Red Hat Linux 7.3

The Official Red Hat Linux x86 Installation Guide

redhat.

Red Hat Linux 7.3: The Official Red Hat Linux x86 Installation Guide
Copyright © 2002 by Red Hat, Inc.

Red Hat, Inc.

1801 Varsity Drive
Raleigh NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588
Research Triangle Park NC 27709 USA

The Red Hat Linux Product Documentation Team consists of the following people:

Sandra A. Moore, Product Documentation Manager — Primary Writer/Maintainer of the *Official Red Hat Linux x86 Installation Guide*; Contributing Writer to the *Official Red Hat Linux Getting Started Guide*

Tammy Fox, Product Documentation Technical Lead — Primary Writer/Maintainer of the *Official Red Hat Linux Customization Guide*; Contributing Writer to the *Official Red Hat Linux Getting Started Guide*; Writer/Maintainer of custom DocBook stylesheets and scripts

Edward C. Bailey, Technical Writer — Contributing Writer to the *Official Red Hat Linux x86 Installation Guide*

Johnray Fuller, Technical Writer — Primary Writer/Maintainer of the *Official Red Hat Linux Reference Guide*

John Ha, Technical Writer — Contributing Writer to the *Official Red Hat Linux Getting Started Guide*

Table of Contents

Introduction

Welcome to the *Official Red Hat Linux x86 Installation Guide*. This guide contains useful information to assist you during the installation of Red Hat Linux. From fundamental concepts such as installation preparation to the step-by-step installation procedure, this book will be a valuable resource as you install Red Hat Linux.

This manual will walk you through a typical Red Hat Linux CD-ROM installation. Once you have completed the installation, you will have a fully functioning Red Hat Linux desktop system.

1. Document Conventions

When you read this manual, you will see that certain words are represented in different fonts, typefaces, sizes, and weights. This highlighting is systematic; different words are represented in the same style to indicate their inclusion in a specific category. The types of words that are represented this way include the following:

`command`

    Linux commands (and other operating system commands, when used) are represented this way. This style should indicate to you that you can type the word or phrase on the command line and press [Enter] to invoke a command. Sometimes a command contains words that would be displayed in a different style on their own (such as filenames). In these cases, they are considered to be part of the command, so the entire phrase will be displayed as a command. For example:

    Use the `cat testfile` command to view the contents of a file, named `testfile`, in the current working directory.

`filename`

    Filenames, directory names, paths, and RPM package names are represented this way. This style should indicate that a particular file or directory exists by that name on your Red Hat Linux system. Examples:

    The `.bashrc` file in your home directory contains bash shell definitions and aliases for your own use.

    The `/etc/fstab` file contains information about different system devices and filesystems.

    Install the `webalizer` RPM if you want to use a Web server log file analysis program.

application

    This style should indicate to you that the program named is an end-user application (as opposed to system software). For example:

    Use Netscape **Navigator** to browse the Web.

[key]

    A key on the keyboard is shown in this style. For example:

To use [Tab] completion, type in a character and then press the [Tab] key. Your terminal will display the list of files in the directory that start with that letter.

[key]-[combination]

    A combination of keystrokes is represented in this way. For example:

The [Ctrl]-[Alt]-[Backspace] key combination will exit your graphical session and return you to the graphical login screen or the console.

text found on a GUI interface

    A title, word, or phrase found on a GUI interface screen or window will be shown in this style. When you see text shown in this style, it is being used to identify a particular GUI screen or an element on a GUI screen (such as text associated with a checkbox or field). Example:

Select the **Require Password** checkbox if you would like your screensaver to require a password before stopping.

top level of a menu on a GUI screen or window

    When you see a word in this style, it indicates that the word is the top level of a pull-down menu. If you click on the word on the GUI screen, the rest of the menu should appear. For example:

Under **Settings** on a GNOME terminal, you will see the following menu items: **Preferences**, **Reset Terminal**, **Reset and Clear**, and **Color selector**.

If you need to type in a sequence of commands from a GUI menu, they will be shown like the following example:

Click on **Programs**=>**Applications**=>**Emacs** to start the Emacs text editor.

button on a GUI screen or window

    This style indicates that the text will be found on a clickable button on a GUI screen. For example:

Click on the **Back** button to return to the webpage you last viewed.

```
computer output
```

    When you see text in this style, it indicates text displayed by the computer on the command line. You will see responses to commands you typed in, error messages, and interactive prompts for your input during scripts or programs shown this way. For example:

Use the `ls` command to display the contents of a directory:

```
$ ls
Desktop            axhome        logs        paulwesterberg.gif
Mail               backupfiles   mail        reports
```

The output returned in response to the command (in this case, the contents of the directory) is shown in this style.

```
prompt
```

A prompt, which is a computer's way of signifying that it is ready for you to input something, will be shown in this style. Examples:

```
$
```

```
#
```

```
[stephen@maturin stephen]$
```

```
leopard login:
```

user input

Text that the user has to type, either on the command line, or into a text box on a GUI screen, is displayed in this style. In the following example, **text** is displayed in this style:

To boot your system into the text based installation program, you will need to type in the **text** command at the boot: prompt.

Additionally, we use several different strategies to draw your attention to certain pieces of information. In order of how critical the information is to your system, these items will be marked as note, tip, important, caution, or a warning. For example:

Note

Remember that Linux is case sensitive. In other words, a rose is not a ROSE is not a rOsE.

Tip

The directory /usr/share/doc contains additional documentation for packages installed on your system.

Important

If you modify the DHCP configuration file, the changes will not take effect until you restart the DHCP daemon.

Caution

Do not perform routine tasks as root — use a regular user account unless you need to use the root account for system administration tasks.

⚠️Warning

> If you choose not to partition manually, a server installation will remove all existing partitions on all installed hard drives. Do not choose this installation class unless you are sure you have no data you need to save.

2. How to Use This Manual

This manual is ideal for users (both new and old) who want a quick and simple installation solution. It will help you prepare your system, walk you through the installation, and assist you in the configuration of Red Hat Linux.

📝Note

> If you currently use Red Hat Linux 4.2 (or greater), you can perform an upgrade. Skim Chapter 1 to review the basics, then read Chapter 3, following the directions as you go. Once you have chosen to perform an upgrade in the installation program, refer to Appendix A.

If you are an experienced user who wants to perform a Red Hat Linux CD-ROM installation, and you do not need a review of the basics, you can skip ahead to Chapter 3 to begin the installation process.

💡Tip

> Refer to the **Red Hat Frequently Asked Questions** for answers to questions and problems that may occur before, during, or after the installation. You will find the FAQ online at:
>
> `http://www.redhat.com/support/docs/faqs/rhl_general_faq/`

2.1. We Need Feedback!

If you spot a typo in the *Official Red Hat Linux x86 Installation Guide*, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report against the component *rhl-ig* in Bugzilla at:

`http://bugzilla.redhat.com/bugzilla/`

When submitting a bug report, be sure to mention the manual's identifier:

`rhl-ig-x86(EN)-7.3-HTML-RHI (2002-04-05T13:43-0400)`

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, please include the section number and some of the surrounding text so we can find it easily.

If you have a support question (for example, if you need help configuring X, or if you are not sure how to partition your hard drive[s]), use the online support system by registering your product at:

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

Chapter 1.
Steps to Get You Started

Note

Although this manual reflects the most current information possible, you should read the Red Hat Linux Release Notes for information that may not have been available prior to our documentation being finalized. The Release Notes can be found on the Red Hat Linux CD #1 and online at:

```
http://www.redhat.com/docs/manuals/linux
```

Before you install Red Hat Linux, you should perform the following steps:

1.1. Do You Have the Right Components?

If you have purchased the Official Red Hat Linux boxed set, you are ready to go. However, mistakes occasionally happen, so now is a good time to double-check the contents of your product.

A black, red, and white Registration Information card is included with your product. A list of the contents of your boxed set version is on the back of the card. Please read over the list and check to make sure that you have all the CDs and manuals that are included with the version of Red Hat Linux that you purchased.

If you have purchased the Official Red Hat Linux boxed set from Red Hat, Inc. (or one of our distributors), and you are missing one or more of the items listed, please let us know. Contact information is also available on the Registration Information card.

How to identify our official boxed set: The bottom of our box has an ISBN number next to one of the bar codes. That ISBN number should be in this form:

```
1-58569-x-y
```

(The x and y will be unique numbers.)

Red Hat partners with companies (international and domestic) so that we can make Red Hat Linux available to you in the most convenient form. Because of these partnerships, you might find that your Red Hat Linux product may not have been actually produced by Red Hat.

If your product has a different ISBN number (or none at all), you will need to contact the company that produced it. Normally, third-party producers will include their logo and/or contact information on the outside of their box; an official Red Hat Linux boxed set lists only our name and contact information.

1.1.1. Where to Find Other Manuals

If your particular product did not include all of the printed Red Hat Linux manuals, you can find them online or on the Red Hat Linux Documentation CD included with your official Red Hat Linux product.

To find the manuals in both HTML and PDF formats online, go to:

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

1.1.2. Registering Your Product

If you have purchased the Official Red Hat Linux boxed set, you should register your product. Registration offers many useful services, such as installation support, access to Red Hat Network, and more. To register your product, go to:

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

You will find your *Product ID* on the Registration Information card in your Official Red Hat Linux boxed set. Once registered, you will have access to all the extras that Red Hat provides to its registered users.

For more information on registering and the scope of Red Hat's technical support offerings, see Appendix C.

1.1.3. No Boxed Set? No Problem!

Of course, not everyone purchases a Red Hat Linux boxed set. It is 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.

For information on downloading and installing Red Hat Linux, refer to:

```
http://www.redhat.com/download/howto_download.html
```

For people installing Red Hat Linux from a CD-ROM not from Red Hat, you may need a boot diskette (also referred to as a boot disk), or if you are using a PCMCIA device during the installation (such as a laptop), PCMCIA boot disks. You may also be able to start the installation directly from the CD. We will discuss this in more detail when we outline the various installation methods. For information on making diskettes, see Section 1.4.2.

1.2. Is Your Hardware Compatible?

Hardware compatibility is particularly important if you have an older system or a system that you built yourself. Red Hat Linux 7.3 should be compatible with most hardware in systems that were factory built within the last two years. However, hardware specifications change almost daily, so it is hard to guarantee that your hardware will be 100% compatible.

The most recent list of supported hardware can be found at:

```
http://hardware.redhat.com/hcl/
```

1.3. Do You Have Enough Disk Space?

Nearly every modern-day operating system (OS) uses *disk partitions*, and Red Hat Linux is no exception. When you install Red Hat Linux, you may have to work with disk partitions. If you have not worked with disk partitions before (or need a quick review of the basic concepts) read Appendix E before proceeding.

If Red Hat Linux will share your system with another OS, you will need to make sure you have enough available disk space on your hard drive(s) for this installation. Refer to Appendix G for instructions on creating a dual boot system.

The disk space used by Red Hat Linux must be separate from the disk space used by other OSes you may have installed on your system, such as Windows, OS/2, or even a different version of Linux. At least two partitions (/ and swap) must be dedicated to Red Hat Linux.

Before you start the installation process, one of the following conditions must be met:

- Your computer must have enough *unpartitioned*[1] disk space for the installation of Red Hat Linux.
- You must have one or more partitions that may be deleted, thereby freeing up enough disk space to install Red Hat Linux.

1.3.1. Installation Disk Space Requirements

Note

These recommendations are based on an installation that only installs one language (such as English). If you plan to install multiple languages to use on your system, you should increase the disk space requirements.

Workstation

A workstation installation, installing either GNOME[2] or KDE[3], requires at least 1.5 GB of free space. Choosing both GNOME and KDE requires at least 1.8 GB of free disk space.

Server

A server installation requires 1.3 GB for a minimal installation without X (the graphical environment), at least 1.4 GB of free space if all components (package groups) other than X are installed, and at least 2.1 GB to install all packages including GNOME and KDE.

Laptop

A laptop installation, when you choose to install GNOME or KDE, requires at least 1.5 GB of free space. If you choose both GNOME and KDE, you will need at least 1.8 GB of free disk space.

---

1.   Unpartitioned disk space means that available disk space on the hard drive(s) you are installing to has not been divided into sections for data. When you partition a disk, each partition will behave like a separate disk drive.

2.   Acronym for GNU Network Object Model Environment. GNOME is part of the GNU project and part of the open source movement. GNOME is a Windows-like desktop system and is not dependent on any one window manager. The main objective of GNOME is to provide a user-friendly suite of applications and an easy-to-use desktop.

3.   Acronym for K Desktop Environment. KDE is a network-transparent contemporary desktop environment for Linux and UNIX workstations, and it is part of the open source movement.

Custom

> A Custom installation requires 350 MB for a minimal installation and at least 3.7 GB of free space if every package is selected.

If you are not sure that you meet these conditions, or if you want to know how to create free disk space for your Red Hat Linux installation, please refer to Appendix E.

1.4. Can You Install Using the CD-ROM?

There are several methods that can be used to install Red Hat Linux.

Installing from a CD-ROM requires that you have purchased a Red Hat Linux 7.3 boxed set, or you have a Red Hat Linux CD-ROM, and you have a CD-ROM drive. Most new computers will allow booting from the CD-ROM. If your system will support booting from the CD-ROM, it is an easy way to begin a local CD-ROM installation.

Your BIOS may need to be changed to allow booting from your CD-ROM drive. For more information about editing your BIOS, see Section 3.3.1.

1.4.1. Alternative Boot Methods

If you cannot boot from the CD-ROM drive, the following alternative boot methods are also available:

Local Boot Disk

> If you need a *local boot disk*[4], you must create it. The local boot disk image file, `boot.img`, is located in the `images` directory on your Red Hat Linux CD-ROM. Refer to Section 1.4.2, for more information on making a boot disk.

PCMCIA Boot Disks

> You may need *PCMCIA boot disks* if you are using a PCMCIA device to install Red Hat Linux. If you need PCMCIA boot disks, you must create them. Refer to Section 1.4.2 for those instructions.

> The following checklist can help you determine if you will need to create PCMCIA boot disks:

> - You will install Red Hat Linux from a CD-ROM, and your CD-ROM drive is attached to your computer through a PCMCIA card.

> - You will use a PCMCIA network adapter during the installation.

> The PCMCIA boot disks image files, `pcmcia.img` and `pcmciadd.img`, are located in the `images` directory on your Red Hat Linux/x86 CD-ROM. Refer to Section 1.4.2 for more information on making a boot disk.

---

4.   A boot disk can be a diskette you created to boot (or start) the installation program, or it can be a diskette you create during the installation process that can later be used to boot the operating system. Normally, your computer boots from a hard disk, but if the hard disk is damaged, you can boot the computer from a bootable diskette.

> **Note**
>
> **USB Diskette Drive** — You can also boot the Red Hat Linux installation program using a boot disk in a USB diskette drive (if your system supports booting from a USB Diskette Drive).

> **Note**
>
> Although it is not required to boot your installation, you may occasionally find that a driver disk is needed to continue with the installation. Appendix F explains why a driver disk may be necessary for your installation, and how to obtain one if needed.

1.4.2. Making Installation Diskettes

You may need to create a diskette from an *image file*; for example, you may need to use updated diskette images obtained from the Red Hat Linux errata page:

```
http://www.redhat.com/apps/support/errata/
```

An image file contains an exact copy (or image) of a diskette's contents. Since a diskette contains file system information in addition to the data contained in files, the contents of the image file are not usable until they have been written to a diskette.

To start, you will need a blank, formatted, high-density (1.44MB), 3.5-inch diskette. You will need access to a computer with a 3.5-inch diskette drive. The computer must be able to run either an MS-DOS program or the dd utility found on most Linux-like operating systems.

The images directory on your Red Hat Linux CD-ROM contains the boot images for Red Hat Linux/x86 . Once you have selected the proper image (such as boot.img for a CD-ROM-based installation or bootnet.img for a network installation), transfer the image file onto a diskette using one of the following methods.

1.4.2.1. Using the `rawrite` Utility

To make a diskette using MS-DOS, use the rawrite utility included on the Red Hat Linux CD-ROM in the dosutils directory. First, label a blank, formatted 3.5-inch diskette appropriately (such as "Boot Disk" or "Updates Disk"). Insert it into the diskette drive. Then, use the following commands (assuming your CD-ROM 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-- : [Enter]
D:\dosutils>
```

First, rawrite 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 have selected. After pressing [Enter] to confirm,

`rawrite` copies the image file onto the diskette. If you need to make another diskette, label that diskette, and run `rawrite` again, specifying the appropriate image file.

Note

The `rawrite` utility only accepts 8.3-type file names, such as `filename.img`. If you download an update image from http://www.redhat.com named something similar to `update-anaconda-03292002.img`, you must rename it as `updates.img` before you run `rawrite`.

1.4.2.2. Using the `dd` Command

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 (such as "Boot Disk" or "Updates Disk"). Insert it into the diskette drive (but do not issue a `mount`[5] command). After mounting the Red Hat Linux CD-ROM, change 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
```

To make another diskette, label that diskette, and run `dd` again, specifying the appropriate image file.

1.5. Which Installation Class is Best For You?

Usually, Red Hat Linux is installed on its own disk partition or set of partitions, or over another installation of Linux.

Warning

Installing Red Hat Linux over another installation of Linux (including Red Hat Linux) does *not* preserve any information (files or data) from a prior installation. Make sure you save any important files! To preserve the current data on your existing system, you should back up your data and/or consider performing an upgrade instead.

Red Hat Linux provides five different classes, or types, of installations:

Workstation

A workstation installation is most appropriate if you are new to the world of Linux, and would like to give it a try. A workstation installation will create a system for your home or desktop use. A graphical, Windows-like environment will be installed.

---

5. When you mount a floppy or CD-ROM, you make that device's contents available to you. See the *Official Red Hat Linux Getting Started Guide* for more information.

Server

A server installation is most appropriate if you would like your system to function as a Linux-based server, and you do not want to heavily customize your system configuration.

Laptop

A laptop installation has been designed to make installing Red Hat Linux on laptops even easier. Much like a workstation installation, it will make sure you have the appropriate packages needed, as well as offer you an automated installation environment.

Custom

A custom installation allows you the greatest flexibility during your installation. You choose your boot loader, which packages you want, and more. Custom installations are most appropriate for those users more familiar with Red Hat Linux installations and for those afraid of losing complete flexibility.

Upgrade

If you already have a version of Red Hat Linux (4.2 or greater) running on your system and you want to quickly update to the latest packages and kernel version, then an upgrade is most appropriate for you.

These classes give you the option of simplifying the installation process (with some potential for loss of configuration flexibility), or retaining flexibility with a slightly more complex installation process. Next, take a detailed look at each class, so you can see which one is right for you.

1.5.1. Workstation Installations

Most suitable for new users, the workstation installation will install your choice of the GNOME or KDE desktop environments, or both, and the X Window System (the graphical software on which the desktop environments are based).

Below are the minimum recommended disk space requirements for a workstation installation where only one language (such as English) will be installed.

- Workstation choosing GNOME or KDE : 1.5 GB
- Workstation choosing both GNOME and KDE, and games: 1.8 GB

If you plan to choose all package groups (for example, GNOME is a group of packages), as well as select additional individual packages, you may want to allow yourself 2.1 GB or more of disk space.

Note

Unlike workstation installations previous to Red Hat Linux 7.0, performing a Red Hat Linux 7.3 workstation installation will not install the network daemon **xinetd** (inet services). When **xinetd** is not installed, you will have a more secure installation. However, in-bound network-related services such as **finger**, **telnet**, **talk**, and **FTP** will not work.[6] If you require these types of services, please go back and choose a server or a custom installation.

---

6. For example, you can `telnet` out to other systems, but other systems cannot `telnet` in to your system.

1.5.1.1. What a Workstation Installation Will Do

If you choose automatic partitioning, a workstation installation will create the following partitions:

- The size of the swap partition is determined by the amount of RAM in your system and the amount of space available on your hard drive. For example, if you have 128 MB of RAM then the swap partition created can be 128 MB – 256 MB (twice your RAM), depending on how much disk space is available.
- A 50 MB partition mounted as `/boot` in which the Linux kernel and related files reside.
- A root partition mounted as `/` in which all other files are stored (the exact size of this partition is dependent on your available disk space).

1.5.2. Server Installations

A server installation is most appropriate for you if you would like your system to function as a Linux-based server, and you do not want to heavily customize your system configuration.

Below are the minimum recommended disk space requirements for a server installation where only one language (such as English) will be installed.

- Server (minimum, no graphical interface): 1.3 GB
- Server (choosing everything, no graphical interface): 1.4 GB
- Server (choosing everything, GNOME and KDE): 2.1 GB

If you plan to choose all group packages, as well as select additional individual packages, you may want to allow yourself 2.3 GB or more of disk space.

During the server installation, *the X Window System is not configured and no GUI will be loaded* when the system boots, unless you choose to install the appropriate packages during package selection.

1.5.2.1. What a Server Installation Will Do

- The size of the swap partition is determined by the amount of RAM in your system and the amount of space available on your hard drive. For example, if you have 128 MB of RAM then the swap partition created can be 128 MB - 256 MB (twice your RAM), depending on how much disk space is available.
- A 384 MB root partition mounted as `/`.
- A partition mounted as `/usr` (the exact size of this partition is dependent on your available disk space).
- A partition mounted as `/home` (the exact size of this partition is dependent on your available disk space).
- A 256 MB partition mounted as `/var`.
- A 50 MB partition mounted as `/boot` in which the Linux kernel and related files are kept.

This disk partitioning scheme results in a reasonably flexible file system configuration for most server tasks.

1.5.3. Laptop Installations

The laptop installation will install your choice of the GNOME or KDE desktop environments (or both) and the X Window System.

Below are the minimum recommended disk space requirements for a laptop installation where only one language (such as English) will be installed.

- Laptop choosing GNOME or KDE: 1.5 GB
- Laptop choosing both GNOME and KDE: 1.8 GB

If you plan to choose all package groups (for example, GNOME is a group of packages), as well as select additional individual packages, you may want to allow yourself 1.7 GB or more of disk space. If you provide this extra space, you will have room for additional data, if needed.

1.5.3.1. What a Laptop Installation Will Do

If you choose automatic partitioning, a laptop installation will create the following partitions:

- The size of the swap partition is determined by the amount of RAM in your system and the amount of space available on your hard drive. For example, if you have 128 MB of RAM then the swap partition created can be 128 MB - 256 MB (twice your RAM), depending on how much disk space is available.
- A 50 MB partition mounted as /boot in which the Linux kernel and related files reside.
- A root partition mounted as / in which all other files are stored (the exact size of this partition is dependent on your available disk space).

Note

Performing a Red Hat Linux 7.3 laptop installation will not install the network daemon **xinetd** (inet services). Withholding **xinetd** results in a more secure installation[7]; network-related services such as **finger**, **telnet**, **talk**, and **FTP** will not work. If you require these types of services, choose a server or a custom installation.

1.5.4. Custom Installations

The *custom installation* allows you the most flexibility during your installation. The workstation and server installations automatically go through the installation process for you and omit certain steps. During a custom installation, you have complete control over the packages that will be installed on your system.

The recommended disk space requirements for a custom installation are as follows:

- Custom (minimum): 350 MB
- Custom (choosing everything): 3.7 GB

7.   For example, you can `telnet` out to other systems, but other systems cannot `telnet` in to your system.

1.5.4.1. What a Custom Installation Will Do

As you might guess from the name, a custom installation puts the emphasis on flexibility. You have complete control over which packages will be installed on your system.

If you choose automatic partitioning, a custom installation will create the following partitions:

- The size of the swap partition is determined by the amount of RAM in your system and the amount of space available on your hard drive. For example, if you have 128 MB of RAM then the swap partition created can be 128 MB - 256 MB (twice your RAM), depending on how much disk space is available.
- A 50 MB partition mounted as `/boot` in which the Linux kernel and related files reside.
- A root partition mounted as `/` in which all other files are stored (the exact size of this partition is dependent on your available disk space).

1.5.5. Upgrading Your System

Upgrading Red Hat Linux 4.2 (or greater) will not delete any existing data. The installation program updates the modular kernel and all currently installed software packages. See Chapter 3 and Appendix A for those instructions.

redhat.

Chapter 2.
Hardware Information and System Requirements Tables

This chapter provides instructions for learning about your hardware and a system requirements table which will help you keep a record of your current system settings and requirements.

2.1. Learning About Your Hardware with Windows

If your computer is already running Windows 9*x*, you can use the following steps to get additional configuration information:

Figure 2-1. Windows 9*x* System Properties

- In Windows, click on the **My Computer** icon using the secondary (normally the right) mouse button. A pop-up menu should appear.
- Select **Properties**. The **System Properties** window should appear. Note the information listed under **Computer** — in particular the amount of RAM listed.
- Click on the **Device Manager** tab. You will then see a graphical representation of your computer's hardware configuration. Make sure the **View devices by type** radio button is selected.

