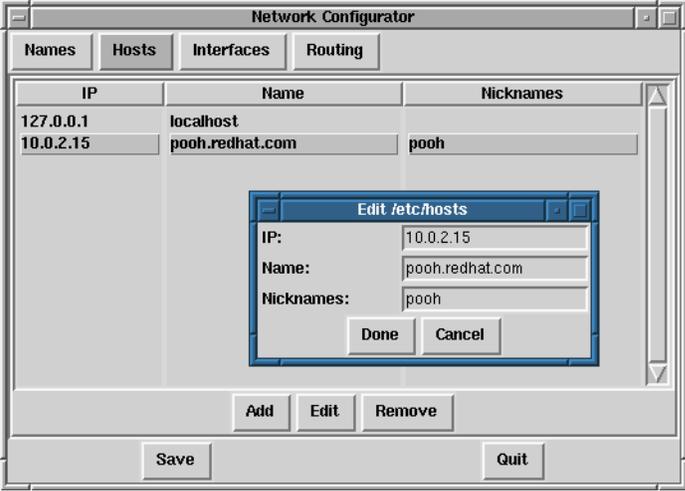


Figure 9.12: Adding/Editing Hosts

## 9.4.2 Managing Hosts

In the **Hosts** management panel you have the ability to add, edit, or remove hosts from the `/etc/hosts` file. Adding or editing an entry involves identical actions. An edit dialog box will appear, simply type the new information and click **Done** when you are finished. See Figure 9.12 for an example.

## 9.4.3 Adding a Networking Interface

If you have added a networking interface to your machine since installing Red Hat Linux, or you didn't configure your ethernet card at install time, you can configure it with a few clicks of a mouse.

**Please Note:** You may need to configure `kernel.d` to load a driver for the network interface you are adding (e.g., `eth0`); see Section 9.6 for more information.

Begin adding an interface by clicking on **Interfaces** in the main panel. This will bring up a window of configured devices with a row of available options, see figure 9.13.

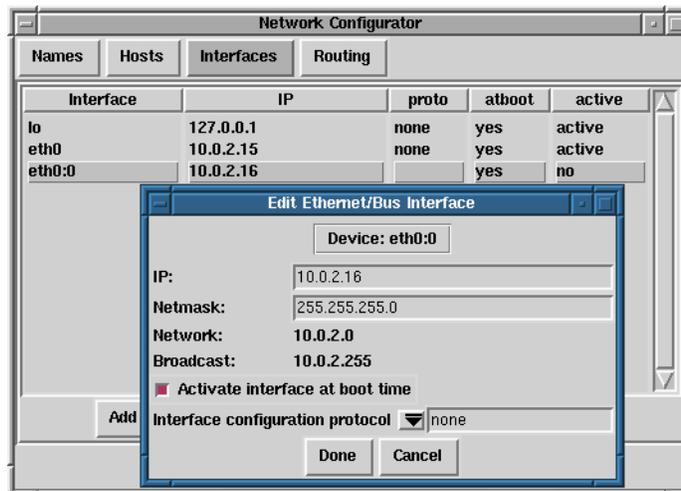


Figure 9.13: Configured Interfaces

To add a device, first click the **Add** button then select the type of interface you

want to configure from the box that appears (See Figure 9.14).



Figure 9.14: Choose Interface Type

### PPP Interface

Adding a PPP interface can be as simple as supplying the phone number, login name and password in the **Create PPP Interface** dialog shown in Figure 9.15. If you need to use PAP authentication for your PPP connection, choose **Use PAP authentication**. In many cases some degree of customization will be needed to establish a PPP connection. Choosing the **Customize** button will allow you to make changes to the hardware, communication, and networking settings for the PPP interface.

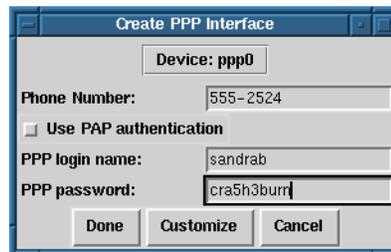


Figure 9.15: Create PPP Interface

### SLIP Interface

In order to configure a SLIP interface you must first supply a phone number, login name, and password. This will supply the initial parameters for the chat script needed to establish a SLIP connection. When you choose **Done**, a dialog titled **Edit SLIP Interface** appears that enables you to further customize the hardware, communication and networking parameters for your SLIP interface.

### PLIP Interface

To add a PLIP interface to your system you only have to supply the IP address, the remote IP address, and the Netmask. You can also select if you want to activate the interface at boot time.

### Ethernet, Arcnet, Token Ring and Pocket Adaptor Interfaces

If you are adding an ethernet, arcnet, token ring or pocket adapter to your computer you will need to supply the following information:

- **Device:** This is determined by netconfig based on the devices already configured.
- **IP Address:** Enter an IP address for your network device.
- **Netmask:** Enter the network mask for your network device.  
The network and broadcast addresses are calculated automatically based on the IP address and netmask you enter.
- **Activate interface at boot time:** If you want the device to be configured automatically when your machine boots select this by clicking on the box.
- **Allow any user to (de)activate interface:** Check this if you want any user to be able to activate or deactivate the interface.
- **Interface configuration protocol:** If you have a BOOTP or DHCP server on your network and would like to use it to configure the interface, choose the appropriate option; otherwise, choose **none**.

After providing the configuration information for your new device, click **Done**. The device should appear in your **Interfaces** list as an inactive device. (The active column should have a label of **no**). To activate the new device, first select it with a mouse click and then choose on the **Activate** button. If it does not come up properly, you may need to reconfigure it by choosing on **Edit**.

#### 9.4.4 Managing Routes

In the Routes management screen you have the ability to add, edit, or remove static networking routes. Adding or editing an entry involves identical actions, just like the Hosts panel. An edit dialog box will appear, simply type the new information and click **Done** when you are finished. See figure 9.16 for an example.

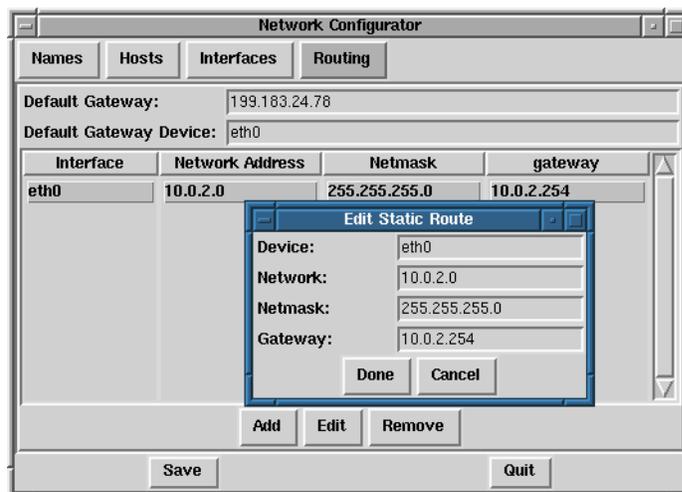


Figure 9.16: Adding/Editing Routes

## 9.5 Time and Date

The time machine allows you to change the time and date by clicking on the appropriate 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.

**Please Note:** Changing the time can seriously confuse programs that depend on the normal progression of time, and could possibly cause problems. Try to quit as many applications and processes as possible before changing the time or date.

## 9.6 Kernel Daemon Configuration

As discussed in Chapter 6, Red Hat Linux includes `kerneld`, the Kernel Daemon, which automatically loads some software and hardware support into memory as it is needed, and unloads it when it is no longer being used.

The tool shown in Figure 9.17 manages the configuration file for `kerneld`. While `kerneld` can load some things, such as filesystems, without explicit configuration, it needs to be told what hardware support to load when it is presented with a generic hardware request.

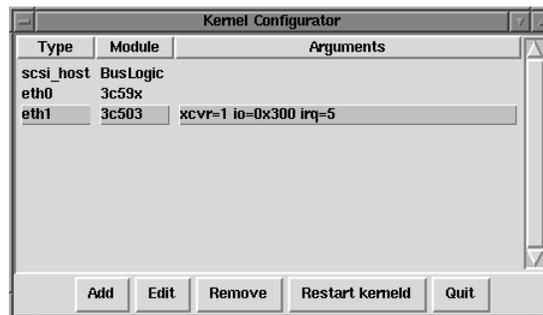


Figure 9.17: Kernel Module Management

For instance, when the kernel wants to load support for ethernet, `kerneld` needs to know which ethernet card you have, and if your ethernet card requires special configuration, it needs to know about that, too.

### 9.6.1 Changing Module Options

To change the options being given to a module when it is loaded, click on the line to select it, then click the **Edit** button. `kernelcfg` will bring up a window which looks like Figure 9.18. The options which `kernelcfg` knows about them (normally all available options) each have their own field. Normally, you will want to ignore the **Other arguments** field. Some modules normally take no arguments; just in case, they have an **Arguments** field which allows you to enter configuration information.

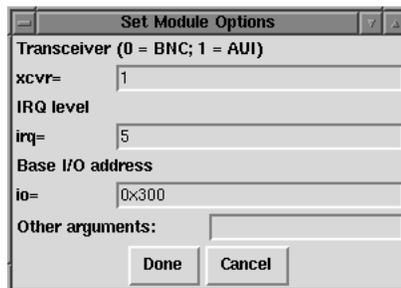


Figure 9.18: Editing Module Options

### 9.6.2 Changing Modules

To change which module gets invoked to provide a generic service, such as an ethernet card or SCSI host adapter module, you need to delete the old one and add a new one. To delete a module, select it by clicking on it, then click on **Delete**. Then click on **Add** to add the new module, as explained in the following section.

If you have changed your SCSI controller (`scsi_hostadapter`), remember to make a new initial ramdisk with the `/sbin/mkinitrd` command as documented in section 6.1.2.

### 9.6.3 Adding Modules

To add a module of any type, click on the **Add** button. You will be presented with a dialog box (Figure 9.19) asking you to choose a module type. Ethernet is `eth`,

Token Ring is `tr`, SCSI controllers are `scsi_hostadapter`, and so on. Click **OK** to continue to

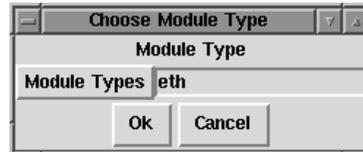


Figure 9.19: Adding a module

the next dialog box. If there is more than one module which can be used for the module type you have chosen, you will be presented with a dialog box (Figure 9.20) which asks which module you want to use, and may also ask for specifics about the type of module; for ethernet, for example, you need to choose from `eth0`, `eth1`, etc. When you are done, click **OK** again to continue to specify any module options in the next dialog box (Figure 9.20), which is the same as the dialog for editing a module.

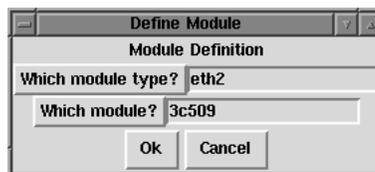


Figure 9.20: Selecting from available modules

#### 9.6.4 Restarting Kerneld

The changes that you make with the Kernel Daemon Configuration tool will be made in the `/etc/conf.modules` file, which `kerneld` reads whenever it is started. Once you have made changes, you can restart `kerneld` by clicking on the **Restart kerneld** button. This will **not** cause any modules which are currently in use to be reloaded, it will only notify `kerneld` to use the configuration when it loads more modules in the future.

# 10

## Glint

Red Hat provides a graphical tool to aid in package installation and removal. It's called glint (Graphical Linux INstallation Tool) and runs under the X Windows System. It allows easy installation, uninstallation, upgrading, querying, and verification of packages. The interface is similar to the one found in many popular file managers and is simple to use.

Operations are performed in glint by selecting the packages to operate on and then selecting the operation to perform via pushbuttons. Installing a package places all of the components of that package on your system. Uninstalling one removes all traces of the package except for configuration files you have modified. Upgrading a package installs the newly available version and uninstalls all other versions that were previously installed. This allows quick upgrading to the latest releases of packages.

The query operation lets you examine the details of both installed or available package. You can view the description of the package, where and when it was built, the files in the package, and other attributes. All of the configuration and documentation components of each package are clearly marked as such to reduce the time you spend looking for them.

Using glint to perform all of these operations is the same as using rpm to do them

from the command line. However, the graphical nature of glint often makes these operations easier to perform.

The normal way to handle glint is to display the available packages and files, select the ones you want to operate on, and then press a button or choose a menu item that performs the operation. For instance, you can install several packages with a few button clicks.

## 10.1 Starting glint

To start glint, simply run `glint &` from any X terminal window. That will bring up a window that looks like the one in figure 10.1. Any user can use glint to query and verify packages, but if you need to install, uninstall, or upgrade packages be sure to run glint as root.

There are two main parts to the glint interface. The first, on the left, allows you to browse and select the packages installed on your system. The right side contains buttons that manipulate the selected packages.

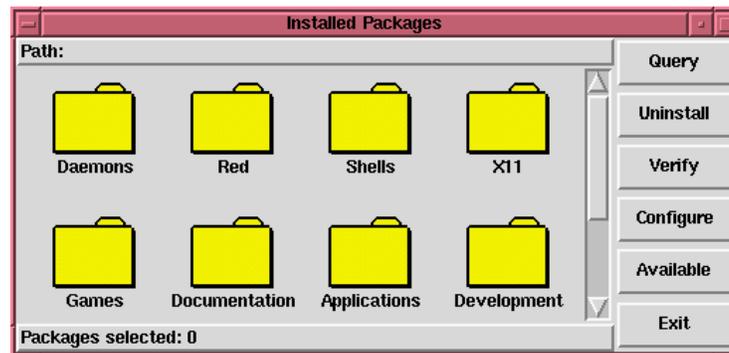


Figure 10.1: Main Glint Window

## 10.2 The Package Display

Each folder icon in glint represents a group of packages. Each group can contain other groups, which allows for flexible locations of packages. Groups are used to

place packages that perform similar functions in similar locations. For example, Red Hat includes many application programs such as editors and spreadsheets. All of the text based ones appear in the “Applications” group. Inside of that group, there is another grouping for all of the editors that are shipped.

By convention, groups are written in the same way as UNIX paths. The top most group is written first, and subsequent groups follow with a slash separating the group names. This means that an X-based drawing program appears in the X11/Applications/Graphics group.

To view the packages and subgroups within a group, double click the left mouse button on a group’s folder icon. The window then changes to show what that package contains. The top line of the package display shows which group you’re currently looking in, as well as the groups leading to the current one. To return to the previous group you were looking at, double click on the “Back” folder, which is always in the upper left hand corner of the folder area (though it often gets scrolled away).

If you’d like to examine a subgroup in a new window, double click the middle mouse button on its folder. If your mouse has only two buttons, click both. This will create a new window with that group in it.

### 10.2.1 Context Sensitive Menus

Pressing the right mouse button on any icon in the package window brings up a small, *context sensitive* menu. The exact items it contains depends on exactly where you press it. They all contain options to select or deselect the item, and many let you install, uninstall, query, upgrade, or verify the item you clicked on. There’s more information on how to do these things later.

To choose an item from a context sensitive menu, press and hold the right mouse button on a icon. While still holding the right button down, move the mouse pointer over the item you’d like to select (which will then become highlighted). Release the right mouse button to select that item and make the menu disappear.

### 10.2.2 Selecting Packages

To select a single package, click the left mouse button on it. You’ll notice a thin border appear around the package’s icon (as shown in figure 10.2) which shows that it’s currently selected. To unselect it, click the left mouse button on it and

the border will disappear. The number of packages currently selected is always displayed at the bottom of the window. A group's folder icon displays the number of packages inside of that group have been selected, or All if all of them have been.

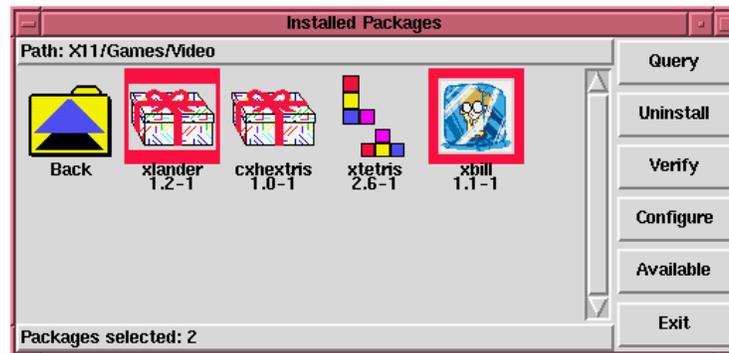


Figure 10.2: Selecting Packages in Glint

The context sensitive menu for a package also allows easy selection and unselection. Using the `select` and `unselect` options on a package's icon selects or unselects that package, while those options on a group's folder icon select and unselect all of the packages in that group. Using these menu options makes selecting groups of packages much quicker than selecting each package individually.

### 10.2.3 Viewing Available Packages

To see what packages are available for you to install, choose the `Available` push-button from any glint window. After a few moments, a new window, like the one shown in figure 10.3, will appear. The differences in the title and buttons indicate that this window is listing packages you may install. Navigating through these packages and selecting them is the same as in the other glint windows.

If you get an error message from glint saying that it can't find any RPMs, see the section below on Configuration.

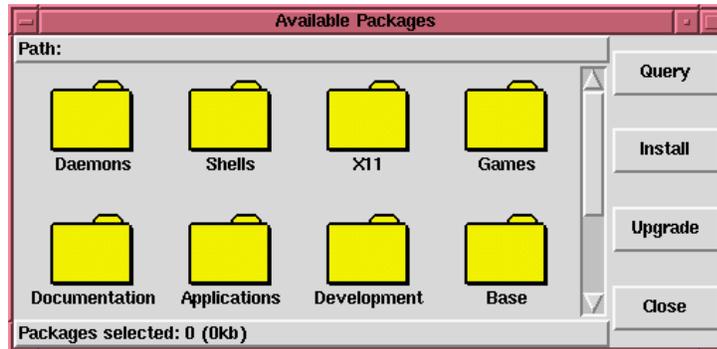


Figure 10.3: Available Window

## 10.3 Configuration

The only configuration information glint needs is the path to new RPMs. When you're using your Red Hat Linux CD-ROM, this will probably be `/mnt/cdrom/RedHat/RPMS`, which is the default path for glint. If you download new RPMs from the Internet or want to install RPMs via a NFS mounted CD-ROM this path will probably be different for you.



Figure 10.4: Configuration Window

To change this path, first be sure to close all of the windows listing available packages you may have open. Then choose the `Configuration` option from one of the remaining windows. This will open a dialog box like the one shown in figure 10.4. Here you can type the full path to the RPMs you'd like to look at. Choosing the `Save` button will save this path, making it the default for future glint sessions. The `Default` button restores the path to the one that glint used when it started.

After changing this path and closing the dialog box, you can use the `Available`

button to view the packages available in the new location.

## 10.4 Package Manipulation

### 10.4.1 Querying Packages

The easiest way to query a single package or group is to use the query option from the icon's context sensitive menu. If you want to query a more diverse set of packages, select them all and use the Query button in one of the windows.

Using either of these methods creates a window like the one shown in figure 10.5. If you choose only one package, it will look a bit different however, so some of this won't apply.

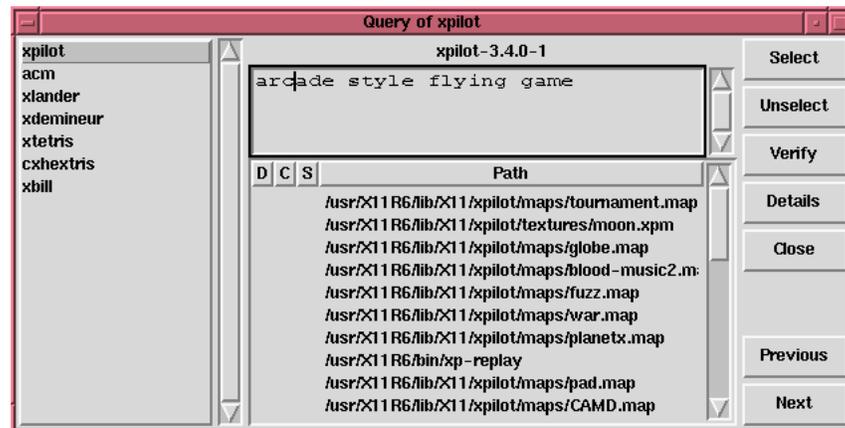


Figure 10.5: Query Window

On the very left of the window is a list of the packages that have been queried. Selecting one of them will change the information in the rest of the window. You may step through them in order by using the Next and Previous buttons on the right side of the window.

The name, version, and release of the current package are in the top middle of the query window. Immediately below this is the description of the package, which can be quite large. A scroll bar is there to let you read the whole thing.

Below the description is a list of the files contained in the package. Along with the full path to the file, the file list tells you a couple of other things. If a `D` appears to the left of the path, that file is a documentation file and would be a good thing to read. If a `C` appears there, then the file is a configuration file. A `*` means that the correct version of that file is not installed on your system. This can occur because a more recent version of a package was installed or because two packages contain different versions of the same file.



Figure 10.6: Query Details

More information on a package can be seen by clicking on the `Details` pushbutton. A window like the one in figure 10.6 will then appear. This lists more information about the package being displayed in the main query window. When you select a new package in the query window, the information in the details window will change to reflect your new choice.

You may also select, unselect, or verify a package while querying it by using the buttons provided. Click on the `Close` button when you are finished looking at the packages.

### 10.4.2 Verifying Packages

Verifying a package checks all of the files in the package to insure they match the ones present on your system. The checksum, file size, permissions, and owner attributes are all checked against the database. This check can be used when you suspect that one of programs files has become corrupted due to the installation of new programs.

Choosing the packages to verify is the same as choosing the packages to query.

Select the packages and use the `Verify` button or choose the `Verify` entry from a context sensitive menu. A window opens like the one in figure 10.7.



Figure 10.7: Verify Window

The three columns in this window describe the package with a problem in it, the file that has the problem, and a brief description of the discrepancies that were found. While the check is running, the current file being checked appears as the last element in the list, and the problem is listed as `(checking)`. A full list of the problems that can be found through verification appears in figure 10.8.

Problem	Description
missing	The file is no longer on your system
mode	permission bits have changed
size	file's size has changed
uid	owner's uid has changed
gid	owner's gid has changed
md5	the md5 checksum has changed
link	the file is a symlink to the wrong place

Figure 10.8: Possible Problems found by Verification

To get more information on the problems found with a file, double click on the file's path. A window like the one shown in figure 10.9 lists the expected and current values of the attributes that are amiss.

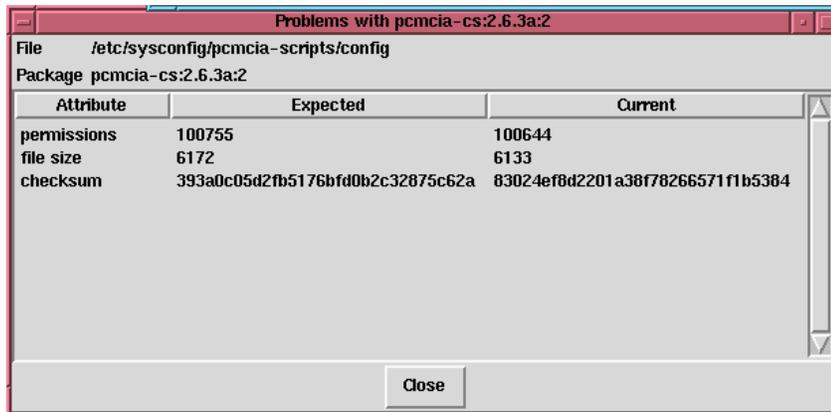


Figure 10.9: Verification Details

### 10.4.3 Installing New Packages

Installing new packages from glint is very simple. First look at the packages available for installation (see the section 10.2.3 for how to do this). You may select any number of these (and query them if you're not sure what they are) for installation (in the same manner you select packages for verification.) If you want to install a single package or group, the context sensitive menus provide a shortcut for doing so. Figure 10.10 shows a window with some packages selected for installation.

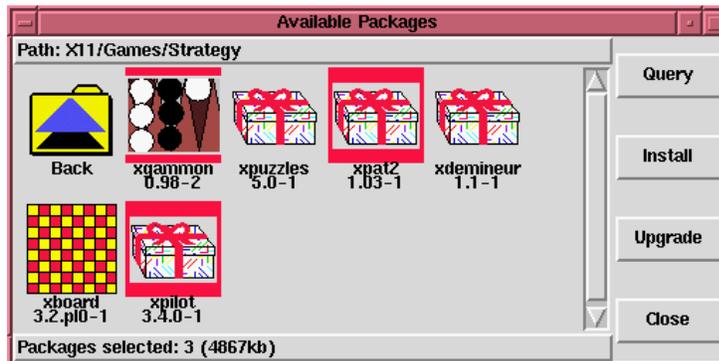


Figure 10.10: Packages Selected for Installation

After you've begun the installation, a window appears like the one in figure 10.11. It tracks the progress of the installation so you'll know something is happening. The top bar shows how much of the current package (whose name is listed inside of it) has been installed while the bottom graph shows how much of the total installation has been finished. The number of packages, package sizes, and time estimates are continually updated.