At this point, you can either double-click on the icons or single-click on the plus sign **+** to look at each entry in more detail. Look under the following icons for more information:

Figure 2-2. Windows 9x System Properties

- **Disk drives** — the type (IDE or SCSI) of hard drive will be found here. (IDE drives will normally include the word "IDE," while SCSI drives will not.)
- **Hard disk controllers** — information about your hard drive controller.
- **CDROM** — information about any CD-ROM drives connected to your computer.

Note

> In some cases, there may be no CD-ROM icon, yet your computer has a functioning CD-ROM drive. This is normal, depending on how Windows was originally installed. In this case, you may be able to learn additional information by looking at the CD-ROM driver loaded in your computer's `config.sys` file.

- **Mouse** — the type of mouse present on your computer.
- **Display adapters** — if you are interested in running the X Window System, you should write down the information you find here.
- **Sound, video and game controllers** — if your computer has sound capabilities, you will find more information about them here.
- **Network adapters** — information on your computer's network card (if you have one).
- **SCSI controllers** — if your computer uses SCSI peripherals, you will find additional information on the SCSI controller here.

This method is not a complete substitute for opening your computer's case and physically examining each component. However, in many cases it can provide sufficient information to continue with the installation.

Note

> This information can also be printed by clicking on the **Print** button. A second window will appear, allowing you to choose the printer, as well as the type of report. The **All Devices and System Summary** report type is the most complete.

If your computer is already running Windows 2000, you can use the following steps to get additional configuration information:
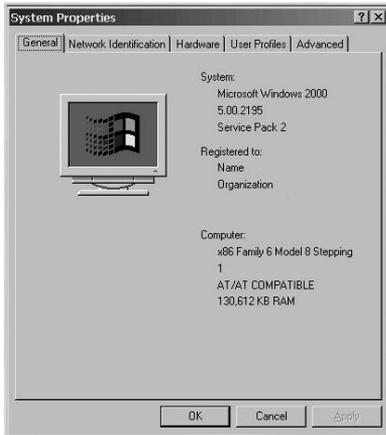
Figure 2-3. Windows 2000 System Properties

- In Windows, click on the **My Computer** icon using the secondary (normally the right) mouse button. A pop-up menu should appear.
- Select **Properties**. The **System Properties** window should appear. Note the information listed under **Computer** — in particular the amount of RAM listed.
- Click on the **Hardware** tab. You will then see your computer's hardware configuration options.
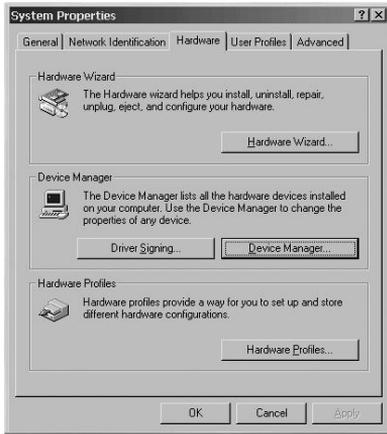
Figure 2-4. Windows 2000 System Properties — Hardware

- Click on the **Device Manager** tab. You will then see a graphical representation of your computer's hardware configuration. Make sure the **View devices by type** radio button is selected.

At this point, you can either double-click on the icons or single-click on the plus sign **+** to look at each entry in more detail. Look under the following icons for more information:
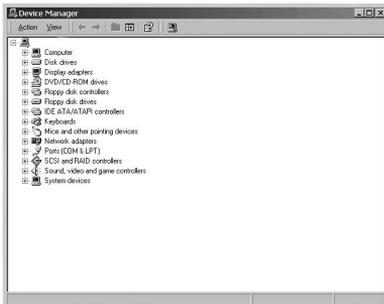
Figure 2-5. Windows 2000 System Properties

- **Disk drives** — the type (IDE or SCSI) of hard drive will be found here. (IDE drives will normally include the word "IDE," while SCSI drives will not.)
- **Hard disk controllers** — information about your hard drive controller.
- **CDROM** — information about any CD-ROM drives connected to your computer.
- **Mouse** — the type of mouse present on your computer.
- **Display adapters** — if you are interested in running the X Window System, you should write down the information you find here.

- **Sound, video and game controllers** — if your computer has sound capabilities, you will find more information about them here.
- **Network adapters** — information on your computer's network card (if you have one).
- **SCSI controllers** — if your computer uses SCSI peripherals, you will find additional information on the SCSI controller here.

This method is not a complete substitute for opening your computer's case and physically examining each component. However, in many cases it can provide sufficient information to continue with the installation.

2.2. Recording Your System's Hardware

Enter information about your system in the table provided as a handy reference to help make your Red Hat Linux installation go more smoothly.

Table 2-1. System Requirements Table

hard drive(s): type, label, size; ex: IDE hda=1.2 GB	
partitions: map of partitions and mount points; ex: `/dev/hda1=/home`, `/dev/hda2=/` (fill this in once you know where they will reside)	
memory: amount of RAM installed on your system; ex: 64 MB, 128 MB	
CD-ROM: interface type; ex: SCSI, IDE (ATAPI)	
SCSI adapter: if present, make and model number; ex: BusLogic SCSI Adapter, Adaptec 2940UW	
network card: if present, make and model number; ex: Tulip, 3COM 3C590	
mouse: type, protocol, and number of buttons; ex: generic 3 button PS/2 mouse, MouseMan 2 button serial mouse	

monitor: make, model, and manufacturer specifications; ex: Optiquest Q53, ViewSonic G773	
video card: make, model number and size of VRAM; ex: Creative Labs Graphics Blaster 3D, 8MB	
sound card: make, chipset and model number; ex: S3 SonicVibes, Sound Blaster 32/64 AWE	
IP, DHCP, and BOOTP addresses: four numbers, separated by dots; ex: 10.0.2.15	
netmask: four numbers, separated by dots; ex: 255.255.248.0	
gateway IP address: four numbers, separated by dots; ex: 10.0.2.245	
one or more name server IP addresses (DNS): one or more sets of dot-separated numbers; ex: 10.0.2.1	
domain name: the name given to your organization; ex: Red Hat's would be `redhat.com`	
hostname: the name of your computer; your personal choice of names; ex: `cookie`, `southpark`	

If any of these networking requirements or terms are unfamiliar to you, contact your network administrator for assistance.

redhat.

Chapter 3.
Installing Red Hat Linux

This chapter explains how to install Red Hat Linux from the CD-ROM. The following topics are discussed:

- Getting familiar with the installation program's user interface
- Starting the installation program
- Selecting an installation method
- Configuration steps during the installation (language, keyboard, mouse, etc.)
- Finishing the installation

Note

If you already have another operating system installed and want to create a dual boot system so that you can use both Red Hat Linux and the other operating system, please read Appendix G for more information.

3.1. The Graphical Installation Program User Interface

If you have used a *graphical user interface (GUI)* before, you will be familiar with this process; simply use your mouse to navigate the screens, "click" buttons, or enter text fields. You can also navigate through the installation using the [Tab] and [Enter] keys.

Note

If you do not wish to use the GUI installation program, the text mode installation program is also available. To start the text mode installation program, use the following boot command:

```
boot: text
```

Please refer to Section 3.2 for a brief overview of text mode installation instructions.

3.1.1. A Note about Virtual Consoles

The Red Hat Linux installation program offers more than the dialog boxes of the installation process. Several different kinds of diagnostic messages are available to you, in addition to providing a way to enter commands from a shell prompt. The installation program displays these messages on five *virtual consoles*, among which you can switch using a single keystroke combination.

These virtual consoles can be helpful if you encounter a problem while installing Red Hat Linux. Messages displayed on the installation or system consoles can help pinpoint a problem. Please see Table 3-1 for a listing of the virtual consoles, keystrokes used to switch to them, and their contents.

Table 3-1. Console, Keystrokes, and Contents

Console	Keystrokes	Contents
1	[Ctrl]-[Alt]-[F1]	installation dialog
2	[Ctrl]-[Alt]-[F2]	shell prompt
3	[Ctrl]-[Alt]-[F3]	install log (messages from installation program)
4	[Ctrl]-[Alt]-[F4]	system-related messages
5	[Ctrl]-[Alt]-[F5]	other messages
7	[Ctrl]-[Alt]-[F7]	X graphical display

Generally, there is no reason to leave the default console (virtual console #7) unless you are attempting to diagnose installation problems.

3.2. The Text Mode Installation Program User Interface

The Red Hat Linux text mode installation program uses a screen-based interface that includes most of the on-screen "widgets" commonly found on graphical user interfaces. Figure 3-1, and Figure 3-2, illustrate the screens you will see.
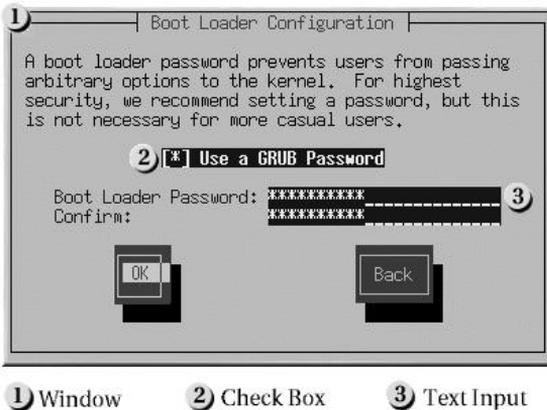
Figure 3-1. Installation Program Widgets as seen in Boot Loader Configuration
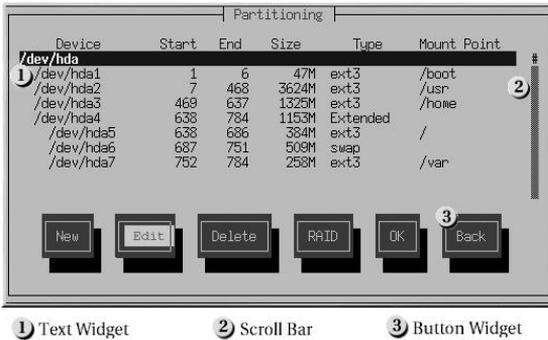
Figure 3-2. Installation Program Widgets as seen in Disk Druid

Here is a list of the most important widgets shown in Figure 3-1, and Figure 3-2:

- Window — windows (usually referred to as *dialogs* in this manual) will appear on your screen throughout the installation process. At times, one window may overlay another; in these cases, you can only interact with the window on top. When you are finished in that window, it will disappear, allowing you to continue working in the window underneath.

- Text Input — text input lines are regions where you can enter information required by the installation program. When the cursor rests on a text input line, you may enter and/or edit information on that line.

- Checkbox — checkboxes allow you to select or deselect a feature. The box displays either an asterisk (selected) or a space (unselected). When the cursor is within a checkbox, press [Space] to select an unselected feature or to deselect a selected feature.

- Text widget — text widgets are regions of the screen for the display of text. At times, text widgets may also contain other widgets, such as checkboxes. If a text widget contains more information than can be displayed in the space reserved for it, a scroll bar appears; if you position the cursor within the text widget, you can then use the [Up] and [Down] arrow keys to scroll through all the information available. Your current position is shown on the scroll bar by a **#** character, which moves up and down the scroll bar as you scroll.

- Button widget — button widgets are the primary method of interacting with the installation program. You progress through the windows of the installation program by navigating these buttons, using the [Tab] and [Enter] keys. Buttons can be selected when they are highlighted.

- Cursor — although not a widget, the cursor is used to select (and interact) with a particular widget. As the cursor is moved from widget to widget, it may cause the widget to change color, or you may only see the cursor itself positioned in or next to the widget. In Figure 3-1, the cursor is positioned on the **OK** button. Figure 3-2, shows the cursor on the **Edit** button.

3.2.1. Using the Keyboard to Navigate

Navigation through the installation dialogs is performed through a simple set of keystrokes. To move the cursor, use [Left], [Right], [Up], and [Down] arrow keys. Use [Tab], and [Alt]-[Tab] to cycle forward or backward through each widget on the screen. Along the bottom, most screens display a summary of available cursor positioning keys.

To "press" a button, position the cursor over the button (using [Tab], for example) and press [Space] or [Enter]. To select an item from a list of items, move the cursor to the item you wish to select and press [Enter]. To select an item with a checkbox, move the cursor to the checkbox 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 equivalent to pressing the **OK** button.

Caution

> Unless a dialog box is waiting for your input, do not press any keys during the installation process (doing so may result in unpredictable behavior).

3.2.2. Displaying Online Help

Once the installation program is loaded into memory, you can obtain information about the installation process and options by pressing [F1] through [F6]. For example, press [F2] to see general information about the online help screens.

3.3. Starting the Installation Program

To start the installation, you must first boot the installation program. Please make sure you have all the resources you will need for the installation. If you have already read through Chapter 1 , and followed the instructions, you should be ready to begin.

Note

> Occasionally, some hardware components require a *driver disk* during the installation. A driver disk adds support for hardware that is not otherwise supported by the installation program. Refer to Appendix F for more information.

3.3.1. Booting the Installation Program

You can boot the Red Hat Linux installation program using any one of the following media (depending upon what your system can support):

- *Bootable CD-ROM* — Your machine supports a bootable CD-ROM drive and you want to perform a local CD-ROM installation.
- *Local boot disk* — Your machine will not support a bootable CD-ROM and you want to install from a local CD-ROM or a hard drive.
- *Network boot disk* — Use a network boot disk to install via NFS, FTP, and HTTP.
- *PCMCIA boot disks* — Use PCMCIA boot disks when you need PCMCIA support, but your machine does not support booting from the CD-ROM drive *or* if you need PCMCIA support in order to make use of the CD-ROM drive on your system. The PCMCIA boot disks can be used for all installation methods (CD-ROM, hard drive, NFS, FTP, and HTTP).

To create a boot disk, refer to Section 1.4.2.

Insert the boot disk into your computer's first diskette drive and reboot (or boot using the CD-ROM, if your computer supports booting from it). Your BIOS settings may need to be changed to allow you to boot from the diskette or CD-ROM.

Tip

> To change your BIOS settings, watch the instructions provided on your display when your computer first boots. You will see a line of text telling you to press the [Del] or [F1] key to enter the BIOS settings.
>
> Once you have entered your BIOS setup program, find the section where you can alter your boot sequence. The default is often C, A or A, C (depending on whether you boot from your hard drive [C] or a diskette drive [A]). Change this sequence so that the CD-ROM is first in your boot order and that C or A (whichever is your typical boot default) is second. This instructs the computer to first look at the CD-ROM drive for bootable media; if it does not find bootable media on the CD-ROM drive, it will then check your hard drive or diskette drive.
>
> Save your changes before exiting the BIOS. For more information, refer to the documentation that came with your system.

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 help screen, press the appropriate function key as listed in the line at the bottom of the screen.

As you boot the installation program, be aware of two issues:

- Once you see the `boot:` prompt, the installation program will automatically begin 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 while the help screen is read from the boot media.

Normally, you only need to press [Enter] to boot. Watch the boot messages to see if the Linux kernel detects your hardware. If your hardware is properly detected, please continue to the next section. If it does not properly detect your hardware, you may need to restart the installation in expert mode.

3.3.1.1. Additional Boot Options

While it is easiest for a user to boot from CD-ROM and perform a graphical installation, sometimes there are installation scenarios where booting in a different manner may be needed. This section discusses addition boot options available for Red Hat Linux.

- If you are having trouble booting into the graphical installation program, you can try to boot using the no framebuffer (`nofb`) boot option.

  At the boot command, enter the following:

  ```
  boot: nofb
  ```

  This option allows you to use the graphical installation program without using a framebuffer.

- If you do not wish to perform a graphical installation, you can start a text mode installation using the following boot command:

  ```
  boot: text
  ```

- ISO images now have an md5sum embedded in them. To test the checksum integrity of an ISO image, at the installation boot prompt, type:

```
boot: linux mediacheck
```

  The installation program will prompt you to insert a CD or select an ISO image to test, and select **OK** to perform the checksum operation. This checksum operation can be performed on any Red Hat Linux CD and does not have to be performed in a specific order (for example, CD #1 does not have the be the first CD you verify). It is strongly recommended to perform this operation on any Red Hat Linux CD that was created from downloaded ISO images. This procedure works with CD-based installations and hard drive and NFS installations using ISO images.

- If the installation program does not properly detect your hardware, you may need to restart the installation in *expert* mode. Enter expert mode using the following boot command:

```
boot: linux noprobe
```

  For text mode installations, use:

```
boot: text noprobe
```

  Expert mode disables most hardware probing, and gives you the option of entering options for the drivers loaded during the installation.

> **Note**
>
> The initial boot messages will not contain any references to SCSI or network cards. This is normal; these devices are supported by modules that are loaded during the installation process.

- If you need to perform the installation in *serial mode*, type the following command:

```
boot: linux console=<device>
```

  For text mode installations, use:

```
boot: linux text console=<device>
```

  In the above command, $<device>$ should be the device you are using (such as ttyS0 or ttyS1). For example, linux console=ttyS0,115200n8.

3.3.1.2. Kernel Options

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

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

For text mode installations, use:

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

After entering any options, press [Enter] to boot using those options.

If you need to specify boot options to identify your hardware, please write them down. The boot options will be needed during the boot loader configuration portion of the installation (please see Section 3.20 for more information).

3.3.1.3. Booting Without Diskettes

The Red Hat Linux CD-ROM can be booted by computers that support bootable CD-ROMs. Not all computers support this feature, so if your system cannot boot from the CD-ROM, there is one other way to start the installation without using a boot disk. The following method is specific to x86-based computers only.

If you have MS-DOS installed on your system, you can boot directly from the CD-ROM drive without using a boot disk. To do this (assuming your CD-ROM is drive `d:`), use the following commands:

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

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.

If your computer cannot boot directly from CD-ROM (and you cannot use a DOS-based autoboot), you will have to use a boot diskette to get things started.

3.4. Selecting an Installation Method

What type of installation method do you wish to use? The following installation methods are available:

CD-ROM

If you have a CD-ROM drive and the Red Hat Linux CD-ROM, you can use this method. You will need a boot disk or a bootable CD-ROM. PCMCIA boot and driver disks may also be used. Refer to Section 3.5, for CD-ROM installation instructions.

Hard Drive

If you have copied the Red Hat Linux ISO images to a local hard drive, you can use this method. You will need a boot disk. PCMCIA boot and driver disks may also be used. Refer to Section 3.6, for hard drive installation instructions.

NFS Image

If you are installing from an NFS server using ISO images or a mirror image of Red Hat Linux, you can use this method. You will need a network boot disk. PCMCIA boot and driver disks may also be used. Refer to Section 3.8, for network installation instructions. Please note that NFS installations may also be performed in GUI mode.

FTP

   If you are installing directly from an FTP server, use this method. You will need a net-
   work boot disk. PCMCIA boot and driver disks may also be used. Refer to Section 3.9,
   for FTP installation instructions.

HTTP

   If you are installing directly from an HTTP (Web) server, use this method. You will need
   a network boot disk. PCMCIA boot and driver disks may also be used. Refer to Section
   3.10, for HTTP installation instructions.

3.5. Installing from CD-ROM

Note

If you already have another operating system installed and want to create a dual boot system so that
you can use both Red Hat Linux and the other operating system, read Appendix G for details.

To install Red Hat Linux from a CD-ROM, choose the **CD-ROM** option from the boot loader
screen and select **OK**. When prompted, insert the Red Hat Linux CD into your CD-ROM
drive (if you did not boot from the CD-ROM). Once the CD is in the CD-ROM drive, select
OK, and press [Enter].

The installation program will then probe your system and attempt to identify your CD-ROM
drive. It will start by looking for an IDE (also known as an ATAPI) CD-ROM drive. If found,
you will continue to the next stage of the installation process (see Section 3.12).

Note

To abort the installation process at this time, reboot your machine and then eject the boot diskette or
CD-ROM. You can safely cancel the installation at any point before the **About to Install** screen. See
Section 3.30 for more information.

If a CD-ROM drive is not detected, you will be asked what type of CD-ROM drive you have.
Choose from the following types:

SCSI

   Select this if your CD-ROM drive is attached to a supported SCSI adapter; the instal-
   lation program will then ask 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 drive is neither an IDE nor a SCSI, it is an *other*. Sound cards with proprietary CD-ROM interfaces are good examples of this type of CD-ROM. The installation program will display a list of drivers for supported CD-ROM drives — choose a driver and, if necessary, specify any driver options. It is recommended that you use the `oldcdrom.img` image file

⬦**Tip**

A partial list of optional parameters for CD-ROM drives can be found in the *Official Red Hat Linux Reference Guide*, in the *General Parameters and Modules* appendix.

3.5.1. What If the IDE CD-ROM Was Not Found?

If you have an IDE (ATAPI) CD-ROM, but the installation program fails to find your IDE (ATAPI) CD-ROM and asks you what type of CD-ROM drive you have, try the following boot command. Restart the installation, and at the `boot:` prompt enter **linux hdX=cdrom**. Replace the **X** with one of the following letters, depending on the interface the unit is connected to, and whether it is configured as master or slave (also known as primary and secondary):

- `a` — first IDE controller, master
- `b` — first IDE controller, slave
- `c` — second IDE controller, master
- `d` — second IDE controller, slave

If you have a third and/or fourth controller, continue assigning letters in alphabetical order, going from controller to controller, and master to slave.

3.6. Installing from a Hard Drive

✎**Note**

Hard drive installations only work from ext2, ext3, or FAT file systems. If you have a file system other than those listed here, such as reiserfs, you will not be able to perform a hard drive installation.

Hard drive installations require the use of the ISO (or CD-ROM) images rather than copying an entire installation tree. After placing the required ISO images (the binary Red Hat Linux CD-ROMs) in a directory, choose to install from the hard drive. You will then point the installation program at that directory to perform the installation.

Verifying that the ISO images are intact before you attempt an installation will help to avoid problems that are often encountered during a hard drive installation. To verify the ISO images are intact prior to performing an installation, use an **md5sum** program (many **md5sum**

programs are available for various operating systems). An **md5sum** program should be available on the same server as the ISO images.

![Note icon] **Note**

> ISO images now have an md5sum embedded in them. To test the checksum integrity of an ISO image, at the installation boot prompt, type:
>
> ```
> boot: linux mediacheck
> ```

The **Select Partition** screen (Figure 3-3) applies only if you are installing from a disk partition (that is, if you selected **Hard Drive** in the **Installation Method** dialog). This dialog allows you to name the disk partition and directory from which you are installing Red Hat Linux.

Enter the device name of the partition containing the Red Hat ISO images. There is also a field labeled **Directory holding images**. If the ISO images are not in the root directory of that partition, enter the path to the ISO images (for example, if the ISO images are in /test/new/RedHat, you would enter /test/new).

After you have identified the disk partition, you will next see the **Welcome** dialog.

Figure 3-3. Selecting Partition Dialog for Hard Drive Installation

3.7. Preparing for a Network Installation

If you are performing a network installation and have booted from a boot disk with the bootnet.img image, the **Configure TCP/IP** dialog appears; for an explanation of this dialog, go to Section 3.22, and then return here.

3.7.1. Setting Up the Server

Because the Red Hat Linux 7.3 installation program is capable of installing Red Hat Linux from multiple CD-ROMs, if you intend to support NFS, FTP, or HTTP installations you must copy the RedHat directory from each CD-ROM comprising Red Hat Linux 7.3 onto a disk drive:

- Insert CD-ROM 1 and execute the following commands:

  - `mount /mnt/cdrom`

  - `cp -var /mnt/cdrom/RedHat /location/of/disk/space`

    Where */location/of/disk/space* is a directory you create such as */export/7.3/*.

  - `umount /mnt/cdrom`

- Insert CD-ROM 2 and execute the following commands:

  - `mount /mnt/cdrom`

  - `cp -var /mnt/cdrom/RedHat /location/of/disk/space`

    Where */location/of/disk/space* is a directory you create such as */export/7.3/*.

  - `umount /mnt/cdrom`

- Next, make */location/of/disk/space* accessible to the installation program (for example, exporting it for NFS installations).

  `Export /location/of/disk/space`

If you are not sure how to do this, refer to the *Official Red Hat Linux Customization Guide* and the *Official Red Hat Linux Reference Guide* for more information.

3.7.1.1. Using ISO Images for NFS Installs

NFS installations can use ISO (or CD-ROM) images rather than copying an entire installation tree. After placing the required ISO images (the binary Red Hat Linux CD-ROMs) in a directory, choose to install via NFS. You will then point the installation program at that directory to perform the installation.

Verifying that the ISO images are intact before you attempt an installation will help to avoid problems that are often encountered during an NFS installation. To verify the ISO images are intact prior to performing an installation, use an **md5sum** program (many **md5sum** programs are available for various operating systems). An **md5sum** program should be available on the same server as the ISO images.

Note

ISO images now have an md5sum embedded in them. To test the checksum integrity of an ISO image, at the installation boot prompt, type:

`boot:` **`linux mediacheck`**

Additionally, if a file called `updates.img` exists in the directory from which you install, then it will be used for **Anaconda** updates. Refer to the file `install-methods.txt` in the `anaconda` RPM package for detailed information on the various ways to install Red Hat Linux, as well as how to apply **Anaconda** updates.

3.8. Installing via NFS

The NFS dialog (Figure 3-4) applies only if you are installing from an NFS server (if you booted from a network or PCMCIA boot disks and selected **NFS Image** in the **Installation Method** dialog).

Figure 3-4. NFS Setup Dialog

Enter the fully-qualified domain name or IP address of your NFS server. For example, if you are installing from a host named `eastcoast` in the domain `redhat.com`, enter `eastcoast.redhat.com` in the **NFS Server** field.

Next, enter the name of the exported directory. If you followed the setup described in Section 3.7, you would enter the directory `/location/of/disk/space/` which contains the `RedHat` directory.

If the NFS server is exporting a mirror of the Red Hat Linux installation tree, enter the directory which contains the `RedHat` directory. (If you do not know this directory path, ask your system administrator.) For example, if your NFS server contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386`.

If the NFS server is exporting the ISO images of the Red Hat Linux CD-ROMs, enter the directory which contains the ISO images.

Next you will see the **Welcome** dialog.

3.9. Installing via FTP

The FTP dialog (Figure 3-5) applies only if you are installing from an FTP server (if you selected **FTP** in the **Installation Method** dialog). This dialog allows you to identify the FTP server from which you are installing Red Hat Linux.

Enter the name or IP address of the FTP site you are installing from, and the name of the directory containing the `RedHat` installation files for your architecture. For example, if the FTP site contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386`.

Figure 3-5. FTP Setup Dialog

If everything has been specified properly, a message box appears indicating that `base/hdlist` is being retrieved.

Next you will see the **Welcome** dialog.

Tip

You can also install Red Hat Linux using ISO images without copying them into a single tree by loopback mounting them as:

```
disc1/,disc2/,disc3/
```

3.10. Installing via HTTP

The HTTP dialog (Figure 3-6) applies only if you are installing from an HTTP server (if you selected **HTTP** in the **Installation Method** dialog). This dialog prompts you for information about the HTTP server from which you are installing Red Hat Linux.

Enter the name or IP address of the HTTP site you are installing from, and the name of the directory there containing the `RedHat` installation files for your architecture. For example, if the HTTP site contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386`.

Figure 3-6. HTTP Setup Dialog

If everything has been specified properly, a message box appears indicating that `base/hdlist` is being retrieved.

Next you will see the **Welcome** dialog.

3.11. Welcome to Red Hat Linux

The **Welcome** screen does not prompt you for any input. Please read over the help text in the left panel for additional instructions and information on where to register your Official Red Hat Linux product.

Please notice the **Hide Help** button at the bottom left corner of the screen. The help screen is open by default. If you do not want to view the help information, click on **Hide Help** to minimize the help portion of the screen.

Click on the **Next** button to continue.

3.12. Language Selection

Using your mouse, select the language you would prefer to use for the installation and as the system default (see Figure 3-7).

Selecting the appropriate language will also help target your time zone configuration later in the installation. The installation program will try to define the appropriate time zone based on what you specify on this screen.

Once you select the appropriate language, click **Next** to continue.

Figure 3-7. Language Selection

3.13. Keyboard Configuration

Choose the keyboard model that best fits your system (see Figure 3-8). If you cannot find an exact match, choose the best **Generic** match for your keyboard type (for example, **Generic 101-key PC**).

Next, choose the correct layout type for your keyboard (for example, U.S. English).

Creating special characters with multiple keystrokes (such as Ñ, Ô, and Ç) is done using "dead keys" (also known as compose key sequences). Dead keys are enabled by default. If you do not wish to use them, select **Disable dead keys**.

Tip

> The following example will help you determine if you need dead keys enabled. An example of a dead key is the backspace (^H) key on a US English 101 Standard Keyboard. Dead keys are not exclusive to non-English keyboards.

To test your keyboard configuration, use the blank text field at the bottom of the screen to enter text.

Once you have made the appropriate selections, click **Next** to continue.

Tip

> To change your keyboard type after you have installed Red Hat Linux, log in as root and use the /usr/sbin/kbdconfig command. Alternatively, you can type setup at the root prompt.
>
> To become root, type su – at the shell prompt in a terminal window and then press [Enter]. Then, enter the root password.
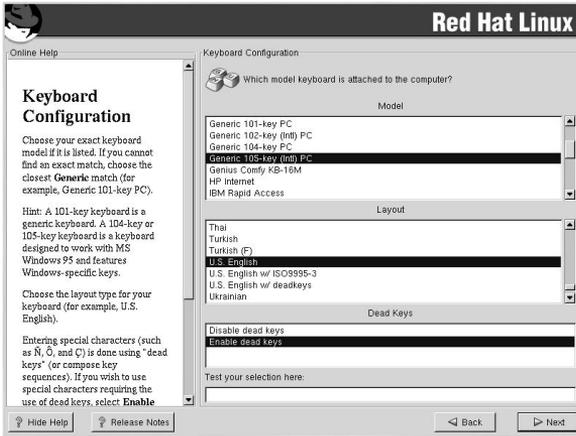
Figure 3-8. Keyboard Configuration

Tip

> To re-enable dead keys (assuming you chose to disable them during the installation), you must comment out the line disabling dead keys in the `XF86Config-4` file (or, it you are using XFree86 version 3, the `XF86Config` file) in `/etc/X11`.
>
> An **InputDevice** section, that would disable dead keys on a keyboard layout that used dead keys (for example, German), would look similar to the following:
>
> ```
> Section "InputDevice"
>         Identifier  "Keyboard0"
>         Driver      "keyboard"
>         Option      "XkbRules" "xfree86"
>         Option      "XkbModel" "pc101"
>         Option      "XkbLayout" "de"
> #       Option      "XkbVariant" "nodeadkeys"
> EndSection
> ```
>
> By default, the keyboard layout for various languages which use dead keys should have them enabled unless the `nodeadkeys` option is present.

3.14. Mouse Configuration

Choose the correct mouse type for your system. If you cannot find an exact match, choose a mouse type that you are sure is compatible with your system (see Figure 3-9).

To determine your mouse's interface, follow the mouse cable back to where it plugs into your system and use the following diagrams. If you are installing Red Hat Linux on a laptop computer, in most cases the pointing device will be PS/2 compatible.

If your mouse is a serial mouse, the port will look similar to .

If your mouse is a PS/2 mouse, the port will look similar to .

If your mouse is a USB mouse, the port will look similar to ⬭ .

If your mouse is a AT (Advanced Technology) mouse, the port will look similar to ⊙ .

If you cannot find a mouse that you are sure is compatible with your system, select one of the **Generic** entries, based on your mouse's number of buttons, and its interface.

Tip

If you have a scroll mouse, select the **MS Intellimouse** entry (with your proper mouse port) as the compatible mouse type.
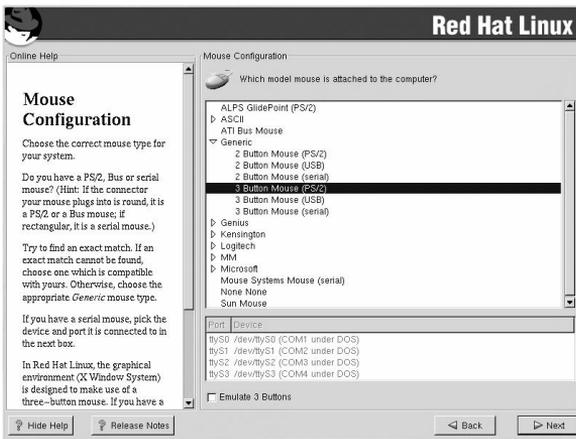
Figure 3-9. Mouse Configuration

If you have a PS/2 or a bus mouse, you do not need to pick a port and device. If you have a serial mouse, you should choose the correct port and device that your serial mouse is on.

The **Emulate 3 Buttons** checkbox allows you to use a two-button mouse as if it had three buttons. In general, the X Window System is easier to use with a three-button mouse. If you select this checkbox, you can emulate a third, "middle" button by pressing both mouse buttons simultaneously.

Tip

To change your mouse configuration after you have completed the installation of Red Hat Linux, become root; then use the /usr/sbin/mouseconfig command from a shell prompt.

To configure your mouse to work as a left-handed mouse, reset the order of the mouse buttons. To do this, after you have booted your Red Hat Linux system, type gpm -B 321 at the shell prompt. You can also change this setting in GNOME by going to the main menu and selecting **Programs** => **Settings** => **Peripherals** => **Mouse**. In KDE, select **Preferences** => **Peripherals** => **Mouse**.

3.15. Install Options

Choose the type of installation you would like to perform (see Figure 3-10). Red Hat Linux allows you to choose the installation type that best fits your needs. Your options are **Workstation**, **Server**, **Laptop**, **Custom**, and **Upgrade**.
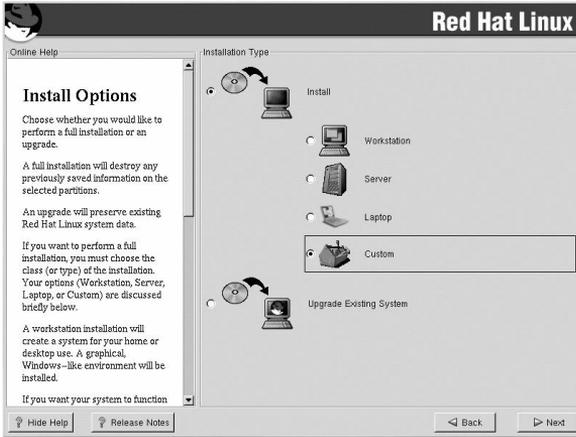
Figure 3-10. Choosing Install or Upgrade

To perform an upgrade, please refer to Appendix A.

For more information about the different installation classes, please refer to Section 1.5.

3.16. Disk Partitioning Setup

Partitioning allows you to divide your hard drive into isolated sections, where each section behaves as its own hard drive. Partitioning is particularly useful if you run more than one operating system. If you are not sure how you want your system to be partitioned, read Appendix E for more information.

On this screen, you can choose to perform automatic partitioning, or manual partitioning using **Disk Druid** or **fdisk** (see Figure 3-11).

Automatic partitioning allows you to perform an installation without having to partition your drive(s) yourself. If you do not feel comfortable with partitioning your system, it is recommended that you *do not* choose to partition manually and instead let the installation program partition for you.

To partition manually, choose either the **Disk Druid** or **fdisk** (recommended for experts only) partitioning tool.
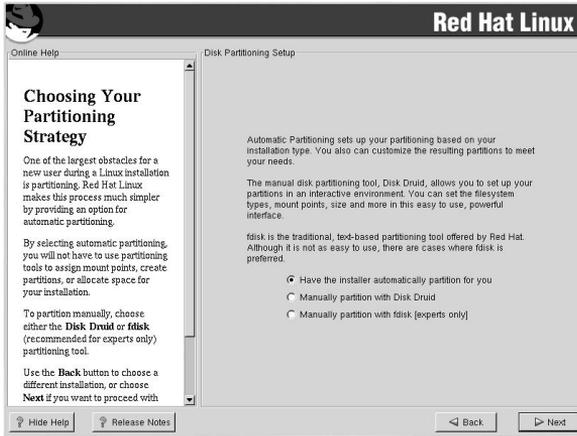
Figure 3-11. Disk Partitioning Setup

If you chose to manually partition using **Disk Druid**, refer to Section 3.18.

If you chose to manually partition using **fdisk**, refer to Section 3.19.

3.17. Automatic Partitioning

Automatic partitioning allows you to have some control concerning what data is removed (if any) from your system. Your options are:

- **Remove all Linux partitions on this system** — select this option to remove only Linux partitions (partitions created from a previous Linux installation). This will not remove other partitions you may have on your hard drive(s) (such as VFAT or FAT32 partitions).
- **Remove all partitions on this system** — select this option to remove all partitions on your hard drive(s) (this includes partitions created by other operating systems such as Windows 95/98/NT/2000).

🖐️**Caution**

> If you select this option, all data on the selected hard drive(s) will be removed by the installation program. Do not select this option if you have information that you want to keep on the hard drive(s) where you are installing Red Hat Linux.

- **Keep all partitions and use existing free space** — select this option to retain your current data and partitions, assuming you have enough free space available on your hard drive(s).
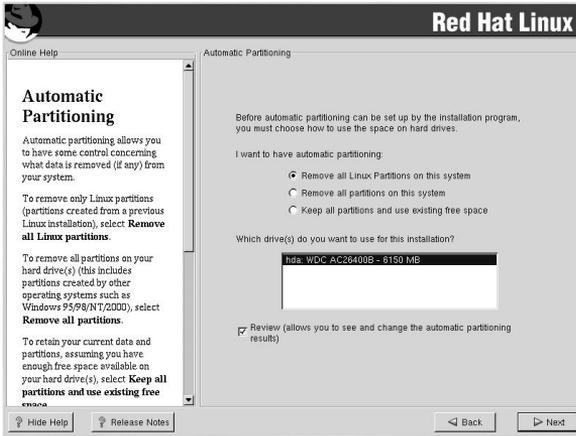
Figure 3-12. Automatic Partitioning

Using your mouse, choose the hard drive(s) on which you want Red Hat Linux to be installed. If you have two or more hard drives, you can choose which hard drive(s) should contain this installation. Unselected hard drives, and any data on them, will not be touched.

To review and make any necessary changes to the partitions created by automatic partitioning, select the **Review** option. After selecting **Review** and clicking **Next** to move forward, you will see the partitions created for you in **Disk Druid**. You will also be able to make modifications to these partitions if they do not meet your needs.

Click **Next** once you have made your selections to proceed.

3.18. Partitioning Your System

If you chose automatic partitioning and did not select **Review**, please skip ahead to Section 3.22.

If you chose automatic partitioning and selected **Review**, you can either accept the current partition settings (click **Next**), or modify the setup using **Disk Druid**, the manual partitioning tool.

If you chose **Manually partition with fdisk**, please skip ahead to Section 3.19.

At this point, you must tell the installation program where to 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 (refer to Figure 3-13).

Note

If you have not yet planned how you will set up your partitions, refer to Appendix E. At a bare minimum, you need an appropriately-sized root partition, and a swap partition equal to twice the amount of RAM you have on the system.
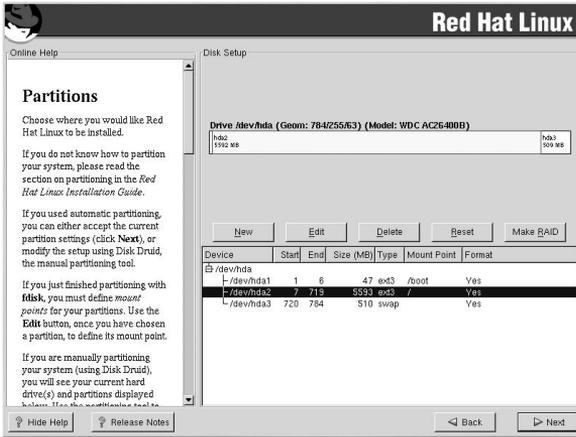
Figure 3-13. Partitioning with Disk Druid

The partitioning tool used in Red Hat Linux 7.3 is **Disk Druid**. With the exception of certain esoteric situations, **Disk Druid** can handle the partitioning requirements for a typical Red Hat Linux installation.

3.18.1. Graphical Display of Hard Drive(s)

Disk Druid offers a graphical representation of your hard drive(s).

Using your mouse, click once to highlight a particular field in the graphical display. Double-click to edit an existing partition or to create a partition out of existing free space.

Above the display, you will see the **drive** name (such as /dev/hda), the **geom** (which shows the hard disk's geometry and consists of three numbers representing the number of cylinders, heads, and sectors as reported by the hard disk), and the **model** of the hard drive as detected by the installation program.

3.18.2. Disk Druid's Buttons

These buttons control **Disk Druid's** actions. They are used to change the attributes of a partition (for example the file system type and mount point) and also to create RAID devices. Buttons on this screen are also used to accept the changes you have made, or to exit **Disk Druid**. For further explanation, take a look at each button in order:

- **New**: Used to request a new partition. When selected, a dialog box appears containing fields (such as mount point and size) that must be filled in.
- **Edit**: Used to modify attributes of the partition currently selected in the **Partitions** section. Selecting **Edit** opens a dialog box. Some or all of the fields can be edited, depending on whether the partition information has already been written to disk.

  You can also edit free space as represented in the graphical display to create a new partition within that space. Either highlight the free space and then select the **Edit** button, or double-click on the free space to edit it.

- **Delete**: Used to remove the partition currently highlighted in the **Current Disk Partitions** section. You will be asked to confirm the deletion of any partition.

- **Reset**: Used to restore **Disk Druid** to its original state. All changes made will be lost if you **Reset** the partitions.

- **Make RAID**: **Make RAID** can be used if you want to provide redundancy to any or all disk partitions. *It should only be used if you have experience using RAID.* To read more about RAID, please refer to *RAID (Redundant Array of Independent Disks)* in the *Official Red Hat Linux Customization Guide*.

  To make a RAID device, you must first create software RAID partitions. Once you have created two or more software RAID partitions, select **Make RAID** to join the software RAID partitions into a RAID device.

3.18.3. Partition Fields

Above the partition hierarchy are labels which present information about the partitions you are creating. The labels are defined as follows:

- **Device**: This field displays the partition's device name.

- **Start**: This field shows the sector on your hard drive where the partition begins.

- **End**: This field shows the sector on your hard drive where the partition ends.

- **Size**: This field shows the partition's size (in MB).

- **Type**: This field shows the partition's type (for example, ext2, ext3, or vfat).

- **Mount Point**: A mount point is the location within the directory hierarchy at which a volume exists; the volume is "mounted" at this location. This field indicates where the partition will be mounted. If a partition exists, but is not set, then you need to define its mount point. Double-click on the partition or click the **Edit** button.

- **Format**: This field shows if the partition being created will be formatted.

3.18.4. Recommended Partitioning Scheme

Unless you have a reason for doing otherwise, we recommend that you create the following partitions:

- A swap partition (at least 32 MB) — swap partitions are used to support virtual memory. In other words, data is written to a swap partition when there is not enough RAM to store the data your system is processing. The size of your swap partition should be equal to twice your computer's RAM, or 32 MB, whichever amount is larger, but no more than 2048 MB (or 2 GB). In **Disk Druid**, the partition field for swap should look similar to the following:

  ```
  <Swap>  hda6  64M   64M  Linux swap
  ```

  For example, if you have 1 GB of RAM or less, your swap partition should be at least equal to the amount of RAM on your system, up to two times the RAM. For more than 1 GB of RAM, 2 GB of swap is recommended. Creating a large swap space partition will be especially helpful if you plan to upgrade your RAM at a later time.

- A /boot partition (50 MB) — the partition mounted on /boot contains the operating system kernel (which allows your system to boot Red Hat Linux), along with files used during the bootstrap process. Due to the limitations of most PC BIOSes, creating a small partition to hold these files is a good idea. For most users, a 50 MB boot partition is sufficient. In **Disk Druid**, the partition field for /boot should look similar to:

```
/boot  hda1  50M   50M  Linux native
```

🤚**Caution**

> If your hard drive is more than 1024 cylinders (and your system was manufactured more than two years ago), you may need to create a /boot partition if you want the / (root) partition to use all of the remaining space on your hard drive.

- A root partition (1.5-3.7 GB) — this is where "/" (the root directory) will be located. In this setup, all files (except those stored in /boot) are on the root partition. A 1.5 GB root partition will permit the equivalent of a workstation installation (with *very* little free space), while a 3.7 GB root partition will let you install every package. In **Disk Druid**, the partition field for / should look similar to:

```
/  hda5  3734M  3734M   Linux native
```

3.18.5. Adding Partitions

To add a new partition, select the **New** button. A dialog box appears (see Figure 3-14).

✏️**Note**

> You must dedicate at least one partition to Red Hat Linux, and optionally more. For more information, see Appendix E.

Figure 3-14. Creating a New Partition

- **Mount Point**: Enter the partition's mount point. For example, if this partition should be the root partition, enter **/**; enter **/boot** for the **/boot** partition, and so on. You can also use the pull-down menu to choose the correct mount point for your partition.

- **File System Type**: Using the pull-down menu, select the appropriate file system type for this partition. For more information on file system types, see Section 3.18.5.1.

- **Allowable Drives**: This field contains a list of the hard disks installed on your system. If a hard disk's box is highlighted, then a desired partition can be created on that hard disk. If the box is *not* checked, then the partition will *never* be created on that hard disk. By using different checkbox settings, you can have **Disk Druid** place partitions as you see fit, or let **Disk Druid** decide where partitions should go.

- **Size (Megs)**: Enter the size (in megabytes) of the partition. Note, this field starts with a "1" (one); unless changed, only a 1 MB partition will be created.

- **Additional Size Options**: Choose whether to keep this partition at a fixed size, to allow it to "grow" (fill up the available hard drive space) to a certain point, or to allow it to grow to fill any remaining hard drive space available.

  If you choose **Fill all space up to (MB)**, you must give size constraints in the field to the right of this option. This allows you to keep a certain amount of space free on your hard drive for future use.

- **Force to be a primary partition**: Select whether the partition you are creating should be one of the first four partitions on the hard drive. If unselected, the partition created will be a logical partition. See Section E.1.3, for more information.

- **Check for bad blocks**: Checking for bad blocks can help prevent data loss by locating the bad blocks on a drive and making a list of them to prevent using them in the future. If you wish to check for bad blocks while formatting each file system, please make sure to select this option.

  Selecting **Check for bad blocks** may dramatically increase your total installation time. Since most newer hard drives are quite large in size, checking for bad blocks may take a long time; the length of time depends on the size of your hard drive. If you choose to check for bad blocks, you can monitor your progress on virtual console #6.

- **Ok**: Select **Ok** once you are satisfied with the settings and wish to create the partition.

- **Cancel**: Select **Cancel** if you do not want to create the partition.

3.18.5.1. File System Types

Red Hat Linux allows you to create different partition types, based on the file system they will use. The following is a brief description of the different file systems available, and how they can be utilized.

- **ext2** — An ext2 file system supports standard Unix file types (regular files, directories, symbolic links, etc). It provides the ability to assign long file names, up to 255 characters. Versions prior to Red Hat Linux 7.2 used ext2 file systems by default.

- **ext3** — The ext3 file system is based on the ext2 file system and has one main advantage — journaling. Using a journaling file system reduces time spent recovering a file system

after a crash as there is no need to **fsck**[1] the file system. The ext3 file system will selected by default and is highly recommended.

- **software RAID** — Creating two or more software RAID partitions allows you to create a RAID device. For more information regarding RAID, refer to the chapter *RAID (Redundant Array of Independent Disks)* in the *Official Red Hat Linux Customization Guide*.
- **swap** — Swap partitions are used to support virtual memory. In other words, data is written to a swap partition when there is not enough RAM to store the data your system is processing.
- **vfat** — The VFAT file system is a Linux file system that is compatible with Windows 95/NT long filenames on the FAT file system.

3.18.6. Editing Partitions

To edit a partition, select the **Edit** button or double-click on the existing partition.

Note

If the partition already exists 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.

3.18.7. Deleting a Partition

To delete a partition, highlight it in the **Partitions** section and click the **Delete** button. You will be asked to confirm the deletion.

Skip to Section 3.20 for further installation instructions.

3.19. Partitioning with fdisk

This section applies only if you chose to use **fdisk** to partition your system.

To partition your system without using **fdisk**, please skip to Section 3.17 for automatic partitioning or Section 3.18 for partitioning with **Disk Druid**.

If you have already completed disk partitioning, skip to Section 3.20 for further installation instructions.

Caution

Unless you have previously used **fdisk** and understand how it works, we do not recommend that you use it. It is much easier for new users to accidentally corrupt or lose data using **fdisk**.

Disk Druid is easier to understand than **fdisk**. To exit **fdisk**, click **Back** to return to the previous screen, deselect **fdisk**, and then click **Next**.

---

1.   The **fsck** application is used to check the file system for metadata consistency and optionally repair one or more Linux file systems.