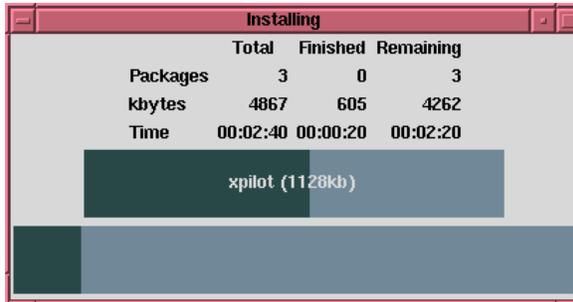


Figure 10.11: Installation Progress

If a problem occurs during the installation, a window will appear listing any errors that occurred. If this happens, you should correct the problems and then try again.

After the installation has completed, the package and groups that have been installed are moved from the available window to the main glint window to show you that they have been successfully installed.

### Upgrading Packages

When a new version of a package has been released, it is easy to install it on your system. Select the packages from the window of available packages in the same way you select packages for installation. Both the Upgrade button and the context sensitive menus will begin the upgrade.

During the upgrade, you'll see a progress indicator like the one for installing packages. When it's finished, the installed packages will appear in the the main glint windows and any old versions of the packages will be removed.

It is much better to use the upgrade option than to uninstall the old versions of a package and then install the new one. Using upgrade ensures that any changes you made to package configuration files get preserved properly, while doing it

manually could cause those changes to be lost.

If you run out of disk space during an installation, the install will fail. However, the package which was being installed when the error occurred may leave some files around. To clean this up, reinstall the package after you've made more disk space available.

#### 10.4.4 Uninstalling Packages

Uninstalling a package is not the same as upgrading one. When a package is uninstalled, any files it uses that are not needed by other packages on your system are removed. Changed configuration files are copied to `<filename>.rpmsave` so you can reuse them later.

Like verifying and querying packages, you can remove a package through the buttons on the right of the glint window or through a context sensitive menu. Remember that when you make a choice from a group's menu, the operation is performed on all of the packages in that group, so be careful!

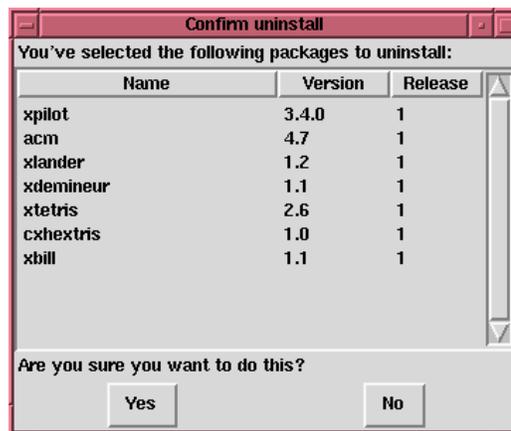


Figure 10.12: Uninstall Window

Once you've begun the uninstall, glint asks for confirmation, showing a window like the one in figure 10.12. All of the packages that are about to be uninstalled are listed. You should look at them all to ensure you're not about to remove something you want to keep. Clicking the `Yes` button will start the uninstallation process. Af-

ter it completes, the packages and groups that have been removed will disappear from any windows they were in.

# 11

## System Administration

This chapter is an overview of the Red Hat Linux system. It will illustrate things that you may not know about the system and things that are somewhat different from other UNIX systems. Note that most system administration tasks are performed with the control-panel, which is covered in chapter 9.

### 11.1 Filesystem Structure

Red Hat Software is committed to the Linux File System Standard, a collaborative document that defines the names and locations of many files and directories. We will continue to track the standard to keep Red Hat compliant.

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 `/usr` partition contains common executables and is not meant to be changed by users. Because of this, the `/usr` partition can be mounted from the CD-ROM or from another machine via read-only NFS. The current Linux Filesystem Standard (FSSTND) document is the authoritative reference to any FSSTND compliant filesystem, but the standard leaves many areas unde-

finer or extensible. In this section we provide an overview of the standard and a description of the parts of the filesystem not covered by the standard.

The complete standard can be viewed at:

`http://www.pathname.com/fhs/`

### 11.1.1 Overview of the FSSTND

The directories and files noted here are a small subset of those specified by the FSSTND document. Check the latest FSSTND document for the most up to date and complete information.

#### The `/etc` Directory

The `/etc` directory is reserved for configuration files that are local to your machine. No binaries are to put in `/etc`. Binaries that were in the past put in `/etc` should now go into `/sbin` or possibly `/bin`.

The `x11` and `skel` directories should be subdirectories of `/etc`:

```
/etc
|- x11
+- skel
```

The `x11` directory is for X11 configuration files such as `XF86Config`. The `skel` directory is for “skeleton” user files, which are files used to populate a home directory when a user is first created.

#### The `/lib` Directory

The `/lib` directory should contain only those libraries that are needed to execute the binaries in `/bin` and `/sbin`.

### The /sbin Directory

The /sbin directory is for executables used only by the root user, and only those executables needed to boot and mount /usr and perform system recovery operations. The FSSTND says:

“/sbin typically contains files essential for booting the system in addition to the binaries in /bin. Anything executed after /usr is known to be mounted (when there are no problems) should be placed in /usr/sbin. Local-only system administration binaries should be placed into /usr/local/sbin.”

At a minimum, the following programs should be in /sbin:

```
clock, getty, init, update, mkswap, swapon,  
swapoff, halt, reboot, shutdown, fdisk, fsck.*,  
mkfs.*, lilo, arp, ifconfig, route
```

### The /usr Directory

The /usr directory is for files that are shareable across a whole site. The /usr directory usually has its own partition, and it should be mountable read only. The following directories should be subdirectories of /usr:

```
/usr  
|- X11R6  
|- bin  
|- dict  
|- doc  
|- etc  
|- games  
|- include  
|- info  
|- lib  
|- local  
|- man  
|- sbin
```

```
| - share  
+- src
```

The `X11R6` directory is for the X Window System (XFree86 on Red Hat Linux), `bin` is for executables, `doc` is for random non-man-page documentation, `etc` is for site-wide configuration files, `include` is for C header files, `info` is for GNU info files, `lib` is for libraries, `man` is for man pages, `sbin` is for system administration binaries (those that do not belong in `/sbin`), and `src` is for source code.

### The `/usr/local` Directory

The FSSTND says:

*“The `/usr/local` hierarchy is for use by the system administrator when installing software locally. It needs to be safe from being overwritten when the system software is updated. It may be used for programs and data that are shareable amongst a group of machines, but not found in `/usr`.”*

The `/usr/local` directory is similar in structure to the `/usr` directory. It has the following subdirectories, which are similar in purpose to those in the `/usr` directory:

```
/usr/local  
| - bin  
| - doc  
| - etc  
| - games  
| - include  
| - info  
| - lib  
| - man  
| - sbin  
+- src
```

### The `/var` directory

Since the FSSTND requires that you be able to mount `/usr` read-only, any programs that write log files or need spool or lock directories probably should write them to the `/var` directory. The FSSTND says `/var` is for

“... variable data files. This includes spool directories and files, administrative and logging data, and transient and temporary files.”

The following directories should be subdirectories of `/var`:

```
/var
|- log
|- catman
|- lib
|- local
|- named
|- nis
|- preserve
|- run
|- lock
|- tmp
+- spool
    |- at
    |- cron
    |- lpd
    |- mail
    |- mqueue
    |- rwho
    |- smail
    |- uucp
+- news
```

System log files such as `wtmp` and `lastlog` go in `/var/log`. The `/var/lib` directory also contains the RPM system databases. Formatted man pages go in `/var/catman`, and lock files go in `/var/lock`. The `/var/spool` directory has subdirectories for various systems that need to store data files.

### 11.1.2 /usr/local in Red Hat Linux

In Red Hat Linux, the intended use for `/usr/local` is slightly different from that specified by the FSSTND. The FSSTND says that `/usr/local` should be where you store software that is to remain safe from system software upgrades. Since system upgrades from Red Hat Software are done safely with the RPM system and Glint, you don't need to protect files by putting them in `/usr/local`. Instead, we recommend you use `/usr/local` for software that is local to your machine.

For instance, let's say you have mounted `/usr` via read-only NFS from *beavis*. If there is a package or program you would like to install, but you are not allowed to write to *beavis*, you should install it under `/usr/local`. Later perhaps, if you've managed to convince the system administrator of *beavis* to install the program on `/usr`, you can uninstall it from `/usr/local`.

## 11.2 Special Red Hat File Locations

In addition to the files pertaining to the RPM system that reside in `/var/lib/rpm` (see chapter 8), there are two other special locations that are reserved for Red Hat 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`, stores 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.

## 11.3 Users, Groups and User-Private Groups

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

The easiest way to manage users and groups is through the Users and Groups module of the control-panel (see section 9 for details on the control-panel and sec-

tion 9.1 for details on the Users and Groups module).

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

### 11.3.1 Standard Users

Table 11.3.1 lists the standard 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 11.3.3 for details on how groups are used.

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

Figure 11.1: Standard Users

### 11.3.2 Standard Groups

Table 11.3.2 lists the standard groups as set up by the installation process (this is essentially the `/etc/group` file).

<b>Group</b>	<b>GID</b>	<b>Members</b>
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	ftp
nobody	99	
users	100	

Figure 11.2: Standard Groups

### 11.3.3 User Private Groups

Red Hat Linux uses a user private group (UPG) scheme, which makes UNIX groups much easier to use. The UPG scheme does not add or change anything in the standard UNIX way of handling groups. It simply offers a new convention for handling groups. Whenever you create a new user, by default, he or she has a unique group. The scheme works as follows:

**User Private Group** Each user has its own primary group, to which only it is a member.

**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. The umask is set in `/etc/profile`.

**SGID bit on Directories** If you set 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.

Most computing sites like to create a group for each major project and assign people to the groups they need to be in. Managing files traditionally has been difficult, though, because when someone creates a file it is owned by the primary group he or she belongs to. When a single person works on multiple projects, it becomes hard to make the files owned by the group that is associated with that project. In the UPG scheme, groups are automatically assigned to files on a project-by-project basis, which makes managing group projects very simple.

Let's say you have a big project called *devel*, with many people editing the devel files in a `devel` directory. Make a group called `devel`, `chgrp` the `devel` directory to `devel`, and add all the `devel` users to the `devel` group. Now, all the `devel` users will be able to edit the `devel` files and create new files in the `devel` directory, and these files will always retain their `devel` group. Thus, they will always be edit-able by other `devel` users.

If you have multiple projects like *devel*, 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 the SGID bit would be redundant.

### User Private Group Rationale

Since the UPG scheme is new, many people have questions about it, and they wonder why it is necessary. The following is the rationale for the scheme.

- 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 enter:

```
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 enter:

```
chmod 775 /usr/lib/emacs/site-lisp
```

- But when a user creates a new file it is assigned the group of the users default group (usually `users`). To prevent this you enter

```
chmod 2775 /usr/lib/emacs/site-lisp
```

which causes everything in the directory to be created with the "emacs" group.

- But the new file needs to be mode 664 for another user in the emacs group to be able to edit it. To do this you make the default umask 002.
- Well, this all works fine, 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 that 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.

## 11.4 User Authentication with PAM

Programs which give users access to privileges of any sort need to be able to authenticate the users. When you log into a system, you provide your name and password, and the login process uses those to authenticate the login—to verify that you are who you say you are. Other forms of authentication than passwords are possible, and it is possible for the passwords to be stored in different ways.

PAM, which stands for “Pluggable Authentication Modules”, is a way of allowing the system administrator to set authentication policy without having to recompile programs which do authentication. With PAM, you control how the modules are plugged into the programs by editing a configuration file.

Most Red Hat Linux users will never need to touch this configuration file. When you use RPM to install programs that need to do authentication, they automatically make the changes that are needed to do normal password authentication. However, you may want to customize your configuration, in which case you need to understand the configuration file.

### 11.4.1 Modules

There are four types of modules defined by the PAM standard. `auth` modules provide the actual authentication, perhaps asking for and checking a password, and set “credentials” such as group membership or kerberos “tickets”. `account` modules check to make sure that the authentication is allowed (the account has not expired, the user is allowed to log in at this time of day, etc.). `password` modules are used to set passwords. `session` modules are used once a user has been authenticated to make it possible for them to use their account, perhaps mounting the user’s home directory or making their mailbox available.

These modules may be *stacked*, so that multiple modules are used. For instance, `rlogin` normally makes use of at least two authentication methods: if “`rhosts`” authentication succeeds, it is sufficient to allow the connection; if it fails, then standard password authentication is done.

New modules can be added at any time, and PAM-aware applications can then be made to use them. For instance, if you have a one-time-password calculator system, and you can write a module to support it (documentation on writing modules is included with the system), PAM-aware programs can use the new module and work with the new one-time-password calculators without being recompiled or

otherwise modified in any way.

## 11.4.2 Services

Each program which uses PAM defines its own “service” name. The login program defines the service type `login`, `ftpd` defines the service type `ftp`, etc. In general, the service type is the name of the program used to **access** the service, not (if there is a difference) the program used to **provide** the service.

## 11.4.3 The Configuration Files

The directory `/etc/pam.d` is used to configure all PAM applications. (This used to be `/etc/pam.conf` in earlier PAM versions; while the `pam.conf` file is still read if no `/etc/pam.d/` entry is found, its use is deprecated.) Each application (really, each **service**) has its own file. A file looks like this:

```
##PAM-1.0
auth      required /lib/security/pam_securetty.so
auth      required /lib/security/pam_pwdb.so shadow nullok
auth      required /lib/security/pam_nologin.so
account   required /lib/security/pam_pwdb.so
password  required /lib/security/pam_cracklib.so
password  required /lib/security/pam_pwdb.so shadow
                                     nullok use_authok
session   required /lib/security/pam_pwdb.so
```

The first line is a comment. Any line that starts with a # character is a comment. The next three lines stack up three modules to use for login authorization. The first line makes sure that *if* the user is trying to log in as root, the tty on which they are logging in is listed in the `/etc/securetty` file *if* that file exists. The second line causes the user to be asked for a password and the password checked. The third line checks to see if the file `/etc/nologin` exists, and if it does, displays the contents of the file, and if the user is not root, does not let him or her log in.

Note that all three modules are checked, *even if the first module fails*. This is a security decision—it is designed to not let the user know why their authentication was disallowed, because knowing why it was disallowed might allow them to break the authentication more easily. You can change this behavior by changing

required to requisite; if any requisite module returns failure, PAM fails immediately without calling any other modules.

The fifth line causes any necessary accounting to be done. For example, if shadow passwords have been enabled, the `pam_pwdb.so` module will check to see if the account has expired, or if the user has not changed his or her password and the grace period for changing the password has expired.

The sixth line (which we've had to wrap) specifies that if the login program changes the user's password, it should use the `pam_pwdb.so` module to do so. (It will do so only if an `auth` module has determined that the password needs to be changed—for example, if a shadow password has expired.)

The final line specifies that the `pam_pwdb.so` module should be used to manage the session. Currently, that module doesn't do anything; it could be replaced (or supplemented by stacking) by any necessary module.

Note that the order of the lines within each file matters. While it doesn't really matter much in which order `required` modules are called, there are other *control flags* available. While `optional` is rarely used, and never used by default on a Red Hat Linux system, `sufficient` and `requisite` cause order to become important.

Let's look at the `auth` configuration for `rlogin`:

```
auth required /lib/security/pam_securetty.so
auth sufficient /lib/security/pam_rhosts_auth.so
auth required /lib/security/pam_pwdb.so shadow nullok
auth required /lib/security/pam_nologin.so
```

That looks *almost* like the `login` entry, but there's an extra line specifying an extra module, and the modules are specified in a different order.

First, `pam_securetty.so` keeps root logins from happening on insecure terminals. This effectively disallows all root `rlogin` attempts. If you wish to allow them (in which case we recommend that you either not be internet-connected or be behind a good firewall), you can simply remove that line.

Second, `pam_nologin.so` checks `/etc/nologin`, as specified above.

Third, if `pam_rhosts_auth.so` authenticates the user, PAM immediately returns success to `rlogin` without any password checking being done. If `pam_rhosts_auth.so` fails to authenticate the user, that failed authentication is ignored.

Finally (if `pam_rhosts_auth.so` has failed to authenticate the user), the `pam_pwdb.so` module performs normal password authentication.

Note that if you do not want to prompt for a password if the security check fails, you can change the `pam_securetty.so` module from `required` to `requisite`

#### 11.4.4 Shadow Passwords

Shadow password support has changed significantly for Red Hat Linux 5.0. Please refer to Section 1.3.2 on page 7 for up-to-date information.

The `pam_pwdb.so` module will automatically detect that you are using shadow passwords and make all necessary adjustments.

#### 11.4.5 More Information

This is just an introduction to PAM. More information is included with the system in `/usr/doc/pam*`, including a *System Administrators' Guide*, a *Module Writers' Manual*, an *Application Developers' Manual*, and the PAM standard, DCE-RFC 86.0. In addition, documentation is available from the Red Hat web site, at <http://www.redhat.com/linux-info/pam/>.

### 11.5 The Boot Process, Init, and Shutdown

#### 11.5.1 System V Init

This section is a brief description of the internals of the boot process. It basically covers in detail how the machine boots using SysV Init and the differences between the original init used in older Linux releases.

Init is the program that gets run by the kernel at boot time. It is in charge of starting all the normal processes that need to run at boot time. These include the gettys that allow you to log in, NFS daemons, FTP daemons, and anything else you want to run when your machine boots.

SysV Init is fast becoming the standard in the Linux world to control the startup of software at boot time. This is because it is easier to use and more powerful and

flexible than the traditional BSD init.

SysV init also differs from BSD init in that the config files are in a subdirectory of `/etc` instead of residing directly in `/etc`. This directory is called `rc.d`. In there you will find `rc.sysinit` and the following directories:

```
init.d
rc0.d
rc1.d
rc2.d
rc3.d
rc4.d
rc5.d
rc6.d
```

`init.d` contains a bunch of scripts. Basically, you need one script for each service you may need to start at boot time or when entering another runlevel. Services include things like networking, `nfs`, `sendmail`, `httpd`, etc. Services do not include things like `setserial` that must only be run once and then exited. Things like that should go in `rc.local` or `rc.serial`.

If you want `rc.local`, it should be in `/etc/rc.d`. Most systems include one even though it doesn't do much. You can also include an `rc.serial` in `/etc/rc.d` if you need to do serial port specific things at boot time.

The chain of events is as follows:

- The kernel looks in several places for `init` and runs the first one it finds
- `init` runs `/etc/rc.d/rc.sysinit`
- `rc.sysinit` does a bunch of necessary things and then runs `rc.serial` (if it exists)
- `init` runs all the scripts for the default runlevel.
- `init` runs `rc.local`

The default runlevel is decided in `/etc/inittab`. You should have a line close to the top like:

```
id:3:initdefault:
```

From this, you'd look in the second column and see that the default runlevel is 3, as should be the case for most systems. If you want to change it, you can edit `/etc/inittab` by hand and change the 3. Be very careful when you are messing with the inittab. If you do mess up, you can fix it by rebooting and doing:

```
LILO boot: linux single
```

This *should* allow you to boot into single user mode so you can fix inittab.

Now, how does it run all the right scripts? If you enter `ls -l` on `rc3.d`, you might see something like:

```
lrwxrwxrwx 1 root root 17 3:11 S10network -> ../init.d/network
lrwxrwxrwx 1 root root 16 3:11 S30syslog -> ../init.d/syslog
lrwxrwxrwx 1 root root 14 3:32 S40cron -> ../init.d/cron
lrwxrwxrwx 1 root root 14 3:11 S50inet -> ../init.d/inet
lrwxrwxrwx 1 root root 13 3:11 S60nfs -> ../init.d/nfs
lrwxrwxrwx 1 root root 15 3:11 S70nfsfs -> ../init.d/nfsfs
lrwxrwxrwx 1 root root 18 3:11 S90lpd -> ../init.d/lpd.init
lrwxrwxrwx 1 root root 11 3:11 S99local -> ../rc.local
```

What you'll notice is that there are no real "files" in the directory. Everything there is a link to one of the scripts in the `init.d` directory. The links also have an "S" and a number at the beginning. The "S" means to start this particular script and a "K" would mean to stop it. The number is there just for ordering purposes. Init will start all the services based on the order they appear. You can duplicate numbers, but it will only confuse you somewhat. You just need to use a two digit number only, along with an upper case "S" or "K" to start or stop the services you need to.

How does init start and stop services? Simple. Each of the scripts is written to accept an argument which can be "start" and "stop". You can execute those scripts by hand in fact with a command like:

```
/etc/rc.d/init.d/httpd.init stop
```

to stop the httpd server. Init just reads the name and if it has a "K", it calls the script with the "stop" argument. If it has an "S" it calls the script with a "start" argument. Why all these runlevels? Some people want an easy way to set up machines to be multi-purpose. I could have a "server" runlevel that just runs httpd, sendmail, networking, etc. Then I could have a "user" runlevel that runs xdm, networking, etc.

### 11.5.2 Init Runlevels

Generally, Red Hat Linux runs in run level 3—full multiuser mode. The following runlevels are used in Red Hat 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 log in because you have a corrupted `/etc/passwd` or have simply forgotten your password, boot into single user mode by typing `linux 1` at the LILO boot prompt. A very bare system will come up and you will be given a shell from which you can fix things.

### 11.5.3 Shutting Down

To shut down Red Hat Linux, issue the `shutdown` command. You can read the `shutdown` man page for complete details, but the two most common usages are:

```
shutdown -h now
shutdown -r now
```

Each will cleanly shutdown the system. After shutting everything down, the first will halt the machine, and the second will reboot.

Do not run the `reboot` or `halt` commands directly in order to prevent damage to your filesystem.

## 11.6 Rescue Modes

When things go wrong, there are several ways to work on fixing them. However, they require that you understand the system well. This manual can't teach you what to do, but we will present the ways that you can use our products to get into rescue modes where you can use your own knowledge to rescue the system.

### 11.6.1 Through LILO

If your system boots, but does not allow you to log in when it has completed booting, you can use the `single` or `emergency` boot option. At the `LILO boot:` prompt, type `linux single` in order to boot in single-user mode. In single-user mode, your local filesystems will be mounted, but your network will not be activated. In emergency mode, almost nothing will be set up. Only the root filesystem will be mounted, and it will be mounted read-only.

### 11.6.2 Emergency Boot Diskettes

The installation diskette set can also be used as a rescue disk set. When you boot the boot diskette, type `rescue` at the boot prompt (on non-Intel platforms, add `rescue` as an option to the kernel you are booting, for example, `linux rescue`), and the install process will ask a few questions, request that you insert the second disk, and set up shells on VC 1 and VC 2. The shell that is running is `ash`, a minimal bourne shell with no history or command-line editing.

The `PATH` environment variable is set so that once you have mounted your root partition on `/mnt`, binaries from your root partition will be run if they do not exist on the diskette. Also, the libraries in `/mnt/lib` and `/mnt/usr/lib` will be used if needed. You may wish to change your `PATH` to not include `/bin` and `/usr/bin`, leaving `/mnt/bin`, `/mnt/usr/bin`, etc., while you are working on repairing your system. If you do that, you will need to use `/bin/umount` to unmount `/mnt`.

Limited versions of important utilities are available before you mount `/mnt`:

`mount` The `mount` command does not require that the device node being mounted already exist, and doesn't require the full name of the device. It assumes `ext2` filesystems unless told otherwise.