If you have chosen to use **fdisk**, the next screen will prompt you to select a drive to partition using **fdisk**. Once you have chosen which drive to partition, you will be presented with the **fdisk** command screen. If you do not know what command to use, type [m] at the prompt for help.

When you are finished making partitions, type [w] to save your changes and quit. You will be taken back to the original **fdisk** screen where you can partition another drive or continue the installation.

Note

> None of the changes you make take effect until you save them and exit **fdisk** using the w command. You can quit **fdisk** at any time without saving changes using the q command.

After you have partitioned your drive(s), click **Next**. You will need to use **Disk Druid** to assign mount points to the partitions you just created with **fdisk**.

You will not be able to add new partitions using **Disk Druid**, but you can edit mount points for the partitions you have already created. For each partition created with **fdisk**, click on the **Edit** button, choose the appropriate mount point for that partition from the pull-down menu, and click on **OK**.

3.20. Boot Loader Installation

In order to boot your Red Hat Linux system without a boot disk, you usually need to install a boot loader. You can choose to install either GRUB (selected by default), LILO, or you can choose not to install a boot loader at all.

First, select which boot loader you want to install. If you do not want to overwrite your current boot loader, choose **Do not install a boot loader**.

Caution

> If you choose not to install GRUB or LILO for any reason, you will not be able to boot your Red Hat Linux system directly, and you will need to use another boot method (such as a boot diskette). Use this option only if you are sure you have another way of booting your Red Hat Linux system!
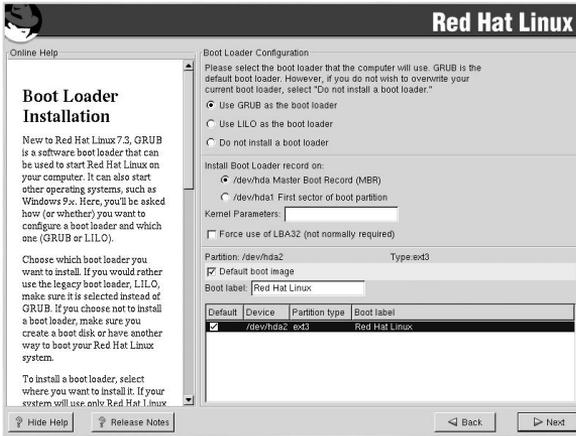
Figure 3-15. Boot Loader Installation

Assuming you chose GRUB or LILO, you must now determine where you want the boot loader to be installed. You may install the boot loader in one of two places:

The master boot record (MBR)

> This is the recommended place to install a boot loader, unless the MBR already starts another operating system loader, such as System Commander. The MBR is a special area on your hard drive that is automatically loaded by your computer's BIOS, and is the earliest point at which the boot loader can take control of the boot process. If you install it in the MBR, when your machine boots, GRUB (or LILO) will present a boot prompt. You can then boot Red Hat Linux or any other operating system that you have configured the boot loader to boot.

The first sector of your boot partition

> This is recommended if you are already using another boot loader on your system. In this case, your other boot loader will take control first. You can then configure that boot loader to start GRUB (or LILO), which will then boot Red Hat Linux.

If your system will use only Red Hat Linux, you should choose the MBR. For systems with Windows 95/98, you should also install the boot loader to the MBR so that it can boot both operating systems.

If you wish to add default options to GRUB or LILO's boot command, enter them into the **Kernel parameters** field. Any options you enter will be passed to the Linux kernel every time it boots.

The **Force use of LBA32 (not normally required)** option allows you to exceed the 1024 cylinder limit for the /boot partition. If you have a system which supports the LBA32 extension for booting operating systems above the 1024 cylinder limit, and you want to place your /boot partition above cylinder 1024, you should select this option.

Every bootable partition is listed, including partitions used by other operating systems. The partition holding your Red Hat Linux system's root file system will have a **Boot label** of Red Hat Linux (for GRUB) or linux (for LILO). Other partitions may also have boot labels. If you would like to add boot labels for other partitions (or change an existing boot label), click once on the partition to select it. Once selected, you can change the boot label by editing the name in the **Boot label** text field.

Note

The **Boot label** column lists what you must enter at the boot prompt, in non-graphical boot loaders, in order to boot the desired operating system.

Once you have loaded the GRUB boot screen, use the arrow keys to choose a boot label or type e for edit. You will be presented with a list of items in the configuration file for the boot label you have selected.

At the graphical LILO screen, press [Ctrl]-[x] to exit to the boot: prompt. If you forget the boot labels defined on your system, you can always press [Tab] at the prompt to display a list of defined boot labels.

3.20.1. Rescue Mode

If you need to use rescue mode, there are several options available to you.

- Using the CD-ROM to boot, type **linux rescue** at the boot: prompt.

- By booting your system from an installation boot diskette made from the boot.img image. This method requires that the Red Hat Linux CD-ROM #1 be inserted as the rescue image or that the rescue image be on the hard drive as an ISO image. Once you have booted using this diskette, type **linux rescue** at the boot: prompt.

- By booting from a network disk made from the bootnet.img or PCMCIA boot disk made from pcmcia.img. Once you have booted using this diskette, type **linux rescue** at the boot: prompt. You can only do this if your network connection is working. You will need to identify the network host and transfer type. For an explanation of how to specify this information, see Section 3.7.

For more information, refer to the *Official Red Hat Linux Customization Guide*.

3.20.2. Alternative Boot Loaders

If you do not wish to use a boot loader, you have several alternatives:

Boot disk

You can use the boot disk created by the installation program (if you create one).

LOADLIN

You can load Linux from MS-DOS. Unfortunately, this 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 (for example, from a boot disk) and then copy the kernel to an MS-DOS partition. LOADLIN is available from

ftp://metalab.unc.edu/pub/Linux/system/boot/dualboot/

and associated mirror sites.

SYSLINUX

>   SYSLINUX is an MS-DOS program very similar to LOADLIN. It is also available from
>
>   `ftp://metalab.unc.edu/pub/Linux/system/boot/loaders/`
>
>   and associated mirror sites.

Some commercial boot loaders

>   You can load Linux using commercial boot loaders. For example, System Commander
>   and Partition Magic are able to boot Linux (but still require GRUB or LILO to be installed
>   in your Linux root partition).

3.20.3. SMP Motherboards, GRUB, and LILO

This section is specific to SMP motherboards only. SMP, short for Symmetric Multiprocessing, is a computer architecture providing fast performance by making multiple CPUs available to complete individual processes simultaneously (multiprocessing).

If the installation program detects an SMP motherboard on your system, it will automatically create two `/boot/grub/grub.conf` or `/etc/lilo.conf` entries (depending on the boot loader you installed), rather than the usual single entry.

Note

Intel®Pentium®4 systems with hyperthreading will have an SMP kernel installed by default.

The two entries in `grub.conf` will be `Red Hat Linux (`*`kernel version`*`)` and `Red Hat Linux (`*`kernel versions`*`-smp)`. The `Red Hat Linux (`*`kernel version`*`-smp)` will boot by default. However, if you have trouble with the SMP kernel, you can elect to boot the `Red Hat Linux (`*`kernel version`*`)` entry instead. You will retain all the functionality as before, but you will only be operating with a single processor.

The two entries in `lilo.conf` will be `linux` and `linux-up`. The `linux` entry will boot by default. However, if you have trouble with the SMP kernel, you can elect to boot the `linux-up` entry instead. You will retain all the functionality as before, but you will only be operating with a single processor.

3.21. GRUB Password

If you did not select GRUB as your boot loader, or did not install a boot loader, skip to Section 3.22.

GRUB passwords provide a security mechanism in an environment where physical access to your server is available.

If you are installing GRUB as your boot loader, you should create a password to protect your system. Without a GRUB password, users with access to your system can pass options to the kernel which can compromise your system security. With a GRUB password in place, the password must first be entered in order to select any non-standard boot options.
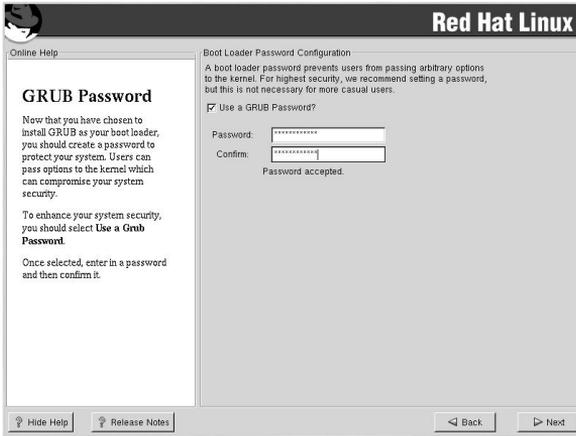
Figure 3-16. GRUB Password

If you choose to use a GRUB password to enhance your system security, be sure to select the checkbox labeled **Use a GRUB Password**.

Once selected, enter a password and confirm it.

3.22. Network Configuration

If you do not have a network device, you will not see this screen. Skip ahead to Section 3.23.

If you have a network device and you have not already configured your networking (such as booting from a network boot disk you created and entering in your network information as prompted), you now have the opportunity (as shown in Figure 3-17) to do so.

If you have multiple devices, you will see a tab for each device. You may switch between devices (for example, between eth0 and eth1) and the information you provide on each tab will be specific to each device.

Indicate if you would like to configure your IP address using DHCP (Dynamic Host Configuration Protocol). If you select **Activate on boot**, your network interface will be started when you boot. If you do not have DHCP client access or you are unsure what to provide here, please contact your network administrator.

Next enter, where applicable, the **IP Address**, **Netmask**, **Network**, and **Broadcast** addresses. If you are unsure about any of these, please contact your network administrator.

If you have a fully qualified domain name for the network device, enter it in the **Hostname** field.

Finally, enter the **Gateway** and **Primary DNS** (and if applicable the **Secondary DNS** and **Ternary DNS**) addresses.
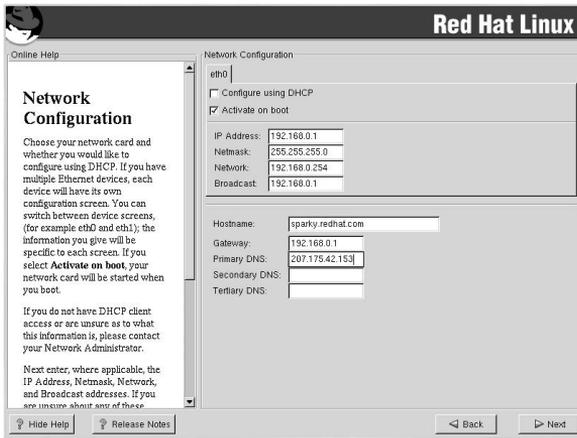
Figure 3-17. Network Configuration

Note

Do not use the numbers as seen in this sample configuration. These values will not work for your own network configuration. If you are not sure what values to enter, contact your network administrator for assistance.

Tip

Even if your computer is not part of a network, you can enter a hostname for your system. If you do not take this opportunity to enter a name, your system will be known as `localhost`.

3.23. Firewall Configuration

Red Hat Linux offers firewall protection for enhanced system security. A firewall exists between your computer and the network, and determines which resources on your computer remote users on the network can access. A properly configured firewall can greatly increase the security of your system.
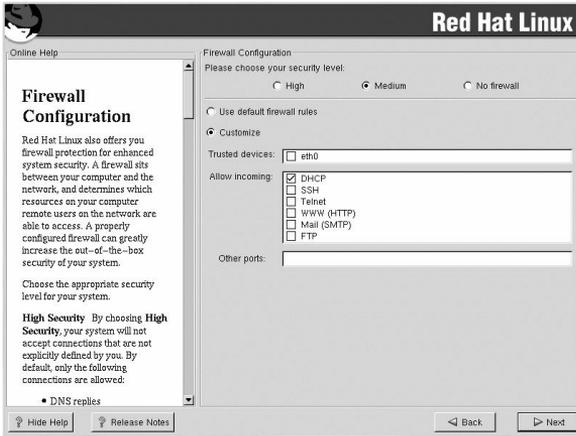
Figure 3-18. Firewall Configuration

Choose the appropriate security level for your system.

High

> If you choose **High**, your system will not accept connections (other than the default settings) that are not explicitly defined by you. By default, only the following connections are allowed:
>
> - DNS replies
> - DHCP — so any network interfaces that use DHCP can be properly configured
>
> If you choose **High**, your firewall will not allow the following:
>
> - Active mode FTP (passive mode FTP, used by default in most clients, should still work)
> - IRC DCC file transfers
> - RealAudio™
> - Remote X Window System clients
>
> If you are connecting your system to the Internet, but do not plan to run a server, this is the safest choice. If additional services are needed, you can choose **Customize** to allow specific services through the firewall.

> **Note**
>
> If you select a medium or high firewall to be setup during this installation, network authentication methods (NIS and LDAP) will not work.

Medium

If you choose **Medium**, your firewall will not allow remote machines to have access to certain resources on your system. By default, access to the following resources are not allowed:

- Ports lower than 1023 — the standard reserved ports, used by most system services, such as **FTP**, **SSH**, **telnet**, **HTTP**, and **NIS**.
- The NFS server port (2049) — NFS is disabled for both remote severs and local clients.
- The local X Window System display for remote X clients.
- The X Font server port (by default, **xfs** does not listen on the network; it is disabled in the font server).

If you want to allow resources such as **RealAudio**™ while still blocking access to normal system services, choose **Medium**. Select **Customize** to allow specific services through the firewall.

Note

If you select a medium or high firewall to be setup during this installation, network authentication methods (NIS and LDAP) will not work.

No Firewall

No firewall provides complete access to your system and does no security checking. Security checking is the disabling of access to certain services. This should only be selected if you are running on a trusted network (not the Internet) or plan to do more firewall configuration later.

Choose **Customize** to add trusted devices or to allow additional incoming services.

Trusted Devices

Selecting any of the **Trusted Devices** allows access to your system for all traffic from that device; it is excluded from the firewall rules. For example, if you are running a local network, but are connected to the Internet via a PPP dialup, you can check **eth0** and any traffic coming from your local network will be allowed. Selecting **eth0** as trusted means all traffic over the Ethernet is allowed, put the ppp0 interface is still firewalled. If you want to restrict traffic on an interface, leave it unchecked.

It is not recommended that you make any device that is connected to public networks, such as the Internet, a **Trusted Device**.

Allow Incoming

Enabling these options allow the specified services to pass through the firewall. Note, during a workstation installation, the majority of these services are *not* installed on the system.

DHCP

If you allow incoming DHCP queries and replies, you allow any network interface that uses DHCP to determine its IP address. DHCP is normally enabled. If DHCP is not enabled, your computer can no longer get an IP address.

SSH

*S*ecure *SH*ell (SSH) is a suite of tools for logging into and executing commands on a remote machine. If you plan to use SSH tools to access your machine through a firewall, enable this option. You need to have the `openssh-server` package installed in order to access your machine remotely, using SSH tools.

Telnet

Telnet is a protocol for logging into remote machines. Telnet communications are unencrypted and provide no security from network snooping. Allowing incoming Telnet access is not recommended. If you do want to allow inbound Telnet access, you will need to install the `telnet-server` package.

WWW (HTTP)

The HTTP protocol is used by Apache (and by other Web servers) to serve webpages. If you plan on making your Web server publicly available, enable this option. This option is not required for viewing pages locally or for developing webpages. You will need to install the `apache` package if you want to serve webpages.

Enabling **WWW (HTTP)** will not open a port for HTTPS. To enable HTTPS, specify it in the **Other ports** field.

Mail (SMTP)

If you want to allow incoming mail delivery through your firewall, so that remote hosts can connect directly to your machine to deliver mail, enable this option. You do not need to enable this if you collect your mail from your ISP's server using POP3 or IMAP, or if you use a tool such as **fetchmail**. Note that an improperly configured SMTP server can allow remote machines to use your server to send spam.

FTP

The FTP protocol is used to transfer files between machines on a network. If you plan on making your FTP server publicly available, enable this option. You need to install the `wu-ftpd` (and possibly the `anonftp`) package for this option to be useful.

Other ports

You can allow access to ports which are not listed here, by listing them in the **Other ports** field. Use the following format: **port:protocol**. For example, if you want to allow IMAP access through your firewall, you can specify **imap:tcp**. You can also explicitly specify numeric ports; to allow UDP packets on port 1234 through the firewall, enter **1234:udp**. To specify multiple ports, separate them with commas.

3.24. Language Support Selection

Red Hat Linux can install and support multiple languages for use on your system.

You must select a language to use as the default language. The default language will be used on your Red Hat Linux system once installation is complete. If you choose to install other languages, you can change your default language after the installation.

If you are only going to use one language on your system, selecting only that language will save significant disk space. The default language is the language you selected to use during the installation. However, if you select only one language, you will only be able to use that specified language after the Red Hat Linux installation is complete.
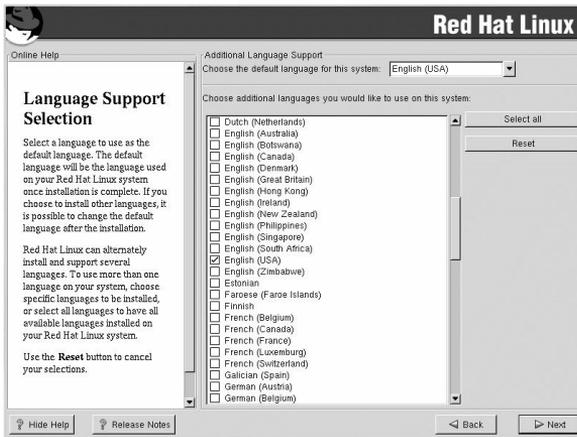

Figure 3-19. Language Support Selection

To use more than one language on your system, choose specific languages to be installed or select all languages to have all available languages installed on your Red Hat Linux system.

Use the **Reset** button to cancel your selections. Resetting will revert to the default; only the language you selected for use during the installation will be installed.

3.25. Time Zone Configuration

You can set your time zone by selecting your computer's physical location or by specifying your time zone's offset from Coordinated Universal Time (UTC).

Notice the two tabs at the top of the screen (see Figure 3-20). The first tab allows you to configure your time zone by your location. You can specify different areas to view: **World**, **North America**, **South America**, **Pacific Rim**, **Europe**, **Africa**, and **Asia**.

On the interactive map, you can also click on a specific city, which is marked by a yellow dot; a red **X** will appear indicating your selection. You can also scroll through a list and choose a time zone.

The second tab allows you to specify a UTC offset. The tab displays a list of offsets to choose from, as well as an option to set daylight saving time.
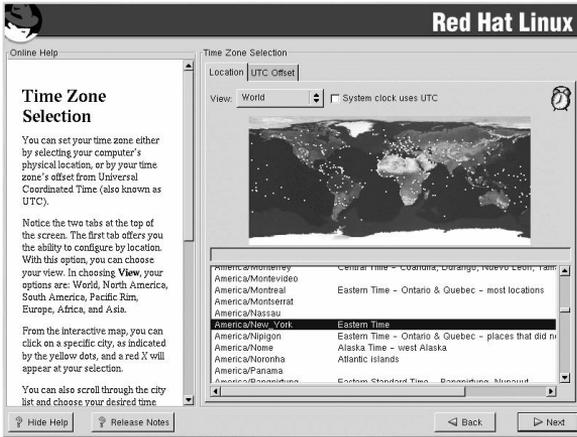
Figure 3-20. Configuring the Time Zone

On both tabs, you can select **System Clock uses UTC**. Please select this if you know that your system is set to UTC.

Tip

If you wish to change your time zone configuration after you have booted your Red Hat Linux system, become root and use the `/usr/sbin/timeconfig` command.

3.26. Account Configuration

The **Account Configuration** screen allows you to set your root password. Additionally, you can set up user accounts for you to log in to once the installation is complete (see Figure 3-21).
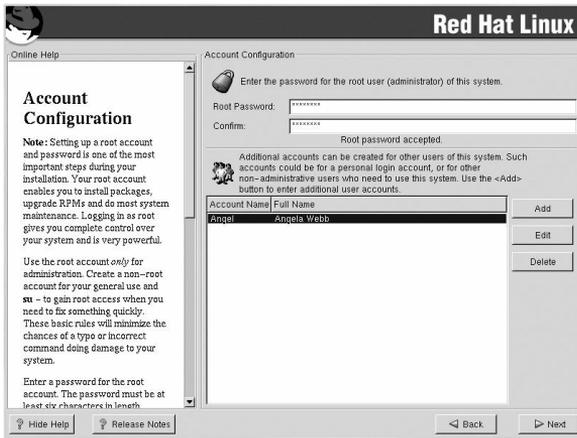
Figure 3-21. Account Creation

3.26.1. Setting the Root Password

Setting up a root account and password is one of the most important steps during your installation. Your root account is similar to the administrator account used on Windows NT machines. The root account is used to install packages, upgrade RPMs, and perform most system maintenance. Logging in as root gives you complete control over your system.

Use the root account only for system administration. Create a non-root account for your general use and su - to root when you need to fix something quickly. These basic rules will minimize the chances of a typo or an incorrect command doing damage to your system.

⧫**Tip**

> To become root, type su - at the shell prompt in a terminal window and then press [Enter]. Then, enter the root password and press [Enter].

The installation program will prompt you to set a root password[2] for your system. You must enter a root password. The installation program will not let you proceed to the next section without entering a root password.

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 should 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 bad passwords. Good passwords mix numerals with upper and lower case letters and do not contain dictionary words: *Aard387vark* or *420BMt-tNT*, for example. Remember that the password is case-sensitive. If you write down your

---

2.   A root password is the administrative password for your Red Hat Linux system. You should only log in as root when needed for system maintenance. The root account does not operate within the restrictions placed on normal user accounts, so changes made as root can have implications for your entire system.

password, keep it in a secure place. However, it is recommended that you do not write down this or any password you create.

Note

Do not use one of the example passwords offered in this manual. Using one of these passwords could be considered a security risk.

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.

3.26.2. Setting Up User Accounts

If you choose to create a user account now, you will have an account to log in to once the installation has completed. This allows you to safely and easily log into your computer without having to be root to create your user account.

Enter an account name. Then enter and confirm a password for that user account. Enter the full name of the account user and press **Add**. Your account information will be added to the account list, and the user account fields will be cleared so that you can add another user.

Figure 3-22. Creating a User Account

Choose **New** to add a new, non-root, user. Enter the user's information and use the **Add** button to add the user to the account list.

You can also **Edit** or **Delete** the user accounts you have created and no longer want.

3.27. Authentication Configuration

If you are performing a workstation, laptop, or server installation, please skip ahead to Section 3.28.

You may skip this section if you will not be setting up network passwords. If you do not know whether you should do this, please ask your system administrator for assistance.

Unless you are setting up **NIS** authentication, you will notice that only **MD5** and **shadow** passwords are selected (see Figure 3-23). We recommend you use both to make your machine as secure as possible.

To configure the NIS option, you must be connected to an NIS network. If you are not sure whether you are connected to an NIS network, please ask your system administrator.
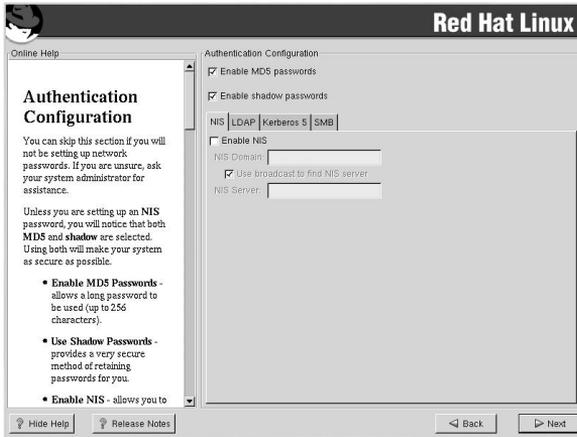

Figure 3-23. Authentication Configuration

- **Enable MD5 passwords** — allows a long password to be used (up to 256 characters), instead of the standard eight letters or less.
- **Enable shadow passwords** — provides a secure method for retaining passwords. The passwords are stored in /etc/shadow, which can only be read by root.
- **Enable NIS** — allows you to run a group of computers in the same Network Information Service domain with a common password and group file. You can choose from the following options:
  - **NIS Domain** — allows you to specify the domain or group of computers your system belongs to.
  - **Use broadcast to find NIS server** — allows you to broadcast a message to your local area network to find an available NIS server.
  - **NIS Server** — causes your computer to use a specific NIS server, rather than broadcasting a message to the local area network asking for any available server to host your system.

> **Note**
>
> If you have selected a medium or high firewall to be setup during this installation, network authentication methods (NIS and LDAP) will not work.

- **Enable LDAP** — tells your computer to use LDAP for some or all authentication. LDAP consolidates certain types of information within your organization. For example, all of the different lists of users within your organization can be merged into one LDAP directory. For more information about LDAP, refer to the *Official Red Hat Linux Reference Guide*, *Lightweight Directory Access Protocol (LDAP)*. You can choose from the following options:
  - **LDAP Server** — allows you to access a specified server (by providing an IP address) running the LDAP protocol.
  - **LDAP Base DN** — allows you to look up user information by its Distinguished Name (DN).
  - **Use TLS** (*Transport Layer Security*) **lookups** — this option allows LDAP to send encrypted user names and passwords to an LDAP server before authentication.

- **Enable Kerberos** — Kerberos is a secure system for providing network authentication services. For more information about Kerberos, see *Using Kerberos 5 on Red Hat Linux* in the *Official Red Hat Linux Reference Guide*. There are three options to choose from here:
  - **Realm** — this option allows you to access a network that uses Kerberos, composed of one or a few servers (also known as KDCs) and a potentially large number of clients.
  - **KDC** — this option allows you access to the Key Distribution Center (KDC), a machine that issues Kerberos tickets (sometimes called a Ticket Granting Server or TGS).
  - **Admin Server** — this option allows you to access a server running **kadmind**.

- **Enable SMB Authentication** — Sets up PAM to use an SMB server to authenticate users. You must supply two pieces of information here:
  - **SMB Server** — Indicates which SMB server your workstation will connect to for authentication.
  - **SMB Workgroup** — Indicates which workgroup the configured SMB servers are in.

3.28. Package Group Selection

After your partitions have been selected and configured for formatting, you are ready to select packages for installation.

Note

Unless you choose a custom installation, the installation program will automatically choose most packages for you. However, you must select either GNOME or KDE (or both) to install a graphical environment.

GNOME and KDE are both graphical desktop environments[3] that handle the overall look and feel of your system. You must choose one of these to have a default graphical setup, but you can also install both to determine for yourself which you prefer.

---

3.   A desktop environment in Linux is similar to the environment you might see in other operating systems. However, environments differ in their look and feel and are easily customized for your individual needs.

If you do not have enough space (at least 1.8 GB) to install both GNOME and KDE, Section 3.28.2 and Section 3.28.3 should make the choice a little easier for you.

You can select components, which group packages together according to function (for example, **C Development**, **Networked Workstation**, or **Web Server**), individual packages, or a combination of the two.

To select a component, click on the checkbox beside it (see Figure 3-24).

Figure 3-24. Package Group Selection

Select each component you wish to install. Selecting **Everything** (at the end of the component list) during a custom installation installs all packages included with Red Hat Linux. If you select every package, you will need approximately 3.7 GB of free disk space.

To select packages individually, check the **Select Individual Packages** box at the bottom of the screen.

3.28.1. Selecting Individual Packages

After selecting the components you wish to install, you can select or deselect individual packages using your mouse (see Figure 3-25).

You can choose to view the individual packages in **tree view** or **flat view**.

Tree view allows you to see the packages grouped by application type.

Flat view allows you to see all of the packages in an alphabetical listing on the right of the screen.
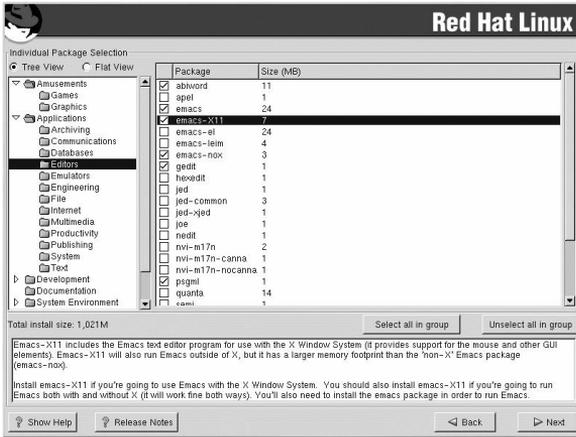
Figure 3-25. Selecting Individual Packages

Using **Tree view**, you see a listing of package groups. When you expand this list (by double-clicking on the folder arrow beside a package group name) and pick one group, the list of packages in that group appears in the panel on the right. **Flat view** allows you to see all of the packages in an alphabetical listing on the right of the screen.

To sort packages alphabetically, click on the **Package** tab. To sort packages by size, click on the **Size (MB)** tab.

To select an individual package, double-click the checkbox beside the package name. A check mark in the box means that a package has been selected.

For more information about a specific package, click on the individual package name. The package information will appear at the bottom of the screen.

You can also select or deselect all packages listed within a particular group, by clicking on the **Select all in group** or **Unselect all in group** buttons.

Note

Some packages (such as the kernel and certain libraries) are required for every Red Hat Linux system and are not available to select or deselect. These base packages are selected by default.

3.28.2. A Brief Introduction to GNOME

GNOME is a powerful graphics-driven desktop environment. GNOME includes a panel (for starting applications and displaying status), a desktop (where data and applications can be placed), multiple window managers (which control the look and feel of your desktop), and a standard set of desktop tools and applications.

GNOME allows you to setup your desktop the way you want it to look and "feel." GNOME's session manager remembers settings and currently running programs. So, once you have set things the way you like, they will stay that way.
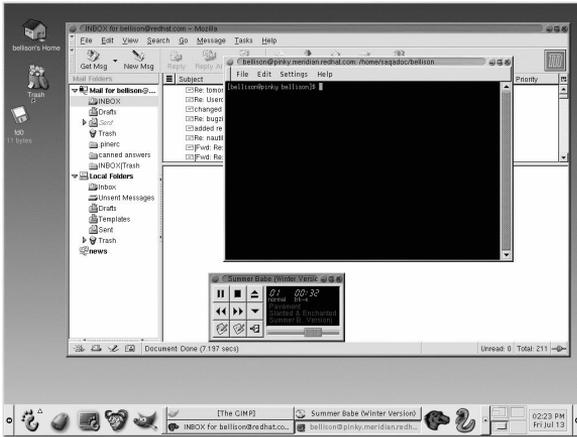
Figure 3-26. Sample GNOME User Screen

Figure 3-26 shows a typical graphical environment using the GNOME desktop environment. A typical graphical environment for KDE would be similar.

Refer to the *Official Red Hat Linux Getting Started Guide* to learn more about GNOME.

3.28.3. A Brief Introduction to KDE

KDE provides a complete desktop environment, including a file manager, a window manager, an integrated help system, a configuration system, numerous tools and utilities, and an ever-increasing number of applications.

KDE offers a contemporary desktop, a searchable help system with convenient access to help on the use of the KDE desktop and its applications, standardized menu and toolbars, keybindings, color schemes, and more.

Refer to the *Official Red Hat Linux Getting Started Guide* to learn more about KDE.

3.28.4. Unresolved Dependencies

Many software packages, in order to work correctly, depend on other software packages 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 these package *dependencies* each time you install or remove software packages.

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 3-27).

The **Unresolved Dependencies** screen appears only if you are missing packages that are needed by the packages you have selected. At the bottom of the screen, under the list of missing packages, an **Install packages to satisfy dependencies** checkbox is selected by default. If you leave this checked, the installation program will resolve dependencies automatically by adding all required packages to the list of selected packages.
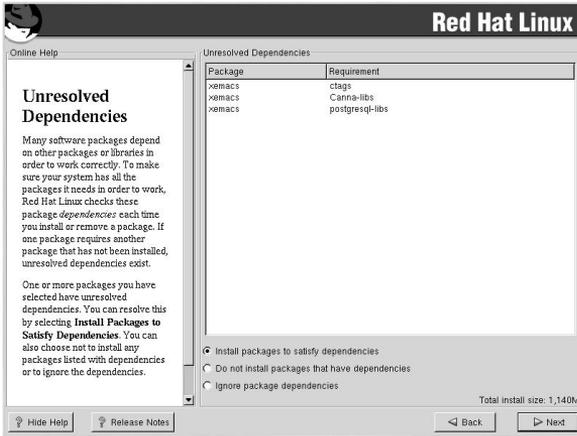
Figure 3-27. Unresolved Dependencies

If you do not wish to install packages that require other packages, select **Do not install packages that have dependencies**.

To install only the packages you have selected and leave the dependencies unresolved, select **Ignore package dependencies**.

3.29. X Configuration — Video Card

The first part of X configuration deals with video card configuration.

3.29.1. Video Card Configuration

Xconfigurator will now present a list of video cards for you to choose from.

If you decided to install the X Window System packages, you now have the opportunity to configure an X server for your system. If you did not choose to install the X Window System packages, skip ahead to Section 3.30.

If your video card does not appear on the list (see Figure 3-28), X may not support it. However, if you have technical knowledge about your card, you may choose **Unlisted Card** and attempt to configure it by matching your card's video chipset with one of the available X servers.
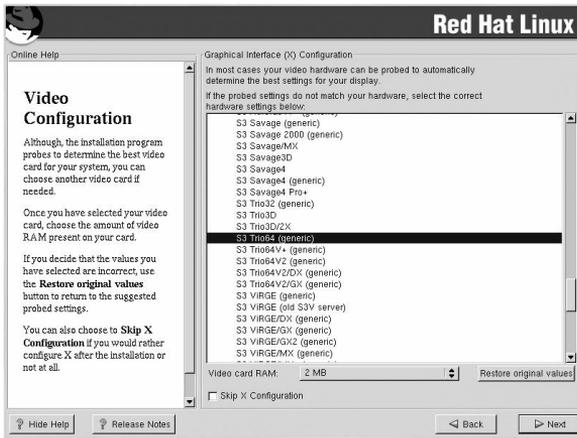
Figure 3-28. Video Card Setup

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

If you decide that the values you have selected are incorrect, you can click the **Restore original values** button to return to the suggested settings.

You can also select **Skip X Configuration** if you would rather configure X after the installation or not at all.

3.30. Preparing to Install

You should now see a screen preparing you for the installation of Red Hat Linux.

For your reference, a complete log of your installation can be found in `/root/install.log` once you reboot your system.

⚠️**Warning**

> If, for some reason, you would rather not continue with the installation process, this is your last opportunity to safely cancel the process and reboot your machine. Once you press the **Next** button, partitions will be written and packages will be installed. If you wish to abort the installation, you should reboot now before any existing information on any hard drive is rewritten.

To cancel this installation process, press your computer's Reset button or use the [Control]-[Alt]-[Delete] key combination to restart your machine.

3.31. Installing Packages

At this point there is nothing left for you to do until all the packages have been installed (see Figure 3-29). How quickly this happens depends on the number of packages you have selected and your computer's speed.

Figure 3-29. Installing Packages

3.32. Boot Disk Creation

To create a boot disk, insert a blank, formatted diskette into your diskette drive (see Figure A-9) and click **Next**.

It is highly recommended that you create a boot disk. If, for some reason, your system were not able to boot properly using GRUB, LILO, or a third-party boot loader, a boot disk would enable you to properly boot your Red Hat Linux system.

After a short delay, your boot disk will be created; remove it from your diskette drive and label it clearly. Note that if you would like to create a boot disk after the installation, you will be able to do so. For more information, please see the mkbootdisk man page, by typing man mkbootdisk at the shell prompt.

If you do not want to create a boot disk, make sure to check the **Skip boot disk creation** checkbox before you click **Next**.

If you boot your system with the boot disk (instead of GRUB or LILO) , make sure you create a new boot disk whenever you make any changes to your kernel (including the installation of a new kernel).

Figure 3-30. Creating Your Boot Disk

3.33. X Configuration — Monitor and Customization

In order to complete X configuration, you must configure your monitor and customize your X settings.

3.33.1. Configuring Your Monitor

If you chose to skip X configuration, go to Section 3.34.

Xconfigurator, the X Window System configuration tool, presents a list of monitors for you to choose from. In the list, you can either use the monitor that is automatically detected for you, or choose another monitor.

If your monitor does not appear on the list, select the most appropriate **Generic** model available. By selecting a **Generic** monitor, **Xconfigurator** will suggest horizontal and vertical sync ranges. These values are generally available in the documentation which accompanies your monitor, or from your monitor's vendor or manufacturer; please check your documentation to make sure these values are set correctly.

Note

If you are installing Red Hat Linux on a laptop with an LCD screen, you should select the most appropriate **Generic** model available.
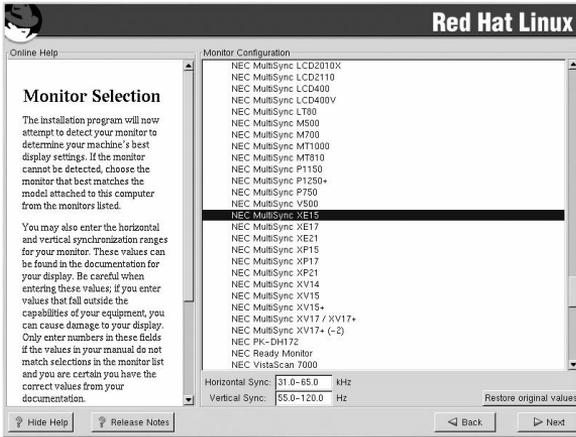
Figure 3-31. Monitor Selection

✋**Caution**

Do not 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. Doing so may overclock your monitor and damage or destroy it.

The horizontal and vertical ranges that **Xconfigurator** suggests for your monitor are also displayed on this screen.

If you decide that the values you have selected are incorrect, you can click the **Restore original values** button to return to the suggested settings.

Click **Next** when you have finished configuring your monitor.

3.33.2. Custom Configuration

Choose the correct color depth and resolution for your X configuration. Click **Test Setting** to try out this configuration. If you do not like what you see during the test, click **No** to choose another resolution.

📝**Note**

If you need to exit out of the X test, use the [Ctrl]-[Alt]-[Backspace] key combination. Also note that this will not work in some test cases.

We recommend that you test your configuration, to make sure the resolution and color settings are usable.

If you installed both GNOME and KDE, you can choose which one to use as your default desktop environment. If you installed one or the other, it will only show GNOME or KDE as the desktop default.

You can also choose whether you want to boot your system into a text or graphical environment once Red Hat Linux is installed. Unless you have special needs, booting into a graphical environment (similar to a Windows environment) is recommended. If you choose to boot into a text environment, you will be presented with a command prompt (similar to a DOS environment).
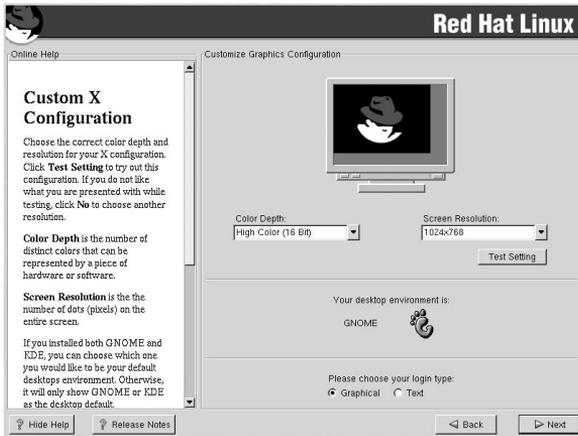

Figure 3-32. X Customization

3.34. Installation Complete

Congratulations! Your Red Hat Linux 7.3 installation is now complete!

The installation program will prompt you to prepare your system for reboot. Remember to remove any installation media (diskette in the diskette drive or CD in the CD-ROM drive) if they are not ejected automatically upon reboot.

If you do not have a boot loader installed and configured, you will need to use the boot disk you created during the installation now.

After your computer's normal power-up sequence has completed, you should see the graphical boot loader prompt, at which you can do any of the following things:

• Press [Enter] — causes the default boot entry to be booted.

• Select a boot label, followed by [Enter] — causes the boot loader to boot the operating system corresponding to the boot label. (Press [?] or [Tab] at LILO's text mode boot loader prompt for a list of valid boot labels.)

• Do nothing — after the boot loader's timeout period, (by default, five seconds) the boot loader will automatically boot the default boot entry.

Do whatever is appropriate to boot Red Hat Linux. You should see one or more screens of messages scroll by. Eventually, you should see a login: prompt or a graphical login screen (if you installed the X Window System and chose to boot into graphical mode by default).

Tip

If you are not sure what to do next, we suggest you begin with the *Official Red Hat Linux Getting Started Guide* which covers topics relating to the basics of your system and is an introduction to using Red Hat Linux. The *Official Red Hat Linux Getting Started Guide* can be found on the Documentation CD or online at:

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

If you are a more experienced user looking for information on administration topics, you may find the *Official Red Hat Linux Reference Guide* to be more helpful.

If you are looking for information on system configuration, you may find the *Official Red Hat Linux Customization Guide* to be helpful.

Appendix A.

Upgrading Your Current System

This appendix walks you through a typical Red Hat Linux 7.3 upgrade.

A.1. What it Means to Upgrade

The installation process for Red Hat Linux 7.3 includes the ability to upgrade from prior versions of Red Hat Linux (version 4.2 and later) which are based on RPM technology.

Upgrading your system installs the modular 2.4.x kernel as well as updated versions of the packages which are currently installed on your system.

The upgrade process preserves existing configuration files by renaming them with an `.rpm-save` extension (for example, `sendmail.cf.rpmsave`). The upgrade process also creates a log of its actions in `/root/upgrade.log`. 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.

> **Note**
>
> Some upgraded packages may require the installation of other packages for proper operation. If you choose to customize your packages to upgrade, you may be required to resolve dependency problems. Otherwise, the upgrade procedure takes care of these dependencies, but it may need to install additional packages which are not on your system.

Depending on how you have partitioned your system, the upgrade program may prompt you to add an additional swap file. If the upgrade program does not detect a swap file that equals twice your RAM, it will ask you if you would like to add a new swap file. If your system does not have a lot of RAM (less than 32 MB), it is recommended that you add this swap file.

A.2. Upgrading Your System

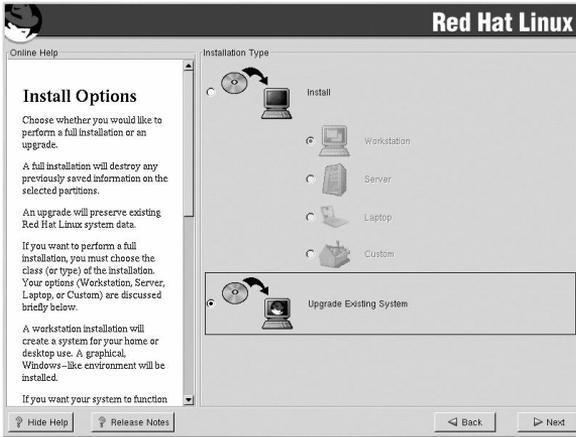At this point, you should have chosen **Upgrade** as your preferred installation type (see Figure A-1).

Figure A-1. Choosing Install or Upgrade

A.3. Upgrading Your File System

Note

This section only pertains to users performing an upgrade from Red Hat Linux version 7.1 or earlier, or from a Red Hat Linux 7.2 installation where ext2 was chosen as the file system.

If the installation program detects the ext2 file system on your Red Hat Linux system, you can choose to retain your current ext2 file system or migrate to the ext3 file system.

The following is a brief description of the ext2 and ext3 file systems, and how they can be utilized.

- **ext2** — An ext2 file system supports standard Unix file types (regular files, directories, symbolic links, etc). It provides the ability to assign long file names, up to 255 characters. Versions prior to Red Hat Linux 7.3 used ext2 file systems by default.
- **ext3** — The ext3 file system is based on the ext2 file system and has one main advantage — journaling. Using a journaling file system reduces time spent recovering a file system after a crash as there is no need to **fsck**[1]the file system.

It is recommended, but not required, that you choose to migrate to the ext3 file system.

If you choose to migrate to the ext3 file system, existing system data will not be modified.

A.4. Customizing Your Upgrade

Do you want to choose the packages to be upgraded or let the installation program perform an automated upgrade (see Figure A-2)?

---

1. The **fsck** application is used to check and optionally repair one or more Linux file systems.

To perform an automated upgrade, click **Next**.

To customize your packages for this upgrade, select the **Customize packages to be upgraded** option and click **Next** (see Section A.8).

Figure A-2. Upgrade Customization

A.5. Boot Loader Configuration

A software boot loader is used to start Red Hat Linux on your computer. It can also start other operating systems, such as Windows. If you are using a Red Hat Linux software boot loader (GRUB or LILO), it will be detected automatically.
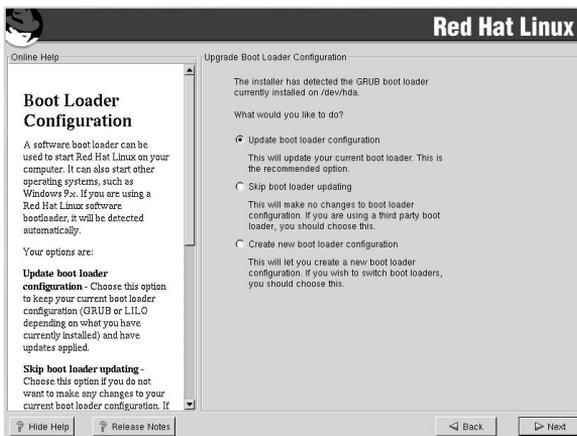

Figure A-3. Upgrade Boot Loader Configuration

On the **Boot Loader Configuration Screen**, your options are:

Update boot loader configuration — Choose this option to keep your current boot loader configuration (GRUB or LILO depending on what you have currently installed) and have updates applied.

Skip boot loader updating — Choose this option if you do not want to make any changes to your current boot loader configuration. If you are using a third party boot loader, you will want to skip updating your boot loader.

Create new boot loader configuration — Choose this option if you want to create a new boot loader for your system. If you currently have LILO and want to switch to GRUB, or if you have been using boot disks to boot your Red Hat Linux system and want to use a software boot loader such as GRUB or LILO, you will want to create a new boot loader configuration (see Section A.5.1 for more information).

Once you have made your selection, click **Next** to continue.

A.5.1. Creating a New Boot Loader Configuration

In order to be able to boot your Red Hat Linux system, you usually need to install a boot loader. In Red Hat Linux 7.3, you can choose to install GRUB or LILO, or you can choose not to install a boot loader at all.

First, select which boot loader you want to install. If you do not want to overwrite your current boot loader, choose **Do not install a boot loader**.

Assuming you chose GRUB or LILO, you must now determine where you want the boot loader to be installed. You may install the boot loader in one of two places:

The master boot record (MBR)

   This is the recommended place to install a boot loader, unless the MBR already starts another operating system loader, such as System Commander. The MBR is a special area on your hard drive that is automatically loaded by your computer's BIOS, and is the earliest point at which the boot loader can take control of the boot process. If you install it in the MBR, when your machine boots, GRUB (or LILO) will present a boot prompt. You can then boot Red Hat Linux or any other operating system that you have configured the boot loader to boot.

The first sector of your boot partition

   This is recommended if you are already using another boot loader on your system. In this case, your other boot loader will take control first. You can then configure that boot loader to start GRUB (or LILO), which will then boot Red Hat Linux.
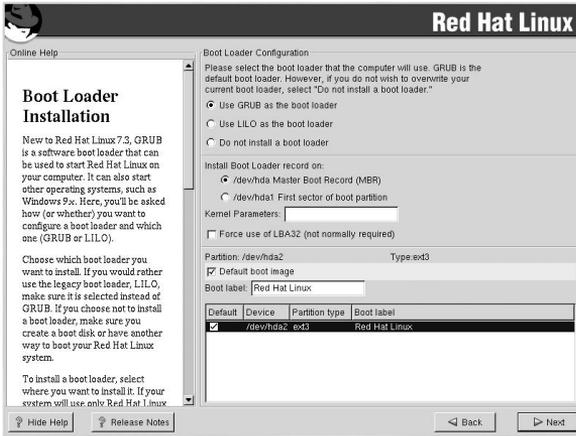
Figure A-4. Boot Loader Configuration

If your system will use only Red Hat Linux, you should choose the MBR. For systems with Windows 95/98, you should also install the boot loader to the MBR so that it can boot both operating systems.

Caution

> If you choose not to install GRUB or LILO for any reason, you will not be able to boot your Red Hat Linux system directly, and you will need to use another boot method (such as a boot diskette). Use this option only if you are sure you have another way of booting your Red Hat Linux system!

If you wish to add default options to GRUB or LILO's boot command, enter them into the **Kernel parameters** field. Any options you enter will be passed to the Linux kernel every time it boots.

The **Force use of LBA32 (not normally required)** option allows you to exceed the 1024 cylinder limit for the /boot partition. If you have a system which supports the LBA32 extension for booting operating systems above the 1024 cylinder limit, and you want to place your /boot partition above cylinder 1024, you should select this option.