```
mount /dev/sda1 /mnt -t ext2
```

is functionally equivalent to

```
mount sda1 /mnt
```

`ash` The full version

`cat` Does not accept any options

---

`chmod` The `chmod` command does not accept symbolic mode names, only octal ones.

`cpio` The full version

`e2fsck` The full version

`fdisk` The full version

`gzip` The full version

`gunzip` The full version

`insmod` The full version

`ls` The full version

`lsmod` Does not accept any options

`mkdir` No GNU long options

`mke2fs` The full version

`mknod` The `mknod` command does not accept symbolic mode names, only octal ones.

`open` The full version

`rm` No GNU long options

`rmmmod` The full version

`sh` Symlink to `ash`

### A Handy Trick

Have you ever rebuilt a kernel and, eager to try out your new handiwork, re-booted before running LILO? And you didn't have an entry for an older kernel in `lilo.conf`? Read on...

Here's a handy trick. In many cases, it's possible to boot your Red Hat Linux/Intel from the Red Hat Linux boot diskette with your root filesystem mounted and ready to go. Here's how:

Enter the following command at the boot diskette's boot : prompt:

```
linux single root=/dev/hdXX initrd=
```

(Replace the XX in /dev/hdXX with the appropriate letter and number for your root partition.)

What does this do? First, it starts the boot in single-user mode, with the root partition set to your root partition. The empty `initrd` specification bypasses the installation-related image on the boot diskette, which will cause you to enter single-user mode immediately.

Is there a downside to this trick? Unfortunately, yes. Because the kernel on the Red Hat Linux boot diskette only has support for IDE built-in, those of you with SCSI-based systems won't be able to use this trick. In that case, you'll have to use the boot diskette's normal rescue mode (which also requires the supplemental diskette).

# A

## Getting Technical Support

This chapter discusses Red Hat Software's support:

- What it is
- How to get it
- Frequently asked questions

### **A.1 An Overview of Our Support Policy**

Red Hat Software provides 30-day installation support for people that have purchased the Official Red Hat Linux product.

Red Hat will provide support to registered purchasers of the Red Hat Linux Boxed Set. This support will be provided by means of electronic mail. In the case that the user doesn't have access to e-mail, fax support will be provided for those requests submitted with complete registration numbers.

In order to receive support it is necessary to register the product via the World

Wide Web at <http://www.redhat.com/support/register> or by sending mail including the registration number to [register@redhat.com](mailto:register@redhat.com).

After registering successfully, support may be obtained by sending a specially formatted message to [support@redhat.com](mailto:support@redhat.com). The message format is described in the support HOW-TO, located at:

<http://www.redhat.com/support/register/support-how2.html>

This e-mail will then be delivered to Red Hat Software support where it will be distributed to a Red Hat Support Engineer.

As the name implies, installation support centers on helping you successfully install Red Hat Linux on your computer. This includes support in three distinct areas:

- Answering questions you may have prior to installation, such as:
  - Hardware compatibility issues.
  - Basic approaches to partitioning your hard drive.
- Helping you get through the installation process:
  - Getting any supported hardware necessary for installation recognized by the installation program.  
(See <http://www.redhat.com/hardware/>)
  - Assisting with the creation of a root and swap partition using the free space available on your hard drive.
  - Using the installation program to configure LILO to boot Red Hat Linux, and one other operating system (such as DOS, Windows 95, or Windows NT) already residing on your hard drive.
- Assisting with final configuration tasks, such as:
  - The successful configuration of the X window system on supported hardware, using either the Metro-X or XFree86 software. (Additional configuration, such as automatically starting X on reboot, or changing/customizing window managers is your responsibility.)
  - Configuring a printer connected directly to the Red Hat Linux system, enabling it to print text.
  - Setting up a mouse to be used with the text-based console, or with the X window system.

- Getting access to the CD-ROM, so that information can be read from it.

Naturally, although our installation support service will get your system running Red Hat Linux, there are many other optional tasks that you might want to undertake, such as compiling a customized kernel, adding support for devices not included in the installation process, and so on.

For assistance with these tasks, please consider the many books on Linux at your local bookstore, or various on-line resources. A starting point in your on-line search for Linux information should always be:

<http://www.redhat.com/support/>

for information specific to Red Hat Linux, or:

<http://www.redhat.com/linux-info/>

for more general Linux information. Another good resource is

<http://www.linux.org/>.

Red Hat Software can only support customers that have purchased the official Red Hat Linux Box set. If you have obtained Red Hat from any other publisher, you must contact them for support. Examples of other publishers would be:

- Macmillan/Sams
- Cheapbytes
- Pacific Hi-Tech (PHT)
- Infomagic
- Linux Systems Labs (LSL)
- ADRAS Computing

Also, RHL-Intel obtained via any of the following methods does not qualify for support from Red Hat Software:

- Red Hat Linux PowerTools Archive
- Downloaded via FTP on the Internet.
- Included in a package such as Applixware or Motif.
- Copied or installed from another user's CD.

## A.2 Getting Support

As of July 1996, Red Hat Software put a new Technical Support system online. The support system automatically routes questions to support technicians. This document describes how to use the Red Hat Software technical support system.

### A.2.1 Registration

In order to receive technical support for your Red Hat Software product, you have to register it. Every official Red Hat product comes with a Registration Card. Your product registration number appears on both the detachable mail-in portion and on the top portion of this card. The product registration number uniquely identifies a product which you have purchased (such as Red Hat Linux) that includes 30 days of free installation technical support by fax or email.

#### Registering via the Web

Registering via the Web

You can register your Red Hat Software product online at Red Hat's World Wide Web site at <http://www.redhat.com/support/register/>. Choose **Register a Product**. Please enter all applicable information, and please be accurate with the system information. This information will aid in solving problems and answering questions more quickly and easily, and incorrect information benefits neither us nor our customers. Also, make sure the electronic mail address you give is correct. All support correspondence will be sent to that address. If this should change, please login to the registration pages and edit this information.

Registration has changed. There will be no more annoying support ID numbers to remember. If you don't have a current login to the support database, you will need to create one. Simply enter a name at the **Support Login** and click on the **Create Login** button to begin registration of your product. If you were already registered for support before August 14, 1997 you can login in by using your email address as your Support Login and your former support/customer ID as the password. The password can be changed once logged in. If you don't remember it, follow the steps as if you never had a login, and re-enter the registration numbers and information.

### Registering via Email

If you don't have access to the World Wide Web, but you do have access to Internet mail, you can send registration details to `register@redhat.com`. Please include the following information:

- Contact Information: name, address, city, state/province, country, zip/postal code, and telephone and fax numbers; also, your support ID if you have one;
- Product Information: name of product you are registering, product registration number (very important!), architecture you're using the product on (Intel, Alpha/AXP, or Sparc), and what you're using the product for (File Server, Internet Server, Workstation, Home Computer, or Other);
- System Information: CPU (i386, i486, i586, i686/PentiumPro, AXP21164, AXP21064, Sparc4C, Sparc4M), CPU speed (MHz), amount of RAM (MB), hard drives (number, size, type, partitions), video card, CD-ROM drive, any network card, and other hardware components (SCSI or ISDN adaptors, etc.).

### Registering via Fax or Snail Mail

Although a mail-in portion is provided with your registration card, we do not recommend registering via fax or snail mail if you have access to either the World Wide Web or Internet mail; if you do so, your registration is likely to be delayed considerably. If you do wish to register via fax, fax a copy of your registration card to Red Hat Software at +1-919-361-2711.

## A.2.2 Support Questions

Technical support is a black art: in most cases, support technicians rely solely on communication with the customer to solve installation problems on hardware the technicians have never physically seen. It is extremely important, therefore, to state your question clearly and concisely, including detailed information such as:

- symptoms of the problem ("Red Hat Linux 5.0 doesn't recognize my CD-ROM drive.")
- when the problem began ("It stopped working yesterday ...")

- what changes you have made to your system around the time of the problem (“... after I disconnected it from the IDE controller.”)
- any diagnostic output specifically related to the problem (“In the bootup messages, it says ‘Cannot find /dev/hdb; device disconnected.’”) however, this can be taken overboard; don’t send us your system logs unless we ask you for them.
- other relevant information (“I’m using the floppy installation method, from the CD.”)

### How to Send Them

Red Hat’s support system is email-based and is partially automated; for this reason, it is important to make sure you send support questions in the correct format, so that your message will be recognized and routed to an appropriate support technician. In order to receive technical support for your Red Hat Software product, you must first register it.

Submitting trouble tickets can be sent in the traditional email way, or can be initiated via the World Wide Web. After providing the information for the registration, you will have a new menu allowing you to edit your registration information or open a trouble ticket. Hurry and login now to open a ticket on the Web! <http://www.redhat.com/support/register/>

To open a support ticket via email, all you need to do is send the support system a message with a Subject: of [registration #nnnn nnnn nnnn nnnn], where nnnn nnnn nnnn nnnn is the registration number of the product for which you are requesting support. For example, if your registration number is fffe 0fff ff00 ff00, the subject line should read:

```
Subject: [registration #fffe 0fff ff00 ff00]
```

The square brackets, the number sign, and the word ‘registration’ must be present. If you wish, you may add explanatory text to the subject line:

```
Subject: [registration #fffe 0fff ff00 ff00] CD-ROM problem
```

Once you’ve opened a ticket, support responses come to you with the support ticket number in the subject line:

Subject: [ticket #12015] CD-ROM problem

To correspond about the same problem, simply send a reply with the same subject:

Subject: Re: [ticket #12015] CD-ROM problem

If you feel that the problem has been solved the ticket can also be closed by your. Simply add close tag. to the same subject:

Subject: Re: [close ticket #12015] CD-ROM problem

Once your problem is solved or your question answered, the technician handling your support ticket can close the ticket. This can also be done by you via the website. Either method will result in a message from the support system stating the ticket is closed and contain a summary of the problem and the solution. You can open a new ticket for your next support question. All past correspondence will be saved in the database under the old ticket number and can be accessed with your account at any time.

### **Where to Send Them**

The address for Red Hat's Technical Support System is [support@redhat.com](mailto:support@redhat.com); all support questions should go there. There are also several related addresses:

- For a copy of the Red Hat FAQ (Frequently Asked Questions, with answers), send mail to [faq@redhat.com](mailto:faq@redhat.com).
- For a copy of the Red Hat PPP Tips (tips on setting up a PPP connection), send mail to [ppp-tips@redhat.com](mailto:ppp-tips@redhat.com).
- To report a bug, send mail to [bugs@redhat.com](mailto:bugs@redhat.com).

## **A.3 Support FAQ (Frequently Asked Questions)**

Here are a few questions that the Red Hat Support Staff see frequently, along with the answers:

**Question** – I've sent several messages to the Red Hat Support System, and I've gotten absolutely no response. Is anybody there?

**Answer** – Check your registration information to make sure your email address is correct. You can check it by logging in at <http://www.redhat.com/support/register/>.

**Question** – I know I have already registered, but I keep getting a message from the support system telling me I'm not registered. Is the system broken?

**Answer** – Make sure to register via the World Wide Web at <http://www.redhat.com/support/register/>. If you didn't please do so, by first accessing the registration page. Next, input a login name and click on **Create Login**. After this you will be prompted to input some personal information. When finished click on the **Submit** button and then enter your registration number which should look something like `ffff 0fff ff00 ff00`. Then be sure to enter all the information about your system. This information can be maintained by logging in and accessing your registration information from <http://www.redhat.com/support/register/>.

**Question** – But I registered via email. Why isn't the support system working?

**Answer** – When you register via email, your message goes to an actual human who registers you manually; this can delay your registration anywhere from several hours to several days.

**Question** – I've done everything perfectly for my registration, but I'm still getting messages telling me I have an invalid registration number.

**Answer** – Please make sure you're using a valid product registration number. If you have not purchased the official box set from Red Hat Software or a Vendor that sells the official version (it will say 'official' on the box) you will not be able to register, as Red Hat Software cannot support products packaged and sold by other software publishers.

**Question** – I see this thing on the registration page that asks me what my login is? What is my login anyway?

**Answer** – If you are a first time user of the the support system or accessed the support system after August 15, 1997 then the Login can be whatever you want it to be. It must be all one word and if there is already a duplicate you will be warned of an error. If you accessed the system before this time and have a valid registration, your login will be the email address you had when

registering and the password will be the old Support ID number. You can change the password if you like; however the login cannot be changed.

**Question** – I've tried entering a login and password but it won't accept it. Why won't it accept the password I'm trying to enter?

**Answer** – The password must be a minimum of 5 alphanumeric characters long and no more than 8. Make sure you typed in the exact same password both times for both password fields on the registration page.

**Question** – I hear that it's possible to change my registration information and view all my old tickets and the past correspondence. How do I do that?

**Answer** – Simply login to the registration page at <http://www.redhat.com/support/register> and all these options will be available to you. If you need to reference an old trouble ticket, change your system information, or just check up on the status of a trouble ticket, visit this page. You can also open a trouble ticket with valid registration numbers directly from the web from these pages.



# B

## Making Installation Diskettes

It is sometimes necessary to create a diskette from an *image file* (for example, you might need to use updated diskette images obtained from the Red Hat Linux Errata).

As the name implies, an image file is a file that contains an exact copy (or image) of a diskette's contents. Since a diskette contains filesystem information in addition to the data contained in files, the image file is not usable until it has been written to a diskette.

To start, you'll need a blank, formatted, high-density (1.44 MB), 3.5-inch diskette. You'll need access to a computer with a 3.5-inch diskette drive, and capable of running a DOS program, or the `dd` utility program found on most Linux-like operating systems.

The image files are found in the following directories on your Red Hat Linux CD:

- `images` – Contains the boot and supplemental images for Red Hat Linux/Intel, and the various kernel and ramdisk images for Red Hat Linux/Alpha.  
On the Red Hat Linux/SPARC CD, this directory contains the boot image, and an image for network booting.

- `mi1o`—Contains the various images for the Red Hat Linux/Alpha miniloader, MILO. This directory exists only on Red Hat Linux/Alpha CDs.

Once you've selected the proper image, it's time to transfer the image file onto a diskette. As mentioned previously, this can be done on a DOS-capable system, or on a system running a Linux-like operating system.

## B.1 Making a Diskette Under MS-DOS

To make a diskette under MS-DOS, use the `rawrite` utility included on the Red Hat Linux CD in the `dosutils` directory. First, label a blank, formatted 3.5-inch diskette appropriately (eg. "Boot Diskette", "Supplemental Diskette", etc). Insert it into the diskette drive. Then, use the following commands (assuming your CD is drive `d:`):

```
C:\> d:
D:\> cd \dosutils
D:\dosutils> rawrite
Enter disk image source file name:  ..\images\boot.img
Enter target diskette drive:  a:
Please insert a formatted diskette into drive A: and
press --ENTER-- :  
D:\dosutils>
```

`rawrite` first asks you for the filename of a diskette image; enter the directory and name of the image you wish to write (for example, `..\images\boot.img`). Then `rawrite` asks for a diskette drive to write the image to; enter `a:`. Finally, `rawrite` asks for confirmation that a formatted diskette is in the drive you've selected. After pressing  to confirm, `rawrite` copies the image file onto the diskette. If you need to make another diskette, label another diskette, and run `rawrite` again, specifying the appropriate image file.

## B.2 Making a Diskette Under a Linux-like O/S

To make a diskette under Linux (or any other Linux-like operating system), you must have permission to write to the device representing a 3.5-inch diskette drive (known as `/dev/fd0` under Linux). First, label a blank, formatted diskette appropriately (eg. "Boot Diskette", "Supplemental Diskette", etc). Insert it into the diskette drive (but don't issue a `mount` command). After mounting the Red Hat Linux CD, change directory to the directory containing the desired image file, and use the following command (changing the name of the image file and diskette device as appropriate):

```
# dd if=boot.img of=/dev/fd0 bs=1440k
```

If you need to make another diskette, label another diskette, and run `dd` again, specifying the appropriate image file.



# C

## Information Specific to Red Hat Linux/SPARC

This appendix describes the differences between Red Hat Linux/SPARC and Red Hat Linux/Intel installations. While it provides a good overview of these differences, it does not completely document the installation process. For that, you will need to read chapters 2, 3, and 6 in order. These chapters will refer you to the appropriate parts of this appendix at the appropriate time.

In addition, there is a Linux/SPARC homepage at [http://www.geog.ubc.ca/s\\_linux.html](http://www.geog.ubc.ca/s_linux.html). It has a wealth of information for people considering Red Hat Linux/SPARC, and is a great resource.

### C.1 Supported Hardware

Red Hat Linux/SPARC supports a variety of hardware based on the SPARC architecture. The most recent list of hardware supported by Red Hat Linux/SPARC can be found at Red Hat Software's World Wide Web site at <http://www.redhat.com/hardware>.

The following list is current as of the time this manual was produced:

- sun4c architecture machines (IPC, SS1, etc)
- sun4m architecture machines (Classic, SS5, SS10, etc)
- bwtwo, cg3, cg6, TCX framebuffer (24 bit on the TCX)
- cg14 framebuffer (in cg3 mode)
- SCSI and Ethernet on all of the above
- type 4 and type 5 keyboards and mice
- external SCSI drives
- CD-ROM drives (external and internal)
- SCSI/Ether SBUS expansion cards
- Any original equipment Sun monitor for the above framebuffers

**Unsupported Hardware** – The following list contains hardware that is currently unsupported:

- VME based sun4m machines, such as the 4/690
- SPARC 5 Model 170 machines
- sun4d (SS1000, SS2000)
- sun4u (UltraSparc1, UltraSparc2, EnterPrise UltraSparc servers)
- sun4 architecture
- Eurocard sun4c machines, called the 'sun4e' under Solaris2.4 and others, are not supported. These are VME-bus sun4c machines, to which the kernel has not been ported yet.
- The following types of CD-ROMs are unsupported (for more information on CD-ROM compatibility issues, please visit <http://saturn.tlug.org/suncdfaq>):
  - All NEC models
  - Some Toshiba XM-4101B revisions

- Toshiba XM-3201B
- AppleCD 300 Plus on older SPARC PROM
- Apple CD600i

## C.2 Installation Overview

Installing Red Hat Linux on a SPARC system is slightly more complex than installing Red Hat Linux/Intel, mostly due to differences in machine architecture. In general, the main steps to a successful installation are:

1. Determine which console commands your SPARC systems supports.
2. Determine whether you have sufficient memory to use a ramdisk-based installation.
3. Determine how you will boot the installation program.
4. Determine what installation method you will use.
5. Determine whether you will install Red Hat Linux/SPARC from a serial terminal.
6. Load and run the Red Hat Linux installation program.

Let's look at each of these steps in more detail.

## C.3 Console Commands

SPARC systems have two different types of boot commands available, depending on the version of the system's console PROM. Systems with a PROM version less than 2.0 use what is known as the "old" style boot command, while systems whose PROM is at version 2.0 or greater can use either the old or the new style boot commands. If your SPARC system's console prompt is "ok", then your system is in new command mode, and can use the new style boot command.

In general, it's preferable to use the new style boot command if your system supports it. Of course, it's possible to use the old style boot command, but you'll need

to know a bit more about your hardware configuration. An excellent reference on older Sun hardware in general (and the old-style boot commands in particular) can be found in *The Sun Hardware Reference*, by James W. Birdsall. It can be found at `ftp://ftp.picarefy.com/pub/Sun-Hardware-Ref/parts/`. The file `part2` contains an in-depth description of the various “ROM monitors” present in older Sun Systems.

In general, the new style boot command is `boot`, followed by a descriptive device name (such as `floppy`, `cdrom`, or `net`).

The old style boot command is `b`, followed by a device specifier in the form:

```
xx(a,b,c)
```

Where `xx` is used to specify the device type (such as `sd` for SCSI disks), `a` is the controller number, `b` is the device’s unit number, and `c` is the partition number. Please refer to *The Sun Hardware Reference* if you have an older SPARC system, and require additional information on its boot command syntax.

## C.4 Ramdisk-based Installation Criteria

When the Red Hat Linux/SPARC installation starts, normally a ramdisk is loaded into memory. This ramdisk contains data and programs required to perform the installation, and is approximately 4 MB in size. Since the ramdisk takes memory away from the Red Hat Linux/SPARC installation program, your SPARC system must have at least 12 MB of memory in order to successfully use a ramdisk-based installation.

For SPARC systems with less than 12 MB of RAM, another approach is available. It is known as an *NFS-mounted root*. As the name implies, instead of using a ramdisk to hold parts of the Red Hat Linux/SPARC installation program, an NFS server is used to export the necessary files to the SPARC system. While this approach uses less memory, it is somewhat more complex to set up.

**Please Note:** Red Hat Linux/SPARC may also be installed by booting from the Red Hat Linux/SPARC CD-ROM. This method uses neither a ramdisk or an NFS-mounted root, as the necessary files are already present on the Red Hat Linux/SPARC CD-ROM. We will discuss booting from CD-ROM in Section C.5.2 on page 202.

### C.4.1 Preparing to Use an NFS-Mounted Root

If you will be using an NFS-mounted root for the installation program, you'll need a network connection for your SPARC system, an NFS server capable of exporting the Red Hat Linux/SPARC CD-ROM (or equivalent files), and the ability to respond to `rarp` requests, so that your SPARC system can obtain its IP address. The necessary commands for `rarp` would be:

```
rarp -s ip.address.of.sparc hw:address:of:sparc
arp -s ip.address.of.sparc hw:address:of:sparc
```

As you might surmise, replace `ip.address.of.sparc` with the IP address of your SPARC system, and replace `hw:address:of:sparc` with the MAC address of your SPARC system's network adapter. The IP address is assigned by your network administrator, and the MAC address is displayed on your SPARC system's console at power-up.

To use an NFS-mounted root during the installation, add the following argument to your boot command:

```
linux nfsroot=nfs.server.ip.address:/path/to/RH/image
```

The word `linux` should always be included in a boot command that includes one or more arguments, and must precede the first argument. Replace `nfs.server.ip.address` with the NFS server's IP address, and `/path/to/RH/image` with the path to the exported directory containing the appropriate Red Hat Linux/SPARC files.

## C.5 Choosing a Boot Method

This section describes how to start (or *boot*) the installation program. Once the installation program is running, you will be able to choose from several installation methods. You can choose from the following installation methods: CD-ROM, NFS, hard disk, and FTP. (Note that if the installation program is booted directly from CD-ROM, the installation will automatically proceed from that CD-ROM.)

There are three different ways a Red Hat Linux/SPARC installation can be started:

**Boot From Diskette** – The installation program is read from a diskette.

**Boot From CD-ROM** – The installation program is read directly from the Red Hat Linux/SPARC CD-ROM.

**Boot From the Network** – The installation program is read from a TFTP server.

### C.5.1 Booting From Diskette

If your SPARC system has a diskette drive, you can boot the Red Hat Linux/SPARC installation program from a diskette. The boot diskette image (known as `boot.img`) is located in the `images/` directory on your Red Hat Linux/SPARC CD-ROM. Please refer to Appendix B on page 193 for instructions on writing the image file to a diskette. Make sure you label the diskette “Boot Diskette”.

#### Diskette Boot Commands

For SPARC systems with a PROM version of 2.0 or greater, the proper boot command (when in new command mode) is:

```
boot floppy
```

On the other hand, SPARC systems with PROM versions less than 2.0 should use the following command at the `>` prompt:

```
b fd()
```

**Please Note:** There have been reports that some systems with pre-2.0 PROMs cannot boot the Red Hat Linux/SPARC installation program from diskette. If you find this to be the case with your SPARC system, you will need to use another boot method.

### C.5.2 Booting From CD-ROM

If your SPARC system has a fully Sun-supported CD-ROM drive, you can boot directly from the Red Hat Linux/SPARC CD-ROM. For SPARC systems with a PROM version of 2.0 or better, use the following command when in new command mode:

```
boot cdrom
```

SPARC systems with PROM versions less than 2.0 may not be able to boot from a CD-ROM at all. Theoretically, if your SPARC system has a CD-ROM at SCSI id 6, the following command should boot the Red Hat Linux/SPARC installation program:

```
b sd(0,6,0)
```

Unfortunately, due to a lack of resources, Red Hat Software has not been able to test this boot command. If one of our readers has been successful getting a pre-2.0 SPARC system booted from the Red Hat Linux/SPARC CD-ROM, please send us mail at [docs@redhat.com](mailto:docs@redhat.com), and we'll update this manual. Thank you!

Note that using an NFS-mounted root after booting from CD-ROM is not supported, as the filesystem on the Red Hat Linux/SPARC CD-ROM performs the same function as an NFS-mounted root. Therefore, no additional boot command arguments should be given for CD-ROM boots.

### C.5.3 Booting From the Network

There are two types of network boots supported by the Red Hat Linux/SPARC installation program:

1. Network boot with NFS-mounted root. This method is required for SPARC systems with less than 12 MB of RAM.
2. Network boot with network-loaded ramdisk. This method can be used by systems with at least 12 MB of RAM.

While booting your SPARC system from the network is fairly straightforward, there are several requirements:

- Your SPARC system must have a network connection.
- Your network must be able to give your SPARC system its IP address via a `rarp`.

- You must have a TFTP server that can download the Red Hat Linux/SPARC kernel and installation program to your SPARC system.
- If you are going to use an NFS-mounted root, an NFS server capable of exporting the Red Hat Linux/SPARC CD-ROM (or equivalent files).

**Setting up RARP** If you are going to use `rarp`, please refer to Section C.4.1 on page 201.

**TFTP Server Setup** If you are going to set up a TFTP server on a Red Hat Linux system, simply install the latest `tftp` package using RPM, and make sure the line in `inetd.conf` that will run `tftp` is uncommented. Don't forget to `kill -HUP inetd` if you needed to make any changes to `inetd.conf`.

Next, you'll need to make a symlink describing the SPARC system to be booted, and pointing to the file from which it should boot. The name of the symlink contains two items:

1. The IP address of the system to be booted, in hexadecimal.
2. A string describing the architecture of the system to be booted.

To convert the more common "dotted decimal" IP address into its hex equivalent, convert each of the address' four groups of numbers into hex. If the resulting hex number is only one digit, add a leading zero to it. Then append all four hex numbers together. For example, take the IP address 10.0.2.254. Convert each set of four numbers into hex, and add a leading zero where necessary:

```
10 = A  or 0A
 0 = 0  or 00
 2 = 2  or 02
254 = FE or FE
```

Therefore, the IP address 10.0.2.254 in hex is: 0A0002FE.

The second part of the symlink name is the SPARC system's architecture. For our example, we'll use `SUN4M`. The IP address and architecture are separated by a dot, resulting in this symlink name:

```
0A0002FE.SUN4M
```

The last step is figuring out what this symlink should point to. There are two choices. If you want to use an NFS-mounted root, use the file `/kernels/vmlinux`. If you would rather use a ramdisk, use the file `/images/tftpboot.img`.

Place the appropriate file in the TFTP server's directory, and create the symlink. In this example, we're using the image that includes a ramdisk:

```
ln -s tftpboot.img 0A0002FE.SUN4M
```

### Network Boot Commands

You're now ready to boot. If you're going to boot `tftpboot.img`, simply use the following command (in new command mode):

```
boot net
```

On the other hand, if you're going to boot from `vmlinux` and use an NFS-mounted root, use this command:

```
boot net linux nfsroot=nfs.server.IP.address:/path/to/RH/image
```

Replace `nfs.server.ip.address` with the NFS server's IP address, and `/path/to/RH/image` with the path to the exported directory containing the appropriate Red Hat Linux/SPARC files.

SPARC systems with PROM versions less than 2.0 should use this boot command, appending the usual NFS root argument if required:

```
b le()
```

## C.6 Choosing an Installation Method

Once your SPARC systems has booted, and the installation program is running, you'll be asked to choose an installation method (unless you've booted directly from CD-ROM, in which case a CD-ROM installation method is assumed). Red Hat Linux/SPARC can be installed by any of the following methods:

- Installing packages from CD-ROM.
- Installing packages from an FTP site.
- Installing packages from an NFS server.
- Installing packages from a locally-attached hard disk.

### **C.6.1 CD-ROM Installation**

This is the most straightforward method. It requires a Red Hat Linux/SPARC CD-ROM, and a Sun-supported CD-ROM connected to your SPARC system.

### **C.6.2 FTP Installation**

This installation method requires a local area network connection and access to an FTP site with the Red Hat Linux/SPARC CD-ROM (or equivalent files).

### **C.6.3 NFS Installation**

Installing via NFS requires a local area network connection and access to an NFS server that can export the contents of the Red Hat Linux/SPARC CD-ROM (or equivalent files).

### **Hard Disk Installation**

This installation method requires that the contents of the Red Hat Linux/SPARC CD-ROM (or equivalent files) have been copied to a hard disk directly attached to your SPARC system. It is important to note that the partition holding these files cannot be used for any other purpose during the installation (ie, it cannot be given a mount point during the installation). In addition, the partition must be in ext2 format.

## C.7 Installation Using a Serial Terminal

You can also install Red Hat Linux/SPARC using a serial terminal attached to your SPARC system. Any terminal which can emulate a VT100 (or a computer with terminal emulation software) will work fine. Boot as you would normally, and at the SILO prompt enter:

```
linux serial
```

The installation program runs on the first serial port at 9600 baud, 8 bits, no parity, 1 stop bit (often called 9600, 8, N, 1). The installation program can run in color, if your serial terminal supports color. Note that a computer running kermit and connected to your SPARC system will display in color.

At any prompt dialog during a serial installation (any dialog with an **OK** button), you can press **Ctrl-Z** to start a subshell. To return to the installation program, enter `exit` at the shell prompt.

When the installation is complete, simply boot normally, and Red Hat Linux should come up on your serial terminal.

## C.8 Installation: Selecting System Components

(For an explanation of system components, please see Section 4.6.1 on page 67.)

**Important!** When selecting system components to install for Red Hat Linux/SPARC, choosing to install **Everything** without choosing **Select individual packages** could result in serious problems, and possibly an unbootable Red Hat Linux/SPARC system. Here's why:

- The last kernel to be installed will be a multi-processor kernel, which will fail to work on sun4c machines.
- The default X server will end up as `XsunMono`, which may not work for most people.

If you do wish to install **Everything**, choose **Select individual packages** as well. Then, deselect all the kernels except for one you wish to install (`kernel-sparc`

works for most) and all X servers except the one for your hardware (XsunMono for bwtwo video cards, Xsun for all 8-bit displays, and Xsun24 for the supported 24-bit displays).

## C.9 SILO Configuration

SILO configuration is nearly identical to LILO configuration. See Section 5.8 on page 82 for more details.

**Please Note:** Unlike LILO, SILO cannot be configured to boot other operating systems during the installation. However, information on configuring SILO to boot other operating systems is available. Please read the file `README` in `/usr/doc/silo*/docs` after installation.

## C.10 Partitioning

There is one additional step required when partitioning a hard drive for Red Hat Linux/SPARC. You must create the third partition of every disk as type `Whole Disk` spanning from cylinder 0 to the end of the disk. It shouldn't be used in any way, but it must exist. You can still create other partitions as you normally would.

Note that this partition will already exist on any disk that has been used under SunOS or Solaris. If you are partitioning a new disk, you can use `fdisk`'s "s" command to create a standard disk label (which includes the whole-disk partition). If you don't care for the size of the other partitions created by "s", you can delete those partitions and recreate them with the sizes you want.

### C.10.1 Swap Partitions and Red Hat Linux/SPARC

Because of the way disk partitioning is done, if the first partition on a disk starts at cylinder 0, and is used for swap space, *it will overwrite the drive's partition table*. This is a bad thing. Therefore, you have two options:

- Do not use the first partition for swap.
- Use the first partition for swap, but when you create it, start it at cylinder 1 instead of cylinder 0.

---

Either approach will save your drive's partition table.

## C.11 X Windows

There is no mouse or X windows config on the SPARC. The 8 bit server is installed by default. If you want to use the 24 bit server (for TCX cards), you need to install the `Xsun24` package. If you only have a bwtwo framebuffer and want to increase performance, install the `XsunMono` package. Installing both of the above requires you to restart X windows. Also, you can uninstall the `Xsun` package if you use one of the above servers.



# D

## Information specific to Red Hat Linux/Alpha

### D.1 Introduction

This appendix describes the differences between Red Hat Linux/Alpha and Red Hat Linux/Intel installations. While it provides a good overview of these differences, you will find it easier to read chapters 2, 3, and 6 in order. These chapters will refer you to the appropriate parts of this appendix at the appropriate time.

In addition, the Linux/Alpha homepage at <http://www.azstarnet.com/~axplinux/> contains a wealth of information for people considering Red Hat Linux/Alpha.

### D.2 Supported Hardware

Red Hat Linux/Alpha supports a variety of hardware based on the Alpha processor and the PCI bus. The most recent list of hardware supported by Red Hat

Linux/Alpha can be found at Red Hat Software's World Wide Web site at <http://www.redhat.com/hardware>. The following list is current as of the time this manual was produced, and contains the system type (in **bold**) followed by systems representative of that type:

**Alcor** – (AlphaStation 600)

**Cabriolet** – (AlphaPC64, AlphaPCI64; Aspen Telluride)

**EB164** – (Aspen Avalanche, Timberline, Summit; Microway Screamer)

**EB64+** – (Aspen Alpine)

**EB66** – (NekoTech Mach 1)

**EB66+**

**Miata** – (Digital Personal Workstation, a-Series and au-Series)

**Mikasa** – (AlphaServer 1000; AlphaServer 1000A not supported)

**Noname** – (XpPCI33, Universal Desktop Box, Multia)

**PC164** – (Aspen Durango)

**PC164LX** – (Aspen Durango II)

**PC164SX**

**Alpha-XLT** – (XL 233, XL 266, XL 300, XL 366, AlphaStation 500; Maverick, Brett, AlphaStation 200, 205, 250, 255, 400)

**Jensen** – (DEC PC 150, 2000 Model 300, 2000 Model 500, Cullean)

**Please Note:**Remember your system type, as we will be referring to it later in the installation process.

All of these platforms except the Jensen include an NCR 810 SCSI controller, although BusLogic PCI SCSI controllers (other than the FlashPoint), the Adaptec AHA2940 SCSI controller, and Qlogic 1020 ISP controllers are also supported. The Jensen design uses an AHA1740 SCSI controller, which is supported.

NE2000, DE422, DE4x5, and DE500 (PCI) Ethernet cards are supported, as is the UDB's internal Ethernet hardware. Token ring support is also included in the kernels.

The X Window System should work on any machine with an S3 based video card (except the Jensen machines). S3 Virge cards, the Matrox Millenium, Millenium II, and the Mystique will work, as will most of the cards from Orchid and Number 9. Most Diamond Stealth cards are also supported. Digital TGA cards (based on the DC21030 chip) are supported in 8 bit mode, and an X server for Mach64 cards is also provided.

**Please Note:** There is an X server available specifically for Jensens with an S3 card. It is available from `ftp://ftp.azstarnet.com/pub/linux/axp/jensen`. Note that this site is mirrored by several of the bigger FTP sites, should you have trouble connecting.

## D.3 Installation Overview

Installing Red Hat Linux on an Alpha system is slightly more complex than installing Red Hat Linux/Intel, mostly due to differences in machine architecture, and the variety of different models supported. In general, the main steps to a successful install are:

1. Create MILO, kernel, and ramdisk diskettes from images available on the Red Hat Linux/Alpha CD.
2. Use the MILO diskette to boot the appropriate Linux kernel.
3. Load and run the Red Hat Linux/Alpha installation program.
4. After the installation completes, install MILO on a small disk partition on your machine.

**Please Note:** While the majority of Alpha systems are supported by MILO, those that are not will need to boot the boot floppy (or CD-ROM) directly from the SRM (System Reference Manual) console. Information on doing this is available from the Red Hat Software web site at

`http://www.redhat.com/linux-info/alpha/faq`. You should also consult your system documentation for the proper boot command syntax, but in general, the proper command would look like this:

```
boot dva0 -file vmlinux.gz -flags "load_ramdisk=1 prompt_ramdisk=1"
```

## D.4 Choosing Diskette Images

The Red Hat Linux/Alpha installation process is diskette-based. The number and contents of the diskettes varies according to the type of Alpha system you have. Therefore, before creating diskettes, it's necessary to pick the appropriate image files. Use the following table as a guide:

System Type	MILO Image (In milo/images)	Boot Image (In images)
<b>Alcor</b> (ARC Console)	N/A	alcor.img
<b>Alcor</b> (SRM Console)	N/A	alcor-s.img
<b>Alpha-XLT</b>	xlt.img	xlt.img
<b>Cabriolet</b>	cab.img	cab.img
<b>EB164</b>	eb164.img	eb164.img
<b>EB64+</b>	eb64p.img	eb64p.img
<b>EB66</b>	eb66.img	eb66.img
<b>EB66+</b>	eb66p.img	eb66p.img
<b>Miata</b>	miata.img	miata.img
<b>Noname</b>	noname.img	noname.img
<b>PC164</b>	pc164.img	pc164.img
<b>PC164LX</b>	lx164.img	pc164.img
<b>PC164SX</b>	sx164.img	sx164.img

**Creating a MILO Diskette** If your system is supported by MILO, you will find the appropriate image in the `milo/images` directory on the Red Hat Linux/Alpha CD-ROM. Refer to Appendix B on page 193 for instructions on writing the image file to a diskette. Make sure you label the diskette "MILO Diskette".

**Please Note:** If you have trouble with MILO, the first thing to try is a newer MILO image. The latest version of MILO is always available at <ftp://ftp.digital.com/pub/DEC/Linux-Alpha/Miniloader/>.

**Creating a Boot Diskette** You will find the appropriate boot diskette image in the `images` directory on the Red Hat Linux/Alpha CD-ROM. Refer to Appendix B on page 193 for instructions on writing the image file to a diskette. Make sure you label the diskette "Boot Diskette".

**Creating a Ramdisk Diskette** All Red Hat Linux/Alpha installations use the same ramdisk image. It is called `ramdisk.img`, and can be found in the `images` directory on the Red Hat Linux/Alpha CD-ROM. Refer to Appendix B on page 193 for instructions on writing the image file to a diskette. Make sure you label the diskette "Ramdisk Diskette".

## D.5 Installing

### D.5.1 Booting the MILO Diskette

(If your Alpha system doesn't support MILO, please skip to Section D.5.2 on the next page.)

Before you can boot your MILO diskette, you'll need to add a new boot selection to your ARC console. You can add a new boot selection for your MILO diskette by following these steps:

1. From the boot menu, select "Supplementary menu".
2. From the supplementary menu, select "Set up the system".
3. From the boot selections menu, select "Add a boot selection".
4. Select "Floppy Disk 0".
5. Enter "`\linload.exe`" as the osloader directory and name after erasing any default information.
6. Indicate that the operating system is in the same partition as the osloader by entering "yes".
7. Enter "\" as the operating system root directory after erasing any default information.
8. Give the boot selection a meaningful name. "Linux" is a nice name.
9. Indicate that you do not want to initialize the debugger at boot time by entering "no".
10. From the boot selection menu, select "Change a boot selection option". Pick the boot selection you just created.

11. Using the arrow keys, select "OSLOADFILENAME", and enter "\milo". Press **Enter**.
12. Using the arrow keys, select "OSLOADOPTIONS", and erase any information that may be there. Press **Enter**.
13. Press **Esc** to leave the "Change a boot selection option" menu.
14. Select "Setup Menu". Select "Supplementary menu, and save changes".
15. Return to the boot menu.

At the end of this process, your boot selection should look something like this one (Not counting differences in boot selection name and disk identification strings:

```
LOADIDENTIFIER =Linux
SYSTEMPARTITION=multi(0)disk(0)fdisk(0)
OSLOADER=multi(0)disk(0)fdisk(0)\linload.exe
OSLOADERPARTITION=multi(0)disk(0)fdisk(0)
OSLOADFILENAME=\milo
OSLOADOPTIONS=
```

At this point your Alpha system has been set to boot MILO from diskette.

## D.5.2 Booting the Kernel Diskette

Now it's time to get things started. We need to start by booting from the kernel diskette you've created. How this is done depends on your Alpha system. If it supports MILO, insert your MILO diskette, and boot from that. At the MILO> prompt, insert your kernel diskette, and enter the following boot command:

```
boot floppy
```

MILO should then read the Linux kernel from your boot disk and start running it.

On the other hand, if you cannot use MILO, and must boot using the SRM console, enter the appropriate boot command for your Alpha system, making sure to add the following arguments to be passed to the kernel:

```
"load_ramdisk=1 prompt_ramdisk=1"
```

### D.5.3 Starting the Installation

After the kernel boots, it prompts for a ramdisk. Remove the kernel diskette, insert the ramdisk diskette, and press **Enter**. After the ramdisk is done loading, you should see the first installation screen.

### D.5.4 Issues Specific to Red Hat Linux/Alpha

Before you return to the installation, you should know about some differences between Red Hat Linux/Intel and its Alpha counterpart. Let's review them:

#### Disk Partitioning

Red Hat Linux/Alpha supports two varieties of hard drive partitioning:

- **Standard `fdisk`-style partitioning:** this has the usual oddities (such as extended partitions) and is done in the same way as on Intel platforms (using Linux `fdisk`). This is the partitioning style you must use if you are using the ARC or AlphaBIOS console.
- **BSD-style disk-labelling:** It is normal to make the first partition on a disk-labelled hard drive the swap partition. We don't recommend using BSD disk-labelling unless it is absolutely necessary (such as in the case of sharing a hard drive between Digital Unix and Red Hat Linux/Alpha). Note that you need to use `aboot` if you use BSD-style disk labels. Note that an RPM package of `aboot` comes with Red Hat Linux/Alpha.

#### Creating a MILO Partition

If your system supports MILO, you should allocate a small (no more than 2 MB) partition to it. The partition type must be DOS. This is where you will copy MILO and `linload.exe` to after the installation completes.

#### XFree86 and TGA

If you have an Alpha with a TGA card and you choose to install XFree86, the installation program prompts you to choose a video card. Choose **Unlisted Server**,

and the installation program will display a list of available servers. Choose the TGA server.

You can now turn back to Section 3.3 on page 39 to complete the installation.

## D.5.5 Finishing Up

After the installation is finished and your system is fully configured, you will be asked to reset your computer. This indicates that your system has been successfully installed.

### Configuring MILO

In order to boot your newly installed system, you'll need to use your MILO diskette. To boot Red Hat Linux/Alpha from MILO, you must use the `boot` command. The command differs slightly depending on where your root partition is. For example, if your root partition is the second partition on your first SCSI hard drive, you would boot as follows:

```
boot sda2:vmlinux.gz root=/dev/sda2
```

However, if your root partition is the third partition on your second IDE drive, you would use this command:

```
boot hdb3:vmlinux.gz root=/dev/hdb3
```

Boot your Alpha system (using the appropriate root partition name for your system, of course). Once Red Hat Linux/Alpha has finished booting, login, and issue the following command:

```
dd if=/dev/fd0 of=/dev/sda1 bs=1440k
```

This will copy MILO (along with `linload.exe`) to the small MILO partition you created during installation.

Finally, you need to create a boot selection that will look for MILO on your MILO partition. Shutdown your Alpha system, and perform the following steps from the ARC console:

1. From the boot menu, select "Supplementary menu".
2. From the supplementary menu, select "Set up the system".
3. From the boot selections menu, select "Add a boot selection".
4. Select the line that describes your MILO partition, if shown. If it's not there, select "New System Partition".
5. Enter the device attributes (ie, SCSI/IDE, bus number, etc.) that describes your MILO partition.
6. Enter "\linload.exe" as the osloader directory and name after erasing any default information.
7. Indicate that the operating system is in the same partition as the osloader by entering "yes".
8. Enter "\" as the operating system root directory after erasing any default information.
9. Give the boot selection a meaningful name. "Linux" is a nice name.
10. Indicate that you do not want to initialize the debugger at boot time by entering "no".
11. From the boot selection menu, select "Change a boot selection option". Pick the boot selection you just created.
12. Using the arrow keys, select "OSLOADFILENAME", and enter "\milo". Press **Enter**.
13. Using the arrow keys, select "OSLOADOPTIONS", and erase any information that may be there. Alternatively, if you would like to have MILO automatically boot your Red Hat Linux/Alpha system, you can enter the boot command that you used to boot Red Hat Linux/Alpha from the MILO floppy. In either case, Press **Enter** when done.
14. Press **Esc** to leave the "Change a boot selection option" menu.
15. Select "Setup Menu". Select "Supplementary menu, and save changes".
16. Return to the boot menu.

Here's a sample boot selection:

```
LOADIDENTIFIER =Linux
SYSTEMPARTITION=scsi(0)disk(0)rdisk(0)partition(1)
OSLOADER=scsi(0)disk(0)rdisk(0)partition(1)\linload.exe
OSLOADERPARTITION=scsi(0)disk(0)rdisk(0)partition(1)
OSLOADFILENAME=\milo
OSLOADOPTIONS=boot sda2:vmlinux.gz root=/dev/sda2
```

Note that in this case we are using the OSLOADOPTIONS line to automatically direct MILO to boot our kernel.

At this point, reboot your system, the boot selection you just created. Once the login: prompt appears, log in as root, and continue with the post-installation tasks in Chapter 6 on page 89.

# E

## Alpha Miniloader Howto

### E.1 Introduction

This document describes the Miniloader for Linux on Alpha AXP (MILO). This firmware is used to initialize Alpha AXP based systems, load and start Linux and, finally, provide PALcode for Linux.

#### E.1.1 Copyright

The Alpha Miniloader (MILO) HOWTO is copyright (C) 1995, 1996, 1997 David A Rusling.

**Copyright.** Like all Linux HOWTO documents, it may be reproduced and distributed in whole or in part, in any medium, physical or electronic, so long as this copyright notice is retained on all copies. Commercial redistribution is allowed and encouraged; however the author would *like* to be notified of such distributions. You may translate this HOWTO into any language whatsoever provided that you leave this copyright statement and disclaimer intact, and that you append a notice stating who translated the document.

**Disclaimer.** While I have tried to include the most correct and up to date information available to me, I cannot guarantee that usage of information in this document does not result in loss of data or equipment. I provide NO WARRANTY about the information in the HOWTO and I cannot be made liable for any consequences resulting from using the information in this HOWTO.

### E.1.2 New Versions of this Document

The latest version of this document can be found in <ftp://gatekeeper.dec.com/pub/Digital/Linux-Alpha/Miniloader/docs> and David Mosberger-Tang is kind enough to include the html form of it in his excellent Linux Alpha FAQ site <http://www.azstarnet.com/~axplinux>.

## E.2 What is MILO?

On Intel based PC systems, the BIOS firmware sets up the system and then loads the image to be run from the boot block of a DOS file system. This is more or less what MILO does on an Alpha based system, however there are several interesting differences between BIOS firmware and MILO, not least of which is that MILO includes and uses standard Linux device drivers unmodified. MILO is firmware, unlike LILO, which relies on the BIOS firmware to get itself loaded. The main functional parts of MILO are:

1. PALcode,
2. Memory set up code (builds page tables and turns on virtual addressing),
3. Video code (BIOS emulation code and TGA (21030)),
4. Linux kernel code. This includes real Linux kernel code (for example, the interrupt handling) and ersatz or mock Linux kernel,
5. Linux block device drivers (for example, the floppy driver),
6. File system support (ext2, MS-DOS and ISO9660),
7. User interface code (MILO),
8. Kernel interface code (sets up the HWRPB and memory map for linux),

### 9. NVRAM code for managing environment variables.

The following paragraphs describe these functional parts in more detail.

PALcode can be thought of as a tiny software layer that tailors the chip to a particular operating system. It runs in a special mode (PALmode) which has certain restrictions but it uses the standard Alpha instruction set with just five extra instructions. In this way, the Alpha chip can run such diverse operating systems as Windows NT, OpenVMS, Digital Unix and, of course, Linux. The PALcode that MILO uses (and therefore Linux itself) is, like the rest of MILO, freeware. It is derived from Digital's Evaluation Board software example Digital Unix PALcode.. The differences between the different PALcodes are because of differences in address mapping and interrupt handling that exist between the Alpha chips (21066 based systems have a different I/O map to 21064+2107x systems) and different Alpha based systems.

For MILO to operate properly it needs to know what memory is available, where Linux will eventually be running from and it must be able to allocate temporary memory for the Linux device drivers. The code maintains a memory map that has entries for permanent and temporary allocated pages. As it boots, MILO uncompresses itself into the correct place in physical memory. When it passes control to the Linux kernel, it reserves memory for the compressed version of itself, the PALcode (which the kernel needs) and some data structures. This leaves most of the memory in the system for Linux itself.

The final act of the memory code is to set up and turn on virtual addressing so that the data structures that Linux expects to see are at the correct place in virtual memory.

MILO contains video code that initialises and uses the video device for the system. It will detect and use a VGA device if there is one, otherwise it will try to use a TGA (21030) video device. Failing that, it will assume that there is no graphics device. The BIOS emulation that the standard, pre-built, images include is Digital's own BIOS emulation which supports most, if not all, of the standard graphics devices available.

Linux device drivers live within the kernel and expect certain services from the kernel. Some of these services are provided directly by Linux kernel code, for example the interrupt handling and some is provided by kernel look-alike routines.

MILO's most powerful feature is that you can embed unaltered Linux device drivers into it. This gives it the potential to support every device that Linux does. MILO includes all of the block devices that are configured into the Linux kernel that it is

built against as well as a lot of the block device code (for example, `llrw_blk()`).

MILO loads the Linux kernel from real file systems rather than from boot blocks and other strange places. It understands MSDOS, EXT2 and ISO9660 filesystems. Gzip'd files are supported and these are recommended, particularly if you are loading from floppy which is rather slow. MILO recognises these by their `.gz` suffix.

Built into MILO is a simple keyboard driver which, together with an equally simple video driver allows it to have a simple user interface. That interface allows you to list file systems on configured devices, boot Linux or run flash update utilities and set environment variables that control the system's booting. Like LILO, you can pass arguments to the Kernel.

MILO must tell the Linux kernel what sort of system this is, how much memory there is and which of that memory is free. It does this using the HWRPB (Hardware Restart Parameter Block) data structure and associated memory cluster descriptions. These are placed at the appropriate place in virtual memory just before control is passed to the Linux kernel.

### E.3 Pre-Built Standard MILO Images.

If you are planning to run Linux on a standard Alpha based system, then there are pre-built "standard" MILO images that you might use. These (along with the sources and other interesting stuff) can be found in <ftp://gatekeeper.dec.com/pub/Digital/Linux-Alpha/Miniloader>.

The `images` subdirectory contains a directory per standard system (eg AlphaPC64) with MILO images having the following naming convention:

1. `MILO` - Miniloader executable image, this image can be loaded in a variety of ways,
2. `fmu.gz` - Flash management utility,
3. `MILO.dd` - Boot block floppy disk image. These should be written using `rawrite.exe` or `dd` on Linux.

The `test-images`, like the `images` subdirectory contains a directory per standard system. These images are somewhat experimental but tend to contain all the latest features.

## E.4 How To Build MILO

You build MILO separately from the Kernel. As MILO requires parts of the kernel to function (for example interrupt handling) you must first configure and build the kernel that matches with MILO that you want to build. Mostly this means building the kernel with the same version number. So, MILO-2.0.25.tar.gz will build against linux-2.0.25.tar.gz. MILO may build against a higher version of the kernel, but there again it may not. Also, now that ELF shared libraries are fully supported, there are two versions of the MILO sources. To build under an ELF system you must first unpack the standard MILO sources and then patch those sources with the same version numbered ELF patch. In the remainder of this discussion, I assume that your kernel sources and object files are stored in the subtree at `/usr/src/linux` and that the linux kernel has been fully built with the command `make boot`

To build MILO, change your working directory to the MILO source directory and invoke `make` with:

```
$ make KSRC=/usr/src/linux config
```

Just like the Linux kernel, you will be asked a series of questions

```
Echo output to the serial port (MINI_SERIAL_ECHO) [y]
```

It's a good idea to echo kernel printk to `/dev/ttyS0` if you can. If you can (and want to), then type "y", otherwise "n". All of the standard, pre-built, MILO images include serial port I/O using COM1.

```
Use Digital's BIOS emulation code (not free) \
(MINI_DIGITAL_BIOS_EMU) [y]
```

This code is included as a library which is freely distributable so long as it is used on an Alpha based system. The sources are not available. If you answer n then the freeware alternative BIOS emulation will be built. It's sources are included with MILO. Note that you cannot right now build choose Digital's BIOS emulation code in an ELF system (the library is not yet ready) and so you must answer no to this question.