Every bootable partition is listed, including partitions used by other operating systems. The partition holding your Red Hat Linux system's root file system will have a **Boot label** of Red Hat Linux (for GRUB) or linux (for LILO). Other partitions may also have boot labels. If you would like to add boot labels for other partitions (or change an existing boot label), click once on the partition to select it. Once selected, you can change the boot label.

Note

> The **Boot label** column lists what you must enter at the boot prompt, in non-graphical boot loaders, in order to boot the desired operating system. However, if you forget the boot labels defined on your system, you can always press [Tab] at the prompt to display a list of defined boot labels.

A.5.1.1. Rescue Mode

If you need to use rescue mode, there are several options available to you.

- Using the CD-ROM to boot, type **`linux rescue`** at the `boot:` prompt.
- By booting your system from an installation boot diskette made from the `boot.img` image. This method requires that the Red Hat Linux CD-ROM #1 be inserted as the rescue image or that the rescue image be on the hard drive as an ISO image. Once you have booted using this diskette, type **`linux rescue`** at the `boot:` prompt.
- By booting from a network disk made from the `bootnet.img` or PCMCIA boot disk made from `pcmcia.img`. Once you have booted using this diskette, type **`linux rescue`** at the `boot:` prompt. You can only do this if your network connection is working. You will need to identify the network host and transfer type. For an explanation of how to specify this information, see Section 3.7.

For more information regarding rescue mode, refer to the *Official Red Hat Linux Customization Guide*.

A.5.1.2. Alternative Boot Loaders

If you do not wish to use GRUB or LILO to boot your Red Hat Linux system, you have several alternatives:

Boot disk

> As previously stated, you can use the boot disk created by the installation program (if you elected to create one).

LOADLIN

> You can load Linux from MS-DOS. Unfortunately, this 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 (for example, from a boot disk) and then copy the kernel to an MS-DOS partition. LOADLIN is available from
>
> `ftp://metalab.unc.edu/pub/Linux/system/boot/dualboot/`
>
> and associated mirror sites.

SYSLINUX

> SYSLINUX is an MS-DOS program very similar to LOADLIN. It is also available from
>
> `ftp://metalab.unc.edu/pub/Linux/system/boot/loaders/`
>
> and associated mirror sites.

Some commercial boot loaders

> You can load Linux using commercial boot loaders. For example, System Commander and Partition Magic are able to boot Linux (but still require GRUB or LILO to be installed in your Linux root partition).

A.5.1.3. SMP Motherboards, GRUB, and LILO

This section is specific to SMP motherboards only. SMP, short for Symmetric Multiprocessing, is a computer architecture providing fast performance by making multiple CPUs available to complete individual processes simultaneously (multiprocessing).

If the installation program detects an SMP motherboard on your system, it will automatically create two `/boot/grub/grub.conf` or `/etc/lilo.conf` entries (depending on the boot loader you installed), rather than the usual single entry.

The two entries in `grub.conf` will be `Red Hat Linux (`*`kernel version`*`)` and `Red Hat Linux (`*`kernel version`*`-smp`). The `Red Hat Linux (`*`kernel version`*`-smp`) will boot by default. However, if you have trouble with the SMP kernel, you can elect to boot the `Red Hat Linux (`*`kernel version`*`)` entry instead. You will retain all the functionality as before, but you will only be operating with a single processor.

The two entries in `lilo.conf` will be `linux` and `linux-up`. The `linux` entry will boot by default. However, if you have trouble with the SMP kernel, you can elect to boot the `linux-up` entry instead. You will retain all the functionality as before, but you will only be operating with a single processor.

A.6. GRUB Password

If you did not install GRUB as your boot loader, skip to Section A.7.

If you are installing GRUB as your boot loader, you should create a password to protect your system. Without a GRUB password, users with access to your system can pass options to the kernel which can compromise your system security. With a GRUB password in place, the password must first be entered in order to select any non-standard boot options.
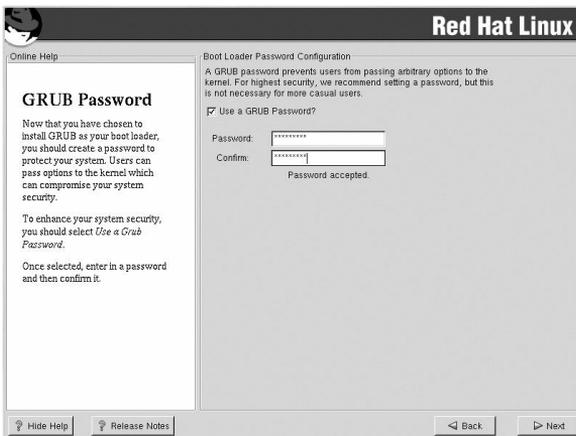

Figure A-5. GRUB Password

To enhance your system security, you should select **Use a Grub Password**.

Once selected, enter a password and confirm it.

A.7. Selecting Packages to Upgrade

On this screen, you can choose which packages you would like to upgrade (see Figure A-6).

You can choose to view the individual packages in **tree view** or **flat view**.

Tree view allows you to see the packages grouped by application type.

Flat view allows you to see all of the packages in an alphabetical listing on the right of the screen.

Using **Tree view**, you see a listing of package groups. When you expand this list (by double-clicking on the folder arrow beside a package group name) and pick one group, the list of packages in that group appears in the panel on the right.

To sort packages alphabetically, click on the **Package** tab. To sort packages by size, click on the **Size (MB)** tab.

To select an individual package, click the checkbox beside the package name. A check mark in the box means that a package has been selected.

For more information about a specific package, click on the individual package name. The package information will appear at the bottom of the screen.

You can also select or deselect all packages listed within a particular group, by clicking on the **Select all in group** or **Unselect all in group** buttons.

Certain packages (for example, the kernel and certain libraries) are required for every Red Hat Linux system and are not available to select or deselect. These base packages are selected by default.
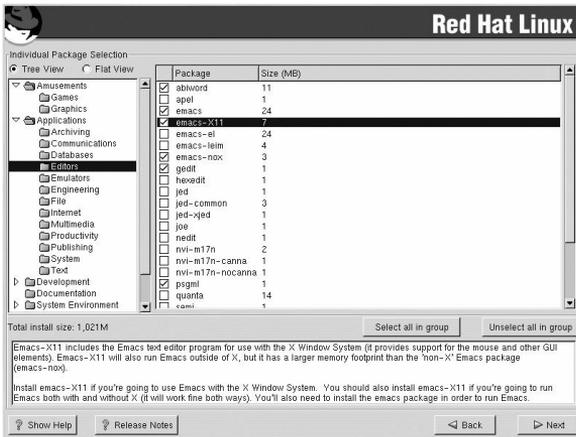

Figure A-6. Individual Package Selection

A.7.1. Unresolved 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 A-7).

The **Unresolved Dependencies** screen appears only if you are missing packages that are needed by your customized package selection.

At the bottom of the screen, under the list of missing packages, an **Install packages to satisfy dependencies** radio button is selected by default. If you leave this checked, the installation program will resolve package dependencies automatically by adding all required packages to the list of selected packages.
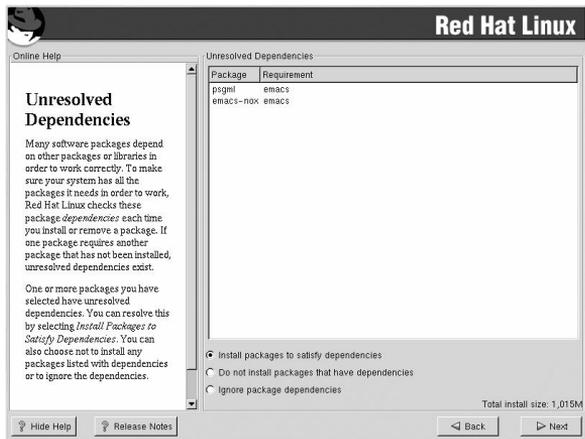

Figure A-7. Unresolved Dependencies

If you do not wish to install packages that require other packages, select **Do not install packages that have dependencies**.

To install only the packages you have selected and leave the dependencies unresolved, select **Ignore package dependencies**.

A.8. Upgrading Packages

At this point there is nothing left for you to do until all the packages have been upgraded or installed (see Figure A-8).
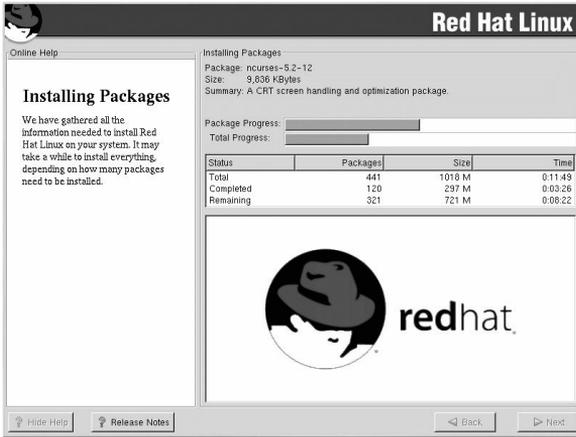
Figure A-8. Installing Packages

A.9. Boot Disk Creation

To create a boot disk, insert a blank, formatted diskette into your diskette drive (see Figure A-9) and click **Next**.

It is highly recommended that you create a boot disk. If, for some reason, your system were not able to boot properly using GRUB, LILO, or a third-party boot loader, a boot disk would enable you to properly boot your Red Hat Linux system.

After a short delay, your boot disk will be created; remove it from your diskette drive and label it clearly. Note that if you would like to create a boot disk after the installation, you will be able to do so. For more information, please see the mkbootdisk man page, by typing man mkbootdisk at the shell prompt.

If you do not want to create a boot disk, make sure to check the **Skip boot disk creation** checkbox before you click **Next**.

If you boot your system with the boot disk (instead of GRUB or LILO) , make sure you create a new boot disk whenever you make any changes to your kernel (including the installation of a new kernel).

Figure A-9. Boot Disk Creation

A.10. Upgrade Complete

Congratulations! Your Red Hat Linux 7.3 upgrade is now complete!

You will now be prompted to prepare your system for reboot. Do not forget to remove any diskette in the floppy drive or CD in the CD-ROM drive if they are not ejected automatically upon reboot. If you do not have a boot loader installed and configured, you will need to use your boot disk now.

Tip

If you need a quick review of some of the basic Linux concepts of Red Hat Linux refer to the *Official Red Hat Linux Getting Started Guide*.

For information dealing with system configuration and administration, refer to the *Official Red Hat Linux Customization Guide* and the *Official Red Hat Linux Reference Guide*.

Appendix B.

Removing Red Hat Linux

To uninstall Red Hat Linux from your system, you will need to remove the GRUB or LILO information from your master boot record (MBR).

In DOS, NT, and Windows 95 you can use **fdisk** to create a new MBR with the *undocumented* flag `/mbr`. This will *ONLY* rewrite the MBR to boot the primary DOS partition. The command should look like the following:

```
fdisk /mbr
```

If you need to remove Linux from a hard drive, and have attempted to do this with the default DOS **fdisk**, you will experience the *Partitions exist but they do not exist* problem. The best way to remove non-DOS partitions is with a tool that understands partitions other than DOS.

You can do this with the installation media by typing **linux expert** at the `boot:` prompt:

```
boot: linux expert
```

Select install (versus upgrade) and at the point when you should partition the drive, choose **fdisk**. In **fdisk**, type [p] to print out the partition numbers, and remove the Linux partitions with the [d] command. When you are satisfied with the changes you have made, you can quit with a [w] and the changes will be saved to disk. If you deleted too much, type [q] and no changes will be made.

Once you have removed the Linux partitions, you can reboot your computer using [Control]-[Alt]-[Delete] instead of continuing with the install.

redhat.

Appendix C.
Getting Technical Support

C.1. Remember to Sign Up

If you have an official edition of Red Hat Linux 7.3 and/or an official Red Hat OEM partner kit, please remember to sign up for the benefits you are entitled to as a Red Hat customer.

You will be entitled to any or all of the following benefits, depending upon the Official Red Hat Linux product you purchased:

- Official Red Hat support — Get help with your installation questions from Red Hat's support team.
- Red Hat Network — Easily update your packages and receive security notices that are customized for your system. For more details, go to:

  `http://rhn.redhat.com`

- Under the Brim: The Official Red Hat E-Newsletter — Every month, get the latest news and product information directly from Red Hat.

To sign up, go to:

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

You will find your *Product ID* on the Registration Information Card in your Official Red Hat Linux boxed set.

C.2. An Overview of Red Hat Support

Note

For more information on how Red Hat's technical support staff can assist you, refer to the service level agreement at:

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

Red Hat provides installation assistance for Official Red Hat Linux boxed set products and covers installation on a single computer. This assistance is intended to help customers successfully install Red Hat Linux. Assistance with installation is offered via telephone and the Web. Note, telephone support is only available with certain Red Hat Linux products. Please check your boxed set to see what types of support are available to you.

Red Hat Support will attempt to answer any questions you may have before the installation process is initiated. Depending on the product purchased, it can include the following:

- Hardware compatibility questions

- Basic hard drive partitioning strategies

Red Hat Support can provide the following assistance during the installation process:

- Getting supported hardware recognized by the Red Hat Linux operating system
- Assistance with hard drive partitioning
- Configuring Red Hat Linux and Windows 9*x*, ME, NT, or 2000 to dual-boot using the Linux boot loaders GRUB or LILO (please note that third party boot loaders and partitioning software are not supported)

We can also help you with basic post-installation tasks, such as:

- Successfully configuring the X Window System using **Xconfigurator**
- Configuring a local parallel port printer to print text
- Configuring a mouse

Our installation assistance service is designed to get Red Hat Linux running on your system as quickly and as easily as possible. However, there are many other things that you may want to do with your Red Hat Linux system, from compiling a custom kernel to setting up a Web server, which are not covered.

For assistance with these tasks, there is a wealth of online information available in the form of HOWTO documents, Linux-related websites, and commercial publications. The various Linux HOWTO documents are included with Red Hat Linux on the Documentation CD in the /HOWTOS directory. These HOWTOS are provided in text files that can easily be read from within Red Hat Linux and other operating systems.

A large number of Linux-related websites are available. The best starting point for finding information on Red Hat Linux is the Red Hat, Inc. website:

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

Many Linux-related books are also available. If you are new to Linux, a book that covers Linux basics will be invaluable. We can recommend several titles: The *Official Red Hat Linux Getting Started Guide*; *Using Linux*, by Bill Ball; *Linux Clearly Explained*, by Bryan Pfaffenberger; *Linux for Dummies*, by Jon "maddog" Hall; and *Learning Red Hat Linux*, by Bill McCarty.

Red Hat also offers various incident-based support plans to assist with configuration issues and tasks that are not covered by installation assistance. Please see the Red Hat Support website for more information at:

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

C.3. Scope of Red Hat Support

Red Hat, Inc. can only provide installation assistance to customers who have purchased an Official Red Hat Linux boxed set and/or an official Red Hat OEM partner kit. If you have obtained Linux from any other company, you must contact that company for support. Other companies include:

- Macmillan
- Sams/Que
- Linux Systems Labs (LSL)

- Mandrake
- CheapBytes

Additionally, Red Hat Linux obtained via any of the following methods does not qualify for support from Red Hat:

- Red Hat Linux PowerTools Archive
- Downloaded via FTP on the Internet
- Included in a package such as Motif or Applixware
- Copied or installed from another user's CD
- A CD-ROM (or CD-ROM set) included in a Linux book or other publication.

C.4. How to Get Technical Support

In order to receive technical support for your Official Red Hat product, you must register your product on Red Hat's website.

Every Official Red Hat product comes with a Product Identification code: a 16-character alphanumeric string. The Product ID for Red Hat Linux 7.3 is located on the Registration Information Card that can be found inside the box.

Note

Do not throw away the card with your Product ID. You must have the Product ID to get technical support. If you lose the certificate, you may not be able to receive support.

The Product ID is the code that will enable technical support and any other benefits or services that you purchased from Red Hat, depending upon which product you purchased.

C.4.1. Signing up for Technical Support

To sign up for technical support, you must:

1. Create a customer profile at:

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

   You may have already completed this step; if you have, continue to the next step.

2. Using the login name and password you created during the customer profile, please log in at the Red Hat Support website at:

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

   If you created a new customer profile, once you activate your product you will see a webpage that shows your registered products. There is also a button, **Access Web Support**, on this page that will take you to the support website.

3. Update your contact information if necessary.

Note

If your email address is not correct, communications regarding your technical support requests CANNOT be delivered to you, and you will not be able to retrieve your login and password by email. Be sure that you give us your correct email address.

If you are worried about your privacy, please see Red Hat's privacy statement at:

```
http://www.redhat.com/legal/privacy_statement.html
```

4. Add a product to your profile. Please enter the following information:

   • The Product ID for your boxed set product

   • The Support Certificate Number or Entitlement Number if the product is a contract

5. Set your customer preferences.

6. Answer the optional customer questionnaire.

7. Submit the form.

If the previous steps were completed successfully, you can now login at

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

and open a new technical service request. However, you must still use your Product ID in order to obtain technical support via telephone (if the product you purchased came with phone support). You will also be asked for your login name when contacting the technical support team via telephone.

C.5. Questions for Technical Support

Technical support is both a science and a mystical art form. In most cases, support technicians must rely on customer observations and communications with the customer in order to diagnose and solve the problem. Therefore, it is extremely important that you are as detailed and clear as possible when you state your questions and report your problems. Examples of what you should include are:

• Symptoms of the problem (for example: "Linux is not able to access my CD-ROM drive. When it tries, I get timeout errors.")

• When the problem began (for example: "My system was working fine until yesterday, when a lightning storm hit my area.")

• Any changes you made to your system (for example: "I added a new hard drive and used **Partition Wizzo** to add Linux partitions.")

• Other information that may be relevant to your situation, such as the installation method (CD-ROM, NFS, HTTP)

• Specific hardware devices that may be relevant to your problem (for example: If you cannot setup networking, what kind of network card do you have?)

C.5.1. How to Send Support Questions

Please login at

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

and open a new service request, or call the phone number for support. If your product came with phone support, or you have purchased a phone support contract, the phone number you will need to call will be provided to you during the sign up process.

For more information on using Red Hat's online support system go to:

```
http://www.redhat.com/support/services/access.html
```

redhat.

<div align="right">

Appendix D.

</div>

Troubleshooting Your Installation of Red Hat Linux

This appendix discusses some common installation problems and their solutions.

D.1. You are Unable to Boot Red Hat Linux

D.1.1. Are You Unable to Boot from the CD-ROM?

Note

There are a few cases where the system BIOS will not allow the Red Hat Linux CD-ROM to boot because of the size of the boot image on the CD-ROM itself. In cases such as these, a boot disk should be made to boot Red Hat Linux. Once booted, the CD-ROMs will work properly for the installation.

If you cannot boot from your Red Hat Linux CD-ROM, you have two options:

1. You can change your BIOS so that booting from the CD-ROM is recognized first in the boot order, or

2. Boot using a boot disk you have created.

To change your BIOS, refer to your system manual for the correct keyboard combination that allows you to access your BIOS, or you can read the key sequence needed while the system begins to boot.

To create a boot disk, follow the instructions in Section 1.4.2.

To boot Red Hat Linux using a boot disk, insert the diskette you have created into your floppy drive and then boot/reboot your computer. Make sure that your BIOS is set to use the floppy or removable disk (A:) to boot.

D.1.2. Are You Unable to Boot from the Local Boot Disk?

If you are experiencing difficulties in getting a local boot disk to boot your system correctly, you may need an updated boot disk.

Check the online errata at

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

for updated diskette images (if available) and follow the instructions in Section 1.4.2, to make an updated boot disk for your system.

D.1.3. Are You Unable to Boot from PCMCIA Boot Disks?

If you are experiencing difficulties in getting the PCMCIA boot disks you made to boot your system correctly, you may need an updated boot disk.

Check the online errata for updated diskette images (if available) and follow the instructions provided to make an updated boot disk for your system.

D.1.4. Is Your System Displaying Signal 11 Errors?

If you receive a fatal signal 11 error during your installation, it is probably due to a hardware error in memory on your system's bus. A hardware error in memory can be caused by problems in executables or with the system's hardware. Like other operating systems, Red Hat Linux places its own demands on your system's hardware. Some of this hardware may not be able to meet those demands, even if they work properly under another OS.

Check to see if you have the latest installation and supplemental boot diskettes from Red Hat. Review the online errata to see if newer versions are available. If the latest images still fail, it may be due to a problem with your hardware. Commonly, these errors are in your memory or CPU-cache. A possible solution for this error is turning off the CPU-cache in the BIOS. You could also try to swap your memory around in the motherboard slots to see if the problem is either slot or memory related.

For more information concerning signal 11 errors, refer to:

```
http://www.bitwizard.nl/sig11/
```

D.1.5. Are You Unable to Boot from a Network Boot Disk?

If you are experiencing difficulties in getting the network boot disk you made to boot your system correctly, you may need an updated boot disk.

Check the online errata for updated diskette images (if available) and follow the instructions provided to make an updated boot disk for your system.

D.2. Trouble Beginning the Installation

D.2.1. Trouble Using PCMCIA Boot Disks?

If you booted using PCMCIA boot disks and want to install via FTP (or NFS or HTTP), but do not see these installation options, you may have a problem with your network card.

If the network card is not initialized during the boot process, the Red Hat Linux installation program will not enable you to configure your system for networking, either during or after the installation itself.

Check the hardware compatibility list at

```
http://hardware.redhat.com/hcl/
```

to see if your network card is compatible and/or supported. If your card is not listed, it may not be compatible with Red Hat Linux.

D.2.2. Is Your Mouse Not Detected?

If the **Mouse Not Detected** screen (see Figure D-1) appears, then the installation program was not able to identify your mouse correctly.

You can choose to continue with the GUI installation or use the text mode installation, which does not require using a mouse. If you choose to continue with the GUI installation, you will need to provide the installation program with your mouse configuration information (see Section 3.14).

Figure D-1. Mouse Not Detected

For an overview of text mode installation instructions, please refer to Section 3.3.

D.2.3. Problems with Booting into the Graphical Installation

The Red Hat Linux installation program uses frame buffers by default. However, there are some video cards that will not work with this setting. The end result will be a problem booting into the graphical installation program.

The installation program will first try to run in frame buffer mode. If that fails, it will try to run in a lower resolution mode. If that still fails, the installation program will run in text mode.

Users who have video cards that will not run at 800 x 600 resolution should type **lowres** at the `boot:` prompt to run the installation program in 640 x 480 resolution.

If this still does not work, you can run the installation program without frame buffers by typing **nofb** at the `boot:` prompt.

D.3. Trouble During the Installation

D.3.1. Partition Creation Problems

If you are having trouble creating a partition (for example, a root (`/`) partition), make sure you are setting its partition type to Linux Native.

Unless your BIOS supports otherwise, make sure /boot does not exceed the 1023 cylinder head. If you do not, the installation program will not allow you to create a /boot or / partition. Some new systems allow you to exceed the 1023 limit (with GRUB and the newer LILO versions that are available), but most machines with older BIOS will not.

D.3.2. Using Remaining Space

You have a swap and a / (root) partition created, and you have selected the root partition to use the remaining space, but it does not fill the hard drive.

If your hard drive is more than 1024 cylinders, you must create a /boot partition if you want the / (root) partition to use all of the remaining space on your hard drive.

D.3.3. Other Partitioning Problems

If you are using **Disk Druid** to create partitions, but cannot move to the next screen, you probably have not created all the partitions necessary for **Disk Druid**'s dependencies to be satisfied.

You must have the following partitions as a bare minimum:

- A /boot partition of type Linux native
- A / (root) partition of type Linux native
- A <swap> partition of type Linux swap

Tip

> When defining a partition's type as Linux swap, you do not have to assign it a mount point. **Disk Druid** automatically assigns the mount point for you.

D.3.4. Are You Seeing Python Errors?

During some upgrades or installations of Red Hat Linux, the installation program (also known as **Anaconda**) may fail with a Python or traceback error. This error may occur after the selection of individual packages or while trying to save the upgrade log in /tmp. The error may look similar to:

```
Traceback (innermost last):
   File "/var/tmp/anaconda-7.1//usr/lib/anaconda/iw/progress_gui.py",
line 20, in run
     rc = self.todo.doInstall ()
   File "/var/tmp/anaconda-7.1//usr/lib/anaconda/todo.py", line 1468, in
doInstall
     self.fstab.savePartitions ()
   File "fstab.py", line 221, in savePartitions
     sys.exit(0)
 SystemExit: 0

 Local variables in innermost frame:
 self: <fstab.GuiFstab instance at 8446fe0>
 sys: <module 'sys' (built-in)>
```

```
ToDo object:
(itodo
ToDo
p1
(dp2
S'method'
p3
(iimage
CdromInstallMethod
p4
(dp5
S'progressWindow'
p6

<failed>
```

This error occurs in some systems where links to /tmp are symbolic to other locations or have been changed since creation. These symbolic or changed links are invalid during the installation process, so the installation program cannot write information and fails.

If you experience such an error, first try to download any available errata for **Anaconda**. Errata can be found at:

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

You can also search for bug reports related to this problem. To search Red Hat's bug tracking system, go to:

```
http://bugzilla.redhat.com/bugzilla
```

Finally, if you are still facing problems related to this error, register your product and contact our support team. To register your product, go to:

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

D.4. Problems After Installation

D.4.1. Trouble With the Graphical GRUB Screen?

If, for some reason, you need to disable the graphical boot screen, you can do so, as root, by editing the /boot/grub/grub.conf file and then rebooting your system.

To do this, comment out the line which begins with splashimage in the grub.conf file. To comment out a line, insert the ; character at the beginning of the line.

Once you reboot, the grub.conf file will be reread and your changes will take place.

You may re-enable the graphical boot screen by uncommenting (or adding) the above line back into the grub.conf file and rebooting.

D.4.2. Trouble With the Graphical LILO Screen?

If, for some reason, you need to disable the graphical boot screen, you can do so, as root, by editing the /etc/lilo.conf file and then rerunning LILO.

First, as root, comment out (or delete) the line which reads `message=/boot/message` in the `/etc/lilo.conf` file. To comment out a line, insert the `;` character at the beginning of the line. Next, rerun LILO by typing `/sbin/lilo -v`. The next time you boot, you will see the text `LILO:` prompt, as used in previous Red Hat Linux releases.

You may re-enable the graphical boot screen by adding (or uncommenting) the above line back into the `lilo.conf` file and rerunning LILO.

D.4.3. Problems with Server Installations and X

If you performed a server installation and you are having trouble getting X to start, you may not have installed the X Window System during your installation.

If you want the X Window System, you can perform an upgrade to install X. During the upgrade, select the X Window System packages, and choose GNOME, KDE, or both.

D.4.4. Problems When You Try to Log In

If you did not create a user account during the installation you will need to log in as root and use the password you assigned to root.

If you cannot remember your root password, you will need to boot your system as `linux single` at the GRUB or LILO `boot:` prompt. Then at the `#` prompt, you will need to type `passwd root`, which will allow you to enter a new password for root. At this point you can type `shutdown -r now` and the system will reboot with your new password.

If you cannot remember your user account password, you must become root. To become root, type `su -` and enter your root password when prompted. Then, type `passwd <username>`. This allows you to enter a new password for the specified user account.

If you selected either the custom or workstation installation and do not see the graphical login screen, check your hardware for compatibility issues. The *Hardware Compatibility List* can be found at:

```
http://hardware.redhat.com/hcl/
```

D.4.5. Does Netscape Navigator Crash on JavaScript Pages?

If Netscape **Navigator** continuously crashes on pages that contain JavaScripts, you may need to edit your `~/.mailcap` file.

Edit the file using **pico** by typing `pico ~/.mailcap` at the prompt in a terminal window. (You may use any text editor.)

Remove the following lines from the file:

```
application/x-javascript;;\
x-mozilla-flags=save
```

You can also turn off JavaScript within Netscape **Navigator** itself. Click on **Edit**=>**Preferences**=>**Advanced** and make sure the **Enable JavaScript** checkbox is not selected.

If these do not help, you can also try to use a newer version of Netscape **Navigator** if available. Check the Red Hat Linux errata website under security advisories for more information.

D.4.6. Your Printer Will Not Work Under X

If you are not sure how to set up your printer or are having trouble getting it to work properly, try using the graphical **printconf** program. Log in as root, open a terminal window, and type `printconf-gui`.

D.4.7. Is Your RAM Not Being Recognized?

Sometimes, the kernel does not recognize all of your memory (RAM). You can check this with the following command:

```
cat /proc/meminfo
```

Find out if the displayed quantity is the same as the known amount of RAM in your system. If they are not equal, add the following line to the `/boot/grub/grub.conf`:

mem=*xx*M

Or, if you used LILO, add the following line to the `/etc/lilo.conf` file:

append="mem=*xx*M"

Replace *xx* with the amount of RAM you have in megabytes. Remember that per-image append lines completely overwrite the global append line. It might be worth adding this to the per-image descriptions, as shown in this example:

mem=128M

append="mem=128M"

In `/boot/grub/grub.conf`, the above example would look similar to the following:

```
#NOTICE: You have a /boot partition. This means that
#        all kernel paths are relative to /boot/
default=0
timeout=30
splashimage=(hd0,0)/grub/splash.xpm.gz
title Red Hat Linux (2.4.6-2)
        root (hd0,0)
        kernel /vmlinuz-2.2.6-2 ro root=/dev/hda3 mem=128M
```

Once you reboot, the changes made to `grub.conf` will be reflected on your system.

In `/etc/lilo.conf`, the above example would look similar to the following:

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

     image=/boot/vmlinuz-2.2.12-20
            label=linux
            root=/dev/sda1
            initrd=/boot/initrd-2.2.12-20.img
            read-only
            append="mem=128M"
```

Remember to run `/sbin/lilo -v` after changing `/etc/lilo.conf`.

Note that you can also produce the same effect by actually passing this option when you are specifying the label/image to use in GRUB or LILO.

Once you have loaded the GRUB boot screen, type `e` for edit. You will be presented with a list of items in the configuration file for the boot label you have selected.

Choose the line that starts with `kernel` and type `e` to edit this boot entry.

At the end of the `kernel` line, add

```
mem=xxM
```

where *xx* equals the amount of RAM in your system.

Press [Enter] to exit edit mode.

Once the GRUB screen has returned, type `b` to boot with your new RAM specifications.

At the graphical LILO screen, press [Ctrl]-[x] to exit to the `boot:` prompt. Next, enter the following at the `boot:` prompt:

```
linux mem=xxM
```

Remember to replace *xx* with the amount of RAM in your system. Press [Enter] to boot.

D.4.8. Problems with Sound Configuration

If you do not have sound after your installation, you may need to run the sound configuration utility. As root, type `sndconfig` in a terminal window.

Note

> `sndconfig` must be run in runlevel 3. More information about runlevels can be found in the *Official Red Hat Linux Reference Guide*, in chapter *Boot Process, Init, and Shutdown*.

If the **sndconfig** application does not help, you may need to select the `Enable sound server startup` option under the **Multimedia**=>**Sound** in the **GNOME Control Center**.

To do this, click on **Panel**=>**Programs**=>**Settings**=>**GNOME Control Center** to launch the **GNOME Control Center**. In the **GNOME Control Center**, select the **Sound** submenu of the **Multimedia** menu. On the right, a **General** sound menu will appear. Select `Enable sound server startup` and then click **OK**.

redhat.

Appendix E.

An Introduction to Disk Partitions

Disk partitions are a standard part of the personal computer landscape and have been for quite some time. However, with many people purchasing computers featuring pre-installed operating systems, relatively few people understand how partitions work. This chapter attempts to explain the reasons for and use of disk partitions so your Red Hat Linux installation will be as simple and painless as possible.

If you are reasonably comfortable with disk partitions, you could skip ahead to Section E.1.4, for more information on the process of freeing up disk space to prepare for a Red Hat Linux installation. This section also discusses the partition naming scheme used by Linux systems, sharing disk space with other operating systems, and related topics.

E.1. Hard Disk Basic Concepts

Hard disks perform a very simple function — they store data and reliably retrieve it on command.

When discussing issues such as disk partitioning, it is important to know a bit about the underlying hardware. Unfortunately, it is easy to become bogged down in details. Therefore, we will use a simplified diagram of a disk drive to help explain what is really happening when a disk drive is partitioned. Figure E-1, shows a brand-new, unused disk drive.

Figure E-1. An Unused Disk Drive

Not much to look at, is it? But if we are talking about disk drives on a basic level, it will do. Say that we would like to store some data on this drive. As things stand now, it will not work. There is something we need to do first. . .

E.1.1. It is Not What You Write, it is How You Write It

Experienced computer users probably got this one on the first try. We need to *format* the drive. Formatting (usually known as "making a *file system*") writes information to the drive, creating order out of the empty space in an unformatted drive.
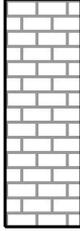
Figure E-2. Disk Drive with a File System

As Figure E-2, implies, the order imposed by a file system involves some trade-offs:

- A small percentage of the drive's available space is used to store file system-related data and can be considered as overhead.
- A file system splits the remaining space into small, consistently-sized segments. For Linux, these segments are known as *blocks*. [1]

Given that file systems make things like directories and files possible, these tradeoffs are usually seen as a small price to pay.

It is also worth noting that there is no single, universal file system. As Figure E-3, shows, a disk drive may have one of many different file systems written on it. As you might guess, different file systems tend to be incompatible; that is, an operating system that supports one file system (or a handful of related file system types) may not support another. This last statement is not a hard-and-fast rule, however. For example, Red Hat Linux supports a wide variety of file systems (including many commonly used by other operating systems), making data interchange between different file systems easy.

---

1.   Blocks really *are* consistently sized, unlike our illustrations. Keep in mind, also, that an average disk drive contains thousands of blocks. But for the purposes of this discussion, please ignore these minor discrepancies.

Figure E-3. Disk Drive with a Different File System

Of course, writing a file system to disk is only the beginning. The goal of this process is to actually *store* and *retrieve* data. Let us take a look at our drive after some files have been written to it.
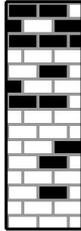
Figure E-4. Disk Drive with Data Written to It

As Figure E-4, shows, 14 of the previously-empty blocks are now holding data. However, by simply looking at this picture, we cannot determine exactly how many files reside on this drive. There may be as few as one or as many as 14 files, as all files use at least one block and some files use multiple blocks. Another important point to note is that the used blocks do not have to form a contiguous region; used and unused blocks may be interspersed. This is known as *fragmentation*. Fragmentation can play a part when attempting to resize an existing partition.

As with most computer-related technologies, disk drives changed over time after their introduction. In particular, they got bigger. Not larger in physical size, but bigger in their capacity to store information. And, this additional capacity drove a fundamental change in the way disk drives were used.

E.1.2. Partitions: Turning One Drive Into Many

As disk drive capacities soared, some people began to wonder if having all of that formatted space in one big chunk was such a great idea. This line of thinking was driven by several issues, some philosophical, some technical. On the philosophical side, above a certain size, it seemed that the additional space provided by a larger drive created more clutter. On the technical side, some file systems were never designed to support anything above a certain capacity. Or the file systems *could* support larger drives with a greater capacity, but the overhead imposed by the file system to track files became excessive.

The solution to this problem was to divide disks into *partitions*. Each partition can be accessed as if it was a separate disk. This is done through the addition of a *partition table*.

Note

While the diagrams in this chapter show the partition table as being separate from the actual disk drive, this is not entirely accurate. In reality, the partition table is stored at the very start of the disk, before any file system or user data. But for clarity, we will keep it separate in our diagrams.
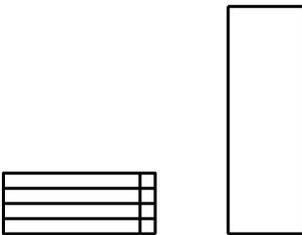
Figure E-5. Disk Drive with Partition Table

As Figure E-5, shows, the partition table is divided into four sections. Each section can hold the information necessary to define a single partition, meaning that the partition table can define no more than four partitions.

Each partition table entry contains several important characteristics of the partition:

- The points on the disk where the partition starts and ends
- Whether the partition is "active"
- The partition's type

Let us take a closer look at each of these characteristics. The starting and ending points actually define the partition's size and location on the disk. The "active" flag is used by some operating systems' boot loaders. In other words, the operating system in the partition that is marked "active" will be booted.

The partition's type can be a bit confusing. The type is a number that identifies the partition's anticipated usage. If that statement sounds a bit vague, that is because the meaning of the partition type is a bit vague. Some operating systems use the partition type to denote a specific file system type, to flag the partition as being associated with a particular operating system, to indicate that the partition contains a bootable operating system, or some combination of the three.

Table E-1, contains a listing of some popular (and obscure) partition types, along with their numeric values.

Table E-1. Partition Types

Partition Type	Value	Partition Type	Value
Empty	00	Novell Netware 386	65
DOS 12-bit FAT	01	PIC/IX	75
XENIX root	02	Old MINIX	80
XENIX usr	03	Linux/MINUX	81
DOS 16-bit <=32M	04	Linux swap	82
Extended	05	Linux native	83
DOS 16-bit >=32	06	Linux extended	85
OS/2 HPFS	07	Amoeba	93
AIX	08	Amoeba BBT	94
AIX bootable	09	BSD/386	a5
OS/2 Boot Manager	0a	OpenBSD	a6
Win95 FAT32	0b	NEXTSTEP	a7
Win95 FAT32 (LBA)	0c	BSDI fs	b7
Win95 FAT16 (LBA)	0e	BSDI swap	b8
Win95 Extended (LBA)	0f	Syrinx	c7
Venix 80286	40	CP/M	db
Novell	51	DOS access	e1
Microport	52	DOS R/O	e3
GNU HURD	63	DOS secondary	f2
Novell Netware 286	64	BBT	ff

By this point, you might be wondering how all this additional complexity is normally used. See Figure E-6, for an example.
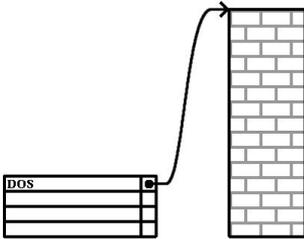
Figure E-6. >Disk Drive With Single Partition

In many cases, there is only a single partition spanning the entire disk, essentially duplicating the method used before partitions. The partition table has only one entry used, and it points to the start of the partition.

We have labeled this partition as being of the "DOS" type. Although it is only one of several possible partition types listed in Table E-1, it is adequate for the purposes of this discussion. This is a typical partition layout for most newly purchased computers with a consumer version of Microsoft Windows™ preinstalled.

E.1.3. Partitions within Partitions — An Overview of Extended Partitions

Of course, over time it became obvious that four partitions would not be enough. As disk drives continued to grow, it became more and more likely that a person could configure four reasonably-sized partitions and still have disk space left over. There needed to be some way of creating more partitions.

Enter the extended partition. As you may have noticed in Table E-1, there is an "Extended" partition type. It is this partition type that is at the heart of extended partitions.

When a partition is created and its type is set to "Extended," an extended partition table is created. In essence, the extended partition is like a disk drive in its own right — it has a partition table that points to one or more partitions (now called *logical partitions*, as opposed to the four *primary partitions*) contained entirely within the extended partition itself. Figure E-7, shows a disk drive with one primary partition and one extended partition containing two logical partitions (along with some unpartitioned free space).
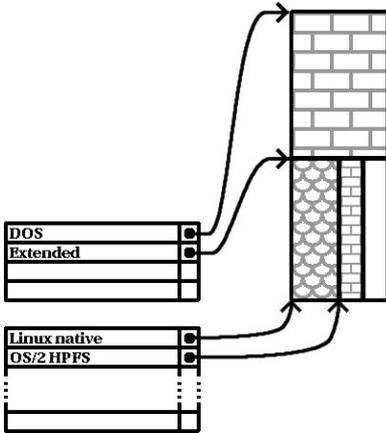
Figure E-7. Disk Drive With Extended Partition

As this figure implies, there is a difference between primary and logical partitions — there can only be four primary partitions, but there is no fixed limit to the number of logical partitions that can exist. (However, in reality, it is probably not a good idea to try to define and use more than 12 logical partitions on a single disk drive.)

Now that we have discussed partitions in general, let us see how to use this knowledge to install Red Hat Linux.

E.1.4. Making Room For Red Hat Linux

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 us look at each scenario in order.

Note

Please keep in mind that the following illustrations are simplified in the interest of clarity and do not reflect the exact partition layout that you will encounter when actually installing Red Hat Linux.

E.1.4.1. Using Unpartitioned Free Space

In this situation, the partitions already defined do not span the entire hard disk, leaving unallocated space that is not part of any defined partition. Figure E-8, shows what this might look like.

Before | *After*
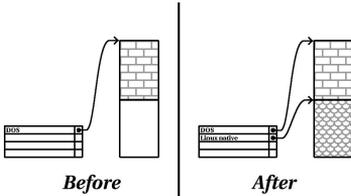
Figure E-8. Disk Drive with Unpartitioned Free Space

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 any case, you can simply create the necessary partitions from the unused space. Unfortunately, this scenario, although very simple, is not very likely (unless you have just purchased a new disk just for Red Hat Linux). Most pre-installed operating systems are configured to take up all available space on a disk drive (see Section E.1.4.3).

Next, we will discuss a slightly more common situation.

E.1.4.2. Using Space from an Unused Partition

In this case, maybe you have one or more partitions that you do not use any longer. Perhaps you have dabbled with another operating system in the past, and the partition(s) you dedicated to it never seem to be used anymore. Figure E-9, illustrates such a situation.

Before | *After*
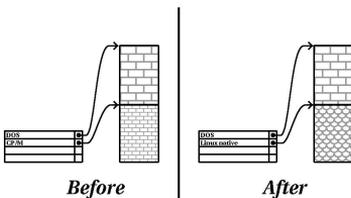
Figure E-9. Disk Drive With an Unused Partition

If you find yourself in this situation, you can use the space allocated to the unused partition. You will first need to delete the partition, and then create the appropriate Linux partition(s) in its place. You can either delete the partition using the DOS `fdisk` command, or you will be given the opportunity to do so during a custom installation.

E.1.4.3. Using Free Space from an Active Partition

This is the most common situation. It is also, unfortunately, the hardest to handle. The main problem is that, even if you have enough free space, it is presently allocated to a partition

that is already 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 the original 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.

🖐**Caution**

If there was an operating system of some type installed on that partition, it will need to be reinstalled as well. Be aware that some computers sold with pre-installed operating systems may not include the CD-ROM media to reinstall the original operating system. The best time to notice if this applies to your system is *before* you destroy your original partition and its operating system installation.

After creating a smaller partition for your existing software, you can reinstall any software, restore your data, and continue your Red Hat Linux installation. Figure E-10 shows this being done.
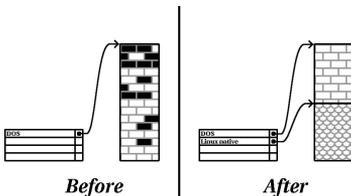
Figure E-10. Disk Drive Being Destructively Repartitioned

🖐**Caution**

As Figure E-10, shows, any data present in the original partition will be lost without proper backup!

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. What software should you use to perform this feat? There are several disk management software products on the market. You will have to do some research to find the one that is best for your situation.

While the process of non-destructive repartitioning is rather straightforward, there are a number of steps involved:

- Compress existing data
- Resize the existing partition
- Create new partition(s)

Next we will look at each step in a bit more detail.

E.1.4.3.1. Compress existing data

As Figure E-11, shows, the first step is to compress the data in your existing partition. The reason for doing this is to rearrange the data such that it maximizes the available free space at the "end" of the partition.

Before   *After*
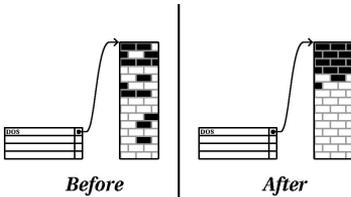
Figure E-11. Disk Drive Being Compressed

This step is crucial. Without it, the location of your data could prevent the partition from being resized to the extent desired. Note also that, for one reason or another, some data cannot be moved. If this is the case (and it severely restricts the size of your new partition(s)), you may be forced to destructively repartition your disk.

E.1.4.3.2. Resize the existing partition

Figure E-12, shows the actual resizing process. While the actual result of the resizing operation varies depending on the software used, in most cases the newly freed space is used to create an unformatted partition of the same type as the original partition.

Before   *After*
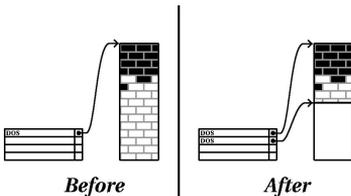
Figure E-12. Disk Drive with Partition Resized

It is important to understand what the resizing software you use does with the newly freed space, so that you can take the appropriate steps. In the case we have illustrated, it would be best to simply delete the new DOS partition and create the appropriate Linux partition(s).

E.1.4.3.3. Create new partition(s)

As the previous step implied, it may or may not be necessary to create new partitions. However, unless your resizing software is Linux-aware, it is likely you will need to delete the partition that was created during the resizing process. Figure E-13, shows this being done.

Before   *After*
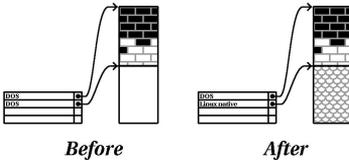
Figure E-13. Disk Drive with Final Partition Configuration

Note

> The following information is specific to Intel-based computers only.

As a convenience to Red Hat Linux users, the DOS `fips` utility is included on the Red Hat Linux/x86 CD 1 in the `dosutils` directory. This is a freely available program that can resize FAT (File Allocation Table) partitions.

Warning

> Many people have successfully used `fips` to resize their hard drive partitions. However, because of the nature of the operations carried out by `fips` and the wide variety of hardware and software configurations under which it must run, Red Hat cannot guarantee that `fips` will work properly on your system. Therefore, no installation support is available for `fips`. *Use it at your own risk.*

That said, if you decide to repartition your hard drive with `fips`, it is *vital* that you do two things:

- *Perform a backup* — Make two copies of all the important data on your computer. These copies should be to removable media (such as tape or diskettes), and you should make sure they are readable before proceeding.

- *Read the documentation* — Completely read the `fips` documentation, located in the `dosutils/fipsdocs` subdirectory on Red Hat Linux/x86 CD 1.

Should you decide to use `fips`, be aware that after `fips` runs you will be left with *two* partitions: the one you resized, and the one `fips` created out of the newly freed space. If your goal is to use that space to install Red Hat Linux, you should delete the newly created partition, either by using `fdisk` under your current operating system or while setting up partitions during a custom installation.

E.1.5. Partition Naming Scheme

Linux refers to disk partitions using a combination of letters and numbers which may be confusing, particularly if you are used to the "C drive" way of referring to hard disks and their partitions. In the DOS/Windows world, partitions are named using the following method:

- Each partition's type is checked to determine if it can be read by DOS/Windows.

- If the partition's type is compatible, it is assigned a "drive letter." The drive letters start with a "C" and move on to the following letters, depending on the number of partitions to be labeled.

- The drive letter can then be used to refer to that partition as well as the file system contained on that partition.

Red Hat Linux uses a naming scheme that is more flexible and conveys more information than the approach used by other operating systems. The naming scheme is file-based, with filenames in the form:

`/dev/xxyN`

Here is how to decipher the partition naming scheme:

`/dev/`

> This string is the name of the directory in which all device files reside. Since partitions reside on hard disks, and hard disks are devices, the files representing all possible partitions reside in `/dev/`.

`xx`

> The first two letters of the partition name indicate the type of device on which the partition resides. You will normally see either `hd` (for IDE disks) or `sd` (for SCSI disks).

`y`

> 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).

`N`

> The final number denotes the partition. The first four (primary or extended) partitions are numbered `1` through `4`. Logical partitions start at `5`. So, for example, `/dev/hda3` is the third primary or extended partition on the first IDE hard disk, and `/dev/sdb6` is the second logical partition on the second SCSI hard disk.

Note

> There is no part of this naming convention that is based on partition type; unlike DOS/Windows, *all* partitions can be identified under Red Hat Linux. Of course, this does not mean that Red Hat Linux can access data on every type of partition, but in many cases it is possible to access data on a partition dedicated to another operating system.

Keep this information in mind; it will make things easier to understand when you are setting up the partitions Red Hat Linux requires.

E.1.6. 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 will 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 Documentation CD in the HOWTO and HOWTO/mini directories. In particular, the Mini-HOWTOs whose names start with Linux+ are quite helpful.

Note

If Red Hat Linux/x86 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.

E.1.7. Disk Partitions and Mount Points

One area that many people new to Linux find confusing is the matter of how partitions are used and accessed by the Linux operating system. In DOS/Windows, it is relatively simple: Each partition gets a "drive letter." You then use the correct drive letter to refer to files and directories on its corresponding partition.

This is entirely different from how Linux deals with partitions and, for that matter, with disk storage in general. The main difference is that each partition is used to form part of the storage necessary to support a single set of files and directories. This is done by associating a partition with a directory through a process known as *mounting*. Mounting a partition makes its storage available starting at the specified directory (known as a *mount point*).

For example, if partition /dev/hda5 were mounted on /usr, that would mean that all files and directories under /usr would physically reside on /dev/hda5. So the file /usr/share/doc/FAQ/txt/Linux-FAQ would be stored on /dev/hda5, while the file /etc/X11/gdm/Sessions/Gnome would not.

Continuing our example, it is also possible that one or more directories below /usr would be mount points for other partitions. For instance, a partition (say, /dev/hda7) could be mounted on /usr/local, meaning that /usr/local/man/whatis would then reside on /dev/hda7 rather than /dev/hda5.

E.1.8. How Many Partitions?

At this point in the process of preparing to install Red Hat Linux, you will need to give some consideration to the number and size of the partitions to be used by your new operating system. The question of "how many partitions" continues to spark debate within the Linux community and, without any end to the debate in sight, it is safe to say that there are probably as many partition layouts as there are people debating the issue.

Keeping this in mind, we recommend that, unless you have a reason for doing otherwise, you should at least create the following partitions:

- *A swap partition* — Swap partitions are used to support virtual memory. In other words, data is written to swap when there is not RAM to hold the data your system is processing. You *must* create a swap partition to correctly use Red Hat Linux. The minimum size of

your swap partition should be equal to twice the amount of your computer's RAM or 32 MB, whichever is larger.

- *A /boot partition* — The partition mounted on /boot contains the operating system kernel (which allows your system to boot Red Hat Linux), along with a few other files used during the bootstrap process.

⚠ Caution

> Make sure you read Section E.1.9 — the information there applies to the /boot partition!

Due to the limitations of most PC BIOSes, creating a small partition to hold these files is a good idea. For most users, a 32 MB boot partition is sufficient.

- *A root partition (/)* — The root partition is where / (the root directory) resides. In this partitioning layout, all files (except those stored in /boot) reside on the root partition. Because of this, it is in your best interest to maximize the size of your root partition. For example, a 1.2 GB root partition *may* permit the equivalent of a workstation installation (with *very* little free space), while a 3.4 GB root partition *may* let you install every package. Obviously, the more space you can give the root partition, the better.

Specific recommendations concerning the proper size for various Red Hat Linux partitions can be found in Section 1.5.

E.1.9. One Last Wrinkle: Using GRUB or LILO

GRUB and LILO are the most commonly used methods to boot Red Hat Linux on Intel-based systems. As operating system loaders, they operate "outside" of any operating system, using only the Basic I/O System (or BIOS) built into the computer hardware itself. This section describes GRUB and LILO's interactions with PC BIOSes and is specific to Intel-compatible computers.

E.1.9.1. BIOS-Related Limitations Impacting GRUB and LILO

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

All the data GRUB and LILO need to access at boot time (including the Linux kernel) is located in the /boot directory. If you follow the partition layout recommended above or are performing a workstation or server install, the /boot directory will be in a small, separate partition. Otherwise, it may reside in the root partition (/). In either case, the partition in which /boot resides must conform to the following guidelines if you are going to use GRUB or 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 GRUB and 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 GRUB and
LILO to work with configurations that do not meet these guidelines. Likewise, some of
GRUB and LILO's more esoteric features may be used to get a Linux system started, even
if the configuration does not meet our guidelines. However, due to the number of variables
involved, Red Hat cannot support such efforts.

Note

   Disk Druid, as well as the workstation and server installations, takes these BIOS-related limitations
   into account.

Appendix F.
Driver Disks

F.1. Why Do I Need a Driver Disk?

While the Red Hat Linux installation program is loading, you may see a screen that asks you for a driver disk. The driver disk screen is most often seen in three scenarios:

- If you run the installation program in `expert mode`
- If you run the installation program by entering **`linux dd`** at the `boot:` prompt
- If you run the installation program on a computer which does not have any PCI devices

F.1.1. So What Is a Driver Disk Anyway?

A driver disk adds support for hardware that is not otherwise supported by the installation program. The driver disk could be produced by Red Hat, it could be a disk you make yourself from drivers found on the Internet, or it could be a disk that a hardware vendor includes with a piece of hardware.

There is no need to use a driver disk unless you need a particular device in order to install Red Hat Linux. Driver disks are most often used for non-standard or very new CD-ROM drives, SCSI adapters, or NICs. These are the only devices used during the installation that might require drivers not included on the Red Hat Linux CD-ROMs (or boot disk, if you created an installation boot disk to begin the install process).

Note

If an unsupported device is not needed to install Red Hat Linux on your system, continue with the installation and add support for the new piece of hardware once the installation is complete.

F.1.2. How Do I Obtain a Driver Disk?

The Red Hat Linux CD-ROM 1 includes driver disk images (`images/drvnet.img` — network card drivers and `images/drvblock.img` — drivers for SCSI controllers) containing many rarely used drivers. If you suspect that your system may require one of these drivers, you should create the driver disk before beginning your Red Hat Linux installation.

Another option for finding specialized driver disk information is on Red Hat's website at

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

under the section called **Bug Fixes**. Occasionally, popular hardware may be made available after a release of Red Hat Linux that will not work with drivers already in the installation program or included on the driver disk images on the Red Hat Linux CD-ROM 1. In such cases, the Red Hat website may contain a link to a driver disk image.

F.1.2.1. Creating a Driver Disk from an Image File

If you have a driver disk image that you need to write to a floppy disk, this can be done from within DOS or Red Hat Linux.

To create a driver disk from a driver disk image using Red Hat Linux:

1. Insert a blank, formatted floppy disk into the first floppy drive.
2. From the same directory containing the driver disk image, such as *drvnet.img*, type `dd if=drvnet.img of=/dev/fd0` as root.

To create a driver disk from a driver disk image using DOS:

1. Insert a blank, formatted floppy disk into the a: drive.
2. From the same directory containing the driver disk image, such as *drvnet.img*, type `rawrite drvnet.img a:` at the command line.

F.1.3. Using a Driver Disk During Installation

Having a driver disk is not enough; you must specifically tell the Red Hat Linux installation program to load that driver disk and use it during the installation process.

Note

A driver disk is different than a boot disk. If you require a boot disk to begin the Red Hat Linux installation, you will still need to create that floppy and boot from it before using the driver disk.

If you do not already have an installation boot disk and your system does not support booting from the CD-ROM, you should create an installation boot disk. For instructions on how make a boot disk, see Section 1.4.2 .

Once you have created your driver disk, begin the installation process by booting from the Red Hat Linux CD-ROM 1 (or the installation boot disk). At the `boot:` prompt, enter either **linux expert** or **linux dd**. Refer to Section 3.3.1 for details on booting the installation program.

The Red Hat Linux installation program will ask you to insert the driver disk. Once the driver disk is read by the installer, it can apply those drivers to hardware discovered on your system later in the installation process.

redhat.

Appendix G.
Configuring a Dual-Boot System

Sharing a computer between two operating systems requires dual booting. You can use either operating system on the computer, but not both at once. Each operating system boots from and uses its own hard drives or disk partitions.

This chapter explains how to configure your system to boot into both Red Hat Linux and another operating system. For clarity, we will assume that the other operating system is Microsoft Windows™. But the general procedures are similar for other operating systems.

Note

> If Red Hat Linux will coexist on your system 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 **fdisk**.

If you do not have any operating systems installed on your computer, install Windows first and then install Red Hat Linux.

- If you are installing Windows 9*x* or Windows ME, you can not define partitions during the Windows installation. Install Windows, and then refer to Section G.3 for instructions on using **fips** to repartition your hard drive and create free space for Linux.

- If you are installing Windows NT or Windows 2000, you can create partitions of a specific size for Windows. Leave enough free space (space that is not partitioned or formatted) on the hard drive to install Linux.

Tip

> While partitioning your hard drive, keep in mind that the BIOS in some systems cannot access more than the first 1024 cylinders on a hard drive. If this is the case, leave enough room for the /boot Linux partition on the first 1024 cylinders of your hard drive to boot Linux. The other Linux partitions can be after cylinder 1024.

Refer to Section 1.3 to determine how much disk space to leave. After installing Windows, refer to Section G.2.

If the computer you want to install Red Hat Linux on is currently running Windows (or some other operating system you have installed), you have an important decision to make. Your choices are:

- Do you want Red Hat Linux to be the only operating system on your computer, despite the fact that you already have Windows on your computer? If yes, you do not have to configure a dual-boot system. Backup any information that you want to save and start the installation. During the installation, if you choose to have the installation program automatically partition your system on the **Disk Partitioning Setup** screen, choose **Remove all partitions on this system**. If you choose manual partitioning with **Disk Druid** or **fdisk**, delete all the existing DOS (Windows) partitions and then create your Linux partitions.

- Do you want to install Red Hat Linux and then have the option of booting either Red Hat Linux or your other operating system? A Red Hat Linux installation can be performed so that Red Hat Linux is installed on your system, but the other operating system is not affected. Since you already have Windows installed, you need to allocate disk space for Linux. Refer to Section G.1, and then refer to Section G.2.

G.1. Allocating Disk Space for Linux

⚠️**Warning**

> Remember to back up all important information before reconfiguring your hard drive. Reconfiguring your hard drive can result in the loss of data if you are not extremely careful. Additionally, be sure to create a boot disk for both operating systems in case the boot loader fails to recognize either of them.

If you already have Windows installed on your system, you must have free hard drive space available on which to install Red Hat Linux. Your choices are as follows:

- Add a new hard drive.
- Use an existing hard drive or partition.
- Create a new partition.

For all three options, be aware that the BIOS in some older systems cannot access more than the first 1024 cylinders on a hard drive. If this is the case, the /boot Linux partition must be located on the first 1024 cylinders of your hard drive to boot Linux.

G.1.1. Add a New Hard Drive

The simplest way to make room for Red Hat Linux is to add a new hard drive to the computer and then install Red Hat Linux on that drive. For example, if you add a second IDE hard drive to the computer, the Red Hat Linux installation program will recognize it as hdb and the existing drive (the one used by Windows) as hda. (For SCSI hard drives, the newly installed Red Hat Linux hard drive would be recognized as sdb and the other hard drive as sda.)

If you choose to install a new hard drive for Linux, all you need to do is start the Red Hat Linux installation program. After starting the Red Hat Linux installation program, just make sure you choose to install Linux on the newly installed hard drive (such as hdb or sdb) rather than the hard drive used by Windows.

G.1.2. Use an Existing Hard Drive or Partition

Another way to make room for Linux is to use a hard drive or disk partition that is currently being used by Windows. For example, suppose that **Windows Explorer** shows two hard drives, C: and D:. This could indicate either that the computer has two hard drives, or a single hard drive with two partitions. In either case (assuming the hard drive has enough disk space), you can install Red Hat Linux on the hard drive or disk partition that Windows recognizes as D:.

This choice is available to you only if the computer has two or more hard drives or disk partitions.

Note

Windows uses letters to refer to removable drives (for example, a ZIP drive) and network storage (virtual drives) as well as for local hard drive space; you cannot install Linux on a removable or network drive.

If a local Windows partition is available in which you want to install Linux, complete the following steps:

1. Copy all data you want to save from the selected hard drive or partition (`D:` in this example) to another location.

2. Start the Red Hat Linux installation program and tell it to install Linux in the designated drive or partition — in this example, in the hard drive or partition that Windows designates as `D:`. Note that Linux distinguishes between hard drives and disk partitions. Thus:

   - If `C:` and `D:` on this computer refer to two separate hard drives, the installation program will recognize them as `hda` and `hdb` (IDE) or `sda` and `sdb` (SCSI). Tell the installation program to install on `hdb` or `sdb`.

   - If `C:` and `D:` refer to partitions on a single drive, the installation program will recognize them as `hda1` and `hda2` (or `sda1` and `sda2`). During the partitioning phase of the Red Hat Linux installation, delete the second partition (`hda2` or `sda2`), then partition the unallocated free space for Linux. You do not have to delete the second partition prior to starting the Red Hat Linux installation.

G.1.3. Create a New Partition

The third way to make room for Linux is to create a new partition for Red Hat Linux on the hard drive being used by the other operating system. If **Windows Explorer** shows only one hard drive (`c:`), and you do not want to add a new hard drive, you must partition the drive. After partitioning, **Windows Explorer** will see a smaller `c:` drive; and, when you run the Red Hat Linux installation program, you can partition the remainder of the drive for Linux.

You can use a destructive partitioning program, such as **fdisk**, to divide the hard drive, but doing so will require you to re-install Windows. (This is probably not your best option.)

A number of non-destructive third-party partitioning programs are available for the Windows operating system. If you choose to use one of these, consult their documentation.

For instructions on how to partition with **FIPS**, a program that is on the Red Hat Linux CD-ROM, refer to Section G.3.

G.2. Installing Red Hat Linux in a Dual-Boot Environment

After Windows is installed and you have free disk space ready for Linux, you can start the Red Hat Linux installation program. Go to Chapter 1 to begin. At this point, the only difference between a Red Hat Linux installation and configuring a dual-boot system during the Red Hat Linux installation is partitioning the hard drive and configuring the boot loader. When you are at the **Disk Partitioning Setup** screen as described in Section 3.16, return to this section.

G.2.1. Disk Partitioning

At the **Disk Partitioning Setup** screen of the installation program, you have a few options. Depending on which option you choose, the steps for configuring a dual-boot system vary. If you do not know how many Linux partitions to create, refer to Section 3.18 for a recommended partitioning scheme. If you choose:

- Automatic partitioning — Choose **Keep all partitions and use existing free space**. This option will leave your Windows partitions on the hard drive and partition the free space or additional hard drive for Red Hat Linux.

- Manual partitioning with **Disk Druid** — Do not delete the existing Windows partitions (they are the partitions of type vfat). Create your Linux partitions on the additional hard drive or in the free space you have reserved for Red Hat Linux.

- Manual partitioning with **fdisk** — Similar to using **Disk Druid** except you will not have the graphical interface. The basic procedure is the same. Do not delete the existing partitions of type FAT16, FAT32, or NTFS. Create your Linux partitions on the additional hard drive or in the free space you have reserved for Red Hat Linux.

G.2.2. Configuring the Boot Loader

When you arrive at the **Boot Loader Installation** screen during the Red Hat Linux installation, choose to install the boot loader. You can use a different boot loader to boot both Red Hat Linux and Windows. Red Hat does not support alternate boot loaders. Thus, this section will discuss how to configure GRUB or LILO to boot both operating systems.

The Red Hat Linux installation program will usually detect Windows and automatically configure the boot loader (GRUB or LILO) to boot either Red Hat Linux or Windows. This can be seen on the boot loader screen of the installation program. An entry named **DOS** appears in the list of operating systems to boot.

G.2.3. Post-Installation

After the installation, whenever you start the computer, you can indicate whether you want to start Red Hat Linux or the other operating system from the boot loader screen. Choose **Red Hat Linux** to boot into Red Hat Linux, and choose **DOS** to boot into Windows.

If you did not partition all the free space on your hard drive for Red Hat Linux, you can partition it for Windows after installing Red Hat Linux. You can use **cfdisk** or **fdisk** to create these partitions. **cfdisk** is easier to use than **fdisk** because it is more interactive and has a menu from which to choose.

⚠️**Warning**

It is highly recommended that you use **cfdisk** or **fdisk** to create partitions after installing Red Hat Linux. Other partitioning software has been known to change the partitioning table of the hard drive and move the Linux partitions. If this happens, the boot loader will not be able to find the Linux partitions and will not boot into Red Hat Linux.

To access the files on the Windows partitions while using Red Hat Linux, refer to the *Accessing a Windows Partition* FAQ in the *Official Red Hat Linux Getting Started Guide*. If you formatted the Windows partitions in NTFS format, this method will not work.

G.3. Partitioning with FIPS

As a convenience to our customers, we provide the **FIPS** utility. This is a freely available program that can resize FAT (File Allocation Table) partitions. It is included on the Red Hat Linux CD-ROM in the dosutils directory. If you are using NTFS partitions, **FIPS** will not work.

Note

Many people have successfully used **FIPS** to repartition their hard drives. However, because of the nature of the operations carried out by **FIPS**, and the wide variety of hardware and software configurations under which it must run, Red Hat cannot guarantee that **FIPS** will work properly on your system. Therefore, no installation support whatsoever is available for **FIPS**; use it at your own risk.

That said, if you decide to repartition your hard drive with **FIPS**, it is vital that you do two things:

- *Perform a Backup* — Make two copies of all the important data on your computer. These copies should be to removable media (such as tape, CD-ROM, or diskettes), and you should make sure they are readable before proceeding.
- *Read the Documentation* — Completely read the **FIPS** documentation, located in the dosutils/fipsdocs directory on the Red Hat Linux CD-ROM 1.

Should you decide to use **FIPS**, be aware that after **FIPS** runs you will be left with two partitions: the one you resized, and the one **FIPS** created out of the newly freed space. If your goal is to use that space to install Red Hat Linux, you should delete the newly created partition, either by using **fdisk** under your current operating system, or while setting up partitions during a custom-class installation.

The following instructions are a simplified version of the **FIPS** documentation file, fips.doc, located in the **FIPS** directory (/dosutils/fips20/*). These instructions should apply in most instances. If you encounter any problems, see the documentation file.

1. From Windows:

   - Do a full backup.
   - Run **scandisk** to verify that the hard drive contains no bad clusters.
   - Decide how to distribute the available space on the hard drive between the operating systems. Use **Windows Explorer** to see the free space on the drive. Make a note of the space (in megabytes) that each operating system will have.
   - If you do not have one, create a Windows boot disk.

     Creating a boot disk varies between different versions of Windows. Consult the Windows documentation for instructions on creating a Windows boot disk.

     The diskette will be formatted, and COMMAND.COM, along with the associated hidden files (IO.SYS, MSDOS.SYS, and DBLSPACE.BIN), will be copied to the diskette.

   - Copy the following files on the Red Hat Linux CD-ROM to the DOS boot disk.

     ```
     dosutils/fips20/fips.exe
     dosutils/fips20/restorrb.exe
     dosutils/fips20/errors.txt
     ```