```
Build PALcode from sources (Warning this is dangerous) \
(MINI_BUILD_PALCODE_FROM_SOURCES) [n]
```

You should only do this if you have changed the PALcode sources, otherwise use the standard, pre-built PALcode included with MILO.

You are now all set to build the MILO image itself:

```
$ make KSRC=/usr/src/linux
```

When the build has successfully completed, the MILO image is in the file called `mi10`. There are a lot of images called `mi10.*`, these should be ignored.

## E.5 How To Load MILO

The most commonly supported method of loading MILO is from the Windows NT ARC firmware as most shipping systems support this. However, there are a wide variety of loading MILO. It may be loaded from:

- a failsafe boot block floppy,
- Windows NT ARC firmware,
- Windows NT AlphaBIOS,
- Digital's SRM console,
- an Alpha Evaluation Board Debug Monitor,
- flash/ROM.

### E.5.1 Loading MILO from the Windows NT ARC firmware

Most, if not all, Alpha AXP based systems include the Windows NT ARC firmware and this is the preferred method of booting MILO and thus Linux. Once the Windows NT firmware is running and you have the correct MILO image for your system, this method is completely generic.

The Windows NT ARC firmware is an environment in which programs can run and make callbacks into the firmware to perform actions. The Windows NT OSLoader is a program that does exactly this. `Linload.exe` is a much simpler program which does just enough to load and execute MILO. It loads the appropriate

image file into memory at 0x00000000 and then makes a swap-PAL PALcall to it. MILO, like Linux, uses a different PALcode to Windows NT which is why the swap has to happen. MILO relocates itself to 0x200000 and continues on through the PALcode reset entry point as before.

Before you add a Linux boot option, you will need to copy `linload.exe` and the appropriate MILO that you wish to load to someplace that the Windows NT ARC firmware can read from. In the following example, I assume that you are booting from a DOS format floppy disk:

1. At the boot menu, select "Supplementary menu..."
2. At the "Supplementary menu", select "Set up the system..."
3. At the "Setup menu", select "Manage boot selection menu..."
4. In the "Boot selections menu", choose "Add a boot selection"
5. Choose "Floppy Disk 0"
6. Enter "`linload.exe`" as the osloader directory and name
7. Say "yes" to the operating system being on the same partition as the osloader
8. Enter "\" as the operating system root directory
9. I usually enter "Linux" as the name for this boot selection
10. Say "No" you do not want to initialise the debugger at boot time
11. You should now be back in the "Boot selections menu", choose the "Change a boot selection option" and pick the selection you just created as the one to edit
12. Use the down arrow to get "OSLOADFILENAME" up and then type in the name of the MILO image that you wish to use, for example "`noname.arc`" followed by return.
13. Press ESC to get back to the "Boot Selections menu"
14. Choose "Setup Menu" (or hit ESC again) and choose "Supplementary menu, and save changes" option
15. ESC will get you back to the "Boot menu" and you can attempt to boot MILO. If you do not want Linux as the first boot option, then you can alter the order of the boot options in the "Boot selections menu".

At the end of all this, you should have a boot selection that looks something like:

```
LOADIDENTIFIER=Linux
SYSTEMPARTITION=multi(0)disk(0)fdisk(0)
OSLOADER=multi(0)disk(0)fdisk(0)\linload.exe
OSLOADPARTITION=multi(0)disk(0)fdisk(0)
OSLOADFILENAME=\noname.arc
OSLOADOPTIONS=
```

You can now boot MILO (and then Linux). You can load `linload.exe` and MILO directly from a file system that Windows NT understands such as NTFS or DOS on a hard disk.

The contents `OSLOADOPTIONS` are passed to MILO which interprets it as a command. So, in order to boot Linux directly from Windows NT without pausing in MILO, you could pass the following in `OSLOADOPTIONS`:

```
boot sda2:vmlinux.gz root=/dev/sda2
```

See E.6 for more information on the commands available.

Another (rather sneaky) way of loading of loading MILO via the WNT ARC firmware is to put MILO onto an MS-DOS floppy and call it `fwupdate.exe` and then choose the "Upgrade Firmware" option.

## E.5.2 Loading MILO from the Windows NT AlphaBIOS

With the introduction of the XLT series, Digital changed the system console interface for its NT systems from ARC to AlphaBIOS. AlphaBIOS is a screen-oriented interface which should be more familiar to PC users. This change in console interface necessitates a change in setup procedure for those who wish to run Linux/Alpha on AlphaBIOS-based systems.

The first thing you must do is to install the latest version of AlphaBIOS on your system. This can be obtained from Digital's "System Software and Driver Updates" Web page, <http://www.windows.digital.com/support/sysoft.htm>. Download the ZIP file, unzip it, and install it as follows:

- Copy the files to a FAT-formatted floppy

- Turn on the system and insert the floppy. At the opening screen, press F2 to go into setup mode
- Select "Upgrade AlphaBIOS"
- Follow the directions on the screen

Once your AlphaBIOS is at the latest revision level, you can start bootstrapping your system as follows:

- Create a FAT-formatted floppy with the following files:
  - linload.exe (from this directory)
  - milo (the version appropriate to your system)
- Turn on the system and insert the floppy. At the opening screen, press F2 to go into setup mode
- Select "Utilities->OS Selection Setup..."
- Press INSERT to add a new operating system selection
- For "Boot Name", enter something like "Linux". Press TAB to get over to the next field.
- Press down-arrow until the selection for "Boot File" is "A:". TAB over to the next field
- Enter "linload.exe". TAB *twice* (i.e. skip the OS Path load device – it's irrelevant)
- Enter "\" for the OS Path load file
- Press ENTER to add the selection.

At this point, AlphaBIOS will probably put up a big, unfriendly dialog box labelled "Warning: Operating System Selection not valid!". Ignore this error (it's only a problem for NT) and press ENTER to continue.

- Press F10 to save the changes you just made; press ENTER to confirm the changes.
- Press ESC twice to get back to the opening screen.

- Use the up and down arrows to select the boot selection you just added, and press ENTER to boot it.
- AlphaBIOS will load linload, which will in turn load MILO. When you get to the MILO prompt, proceed as you would for a normal ARC-based system installation.

If you dedicated the first partition of your first disk drive to a small FAT partition for booting (as the installation procedure advises you to), then once Linux is installed you should copy linload.exe and milo to this partition. Once you shut down, you can then modify the Linux menu selection to load MILO from this partition as follows:

- At the opening screen, select F2 to go into setup mode
- Select "Utilities->OS Selection setup"
- Highlight the entry for Linux, then press F6 to edit it.
- TAB over to the device portion of the "Boot File" line. Use the up and down arrow keys to select the hard-disk partition where linload and milo reside (typically "Disk 0 Partition 1" or "Disk 1 Partition 1"). Press ENTER to confirm the selection
- If you wish to auto-boot linux after MILO is loaded, then TAB over to the "OS Options" line and enter the MILO command to boot the system, e.g. "boot sda2:vmlinux.gz" Press ENTER to confirm the selection.
- Press F10 to save the changes. Press ENTER to confirm.

Once you have done this, booting and running Linux on an AlphaBIOS based system should be very similar to doing so on an ARC system.

### E.5.3 Loading MILO from the Evaluation Board Debug Monitor

Evaluation boards (and often designs cloned from them) include support for the Alpha Evaluation Board Debug Monitor. Consult your system document before considering this method of booting MILO. The following systems are *known* to include Debug Monitor support:

- AlphaPC64 (Section E.5.7)

- EB64+ (Section E.5.7)
- EB66+ (Section E.5.7)
- EB164 (Section E.5.7)
- PC164 (Section E.5.7)

Before you consider this method, you should note that the early versions of the Evaluation Board Debug Monitor did not include video or keyboard drivers and so you must be prepared to connect another system via the serial port so that you can use the Debug Monitor. Its interface is very simple and typing help shows a whole heap of commands. The ones that are most interesting include the word `boot` or `load` in them.

The Evaluation Board Debug Monitor can load an image either via the network (`netboot`) or via a floppy (`flboot`). In either case, set the boot address to `0x200000` (`> bootadr 200000`) before booting the image.

If the image is on floppy (and note that only DOS formatted floppies are supported), then you will need to type the following command:

```
AlphaPC64> flboot <MILO-image-name>
```

#### E.5.4 Loading MILO from a Failsafe Boot Block Floppy

Only the AxpPCI33 is *known* to include failsafe boot block floppy support (Section E.5.7).

If you do not have a standard pre-built MILO `.dd` image, then you may need to build an SRM boot block floppy. Once you have built MILO, you need to do the following on Digital Unix box:

```
fddisk -fmt /dev/rfd0a
cat mboot bootm > /dev/rfd0a
disklabel -rw rfd0a 'rx23' mboot bootm
```

Or on a Linux box:

```
cat mboot bootm > /dev/fd0
```

If you have a standard MILO image available (say `MILO.dd`) then you would build a boot block floppy using the following command:

```
dd if=MILO.dd of=/dev/fd0
```

### E.5.5 Loading MILO from Flash

There are a number of systems where MILO can be blown into flash and booted directly (instead of via the Windows NT ARC firmware):

- AlphaPC64 (Section E.5.7)
- Noname (Section E.5.7)
- EB66+ (Section E.5.7)
- EB164 (Section E.5.7)
- PC164 (Section E.5.7)

### E.5.6 Loading MILO from the SRM Console

The SRM (short for System Reference Manual) Console knows nothing about filesystems or disk-partitions, it simply expects that the secondary bootstrap loader occupies a consecutive range of physical disk sectors starting from a given offset. The information describing the secondary bootstrap loader (its size and offset) is given in the first 512 byte block. To load MILO via the SRM you must generate that structure on a device which the SRM can access (such as a floppy disk). This is what `mboot` and `bootm`, `mboot` is the first block (or boot description) and `mboot` is the MILO image rounded up to a 512 byte boundary.

To load MILO from a boot block device, either build `mboot` and `bootm` and push them onto the boot device using the following command:

```
$ cat mboot bootm > /dev/fd0
```

Or, grab the appropriate `MILO.dd` from a web site and write it onto the boot device using either `RAWRITE.EXE` or `dd`.

Once you have done that you can boot the SRM console and use one of its many commands to boot MILO. For example, to boot MILO from a boot block floppy you would use the following command:

```
>>>boot dva0
(block dva0.0.0.0.1 -flags 0)
block 0 of dva0.0.0.0.1 is a valid boot block
reading 621 blocks from dva0.0.0.0.1
bootstrap code read in
base = 112000, image_start = 0, image_bytes = 4da00
initializing HWRPB at 2000
initializing page table at 104000
initializing machine state
setting affinity to the primary CPU
jumping to bootstrap code
MILO Stub: V1.1
Unzipping MILO into position
Allocating memory for unzip
####...
```

The following systems are *known* to have SRM Console support:

- Noname (Section E.5.7)
- AlphaPC64 (Section E.5.7)
- EB164 (Section E.5.7)
- PC164 (Section E.5.7)

## E.5.7 System Specific Information

### AxpPCI33 (Noname)

The Noname board can load MILO from the Windows NT ARC firmware (Section E.5.1), from the SRM Console (Section E.5.6), and from a failsafe boot block floppy (Section E.5.4). A flash management utility, runnable from MILO is available so that once MILO is running, it can be blown into flash (Section E.7). However, be warned that once you have done this you will lose the previous image held there as there is only room for one image.

The way that Noname boots is controlled by a set of jumpers on the board, J29 and J28. These look like:

```

                4
J29           2 x x x 6
              1 x x x 5

J28           2 x x x 6
              1 x x x 5
                3

```

The two options that we're interested in are J28, pins 1-3 which boots the console/loader from flash and J29, pins 1-3 which boots the console/loader from a boot block floppy. The second option is the one that you need to first boot MILO on the Noname board.

Once you've selected the boot from floppy option via the jumpers, put the SRM boot block floppy containing MILO into the floppy and reboot. In a few seconds (after the floppy light goes out) you should see the screen blank to white and MILO telling you what's going on.

If you are really interested in technical stuff, the Noname loads images off of the floppy into physical address 0x104000 and images from flash into 0x100000. For this reason, MILO is built with it's PALcode starting at 0x200000. When it is first loaded, it moves itself to the correct location (see relocate.S).

### AlphaPC64 (Cabriolet)

The AlphaPC64 includes the Windows NT ARC firmware (Section E.5.1), the SRM Console (Section E.5.6) and the Evaluation Debug Monitor (Section E.5.3). These images are in flash and there is room to add MILO so that you can boot MILO directly from flash. A flash management utility, runnable from MILO is available so that once MILO is running, it can be blown into flash (Section E.7). This system supports MILO environment variables.

You select between the boot options (and MILO when it is been put into flash) using a combination of jumpers and a boot option which is saved in the NVRAM of the TOY clock.

The jumper is J2, SP bits 6 and 7 have the following meanings:

- SP bit 6 should always be out. If this jumper is set then the SROM mini-

debugger gets booted,

- SP bit 7 in is boot image selected by the boot option byte in the TOY clock,
- SP bit 7 out is boot first image in flash.

So, with bit 7 out, the Debug Monitor will be booted as it is always the first image in flash. With bit 7 in, the image selected by the boot option in the TOY clock will be selected. The Debug Monitor, the Windows NT ARC firmware and MILO all support setting this boot option byte but you must be very careful using it. In particular, you cannot set the boot option so that next time the system boots MILO when you are running the Windows NT ARC firmware, it only allows you to set Debug Monitor or Windows NT ARC as boot options.

To get MILO into flash via the Evaluation Board Debug Monitor, you will need a flashable image. The build procedures make MILO.rom, but you can also make a rom image using the makerom tool in the Debug Monitor software that comes with the board:

```
> makerom -v -i7 -l200000 MILO -o mini.flash
```

(type makerom to find out what the arguments mean, but 7 is a flash image id used by the srom and -l200000 gives the load address for the image as 0x200000).

Load that image into memory (via the Debug Monitor commands fload, netload, and so on) at 0x200000 and then blow the image into flash:

```
AlphaPC64> flash 200000 8
```

(200000 is where the image to be blown is in memory and 8 is the segment number where you put the image. There are 16 1024\*64 byte segments in the flash and the Debug Monitor is at seg 0 and the Windows NT ARC firmware is at seg 4).

Set up the image that the srom will boot by writing the number of the image into the TOY clock.

```
AlphaPC64> bootopt 131
```

(131 means boot the 3rd image, 129 = 1st, 130 = 2nd and so on).

Power off, put jumper 7 on and power on and you should see the MILO burst into life. If you don't then take jumper 7 back off and reboot the Debug Monitor.

### **EB66+**

The EB66+, like all of the Alpha Evaluation Boards built by Digital contains the Evaluation Board Debug Monitor and so this is available to load MILO (Section E.5.3). Quite often (although not always) boards whose design is derived from these include the Debug Monitor also. Usually, these boards include the Windows NT ARC firmware (Section E.5.1). A flash management utility, runnable from MILO is available so that once MILO is running, it can be blown into flash (Section E.7). This system supports MILO environment variables.

These systems have several boot images in flash controlled by jumpers. The two jumper banks are J18 and J16 and are located at the bottom of the board in the middle (if the Alpha chip is at the top). You select between the boot options (and MILO when it is been put into flash) using a combination of jumpers and a boot option which is saved in the NVRAM of the TOY clock.

Jumper 7-8 of J18 in means boot the image described by the boot option. Jumper 7-8 of J18 out means boot the Evaluation Board Debug Monitor.

Blowing an image into flash via the Evaluation Board Debug Monitor is exactly the same procedure as for the AlphaPC64 (Section E.5.7).

### **EB64+/Aspen Alpine**

This system is quite like the AlphaPC64 except that it does not contain flash which MILO can be loaded from. The EB64+ has two ROMs, one of which contains the Windows NT ARC firmware (Section E.5.1). and the other contains the Evaluation Board Debug Monitor (Section E.5.3).

The Aspen Alpine is a little different in that it only has one ROM; this contains the Windows NT ARC firmware.

### **Universal Desktop Box (Multia)**

This is a very compact pre-packaged 21066 based system that includes a TGA (21030) graphics device. Although you can *just* fit a half height PCI graphics card in the box you are better off waiting for full TGA support in XFree86. It includes the Windows NT ARC firmware and so booting from that is the preferred method (Section E.5.1).

**EB164**

The EB164, like all of the Alpha Evaluation Boards built by Digital contains the Evaluation Board Debug Monitor and so this is available to load MILO (Section E.5.3). Quite often (although not always) boards whose design is derived from these include the Debug Monitor also. Usually, these boards include the Windows NT ARC firmware (Section E.5.1). The SRM console is also available (Section E.5.6). A flash management utility, runnable from MILO is available so that once MILO is running, it can be blown into flash (Section E.7). This system supports MILO environment variables.

These systems have several boot images in flash controlled by jumpers. The two jumper bank is J1 and is located at the bottom of the board on the left (if the Alpha chip is at the top). You select between the boot options (and MILO when it is been put into flash) using a combination of jumpers and a boot option which is saved in the NVRAM of the TOY clock.

Jumper SP-11 of J1 in means boot the image described by the boot option. Jumper SP-11 of J1 out means boot the Evaluation Board Debug Monitor.

Blowing an image into flash via the Evaluation Board Debug Monitor is exactly the same procedure as for the AlphaPC64 (Section E.5.7).

**PC164**

The PC164, like all of the Alpha Evaluation Boards built by Digital contains the Evaluation Board Debug Monitor and so this is available to load MILO (Section E.5.3). Quite often (although not always) boards whose design is derived from these include the Debug Monitor also. Usually, these boards include the Windows NT ARC firmware (Section E.5.1). The SRM console is also available (Section E.5.6). A flash management utility, runnable from MILO is available so that once MILO is running, it can be blown into flash (Section E.7). This system supports MILO environment variables.

These systems have several boot images in flash controlled by jumpers. The main jumper block, J30, contains the system configuration jumpers and jumper CF6 in means that the system will boot the Debug Monitor, the default is out.

Blowing an image into flash via the Evaluation Board Debug Monitor is exactly the same procedure as for the AlphaPC64 (Section E.5.7).

## XL266

The XL266 is one of a family of systems that are known as Avanti. It has a riser card containing the Alpha chip and cache which plugs into the main board at right angles. This board can replace the equivalent Pentium board.

Some of these systems ship with the SRM console but others, notably the XL266 ship with only the Windows NT ARC firmware (Section E.5.1).

Here is my list of compatible systems:

- AlphaStation 400 (Avanti),
- AlphaStation 250,
- AlphaStation 200 (Mustang),
- XL. There are two flavours, XL266 and XL233 with the only difference being in processor speed and cache size.

**Note** The system that I use to develop and test MILO is an XL266 and so this is the only one that I can guarantee will work. However, technically, all of the above systems are equivalent; they have the same support chipsets and the same interrupt handling mechanisms.

## Platform2000

This is a 233Mhz 21066 based system.

## E.6 MILO's User Interface

Once you have correctly installed/loaded/run MILO you will see the MILO (for MIniLOader) prompt displayed on your screen. There is a very simple interface that you must use in order to boot a particular Linux kernel image. Typing "help" is a good idea as it gives a useful summary of the commands.

### E.6.1 The "help" Command

Probably the most useful command that MILO has:

```
MILO> help
MILO command summary:

ls [-t fs] [dev:[dir]]
    - List files in directory on device
boot [-t fs] [dev:file] [boot string]
    - Boot Linux from the specified device and file
run [-t fs] dev:file
    - Run the standalone program dev:file
show
    - Display all known devices and file systems
set VAR VALUE
    - Set the variable VAR to the specified VALUE
unset VAR
    - Delete the specified variable
reset
    - Delete all variables
print
    - Display current variable settings
help [var]
    - Print this help text
```

Devices are specified as: fd0, hda1, hda2, sda1...  
Use the '-t filesystem-name' option if you want to use anything but the default filesystem ('ext2').  
Use the 'show' command to show known devices and filesystems.  
Type 'help var' for a list of variables.

**Note** that the `bootopt` command only appears on AlphaPC64 (and similar) systems. Refer to the board's documentation to find out just what it means.

**Devices.** Until you use a command that needs to make use of a device, no device initialisation will take place. The first `show`, `ls`, `boot` or `run` commands all cause the devices within MILO to be initialised. Devices are named in the same way (exactly) that Linux itself will name them. So, the first IDE disk will be called 'hda' and its first partition will be 'hda1'. Use the `show` command to show what devices are available.

**File Systems.** MILO supports three file systems, MSDOS, EXT2 and ISO9660. So long as a device is available to it, MILO can `listboot` or `run` an image stored on one of these file systems. MILO's default file system is EXT2 and so you have tell MILO that the file system is something other than that. All of the commands that use filenames allow you to pass the file system using the `-t [filesystem]` option. So, if you wanted to list the contents of a SCSI CD ROM, you might type the following:

```
MILO> ls -t iso9660 scd0:
```

**Variables.** MILO contains some settable variables that help the boot process. If you are loading via the Windows NT ARC firmware, then MILO makes use of the boot option environment variables set up by that firmware. For some systems, MILO (for example, the AlphaPC64) maintains its own set of environment variables that do not change from boot to boot. These variables are:

```
MILO> help var
Variables that MILO cares about:
MEMORY_SIZE - System memory size in megabytes
BOOT_DEV    - Specifies the default boot device
BOOT_FILE   - Specifies the default boot file
BOOT_STRING - Specifies the boot string to pass to the kernel
SCSIn_HOSTID - Specifies the host id of the n-th SCSI controller.
PCI_LATENCY - Specifies the PCI master device latency
AUTOBOOT    - If set, MILO attempts to boot on powerup
              and enters command loop only on failure.
AUTOBOOT_TIMEOUT - Seconds to wait before auto-booting on powerup.
```

**BEWARE** of setting AUTOBOOT without setting the timeout, you can end up automatically booting after a timeout of 0 seconds which may not be what you want.

PCILATENCY is the number of PCI bus cycles that PCI bus mastering device hangs onto the PCI bus for once it has control. The default value for this is 32 and the maximum value is 255. Setting a high value means that each time a device gets control of the PCI bus it can transfer more data. However, setting it high also means that devices may have to wait longer to get control.

## E.6.2 Booting Linux

The `boot` command boots a linux kernel from a device. You will need to have a linux kernel image on an EXT2 formatted disk (SCSI, IDE or floppy) or an ISO9660 formatted CD available to MILO. The image can be gzip'd and in this case MILO will automatically gunzip it. Early versions of MILO recognised a gzip'd file by the `.gz` suffix but later MILOs look for magic numbers in the image.

You should note that the version of MILO does not usually have to match the version of the Linux kernel that you are loading. You boot Linux using the following command syntax:

```
MILO> boot [-t file-system] device-name:file-name \  
          [[boot-option] [boot-option] ...]
```

Where `device-name` is the name of the device that you wish to use and `file-name` is the name of the file containing the Linux kernel. All arguments supplied after the file name are passed directly to the Linux kernel.

If you are installing Red Hat, then you will need to specify a root device and so on. So you would use:

```
MILO> boot fd0:vmlinux.gz root=/dev/fd0 load_ramdisk=1
```

MILO will automatically contain the block devices that you configure into your `vmlinux`. I have tested the floppy driver, the IDE driver and a number of SCSI drivers (for example, the NCR 810), and these work fine. Also, it is important to set the host id of the SCSI controller to a reasonable value. By default, MILO will initialize it to the highest possible value (7) which should normally work just fine. However, if you wish, you can explicitly set the host id of the *n*-th SCSI controller in the system by setting environment variable `SCSI $n$ _HOSTID` to the appropriate value. For example, to set the hostid of the first SCSI controller to 7, you can issue the following command at the MILO prompt:

```
setenv SCSI0_HOSTID 7
```

### E.6.3 Rebooting Linux

You may want to reboot a running Linux system using the `shutdown -r now` command. In this case, the Linux kernel returns control to MILO (via the HALT CallPAL entrypoint). MILO leaves a compressed copy of itself in memory for just this reason and detects that the system is being rebooted from information held in the HWRPB (Hardware Restart Parameter Block). In this case it starts to reboot using exactly the same command that was used to boot the Linux kernel the last time. There is a 30 second timeout that allows you to interrupt this process and boot whatever kernel you wish in whatever way you wish.

### E.6.4 The "bootopt" command

For flash based systems such as the AlphaPC64, EB164 and the EB66+, there are a number of possible boot options and these are changed using the `bootopt` com-

mand. This has one argument, a decimal number which is the type of the image to be booted the next time the system is power cycled or reset:

**0** Boot the Evaluation Board Debug Monitor,

**1** Boot the Windows NT ARC firmware.

In order to tell the boot code to boot the MILO firmware from flash then you need a boot option that means boot the N'th image. For this, you need to 128 plus N, so if MILO is the third image, you would use the command:

```
MILO> bootopt 131
```

**Note:** Be very careful with this command. A good rule is never to set bootopt to 0 (the Evaluation Board Debug Monitor), but instead use the system's jumpers to achieve the same thing.

## E.7 Running the Flash Management Utility

The run command is used to run the flash management utility. Before you start you will need a device available to MILO that contains the updateflash program. This (like vmlinux) can be gzip'd. You need to run the flash management utility program from the MILO using the (run) command:

```
MILO> run fd0:fmv.gz
```

Once it has loaded and initialised, the flash management utility will tell you some information about the flash device and give you a command prompt. Again the help command is most useful.