```
dosutils/fips20/fips.doc
dosutils/fips20/fips.faq
```

- Defragment the hard drive so that all the data on the hard drive is located at the beginning of the drive.

2. Insert the Windows boot disk into the floppy drive and reboot the system.

3. Start **FIPS** (type `fips` at the prompt).

   When **FIPS** begins, you'll find a welcome screen similar to the following:

   ```
   FIPS version 2.0, Copyright (C) 1993/4 Arno Schaefer
   FAT32 Support, Copyright (C) 1997 Gordon Chaffee

   DO NOT use FIPS in a multitasking environment like Windows, OS/2, Desqview,
   Novell Task manager or the Linux DOS emulator; boot from a DOS boot disk first.

   If you use OS/2 or a disk compressor, read the relevant sections in FIPS.DOC.

   FIPS comes with ABSOLUTELY NO WARRANTY, see file COPYING for details.

   This is free software, and you are welcome to redistribute it
   under certain conditions; again, see file COPYING for details.

   Press any key.
   ```

   When you press a key, a root partition screen will appear. (Note that, if the computer has more than one hard drive, you'll be asked to select which one you want to partition.)

   When you press a key, details about the hard drive, such as the following, will appear.

   ```
   Boot sector:
   Bytes per sector: 512
   Sectors per cluster: 8
   Reserved sectors: 1
   Number of FATs: 2
   Number of rootdirectory entries: 512
   Number of sectors (short): 0
   Media descriptor byte: f8h
   Sectors per FAT: 145
   Sectors per track: 63
   Drive heads: 16
   Hidden sectors: 63
   Number of sectors (long): 141057
   Physical drive number: 80h
   Signature: 29h

   Checking boot sector ... OK
   Checking FAT ... OK
   Searching for free space ... OK

   Do you want to make a backup copy of your root and boot sector before
   proceeding? (y/n)
   ```

   You should select [y], for `yes`, to make a backup copy of your root and boot sector before proceeding with **FIPS**.

   Next, you will be presented with the following message:

   ```
   Do you have a bootable floppy disk in drive A: as described in the
   documentation? (y/n)
   ```

Verify that a DOS boot disk is in the floppy drive, and type [y], for `yes`. A screen similar to the following will appear, allowing you to resize the partition.

```
Writing file a:\rootboot:000

Enter start cylinder for new partition (33-526)

Use the cursor keys to choose the cylinder, <enter> to continue

Old partition          Cylinder              New partition
 258.9 MB                 33                    3835.8 MB
```

Figure G-1. Partition Resizing Screen

The initial values allocate *all* free space on the disk to the new partition. This is not what you want, because this setting would leave no free space on your Windows partition. Press the [right arrow] to increase the size of the Windows partition and decrease the size of the new (Linux) partition; press the [left arrow] to decrease the size of the Windows partition and increase the size of the Linux partition. When the sizes are what you want, press [Enter]. A verification screen will appear.

If you type **r** (to re-edit the partition tables), Figure G-1 reappears, allowing you to change the partition sizes. If you answer **c**, a confirmation screen, Figure G-2, appears:

```
New boot sector:

Boot sector:
Bytes per sector: 512
Sectors per cluster: 8
Reserved sectors: 1
Number of FATs: 2
Number of rootdirectory entries: 512
Number of sectors (short): 0
Media descriptor byte: f8h
Sectors per FAT: 145
Sectors per track: 63
Drive heads: 16
Hidden sectors: 63
Number of sectors (long): 141057
Physical drive number: 80h
Signature: 29h

Checking boot sector ... OK

Ready to write new partition scheme to disk
Do you want to proceed (y/n)?
```

Figure G-2. FIPS Confirmation Screen

Answering **y** completes the resizing operation. A harmless error message may occur, stating in effect that **FIPS** cannot reboot the system.

After a successful operation, the disk will have two partitions. The first partition (`hda1` or `sda1`) will be used by Windows. We recommend that you start Windows (remember to remove the boot disk from drive A:) and run **scandisk** on drive C:.

If you encounter any problems (for example, Windows will not boot), you can reverse the **FIPS** resizing operation with the **restorrb.exe** command, which you copied to your DOS boot disk. In case of any errors, read the **FIPS** documentation files (`fips.doc` and `fips.faq`),

which describe a number of factors that could cause the resizing operation to fail. If all else fails, you can restore Windows with the backup you made.

The second partition (`hda2` or `sda2`) contains the space that the Red Hat Linux installation program will use. When the **Disk Druid** screen appears during installation, delete this partition (the installation manual explains how), then proceed with Linux partitioning.

Index

Symbols

/boot, 52
/boot partition
  (See partition, /boot)
/root/install.log
  install log file location, 75

A

adding partitions, 53
  file system type, 54
ATAPI CD-ROM
  unrecognized, problems with, 39
authentication
  configuration, 68
    Kerberos, 68
    LDAP, 68
    MD5 passwords, 68
    NIS, 68
    shadow passwords, 68
autoboot, 37
automatic partitioning, 48, 49

B

BIOS, issues related to GRUB, 122
BIOS, issues related to LILO, 122
boot disk creation, 76, 90
boot loader, 56
  alternatives to, 58, 86
    boot disk, 58, 86
    commercial products, 59, 86
    LOADLIN, 58, 86
    SYSLINUX, 59, 86
  configuration, 56
  GRUB, 56
  installation, 56
  installing on boot partition, 57, 84
  LILO, 56
  MBR, 57
boot methods
  local boot disk, 18
  PCMCIA boot disks, 18
  USB diskette drive, 19
boot options, 35
  expert mode, 36
  mediacheck, 36
  no framebuffer mode, 35
  serial mode, 36
  text mode, 35
bootable CD-ROM, 37

booting
  installation program, 34

C

canceling the installation, 38
CD-ROM
  ATAPI, 38
    unrecognized, problems with, 39
  bootable, 37
  IDE, 38
    unrecognized, problems with, 39
  installation from, 38
  other, 39
  SCSI, 38
class
  installation, 48
clock, 65
configuration
  clock, 65
  GRUB, 56
  hardware, 29
    finding with Windows, 25
  LILO, 56
  network, 60
  time, 65
  time zone, 65
consoles, virtual, 31
conventions
  document, ix
Custom
  disk space, 18

D

dd
  creating installation diskette, 20
deleting partitions, 55
dependencies
  installing packages, 73
  upgrading packages, 88
disk
  driver, 125
Disk Druid
  adding partitions, 53
    file system type, 54
  buttons, 51
  deleting partitions, 55
  editing partitions, 55
  partitions, 50
disk partitioning, 48
disk space, 16
disk space requirements, 17
  custom, 23
  laptop, 23

W

Windows

    finding hardware configuration with, 25

X

X

    configuration, 74

X Window System, 74

Xconfigurator

    monitor setup, 77

    video card setup, 74