```
Linux MILO Flash Management Utility V1.0
```

```
Flash device is an Intel 28f008SA
 16 segments, each of 0x10000 (65536) bytes
Scanning Flash blocks for usage
Block 12 contains the environment variables
FMU>
```

Note that on systems where environment variables may be stored and where there is more than one flash block (for example, the AlphaPC64) the flash management utility will look for a block to hold MILO's environment variables. If such a block already exists, the flash management utility will tell you where it is. Otherwise, you must use the `environment` command to set a block and initialise it. In the above example, flash block 12 contains MILO's environment variables.

### E.7.1 The "help" command

```
FMU> help
FMU command summary:

list          - List the contents of flash
program       - program an image into flash
quit          - Quit
environment   - Set which block should contain the environment \
                variables
bootopt num   - Select firmware type to use on next power up
help          - Print this help text
FMU>
```

*Note* that the `environment` and `bootopt` commands are only available on the EB66+, the AlphaPC64, EB164 and PC164 systems (and their clones).

### E.7.2 The "list" command

The "list" command shows the current usage of the flash memory. Where there is more than one flash block, the usage of each flash block is shown. In the example below you can see that Windows NT ARC is using blocks 4:7 and block 15.

```
FMU> list
Flash blocks:  0:DBM  1:DBM  2:DBM  3:WNT  4:WNT  5:WNT  6:WNT
              7:WNT  8:MILO  9:MILO 10:MILO 11:MILO 12:MILO 13:U 14:U
              15:WNT
Listing flash Images
Flash image starting at block 0:
  Firmware Id:  0 (Alpha Evaluation Board Debug Monitor)
  Image size is 191248 bytes (3 blocks)
  Executing at 0x300000
```

```
Flash image starting at block 3:
  Firmware Id:  1 (Windows NT ARC)
  Image size is 277664 bytes (5 blocks)
  Executing at 0x300000
Flash image starting at block 8:
  Firmware Id:  7 (MILO/Linux)
  Image size is 217896 bytes (4 blocks)
  Executing at 0x200000
FMU>
```

### E.7.3 The "program" command

The flash management utility contains a compressed copy of a flash image of MILO. The "program" command allows you to blow this image into flash. The command allows you to back out, but before you run it you should use the "list" command to see where to put MILO. If MILO is already in flash, then the flash management utility will offer to overwrite it.

```
FMU> program
Image is:
  Firmware Id:  7 (MILO/Linux)
  Image size is 217896 bytes (4 blocks)
  Executing at 0x200000
Found existing image at block 8
Overwrite existing image? (N/y)? y
Do you really want to do this (y/N)? y
Deleting blocks ready to program: 8 9 10 11
Programming image into flash
Scanning Flash blocks for usage
FMU>
```

Wait until it has completed before powering off your system.

Note: I cannot emphasise just how careful you must be here not to overwrite an existing flash image that you might need or render your system useless. A very good rule is never to overwrite the Debug Monitor.

### E.7.4 The "environment" command

This selects a flash block to contain MILO's environment variables.

### E.7.5 The "bootopt" command

This is just the same as MILO's "bootopt" command, see (Section E.6.4).

### E.7.6 The "quit" command

This is really pretty meaningless. The only way back to MILO (or anything else) once the flash management utility has run is to reboot the system.

## E.8 Restrictions.

Unfortunately this is not a perfect world and there, as always, some restrictions that you should be aware of.

MILO is not meant to load operating systems other than Linux, although it can load images linked to run at the same place in memory as Linux (which is 0xFFFFFC0000310000). This is how the flash management utilities can be run.

The PALcode sources included in `miniboot/palcode/blah` are correct, however there are problems when they are built using the latest `gas`. They *do* build if you use the ancient `a.out gas` that's supplied in the Alpha Evaluation Board toolset (and that's how they were built). I'm trying to get someone to fix the new `gas`. Meanwhile, as a workaround, I have provided pre-built PALcode for the supported boards and David Mosberger-Tang has a fixed `gas` on his ftp site.

## E.9 Problem Solving.

Here are some common problems that people have seen, together with the solutions.

### Reading MS-DOS floppies from the Evaluation Board Debug Monitor.

Some of the older versions of the Evaluation Board Debug Monitor (pre-version 2.0) have a problem with DOS format floppies generated from Linux. Usually, the Debug Monitor can load the first few sectors all right, but then goes into an endless loop complaining about "bad sectors." Apparently, there is an incompatibility

between the DOS file system as expected by the Debug Monitor and the Linux implementation of DOSFS. To make the long story short: if you run into this problem, try using DOS to write the floppy disk. For example, if loading the file `MILO.cab` doesn't work, use a DOS machine, insert the floppy and then do:

```
copy a:MILO.cab c:
copy c:MILO.cab a:
del c:MILO.cab
```

Then try booting from that floppy again. This normally solves the problem.

#### **MILO displays a long sequence of O> and does not accept input.**

This usually happens when MILO was built to use COM1 as a secondary console device. In such a case, MILO echo output to COM1 and accepts input from there also. This is great for debugging but not so great if you have a device other than a terminal connected. If this happens, disconnect the device or power it down until the Linux kernel has booted. Once Linux is up and running, everything will work as expected.

#### **MILO complains that the kernel image has the wrong magic number**

Older versions of MILO did not support the ELF object file format and so could not recognise an ELF image and this might be your problem. If this is reported, upgrade to the latest MILO that you can find. All 2.0.20 and beyond MILOs support ELF. On the other hand it could be that the image is indeed damaged. You should also note that MILO does not yet automatically distinguish between GZIP'd and non-GZIP'd images; you need to add the ".gz" suffix to the file name.

#### **MILO prints "...turning on virtual addressing and jumping to the Linux Kernel" and nothing else happens**

One obvious problem is that the kernel image is wrongly built or is built for another Alpha system altogether. Another is that the video board is a TGA (Zlxp) device and the kernel has been built for a VGA device (or vice versa). It is worth building the kernel to echo to COM1 and then connecting a terminal to that serial port or retrying the kernel that came with the Linux distribution that you installed.

#### **MILO does not recognise the SCSI device**

The standard MILO images include as many device drivers as are known to be stable for Alpha (as of now that includes the NCR 810, QLOGIC ISP, Buslogic and Adaptec 2940s and 3940 cards). If your card is not included, it may be that the driver is not stable enough on an Alpha system yet. Again, the latest MILO images

---

are worth trying. You can tell which SCSI devices a MILO image has built into it by using the "show" command.

## **E.10 Acknowledgements.**

I would like to thank:

- Eric Rasmussen and Eileen Samberg the authors of the PALcode,
- Jim Paradis for the keyboard driver, the original MILO interface and the AlphaBIOS work,
- Jay Estabrook for his help and bugfixes,
- David Mosberger-Tang for the freeware BIOS emulation and his support and encouragement,
- Last (and not least) Linus Torvalds for the timer code and his kernel.

There are a number of things that still need doing to MILO, if you want to add something yourself, then do let me know  `david.rusling@reo.mts.dec.com`  so that we do not duplicate our efforts.

Finally, a big thank you to Digital for producing such a wonderful chip (and paying me to do this).



# F

## General Parameters

This appendix is provided to illustrate *some* of the possible parameters that may be needed by certain drivers. It should be noted that, in most cases, these additional parameters are unnecessary.

Please keep in mind that if a device you are attempting to use requires one of these parameters, and support for that device is *not* compiled into the kernel, the traditional method of adding the parameter to the LILO boot command will not work. Drivers loaded as modules require that these parameters are specified when the module is loaded. The Red Hat Linux installation program gives you the option to specify module parameters when a driver is loaded.

For more information concerning the device support compiled into the kernel used by the Red Hat Linux installation program, please refer to Section 2.6 on page 32.

One of the more commonly used parameters, the `hdX=cdrom` parameter, *can* be entered at the boot prompt, as it deals with support for IDE/ATAPI CD-ROMs, which is part of the kernel.

## F.1 CD-ROM parameters

Hardware	Parameter
Mitsumi CD-ROM	<code>mcd=port,irq</code>
Sony CDU 31 or 33 CD-ROM	<code>cdu31a_port=base_addr cdu31a_irq=irq</code>
Aztech CD268	<code>aztcd=port</code>
SB Pro or 16 compatible	<code>sbpcd=io_addr,sb_pro_setting</code>
ATAPI/IDE CD-ROM Drives	<code>hdx=cdrom</code>

Examples of the above would be:

Configuration	Example
non-IDE Mitsumi CD-ROM on port 340, IRQ 11	<code>mcd=0x340,11</code>
Sony CDU 31 or 33 at port 340, no IRQ	<code>cdu31a_port=0x340 cdu31a_irq=0</code>
Aztech CD-ROM at port 220	<code>aztcd=0x220</code>
ATAPI CD-ROM, jumpered as master on 2nd interface	<code>hdc=cdrom</code>
Panasonic-type CD-ROM on a SoundBlaster at port 230	<code>sbpcd=0x230,1</code>

**Please Note:** Most newer Sound Blaster cards come with IDE interfaces. You do not need to use `sbpcd` parameters, only use `hdx` parameters.

## F.2 SCSI parameters

Hardware	Parameter
SCSI Seagate ST0x	<code>controller.type=1</code> <code>base.address=shmemaddr irq=irq</code>
SCSI Future Domain TMC-8xx	<code>controller.type=2</code> <code>base.address=shmemaddr irq=irq</code>
SCSI Trantor T128	<code>t128=base_addr,irq</code>
SCSI NCR-5380 Based	<code>ncr5380=port,irq,dma_channel</code>
Adaptec 152x, SB SCSI	<code>aha152x=port,irq,scsi_id,reconnect,parity</code>
Adaptec 1542	<code>bases=base_addr</code>
Buslogic	<code>buslogic=base_addr</code>
Pro Audio Spectrum / Studio 16	<code>pas16=port,irq</code>

Some examples would be:

Configuration	Example
Adaptec AHA1522 at port 330, IRQ 11, SCSI ID 7	<code>aha152x=0x330,11,7</code>
Adaptec AHA1542 at port 330	<code>bases=0x330</code>
Future Domain TMC-800 at CA000, IRQ 10	<code>controller.type=2</code> <code>base.address=0xca000 irq=10</code>

When a parameter has commas, make sure you do not put a space after a comma.

## F.3 Ethernet parameters

Most ethernet drivers accept parameters to specify a base IO address and an IRQ as follows:

```
io=base_addr irq=irq
```

For example, for a 3com 3c509 ethernet card located at IO address 210 (IO addresses are usually in hexadecimal) and IRQ 10, use the following parameters for the 3c509 driver:

```
io=0x210 irq=10
```

You can use multiple ethernet cards in one machine . If each card uses a different driver (e.g., 3c509 and a DE425), you simply need to add aliases (and possibly options) for each card to `/etc/conf.modules`; for example:

```
alias eth0 3c509
options 3c509 io=0x210 irq=10
alias eth1 de4x5
options de4x5 io=0
```

See Section 9.6 on page 148 for more information.

However, if any two ethernet cards use the same driver (e.g., two 3c509's or a 3c595 and a 3c905), you will need to compile a custom kernel with the ethernet driver built in. In that case, you can use the "classic" LILO boot : parameters of the form:

```
ether=irq,base_addr,interface
```

For example:

```
LILO boot: linux ether=10,0x210,eth0 ether=11,0x300,eth1
```

(For more information about using more than one ethernet card, see the *Multiple-Ethernet* mini-HOWTO.)

# G

## Metro X

### G.1 Special Features

#### G.1.1 Hot-Key Exit

Metro-X supports a hot-key exit. This is useful for terminating the X server if the last client in your `.xinitrc` or `.xsession` does not provide a way to exit the session. The exit key combination is `Ctrl-Alt-Backspace`.

#### G.1.2 Hot-Key Resolution Switching

When configuring Metro-X you can select several screen resolutions. You can then change resolutions on the fly while running your X server by pressing the appropriate key combination. To switch to the next resolution, press `Ctrl-Alt-+`. To switch to the previous resolution, press `Ctrl-Alt-`.

### G.1.3 Hardware Panning

Metro-X supports hardware panning. This lets you set the virtual resolution higher than the monitor's physical resolution. Windows can then extend beyond the screen's physical boundaries and can be viewed by scrolling into the extended area.

### G.1.4 Multiple Screens

Metro-X supports up to four screens which can all be controlled simultaneously with a single keyboard and mouse. This allows you to run many applications without overlapping windows, and is particularly useful for image-processing work.

In a multiple screen configuration there is one X process, one mouse, and one keyboard - but several screens. The mouse/cursor can be moved between the different screens. Clients connect to the different screens with the syntax `-display :0.X` where X is the screen number.

#### Supported Cards

Use up to four of any one card from the following list:

- One or more Matrox MGA Impression Plus
- One or more Matrox MGA Ultima Plus
- One or more Matrox MGA Millennium

### G.1.5 International Keyboards

Metro-X uses the X Keyboard extension to support various keyboards. The keyboard description contains information about the physical layout of a keyboard, the key codes that come from the keyboard, and the symbol information needed to map the keycodes into keysyms.

Metro-X provides many predefined configurations, sorted by language and country. These configurations are selected through `configX` (section 6.3.3 on page 98).

## G.2 Metro X Configuration File

### G.2.1 Warning

The primary configuration file for the X server is `Xconfig`. If inappropriate changes are made to this file, the proper functioning of Metro-X will be impaired! Do not edit the `Xconfig` file unless absolutely necessary.

### G.2.2 Xconfig

Metro-X looks for the `Xconfig` file in the following locations and in the following order:

```
$HOME/Xconfig
$HOME/.Xconfig
/etc/X11/Xconfig (on Linux only)
/usr/lib/X11/Metro/Xconfig
```

Most options are set with the configuration utility (`ConfigX`, section 6.3.3 on page 98). Additional options can be added using an editor, and `ConfigX` will preserve these.

### G.2.3 Keywords

Descriptions of the `Xconfig` keywords are listed below. They are alphabetized here for your convenience and any order dependencies are noted within the individual entries.

#### Options Set In Configuration Utility

The following list of `Xconfig` keywords are options controlled via the Metro-X Configuration Utility. Any changes to these options should be made through the utility to avoid accidental damage to the function of Metro-X.

`BitsPerPixel <number>` Some X drivers support packed 24 bits per pixel memory organization. This means that each pixel takes up three bytes of video memory. In this case, `configX` inserts the line: `BitsPerPixel 24`

**Buttons** <number> This specifies the number of buttons on the mouse. *This entry must immediately follow the Mouse entry in the Xconfig file.*

**Card** <string> This designates the model of the graphics card in use. This is only used by configX.

**Chipset** <string> This displays the chipset used by the server. Normally configX inserts the correct chipset into the Xconfig file.

**Clocks** <number> [<number>...] This contains a list of numbers specifying the frequencies associated with the clock generator on the card. The numbers are specified in MHz.

**Depth** <number> This references the color depth of the screen, which controls how many colors can be displayed. Different cards and drivers support different depth values: 4, 8, 16, 24, or 32. Depth 4 supports 16 colors and depth 32 supports 16 million colors.

**Dimensions** <width> <height > This keyword represents the dimensions of the screen in millimeters, which is passed to X clients upon request.

**DPMSoffDelay** <number> This controls the number of seconds after the last user input before the screen enters DPMS off mode. Off mode provides maximum power reduction. To activate this mode, the X server disables both the horizontal and vertical sync signals to the monitor.

**DPMSstandbyDelay**<number> This controls the number of seconds after the last user input before the screen enters DPMS standby mode. Standby mode provides minimal power reduction. To activate this mode, the X server disables the horizontal sync signal to the monitor.

**DPMSsuspendDelay**<number> This controls the number of seconds after the last user input before the screen enters DPMS suspend mode. Suspend mode provides significant power reduction. To activate this mode, the X server disables the vertical sync signal to the monitor.

**Driver**<string> This references the name of the driver, which usually corresponds to the name of the X server. For example:

```
Driver 8514
```

initiates the section for defining the 8514 driver.

**Emulate3Buttons** This entry is used when a two-button mouse is emulating a three-button mouse. This entry must immediately follow the Mouse entry in the

Xconfig file.

**FontPath** *<string>* FontPath is followed by a comma-delimited string of font directories. A font directory is any directory which contains fonts and a fonts.dir file.

**HotKeyExit** This allows you to shut down the X server with the key sequence **Ctrl-Alt-Backspace**. If this keyword is not present, **Ctrl-Alt-Backspace** does nothing special.

**KbdDesc** *<string>* The quoted string must match one of the entries in the KeyboardInfo file.

**KbdDescFile** *<string>* The string describes the keyboard.

**LogFile** *<filename>* Errors from the X server are logged here.

**Modes** The entries following the Modes give display mode timing information. If more than one mode is specified, the server can switch between modes. See Modes ( section G.2.4 on page 261 for more details.

**Monitor** "*<string>*" The quoted string must match one of the entries in the MonitorInfo file. **Mouse***<type>* [*device*] This describes the type of mouse in use. Type is a keyword that is OS dependant; the device parameter is usually required and is a file in the /dev directory.

**MouseDesc** "*<string>*" The quoted string must match one of the entries in the MouseInfo file.

**ScreenLayout***<screen>* *<left\_neighbor>* *<right\_neighbor>* *<top\_neighbor>* *<bottom\_neighbor>* *<left\_edge>* *<right\_edge>* *<top\_edge>* *<bottom\_edge>* This defines how the mouse cursor behaves when it gets to the edge of the screen. There is one ScreenLayout entry for each configured screen. The \*\_neighbor fields tell which screen to move the mouse to. A value of -1 indicates that the mouse should stop at that edge. The \*\_edge fields tell at which edge of the corresponding \*\_neighbor screen the mouse will appear. Screens are numbered from 0 to 4. Edges are numbered as follows:

**NOTE:** Any changes made to this entry will be lost if configX is run.

**ScreenLayout Index***<number>* This tells which of the predefined screen layouts was chosen in configX.

**Virtual***<width>* *<height>* This designates the virtual width and height of the screen. The virtual width and height may be larger than the physical width and

EDGES	NUMBER
ENTER_LEFT	1
ENTER_RIGHT	2
ENTER_UP	3
ENTER_DOWN	4

**Figure G.1:** ScreenLayout Edge Numbering

height if hardware panning is used.

`weight<string>` This designates the RGB color specification, which is dependent on the depth value. When the depth value is 16, the options are:

```
R5:G5:B5 (default) or
R5:G6:B5
```

When the depth value is 32, the options are:

```
R8:G8:B8 (default) or
B8:G8:R8
```

When the depth value is 4 or 8, the weight value is ignored. `XscreenSaverDelay` This controls the number of seconds (with no input) before the screen is blanked.

### Options Not Set in Configuration Utility

The following keywords control options not set through the configuration utility, and require manual editing of the Xconfig file.

`BackBuffer` This entry tells the X server to allocate off-screen memory for a Back-Buffer.

`CursorAdjustX`

`CursorAdjustY` If the hardware cursor does not position correctly over the X Window screen, `CursorAdjustX` and `CursorAdjustY` allow you to specify a delta (number of pixels) to adjust the cursor left or right.

`CursorScreenOne` In a multiple screen configuration, this setting constrains the cursor to screen 0.

`NoBackingStore` This turns off backing store. Backing store is a server feature that allows programs to create complex drawings and have the server automatically handle exposure requests, thus preventing lengthy re-draw cycles. Backing store uses system memory to hold the window contents and can actually slow down the system if there is limited memory.

`NoCursorWrap` In a multiple screen configuration, this setting prevents the cursor from wrapping from screen 1 to screen 0.

`NoFontCache` This turns off font caching.

`NoKeyboard` This setting disallows keyboard input.

`NoLEDs` No LED processing is done by the server.

`NoPixmapCache` This turns off pixmap and stipple caching.

`NotPrimaryMonitor` No console processing is done for Virtual Terminal Support. This is normally used when X is run on a graphics card that is not driving the console.

`Option <string>` This specifies a special option passed to the server. It might specify the card model or a specific feature not typically used. Normally `configX` inserts the correct options in the Xconfig file.

`Option HardwarePex` Add this line to the Xconfig file to use the Hardware Z Buffer on the Matrox MGA Impression Plus graphics card. This file entry must immediately follow the Driver MGA entry.

`Option NoHardwareCursor` Add this line to disable hardware cursor support. This file entry must immediately follow the Driver entry.

`RemapCursorPad` This configuration causes the X server to send back different keycodes for the keypad arrow keys and cursorpad arrow keys. This is a minor violation of the X protocol, but it allows clients to differentiate between these keys.

`SnfFormat <number> <number> <number> <number>` Metro-X X11.5 can read X11.4 SNF (Server Natural Format) fonts. Since there is no standard for SNF fonts, the following four numbers can be adjusted to read most SNF fonts. The four number parameters represent the following values, normally built into the server at compile time:

```
BITMAP_BIT_ORDER
IMAGE_BYTE_ORDER
GLYPHPADBYTES
```

```
BITMAP_SCANLINE_UNIT
```

For example, to read SCO SNF fonts, add this line to the Xconfig file:

```
SnfFormat 1 0 1 1
```

Or, to read SNF fonts compatible with X11 R4 from Interactive Systems Corporation, add this line to the Xconfig file:

```
SnfFormat 1 1 4 4
```

`Terminate` `Terminate` causes the server to exit when the last client disconnects. Normally `xinit` kills the X server when the last client disconnects, before the server can reset.

`Trace <number>` This prints additional debugging information.

`VideoRam <number>` This specifies the amount of video memory on the graphics card, specified in terms of 1024-byte units (with one megabyte equal to a value of 1024). The X server usually detects how much memory is on the card.

`VisualClass <number>` This uses the default `VisualClass` from `/usr/include/X11/X.h`, and is based on the following table:

VISUAL CLASS	NUMBER
StaticGray	0
GrayScale	1
StaticColor	2
PseudoColor	3
TrueColor	4
DirectColor	5

**Figure G.2:** Visual Classes

On monochrome monitors the line

```
VisualClass 1
```

should be added to the Xconfig file.

### G.2.4 Modes

The last section of the `xconfig` file is the Mode database (`Modes`). It contains the VESA timing specifications for various resolutions. The default values work for most graphics cards, but they can be tweaked to center the image and adjust its size on the monitor. Most of the Metro-X servers use this database to configure the card for graphics mode. The numbers indicate which clock to use from the `DisplayConfig` file's `CLOCKS` line, and also show the horizontal and vertical timing values. More specifically, an example of the numbers are:

```
''1024x768'' 65 1024 1092 1220 1344 768 786 791 810
```

The following table shows the definition of each Modes number from the prior example.

VALUE	Meaning
"1024X768"	Screen Resolution in Pixels
65	Pixel Clock (MHz)
1024	Pixels Displayed (H)
1092	Sync Start (H)
1220	Sync End (H)
1344	Total (H)
768	Pixels Displayed (V)
786	Sync Start (V)
791	Sync End (V)
810	Total (V)

**Figure G.3:** Modes Definitions

These numbers represent the ranges the electron beam traverses per one vertical refresh of the screen. They determine where the image is positioned on the screen, and what the refresh rate is.

#### Adjusting the Image Position

You can edit the Modes numbers if the screen image does not properly center or fill the monitor screen. Be sure to center the image before stretching or shrinking it. Any adjustments you make will go into effect after restarting the X server.

**Centering the Image:** The `Pixels Displayed` numbers must match the desired physical screen resolution. The `Sync Start` and `Sync End` numbers must be increased or decreased (to move left or right; up or down) at the same time and by the same amount, and stay within the range of the `Displayed` and `Total` values.

- **Image too far right:** To adjust an image that is too far to the right, increase the `Horizontal Sync Start` and `Horizontal Sync End` values equally.
- **Image too far left:** To adjust an image that is too far to the left, decrease the `Horizontal Sync Start` and `Horizontal Sync End` values equally. For example, to adjust an image that is too far to the left, you could begin the process by adding 8 to the two horizontal timing numbers (`Sync Start` and `Sync End`). Make sure that the `Sync End` number does not exceed the `Horizontal Total` number. The same technique applies for vertical positioning:
- **Image too low:** To adjust an image that is too far to the bottom, decrease the `Vertical Sync Start` and `Vertical Sync End` values equally.
- **Image too high:** To adjust an image that is too far to the top, increase the `Vertical Sync Start` and `Vertical Sync End` values equally. For example, to adjust an image that is too high, you could begin the process by adding 8 to the two vertical timing numbers (`Sync Start` and `Sync End`). Make sure that the `Sync End` number does not exceed the `Vertical Total` number.

**Stretching or Shrinking the Image:** The `Pixels Displayed` numbers must match the desired physical screen resolution. The `Total` value is increased or decreased to stretch or shrink the image. In addition, to hold the center position, the `Sync Start` and `Sync End` numbers must be increased or decreased equally, based on one-half of the difference in the changed `Total` value. For instance, if you added 10 to the `Horizontal Total` value, you would also add 5 to the `Horizontal Sync Start` value and 5 to the `Horizontal Sync End` value.

- **Image too narrow horizontally:** To stretch an image that is too narrow horizontally (extra space appears on the left and right), increase the `Horizontal Total` by amount  $x$ , then add one-half of  $x$  to both the `Horizontal Sync Start` and `Horizontal Sync End` values.
- **Image too narrow vertically:** To stretch an image that is too narrow vertically (extra space appears on the top and bottom), increase the `Vertical Total`

by amount  $x$ , then add one-half of  $x$  to both the Vertical Sync Start and Vertical Sync End values.

- Image too wide horizontally: To shrink an image that is too wide horizontally (the image extends beyond the left and right edges), decrease the Horizontal Total by amount  $x$ , then subtract one-half of  $x$  from both the Horizontal Sync Start and Horizontal Sync End values.
- Image too tall vertically: To shrink an image that is too tall vertically (the image extends beyond the top and bottom edges), decrease the Vertical Total by amount  $x$ , then subtract one-half of  $x$  from both the Vertical Sync Start and Vertical Sync End values.

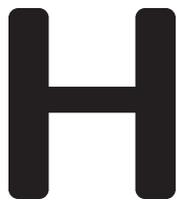
### Extensions

Extensions are configured in the `/usr/lib/X11/Metro/Extensions` file. Most extensions are dynamically loaded; that is, if they are needed they are loaded automatically without any additional manual configuration. An extension that is "ON" is loaded at server startup and never unloaded. An extension that is "DYN" is dynamically loaded on request by an X client `XQueryExtension` call. These extensions display with `XListExtensions` or `xdpyinfo`. An extension that is "OFF" can not be used, nor is it displayed with `XListExtensions` or `xdpyinfo`.

### Special Notes

- Multi-Buffering and DOUBLE-BUFFER can not be dynamically loaded. They must be set to "ON" or "OFF".
- GLX (OpenGL) can be dynamically loaded, however GLX depends on DOUBLE-BUFFER being "ON". Therefore, to use GLX, set DOUBLE-BUFFER to "ON" and GLX to "DYN".
- X3D-PEX (PEX) can be dynamically loaded, however X3D-PEX depends on Multi-Buffering being "ON". Therefore, to use X3D-PEX, set Multi-Buffering to "ON" and X3D-PEX to "DYN".





# Red Hat Linux Frequently Asked Questions

## H.1 Introduction

This is the Official Red Hat Linux FAQ. It is maintained by [faq-maintainer@redhat.com](mailto:faq-maintainer@redhat.com); all comments or suggestions for this FAQ should be sent to that address. A recent version of this FAQ is available from <http://www.redhat.com>, or by email from [faq@redhat.com](mailto:faq@redhat.com).

## H.2 General Questions

### H.2.1 Where can I get information on Red Hat Linux?

- Via the Web: <http://www.redhat.com>
- Via FTP: <ftp://ftp.redhat.com>

## H.2.2 How Can I Get Red Hat Linux?

Red Hat Linux is available on CD from Red Hat Software or various Red Hat Linux resellers; it is also available via FTP from `ftp://ftp.redhat.com` and its many mirror sites.

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

There are several mirrors of the Red Hat FTP site:

FTP Site	Directory
=====	=====
UK	
--	
<code>sunsite.doc.ic.ac.uk</code>	<code>/packages/linux/redhat</code>
<code>nfs-uxsup.csx.cam.ac.uk</code>	<code>/pub/redhat/</code>
<code>ftp.ox.ac.uk</code>	<code>/pub/linux/redhat</code>
Europe	
-----	
<code>ftp.osb.hu</code>	<code>/pub/mirrors/linux/redhat</code>
<code>sunsite.icm.edu.pl</code>	<code>/pub/Linux/redhat</code>
<code>ftp.vma.bme.hu</code>	<code>/pub/linux/redhat</code>
<code>ftp.sfos.ro</code>	<code>/mirror/RedHat</code>
<code>ftp.u-bordeaux.fr</code>	<code>/pub/Linux/redhat</code>
<code>ftp.tu-chemnitz.de</code>	<code>/pub/linux/redhat-mirror/</code>
<code>ftp.wcfauna.ee</code>	<code>/pub/Linux/RedHat/</code>
<code>ftp.uni-bayreuth.de</code>	<code>/pub/linux/ftp.redhat.com/</code>
<code>ftp.pk.edu.pl</code>	<code>/pub/linux/redhat</code>
<code>sunsite.uio.no</code>	<code>/pub/unix/linux/redhat/</code>
<code>ftp.etud.insa-tlse.fr</code>	<code>/pub/redhat</code>
<code>ftp.arch.pwr.wroc.pl</code>	<code>/mirror/linux/redhat/</code>
<code>ftp.ras.ru</code>	<code>/pub/Linux/redhat</code>
<code>ftp.u-bordeaux.fr</code>	<code>/pub/Linux/redhat</code>
<code>ftp.ibp.fr</code>	<code>/pub/linux/distributions/redhat</code>
<code>redhat.netvision.net.il</code>	<code>/pub/mirrors/redhat</code>
<code>ftp.sorosis.ro</code>	<code>/pub/mirrors/ftp.redhat.com</code>
<code>ftp.mpi-sb.mpg.de</code>	<code>/pub/linux/mirror/ftp.redhat.com</code>

```
ftp.jate.u-szeged.hu /pub/linux/redhat
ftp.ibp.fr /pub/linux/distributions/redhat
ftp.gwdg.de /pub/linux/install/redhat
ftp.sgg.ru /mirror/redhat
sunsite.mff.cuni.cz /OS/Linux/Distributions/Redhat
ftp.ton.tut.fi /pub/Linux/RedHat
garbo.uwasa.fi /pub/mirror/RedHat
ftp.funet.fi /pub/Linux/images/RedHat
ftp.arch.pwr.proc.pl /mirror/linux/redhat
ftp.rhi.hi.is /pub/linux/RedHat
ftp.nvg.ntnu.no /pub/linux/redhat
ftp.pk.edu.pl /pub/linux/redhat
ftp.nluug.nl /pub/os/Linux/distr/RedHat
dutepp0.et.tudelft.nl /pub/Unix/Linux/Distributions/redhat
ftp.iol.ie /pub/Unix/Linux/distributions/RedHat
sunsite.auc.dk /pub/os/linux/redhat
ftp.sorosis.ro /pub/mirrors/ftp.redhat.com
```

## Africa

-----

```
ftp.is.co.za /linux/distributions/redhat
ftp.leg.uct.ac.za /pub/linux/distributions/redhat/
ftp.sun.ac.za /pub/linux/redhat/
```

## Australia

-----

```
dopey.gist.net.au /pub/redhat
ftp.cse.unsw.EDU.AU /pub/linux/redhat
www.unimelb.edu.au /pub/unix/linux/redhat
sunsite.anu.edu.au /pub/linux/distributions/redhat/
```

## Asia

-----

```
ftp.jaist.ac.jp /pub/os/linux/redhat
NCTUCCA.edu.tw /OS/Linux/distributions/RedHat
ftp.dacom.co.kr /pub/Linux/redhat
ftp.lab.kdd.co.jp /OS/Linux/packages/redhat
sunsite.ust.hk /pub/Linux/distributions/redhat
ftp.tku.edu.tw /Unix/Linux/RedHat
ftp.wownet.net /LINUX/redhat
ftp.kreonet.re.kr /pub/Linux/redhat
```

```

ftp.iss.nus.sg           /pub/redhat

South America
-----

ftp.sunsite.dcc.uchile.cl /pub/OS/linux/redhat

North America
-----

ftp.kernel.org           /pub/mirrors/redhat
sunsite.unc.edu         /pub/Linux/distributions/redhat
tsx-11.mit.edu          /pub/linux/distributions/redhat/
ftp.redhat.com          /pub/linux/distributions/RedHat/
peyote-asesino.nuclecu.unam.mx /linux/redhat-4.1
ftp.epix.net            /pub/Linux/distributions/redhat
linux.eecs.umich.edu    /pub/linux/redhat
redhat.aiss.uiuc.edu    /pub/redhat
ftp.varesearch.com      /pub/mirrors/redhat
ftp.xtn.net             /pub/linux/redhat
ftp-nog.rutgers.edu     /pub/linux/distributions/redhat
ftp.interpath.net       /pub/linux/redhat
schlitz.cae.wisc.edu    /pub/Linux/RedHat
ftp.engr.uark.edu       /pub/linux/redhat
ftp.infomagic.com       /pub/mirrors/linux/RedHat
ftp.wgs.com             /pub/linux/redhat
ftp.hkstar.com          /pub/Linux/redhat
ftp.pht.com             /pub/linux/redhat
linux.ucs.indiana.edu   /pub/linux/redhat
ftp.ou.edu              /linux/redhat
ftp.cc.gatech.edu       /pub/linux/distributions/redhat
uiarchive.uiuc.edu      /pub/systems/linux/redhat
ftp.caldera.com         /pub/mirrors/redhat
ftp.cms.uncwil.edu      /linux/redhat
ftp.wilmington.net     /linux/redhat
gatekeeper.dec.com      /pub/linux/redhat
ftp.rge.com             /pub/systems/linux/redhat
linuxwww.db.erau.edu    /pub/linux/distrib/redhat
ftp.eit.com             /pub/mirrors/redhat
ftp.real-time.com       /pub/redhat
ftp.priority.net        /pub/redhat
ftp.miranda.org         /pub/mirrors/redhat
ftp.siriuscc.com        /pub/Linux/Redhat
gandalf.ncssm.edu       /pub/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.

### H.2.4 How do I order from Red Hat Software?

- Via the Web: <http://www.redhat.com>
- Via email: [sales@redhat.com](mailto:sales@redhat.com)
- Via fax: +1-919-547-0024
- Via telephone (toll-free): +1-888-RED-HAT1
- Via telephone (toll call): +1-919-547-0012

### H.2.5 How do I contact Red Hat Software?

Red Hat Software  
4201 Research Commons  
Suite 100  
79 T. W. Alexander Drive  
P. O. Box 13588  
Research Triangle Park, NC 27709  
USA  
tel: +1-919-547-0012  
fax: +1-919-547-0024  
<http://www.redhat.com>  
<ftp://ftp.redhat.com>  
<mailto:redhat@redhat.com>

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

Keep your eye on:

<http://www.redhat.com/errata>

for new information we've made available since the current CD was pressed. Also check the `contrib` directory on our FTP mirrors for packages that users have contributed. We also make periodic announcements to the `redhat-announce-list` with updates.

### **H.2.7 What do I do if I have media (manual or CD) problems?**

If you are experiencing problems with the product you received (e.g., your manual is missing pages or your CD is scratched), please contact Red Hat Software at [sales@redhat.com](mailto: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 Red Hat Support, as they are likely isolated and your product will simply need to be replaced.

(Note: This is not a common problem, but it does happen from time to time.)

### **H.2.8 What's the quickest and best way to get support for Red Hat?**

- The Red Hat Software WWW site (<http://www.redhat.com>)
- The Red Hat Linux Users' Mailing Lists; see <http://www.redhat.com/mailling-lists/>.

### **H.2.9 How do I report a bug?**

Send email to [bugs@redhat.com](mailto:bugs@redhat.com).

### **H.2.10 What versions of Red Hat Linux are there?**

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

- Preview (or Beta) - Summer 1994
- Halloween - Fall 1994
- Mother's Day 1.0 - Summer 1995
- Mother's Day 1.1 - Late Summer 1995
- Red Hat Linux 2.0beta - Late Late Summer 1995
- Red Hat Linux 2.0 - Early Fall 1995

- Red Hat Linux 2.1 - Late Fall 1995
- Red Hat Linux/ AXP 2.1-Beta - Late Fall 1995
- Red Hat Linux/ AXP 2.1 - January 1996
- Red Hat Linux 3.0.3 (Picasso) - March 1996
- Red Hat Linux Rembrandt (beta) - August 1996
- Red Hat Linux (Intel/ Alpha/ SPARC) 4.0 - October 1996
- Red Hat Linux (Intel/ Alpha/ SPARC) 4.1 - January 1997
- Red Hat Linux (Intel/ Alpha/ SPARC) 4.2 - May 1997
- Red Hat Linux (Intel/ Alpha/ SPARC) 5.0 - December 1997

### **H.2.11 What hardware does Red Hat Linux support?**

For a current list of supported hardware, please see:

<http://www.redhat.com/hardware>

### **H.2.12 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.

### **H.2.13 What applications are there for Red Hat Linux?**

There are many applications for Linux available. Please see <http://www.redhat.com/linux-info> 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.

### H.2.14 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 with a packaging system called RPM that ONLY uses pristine source (the same as what you'd find at the author's site) and possibly a patch by Red Hat. Complete information on RPM is available on <http://www.redhat.com/>.

### H.2.15 Does Red Hat Linux support UMSDOS installs?

No. It is not a good performance idea.

## H.3 Installation Questions

### H.3.1 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 hdX=cdrom
```

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

### H.3.2 How do I run the install program again?

Reboot your machine using the disk 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.

### H.3.3 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.

### H.3.4 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.

### H.3.5 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.

### H.3.6 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.

### H.3.7 How do I use the NYS/NIS system that comes with Red-Hat?

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

### H.3.8 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.

### H.3.9 Why does my Diamond video card 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).

### H.3.10 Where can I find an list of known problems?

To find solutions to known problems, please see:

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

and look for both a "General Errata" page and a platform-specific Errata page for your version of Red Hat Linux.

### H.3.11 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:

```
rm /etc/X11/X
ln -s ../../usr/X11R6/bin/Xmetro /etc/X11/X
```

## **H.4 Using Red Hat Linux**

### **H.4.1 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 basically 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.

## **H.5 Copyright Notice**

This document is Copyright (C) 1997 by Red Hat Software. Redistribution of this document is permitted as long as the content remains completely intact and unchanged. In other words, you may reformat and reprint or redistribute only.



# Red Hat and Freely Distributable Software.

## I.0.1 Red Hat History

Linux, like Unix itself, is a very modular operating system. The skills required to select, compile, link, and install, the various components that are needed for a complete Linux OS are beyond the experience of most people who might want to use Linux. The various Linux distributions go a long way towards solving this for the average Linux user, but most do not address the problem of how to upgrade your Linux system once you get it successfully installed. Most users found it easier to delete their whole Linux system and reinstall from scratch when they needed to upgrade.

The Red Hat distribution makes Linux easier to install and maintain by providing the user with advanced package management, graphical (point and click!) system installation and control, and system administration tools.

## I.1 The Free Software Model

Probably the best feature of Linux and the GNU utilities in general and Red Hat Linux in particular is that it is distributable under the terms of the GNU Public License (GPL). This feature has allowed research institutions, universities, commercial enterprises, and hackers, to develop and use Red Hat Linux and related technologies co-operatively without fear that their work would someday be controlled and restricted by a commercial vendor.

In short, the GPL changes the model of software development and distribution to one much like the model our Legal system and its industry uses. If a lawyer designs an argument that wins his case in front of the supreme court his reward is not only the fees his client pays him but also the additional clients that his achievement attracts to his practice. The "argument" he used becomes available for any other lawyer to use without restriction, and in fact becomes part of our collective legal heritage.

This new model is already a new industry with companies like Cygnus Support (see <http://www.cygnus.com>) becoming multi-million dollar enterprises on the basis of providing support and services for large commercial users of GPL'd and other "freely distributable" software. Red Hat Software, Inc.'s rapid growth is based, similarly, on our development and support of the Red Hat Linux distribution, a product that we distribute worldwide on CD and over the Internet under the terms of the GPL.

## I.2 The Evolution of "Free Software"

The Unix operating system was also developed on a cooperative basis by AT&T with the active support of many universities and research establishments. Many of the early Unix contributors were disappointed when AT&T began in the late 1970's and 1980's, to enforce their copyright terms over the use and distribution of Unix. This eventually resulted in the long and bitter legal battle between AT&T (and subsequently Novell who bought the Unix copyrights from AT&T in 1991) and the University of California at Berkeley. In a partial response to this problem of individual control of a public software development project, The Massachusetts Institute of Technology (MIT) in Cambridge, MA has, since 1984, supported Richard Stallman's efforts to launch a movement in the software development community to support the development of free software. This resulted in the establishment of The Free Software Foundation (the FSF), and its GNU project. The name GNU

follows in an early Unix tradition of using recursive acronyms as program names, and stands for Gnu's Not Unix. The FSF continues to be head-quartered at MIT, and has enjoyed amazing success over the last decade with products like GNU C, GNU C++, Emacs, and many other products, many of which form substantial and key building blocks of the Linux operating system. Richard Stallman insists, with some justification, that Linux should correctly be known as a GNU system.

There are many other "free software" licenses that some of the code in Red Hat Linux has been licensed under. The most popular of these are the University of California at Berkeley BSD License whereby the copyright holder remains the registered owner but where no restrictions on further distribution are imposed on the user of the code.

Also coming from MIT, but not related to the FSF or the GPL, is the X Consortium, which is a Unix industry sponsored effort to build and maintain the X Window System. In turn, a group called the XFree86 Project has taken on the responsibility to port the X Window System to the Intel platform. We have included copies of each of these licenses in the appendix of this manual for your convenience. Linux is continuing to progress very rapidly because of the contributions of this world wide development effort, to which you and your company or school are encouraged to contribute.

### I.3 The GNU Public License

Red Hat Software, Inc. has chosen to license its code under the GPL for two reasons. One is that this is the easiest thing to do since the Linux kernel is licensed this way and our code interacts with other GPL'd code sufficiently to make any other Licensing scheme difficult. The other is that for the purpose of cooperative development purposes we genuinely feel that the GPL is in fact the most effective license format.

The purpose of the GPL is much like any standard copyright license, to define what you can and cannot use the product for. However it differs substantially from commercial licenses because instead of restricting your use of the software its purpose is to *guarantee* your use of the software.

It in no way restricts the use of standard commercially licensed software with Linux or other GPL'd code. Only the changes that you make to the code that you "borrowed" from someone else that was licensed under the GPL must be made available to others under the same unrestricted rules defined by the GPL. Any

product that simply uses or operates on top of Linux can be licensed under any commercial rules you chose.

There are many leading commercial software developers who are actively supporting their commercial products on Linux for which you must purchase a license, such as Motif from the OSF, Empress RDBMS from Empress Software, WordPerfect, and recently Applixware - the Office Suite from Applix Inc. and Red Hat Software. You will find a complete copy of the GPL in section J.

## I.4 Will Linux Replace Commercial OSes?

The huge development effort and wide distribution of the Linux OS will ensure that it takes its place as a real, viable, and significant alternative to commercially restricted operating systems. The open development model, availability of source, and lack of license restrictions are "features" of the Linux OS that commercial OS developers simply cannot offer. Software development groups who need this model include groups from government affiliated research organizations, to academic research and teaching projects, to commercial software application developers.

The recent rapid increase in new applications becoming available for Linux and the rapidly growing user base of these technologies are causing even the largest computer industry organizations to take Linux seriously. Even Datapro (a McGraw-Hill Company) in their recent 1996 Survey of the UNIX industry concluded that: "Programmers are taking a hard look at the viability of Linux on production platforms now that Linux costs less than Microsoft and has the added benefits of UNIX, such as great performance, inherent power tool sets, and communication capabilities".

The new distribution and development model that Linux, the FSF's GNU code, and other "freely distributable" software represent is a new model that is already successfully competing with the older commercially-restricted model. There is an exciting future for these technologies that we are pleased that you are now a part of.

## I.5 Red Hat Software, Inc.'s Role in Linux Development

The Red Hat Linux distribution is also issued under the GPL. We see our role as making the Linux OS easy to install, configure, and maintain for professional developers, system administrators, and students. Red Hat Linux is available at no cost and few restrictions (see the GPL) on many ftp sites on the Internet. It is also available on CD from Red Hat Software <http://www.redhat.com> or (919) 547-0012 and from most of the Linux archive CD vendors.

Red Hat Software, Inc. provides support for users of the Linux OS, from individuals running a single workstation to large organizations running hundreds of networked machines, in a variety of programs.

We gratefully acknowledge the support that we have received from the Linux development community including Linus Torvalds, everyone on the Linux contributors list, and the thousands of other developers who helped make Linux the success it is today. We would also like to acknowledge the growing contribution of commercial developers to our efforts including Digital Equipment Corporation, O'Reilly & Associates, Metro Linux Inc, and many others. You are encouraged to use Red Hat's Linux distribution, to develop and add features, or simply package programs that you need or that you think others would appreciate. Please do not hesitate to contact us for any help you require.

You can contact Red Hat Software directly at: [info@redhat.com](mailto:info@redhat.com).

## I.6 Further Reading

There are Internet based listings of both the freely distributable, and commercial software available for Linux in documents such as the Linux Journal, the Commercial HOWTO, the Linux Software Map, and from many Linux vendors, including Red Hat Software's retail division. Many third party and commercial products can be purchased through Red Hat's web site, including Motif, database development systems, editors, and more. Info can be obtained by sending a fax request to (919) 572-6726, and the web site is located at <http://www.redhat.com>.

For more information on Linux please see:

- *Running Linux*, M. Welsh, O'Reilly & Associates, 1995.

- *Linux: Network Administrators Guide*, O'Reilly & Associates, 1995.
- Red Hat Software (<http://www.redhat.com> and <ftp://ftp.redhat.com/pub>)
- The Free Software Foundation (<ftp://prep.ai.mit.edu/pub/gnu> or [gnu@prep.ai.mit.edu](mailto:gnu@prep.ai.mit.edu))
- The X Consortium (<http://www.x.org>) and XFree86 (<http://www.xfree86.org>)
- For general Unix questions, please see the full line of O'Reilly & Associates books. They are perhaps the best reference material for all things Unix related. See the O'Reilly web site (<http://www.ora.com>) for more information.

# J

## Licenses

Nearly all of the software components on the Red Hat Linux CD-ROM are freely redistributable. A few of them require special permission to redistribute, and Red Hat Software, Inc. has obtained these permissions. The majority of the software components are distributable under the terms of one of the three licenses in this chapter. Please see each software component for precise distribution terms.

All the software on the CD-ROM produced by Red Hat Software is Copyright © 1995 by Red Hat Software, Inc. Unless otherwise noted, all such software is freely redistributable under the terms of the GNU Public License (GPL).

The terms Red Hat, rpm, and glint are trademarks of Red Hat Software, Inc.

### J.1 The BSD Copyright

Copyright ©1991, 1992, 1993, 1994 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## J.2 X Copyright

Copyright ©1987 X Consortium

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE X CONSORTIUM BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of the X Consortium shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from the X Consortium.

Copyright ©1987 by Digital Equipment Corporation, Maynard, Massachusetts.

All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and

that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Digital not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

DIGITAL DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL DIGITAL BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

## J.3 Gnu Public License

GNU GENERAL PUBLIC LICENSE Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc. 675 Mass Ave, Cambridge, MA 02139, USA  
Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### J.3.1 Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation’s software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author’s protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors’ reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

## J.3.2 GNU GENERAL PUBLIC LICENSE

### TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

#### 0

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

#### 1

You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

#### 2

You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

- b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

### 3

You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a. Accompany it with the complete corresponding machine readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of

the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

#### 4

You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

#### 5

You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

#### 6

Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

#### 7

If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

**8**

If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

**9**

The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

**10**

If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

**11**

NO WARRANTY

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

### J.3.3 Appendix

#### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what
it does.> Copyright (C) 19yy <name of author>
```

```
This program is free software; you can redistribute it and/or
modify it under the terms of the GNU General Public License
as published by the Free Software Foundation; either version
2 of the License, or (at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty
of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
the GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright ©19yy name of author Gnomovision
comes with ABSOLUTELY NO WARRANTY; for details type `show w`.
This is free software, and you are welcome to redistribute
it under certain conditions; type `show c` for details.
```

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than 'show w' and 'show c'; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in
the program 'Gnomovision' (which makes passes at compilers)
written by James Hacker.
<signature of Ty Coon>, 1 April 1989 Ty Coon, President of
Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

# Index

## Symbols

/etc directory .....	164
/etc/hosts file, managing .....	144
/lib directory .....	164
/sbin directory .....	165
/usr directory .....	165
/usr/local directory .....	166, 168
/var directory .....	167
/etc/pam.conf .....	174
/etc/pam.d.....	174

## A

accounts	
adding .....	132
deactivating .....	135
editing.....	132
removing.....	134
acknowledgments.....	xv
adding NFS mounts .....	137
adding user accounts .....	132
administration, system .....	163
Alpha .....	211
boot diskette, creating .....	214
booting kernel diskette.....	216
booting MILO diskette.....	215
diskette images .....	214
installation overview .....	213
installing .....	215
introduction to installing on.....	211
issues specific to .....	217
MILO diskette, creating.....	214

MILO HOWTO .....	221
MILO partition, creating.....	217
MILO, configuring.....	218
partitioning disks.....	217
ramdisk diskette.....	214
SRM console.....	213
starting the installation .....	217
supported hardware .....	211
TGA.....	217
Alpha computer .....	xiii
Alpha-specific information ...	17, 18, 20-22, 29, 31, 37
AMD.....	90
anonymous FTP .....	109
antispam sendmail, new .....	9
Apache.....	110
Apache web server, configuring .	110
arena.....	116
AT&T .....	278
auto probing .....	2
autoboot .....	39
autostart, selecting services for ...	76

## B

backup script, new .....	8
bash .....	131
BIOS, issues related to LILO .....	29
boot diskette for Alpha, creating .	214
boot Manager .....	<i>see</i> LILO
boot method selection for SPARC201	
booting	

rescue mode ..... 179  
 a trick ..... 181  
 using diskettes ..... 180  
 using LILO ..... 180  
 single user ..... 178  
 booting from diskette for SPARC 202  
 booting installation program ..... 38  
 BSD ..... 217  
 BSD Copyright ..... 283

## C

card  
 license ..... 17  
 registration ..... 17  
 CD-ROM  
 ATAPI ..... 42  
 IDE ..... 42  
 install ..... *see* install, CD-ROM  
 other ..... 43  
 SCSI ..... 43  
 CD-ROM boot commands, SPARC 202  
 CD-ROM installation ..... 23  
 CD-ROM installation for SPARC 206  
 CD-ROM module parameters ... 250  
 changing time / date ..... 147  
 clock ..... 76  
 clone network devices ..... 8  
 color ls ..... 112  
 comp.os.linux ..... 119  
 Compact Discs, Red Hat Linux ... 17  
 components  
 selecting ..... 67  
 selecting for SPARC ..... 207  
 components, missing ..... 18  
 configuration  
 anonymous FTP ..... 109  
 clock ..... 76  
 ethernet ..... 146  
 filesystem ..... 136  
 glint ..... 155  
 group ..... 132  
 hosts ..... 144  
 kernel ..... 148  
 adding modules ..... 149  
 changing modules ..... 149  
 module options ..... 149  
 restarting ..... 150  
 LAN manager printer ..... 140  
 local printer ..... 139  
 Metro-X ..... 74, 96  
 MILO ..... 218  
 network ..... 75, 142  
 network device, adding ..... 144  
 network dialogs ..... 76  
 network routes ..... 147  
 NFS ..... 110  
 NFS mounts ..... 137  
 PLIP ..... 146  
 pocket network adapters ... 146  
 post-installation ..... 89  
 PPP ..... 145  
 printer ..... 77, 137  
 printer test page ..... 141  
 remote printer ..... 140  
 selecting nameservers ..... 142  
 SILO ..... 208  
 SLIP ..... 146  
 SMB printer ..... 140  
 time ..... 76  
 time zone ..... 76  
 token ring ..... 146  
 user ..... 132  
 X Windows ..... 73  
 XFree86 ..... 74  
 configuration, hardware ..... 19  
 configuration, video ..... 20  
 configure  
 Apache ..... 110  
 WWW server ..... 110  
 configuring services ..... 3

configX.....98  
 console commands for SPARC... 199  
 consoles, virtual..... 36  
 control panel  
     seecontrolpanel..... 131  
 controlpanel.....131  
 Copyright.....ii  
 Costales, Bryan.....107  
 creating user groups.....135  
 CSLIP.....32  
 Cygnus.....278  
 Cyrix.....90

## D

daemon, kernel.....148  
 date  
     setting.....147  
 dd, creating installation diskette with  
     195  
 deactivating user accounts.....135  
 dependencies, packages.....68  
 destructive partitioning.....27  
 devices  
     network, clone.....8  
 Diamond.....see Video  
     Adapter,Diamond  
 Digital Equipment Corporation..281  
 directories  
     /etc.....164  
     /lib.....164  
     /sbin.....165  
     /usr.....165  
     /usr/local.....166, 168  
     /var.....167  
 Disk Druid.....2, 53  
     aborting.....60  
     adding NFS with.....59  
     adding partitions with.....57  
     buttons.....56  
     current partitions screen.....54  
     deleting partitions with.....58  
     drive summary screen.....55  
     editing partitions with.....59  
     problems adding partitions..58  
     starting over.....60  
     using.....54  
     when finished with.....60  
 disk partitioning, SPARC.....208  
 disk, partitioning.....25  
 disk-based installation for SPARC206  
 diskette  
     boot, creating.....193  
     supplemental, creating.....193  
 diskette boot commands, SPARC 202  
 diskette group, new.....9  
 diskette images for Alpha.....214  
 diskette, making under Linux-like O/S  
     195  
 diskette, making with MS-DOS..194  
 diskettes.....17  
     boot.....17  
     images, updated.....18  
     supplemental.....17  
 documentation.....113  
     FAQs.....116  
     finding.....115  
     HOWTOs.....116  
     HTML.....116  
     info pages.....117  
     locate command.....117  
     makewhatis, searching with.115  
     man pages.....113  
         printing.....114  
         searching.....115  
         section.....114  
     on-line.....113  
     package documentation.....115  
     PAM.....176  
     READMEs.....115  
 documentation, how to improve .xvi

documentation, obtaining additional  
31  
drivers, kernel . . . . . 32

## E

editing user accounts . . . . . 132  
editing user groups . . . . . 136  
editor's acknowledgements . . . . . xvi  
editor's notes . . . . . xv  
emacs . . . . . 118, 279  
Empress . . . . . 280  
errata . . . . . 18  
ET4000 . . . *see* Video Adapter, ET4000  
ethernet . . . . . 146  
ethernet module parameters . . . . . 251  
ethernet, supporting multiple cards 252  
expert installation . . . . . 2  
expert installation mode . . . . . 38  
exporting NFS filesystems . . . . . 111

## F

### FAQ

Adaptec problems . . . . . 271  
applications for Red Hat Linux  
271  
booting Unixware . . . . . 275  
bugs, reporting . . . . . 270  
contacting Red Hat Software 269  
Diamond video cards . . . . . 275  
errata information . . . . . 269  
FTP sites mirroring Red Hat Linux  
266  
Getting Red Hat Linux . . . . . 266  
getting support for Red Hat Linux  
270  
install onto MS-DOS partitions 274  
known problems . . . . . 275  
missing programs . . . . . 273  
mounting CD-ROMs . . . . . 273  
NYS/NIS usage . . . . . 274

obtaining information on Red Hat  
Linux . . . . . 265  
ordering from Red Hat Software  
269  
previous versions of Red Hat Linux  
270  
RAM-related issues . . . . . 274  
re-running installation program  
272  
Red Hat Linux . . . . . 265  
resolving media problems . . 270  
SoundBlaster-16 IDE CD-ROM sup-  
port . . . . . 272  
source code . . . . . 272  
starting Metro-X . . . . . 275  
supported hardware . . . . . 271  
technical support . . . . . 189  
UMSDOS support . . . . . 272  
updates to Red Hat Linux . . 269  
xmh broken . . . . . 276  
FAQs . . . . . 116  
FAT32 partitions . . . . . 28  
fdisk . . . . . 53  
overview of . . . . . 61  
using . . . . . 61  
features, new to 5.0 . . . . . 1  
feedback, how to give . . . . . xvi  
filesystem  
standard . . . . . 164  
structure . . . . . 163  
filesystem configuration . . . . 64, 136  
adding NFS mounts . . . . . 137  
other partitions . . . . . 64  
root partition . . . . . 64  
filesystems  
NFS, exporting of . . . . . 111  
NFS, mounting of . . . . . 111  
finishing installation . . . . . 73  
fips partitioning utility . . . . . 28  
formatting partitions . . . . . 66

framebuffer ..... 209  
 free software ..... 278  
 Free Software Foundation... *see* FSF  
 free software, evolution of ..... 278  
 Frequently Asked Questions ..... *see*  
     FAQs  
 friends, impressing with RPM... 128  
 FSF ..... 278–280, 282  
     email ..... 282  
     ftp ..... 282  
 FSSTND ..... 164  
 fstab ..... 136, 137  
 FTP  
     anonymous ..... 109  
     ftpaccess ..... 109  
     ftphosts ..... 109  
     ftpusers ..... 109  
 FTP installation ..... 23  
 FTP installation for SPARC ..... 206

## G

getting started ..... xiv  
 GLINT ..... *see* glint  
 glint ..... 151  
     available packages ..... 154  
     installing packages with .... 159  
     menus ..... 153  
     package display ..... 152  
     package manipulation ..... 156  
     querying packages ..... 156  
     removing packages with .... 161  
     selecting packages ..... 153  
     starting ..... 152  
     uninstalling packages with.. 161  
     upgrading packages with... 160  
     verifying packages ..... 157  
 glint configuration ..... 155  
 GNU ..... 279, 280  
 GNU general public license ..... 285  
 Gnu General Public License . 278, 279

GPL ..... 278–281  
 group configuration ..... 132  
 groups ..... 168  
     creating ..... 135  
     editing ..... 136  
     standard ..... 169  
     user-private ..... 168, 171  
     rationale ..... 172

## H

halt ..... 165, 179  
 hard disk installation for SPARC . 206  
 hard drive installation ..... 25  
 hardware configuration ..... 19  
 hardware supported by Alpha... 211  
 hardware supported by SPARC.. 197  
 history  
     Red Hat Linux ..... 277  
 hostname ..... 142  
 hosts, managing ..... 144  
 hosts.allow ..... 108  
     example ..... 108  
 hosts.deny ..... 108  
 HOWTOs ..... 116  
 HTML ..... 110

## I

info pages ..... 117  
 information  
     Alpha-specific . 17, 18, 20–22, 29,  
       31, 37  
     Intel-specific ... 16–18, 22–25, 28,  
       30, 39  
     SPARC-specific . 3, 17, 18, 20–22,  
       29, 37, 53  
 information, network ..... 20  
 information, pre-installation ..... 19  
 init, SysV-style ..... 176  
 initrd ..... 93  
 initscript utility, new ..... 8

install

- CD-ROM..... 23, 41
- component selection ..... 67
- continuing ..... 51
- finishing ..... 73, 84
- FTP ..... 23, 42, 43
- getting Red Hat Linux ..... 16
- Hard Drive ..... 25, 42
- hard drive ..... 47
- NFS ..... 24, 41, 45
- NFS server information ..... 47
- package selection ..... 66, 67
- PCMCIA support ..... 22, 40
- selecting network drivers .... 45
- TCP/IP networking ..... 45
- upgrade ..... 51

Install, preparing for ..... 15

installation

- printer ..... 77

installation guide ..... 16

installation method

- CD-ROM ..... 41, 42
- FTP ..... 42, 43
- hard drive ..... 42, 47
- NFS ..... 41, 43
- selecting ..... 41

installation method, choosing for SPARC 205

installation methods ..... 21

installation mode, expert ..... 38

installation overview for Alpha . 213

installation overview for SPARC . 199

installation program

- booting ..... 38
- booting without diskette ..... 39
- keyboard navigation ..... 36
- starting ..... 37
- user interface ..... 35
- virtual consoles ..... 36

installation, continuing ..... 51

installation, starting ..... 35, 39

installing on Alpha ..... 215

installing packages ..... 66

installing packages with glint .... 159

Intel ..... 90

Intel-specific information ..... 16–18, 22–25, 28, 30, 39

## K

kernel ..... 32

- building ..... 89, 93
- custom ..... 89, 93
- initrd image for ..... 93
- modular ..... 89, 90
- monolithic ..... 93

kernel daemon ..... 148

kernel diskette, booting on Alpha 216

kernel drivers ..... 32

kernel options ..... 38

kerneld ..... 148

- adding modules ..... 149
- changing modules ..... 149
- module options ..... 149
- restarting ..... 150

keyboard navigation, installation program ..... 36

keyboard type, selecting ..... 40

keymap

- see keyboard type, selecting .. 40

kickstart installation ..... 2

## L

LAN manager printer, configuring 77, 79

LDP ..... 113, 116

license card ..... 17

licenses ..... 283

LILO ..... 82, 165

- /etc/lilo.conf ..... 92
- Adding options to ..... 83

- alternatives to ..... 83
  - commercial products ..... 84
  - LOADLIN ..... 83
  - SYSLINUX ..... 84
- installing ..... 82
- MBR, installing on ..... 82
- root partition, installing on... 82
- LILO, BIOS-related issues ..... 29
- LILO, partition-related issues ..... 28
- Linux Documentation Project .... *see*
  - LDP, 116
- Linux-like O/S
  - creating installation diskette with 195
- LOADLIN ..... 83
- local printer, configuring ..... 77, 78
- locate command, finding documentation with ..... 117
- lpd printer, configuring ..... 77, 78

## M

- mailing lists
  - xp-list ..... 118
  - redhat-announce-list ..... 118
  - redhat-devel-list ..... 118
  - redhat-digest ..... 118
  - redhat-list ..... 118
  - rpm-list ..... 118
  - sparc-list ..... 118
  - support from ..... 118
- makewhatis ..... 115
- man pages ..... 113
  - printing ..... 114
  - searching ..... 115
  - sections ..... 114
- Matt Welsh ..... *see* Welsh, M.
- Maximum RPM ..... 130
- menus, glint ..... 153
- methods, installation ..... 21
- Metro Link ..... 94, 97

- Metro X
  - font path ..... 105
  - multiple screens ..... 104
- Metro-X ..... 97, 98, 253
  - adjusting image position .... 261
  - configuration ..... 74, 96
  - configuration files ..... 255
  - configuration utility options 255
  - configX ..... 96
  - configX.curses ..... 97
  - configX.ncurses ..... 96
  - extensions ..... 263
  - hardware panning ..... 254
  - hot-key exit ..... 253
  - hot-key resolution switching 253
  - keyboards
    - international ..... 254
    - multiple screens ..... 254
    - supported cards ..... 254
    - special notes ..... 263
  - Xconfig file ..... 255
    - keywords ..... 255
    - modes, VESA ..... 261
- MILO ..... 218
- MILO diskette for Alpha, creating 214
- MILO diskette, booting on Alpha 215
- MILO HOWTO for Alpha ..... 221
- MILO partition, creating for Alpha 217
- missing components ..... 18
- MIT ..... 278
- mkswap ..... 165
- modular sound drivers ..... 13
- module parameters ..... 249
- modules
  - PAM ..... 173
- Motif ..... 280, 281
- mounting NFS filesystems ..... 111
- mouse
  - selection ..... 97, 106
- mouse, configuring ..... 73

MS-DOS  
   creating installation diskette with  
     194

## N

nameserver  
   selecting ..... 142  
 netcfg ..... 8  
 netnews, support from ..... 119  
 Netscape ..... *see* WWW, browsers  
 network  
   interface  
     aliasing ..... 142  
   network adapters, pocket ..... 146  
   network boot commands, SPARC203  
   network boot commands, SPARC in-  
     stall ..... 205  
   network config utility, new ..... 8  
   network configuration ..... 75, 142  
     adding device ..... 144  
   network configuration dialogs ..... 76  
   network devices, clone ..... 8  
   network information ..... 20  
   network routes, managing ..... 147  
   networking ..... 144  
 new features ..... 1  
   anti-spam sendmail ..... 9  
   auto probing ..... 2  
   backup script ..... 8  
   configuring services ..... 3  
   Disk Druid ..... 2  
   diskette group ..... 9  
   expert installation ..... 2  
   filesystem mount tool ..... 5  
   glibc ..... 6  
   initscript utility ..... 8  
   installation-related ..... 1  
   kickstart installation ..... 2  
   miscellaneous ..... 12  
   network config utility ..... 8

network device tool ..... 5  
 shadow utilities ..... 7  
 sound card config tool ..... 13  
 sound drivers ..... 13  
 sysadmin-related ..... 6  
 user info tool ..... 4  
 user password tool ..... 4  
 user-mode tools ..... 3  
 window manager enhancements  
   14  
 Xconfigurator ..... 3  
 XFree86 configuration ..... 3

## NFS

  adding to filesystem configura-  
     tion ..... 137  
   configuration ..... 110  
   exporting ..... 111  
   mounting ..... 111  
 NFS installation ..... 24  
 NFS installation for SPARC ..... 206  
 NFS-mounted root, SPARC install us-  
   ing ..... 201  
 non-destructive partitioning ..... 27  
 notes, editor's ..... xv  
 Novell ..... 278

## O

O'Reilly & Associates .. 107, 110, 112,  
   281  
 on-line documentation ..... 113  
 options, kernel ..... 38  
 OS/2 ..... 28, 82  
 OSF ..... 280  
 overview ..... xiii

## P

package  
   installation screen ..... 69  
 package manipulation with glint 156  
 packages

- dependencies . . . . . 68, 124
- determining file ownership . 128
- documentation . . . . . 115
- finding deleted files from . . . 128
- handy hints . . . . . 128
- installing . . . . . 66, 123
- installing with glint . . . . . 159
- locating documentation for . 129
- obtaining list of files . . . . . 129
- preserving config files . . . . . 125
- querying . . . . . 126
- querying uninstalled . . . . . 129
- removing . . . . . 124
- removing with glint . . . . . 161
- selecting . . . . . 66
- selecting individual . . . . . 67
- uninstalling with glint . . . . . 161
- upgrading . . . . . 125
- upgrading with glint . . . . . 160
- using glint to view available 154
- verifying . . . . . 127
- verifying with glint . . . . . 156, 157
- PAM . . . . . 173
  - additional information . . . . . 176
  - configuration files . . . . . 174
  - modules . . . . . 173
  - services . . . . . 174
- parameters, CD-ROM modules . . 250
- parameters, ethernet modules . . 251
- parameters, module . . . . . 249
- partition
  - /home . . . . . 30
  - /tmp . . . . . 31
  - /usr . . . . . 30
  - /usr/local . . . . . 31
  - /usr/src . . . . . 31
  - MILO . . . . . 31
  - root . . . . . 30
  - swap . . . . . 30
- partitioning . . . . . 25
  - destructive . . . . . 27
  - how many . . . . . 30
  - LILO issues related to . . . . . 28
  - non-destructive . . . . . 27
  - other operating systems . . . . . 28
  - using free space . . . . . 26
  - using in-use partition . . . . . 27
  - using unused partition . . . . . 27
- partitioning disks on Alpha . . . . 217
- partitions
  - changing table of . . . . . 64
  - creating . . . . . 53
  - formatting . . . . . 66
  - numbering of . . . . . 63
- password
  - root, setting . . . . . 81
- passwords
  - shadow . . . . . 176
- PCMCIA support . . . . . 40
- PCMCIA, support during installation
  - 22
- PLIP . . . . . 32, 146
- plip . . . . . *see* PLIP
- Plug and Play . . . . . 13
- pluggable authentication modules *see*
  - PAM
- pocket network adapters . . . . . 146
- post-installation configuration . . . 89
- PPP . . . . . 32, 145
- ppp . . . . . *see* PPP
- pre-installation information . . . . . 19
- printer configuration . . . . . 77, 137
  - finalizing . . . . . 80
  - LAN manager . . . . . 77, 79, 140
  - local . . . . . 77, 78, 139
  - remote . . . . . 140
  - remote lpd . . . . . 77, 78
  - SMB . . . . . 140
  - test page . . . . . 141
  - verifying . . . . . 80

processor  
 AMD ..... 90  
 Cyrix ..... 90  
 Intel ..... 90  
 programs, running at boot time.. 112

## Q

querying packages with glint... 156  
 quick start ..... xiv

## R

ramdisk diskette for Alpha, creating  
 214  
 ramdisk, criteria for use on SPARC  
 200  
 RARP commands, SPARC install 204  
 rawrite, creating installation diskette  
 with ..... 194  
 rc.local, modifying ..... 112  
 README files ..... 115  
 recursion ..... *see* recursion  
 Red Baron ..... 116  
 Red Hat Linux  
 what it is ..... xii  
 Red Hat Linux boxed set, contents of  
 16  
 Red Hat Linux CDs ..... 17  
 Red Hat Linux FAQ ..... 265  
 Red Hat Linux history ..... 277  
 Red Hat Linux, licenses used by 283  
 Red Hat Linux, methods of installing  
 21  
 Red Hat Package Manager .. *see* RPM  
 Red Hat Software  
 contacting  
 email ..... 281  
 FTP ..... 282  
 mailing lists ..... 118  
 axp-list ..... 118  
 redhat-announce-list ..... 118

redhat-devel-list ..... 118  
 redhat-digest ..... 118  
 redhat-list ..... 118  
 rpm-list ..... 118  
 sparc-list ..... 118  
 WWW ..... 118, 281, 282  
 Red Hat Software, contributions to Linux  
 281  
 Red Hat-specific file locations ... 168  
 registration card ..... 17  
 registration sticker ..... 17  
 remote lpd printer, configuring 77, 78  
 removing packages with glint ... 161  
 removing user accounts ..... 134  
 repartitioning strategy ..... 26  
 rescue mode ..... 179  
 using a trick ..... 181  
 using diskettes ..... 180  
 using LILO ..... 180  
 restarting kernel ..... 150  
 rhbackup ..... 8  
 root password ..... 81  
 routes, managing ..... 147  
 RPM ..... *see* rpm  
 rpm ..... 121  
 book written about ..... 130  
 dependencies ..... 124  
 design goals ..... 122  
 determining file ownership with  
 128  
 file conflicts, resolving ..... 124  
 finding deleted files with ... 128  
 handy hints ..... 128  
 installing ..... 123  
 locating documentation with 129  
 mailing list devoted to ..... 130  
 other resources ..... 130  
 preserving config files ..... 125  
 querying ..... 126  
 querying for file list ..... 129

- querying uninstalled packages 129
- uninstalling . . . . . 124
- upgrading . . . . . 125
- using . . . . . 123
- verifying . . . . . 127
- website devoted to . . . . . 130
- ruling the world, linux . . . . . 280
- Running Linux* . . . . . 281

## S

- SCSI . . . . . 32
- SCSI support . . . . . 53
- selecting components . . . . . 67
- selecting packages . . . . . 66
- selecting packages with glint . . . . . 153
- sendmail . . . . . 107
  - aliases . . . . . 107
  - masquerading . . . . . 108
  - with UUCP . . . . . 107
- serial terminal installation for SPARC 207
- services
  - PAM . . . . . 174
- services, configuring . . . . . 3
- services, controlling access to . . . . . 108
- services, selecting for autostart . . . . . 76
- setenv . . . . . 131
- shadow passwords . . . . . 176
- shadow utilities, improved . . . . . 7
- shutdown . . . . . 165, 179
- SILO configuration . . . . . 208
- SLIP . . . . . 32, 146
- slip . . . . . *see* SLIP
- sound card config tool, new . . . . . 13
- sound drivers, new . . . . . 13
- spam, sendmail changes to combat 9
- SPARC
  - boot method, choosing . . . . . 201
  - booting
    - from diskette . . . . . 202

- CD-ROM boot commands . . . . . 202
- CD-ROM installation . . . . . 206
- choosing installation method 205
- console commands . . . . . 199
- disk partitioning . . . . . 208
- diskette boot commands . . . . . 202
- FTP installation . . . . . 206
- hard disk installation . . . . . 206
- installation overview . . . . . 199
- installation with serial terminal 207
  - network boot commands . . . . . 203
- NFS installation . . . . . 206
- NFS-mounted root, installation using . . . . . 201
- overview of information specific to . . . . . 197
- ramdisk criteria . . . . . 200
- selecting components . . . . . 207
- SILO configuration . . . . . 208
- supported hardware . . . . . 197
- swap partitions . . . . . 208
- unsupported CD-ROMs . . . . . 198
- unsupported hardware . . . . . 198
- using network boot commands 205
  - using RARP during install . . . . . 204
  - using TFTP during install . . . . . 204
  - X Windows . . . . . 209
- SPARC computer . . . . . xiii
- SPARC-specific information 3, 17, 18, 20–22, 29, 37, 53, 197
- SRM console for Alpha . . . . . 213
- Stallman, Richard . . . . . 278
- standard groups . . . . . 169
- standard users . . . . . 169
- starting installation . . . . . 39
- starting installation program . . . . . 37
- sticker, registration . . . . . 17
- strategy, repartitioning . . . . . 26

structure, filesystem . . . . . 163  
 support from mailing lists . . . . . 118  
 support from Usenet . . . . . 119  
 support, technical . . . . . 183  
 swap partitions for SPARC . . . . . 208  
 swap space, initializing . . . . . 65  
 swapoff . . . . . 115, 165  
 swapon . . . . . 115, 165  
 SYSLINUX . . . . . 84  
 system administration . . . . . 163  
 System Commander . . . . . 84  
 system shutdown . . . . . 179  
 SysV init . . . . . 176  
   directories used by . . . . . 177  
   runlevels used by . . . . . 179

**T**

TCP wrapper . . . . . 108  
 tcpd . . . . . 108  
 tcsh . . . . . 131  
 technical support . . . . . 183  
   FAQ . . . . . 189  
   how to send questions for . . . 188  
   how to state problem for . . . 187  
   obtaining . . . . . 186  
   overview . . . . . 183  
   registering . . . . . 186  
     via email . . . . . 187  
     via fax . . . . . 187  
     via snail mail . . . . . 187  
     via WWW . . . . . 186  
   where to send support request 189  
 test page, printer . . . . . 141  
 TFTP setup, SPARC install . . . . . 204  
 TGA for Alpha . . . . . 217  
 time  
   setting . . . . . 147  
 time zone, setting . . . . . 76  
 token ring . . . . . 146  
 Torvalds, Linus . . . . . xi, xv, 281

**U**

uninstalling packages with glint . 161  
 University of California at Berkeley  
   278  
 upgrade . . . . . xii  
 upgrade, description of . . . . . 52  
 upgrade, performing . . . . . 51  
 upgrading packages with glint . . 160  
 upgrading, how to . . . . . xv  
 Usenet . . . . . 119  
 user accounts  
   adding . . . . . 132  
   deactivating . . . . . 135  
   editing . . . . . 132  
   removing . . . . . 134  
 user configuration . . . . . 132  
 user groups  
   creating . . . . . 135  
   editing . . . . . 136  
 user interface, installation program 35  
 user-private groups . . . . . 168, 171  
 user-private groups, rationale behind  
   172  
 users . . . . . 168  
   standard . . . . . 169

**V**

verifying packages with glint . . . 157  
 Video Adapter . . . . . 95  
   S3 . . . . . 95  
 video configuration . . . . . 20  
 viewing available packages with glint  
   154  
 virtual consoles . . . . . 36

**W**

Welsh, M. . . . . 281  
 window manager, enhancements . 14  
 WordPerfect . . . . . 280  
 World Wide Web . . . . . *see* WWW

World's Greatest Programmer . . . . *see*  
    Torvalds, Linus  
WWW . . . . . 109, 116  
    browsers . . . . . 109  
    server . . . . . *see* Apache

## X

X Consortium . . . . . 279  
X Windows copyright . . . . . 284  
X windows for SPARC . . . . . 209  
X windows, configuring . . . . . 73  
Xconfigurator . . . . . 74, 95, 96  
    enhancements . . . . . 3  
XFree86 . . . . . 279, 282  
    configuration . . . . . 74, 94  
        Xconfigurator . . . . . 94, 95  
        xf86config . . . . . 94  
    server chart . . . . . 95