

Red Hat Linux 7.1

The Official Red Hat Linux Alpha Installation Guide

ISBN: ISBN: N/A

 Red Hat, Inc.

2600 Meridian Parkway
Durham, NC 27713 USA
+1 919 547 0012 (Voice)
+1 919 547 0024 (FAX)
888 733 4281 (Voice)
P.O. Box 13588
Research Triangle Park, NC 27709 USA

© 2001 Red Hat, Inc.

rhl-ig-alpha(EN)-7.1-Print-RHI (2001-04-20T16:12-0400)

Copyright © 2001 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.

Distribution of the work or derivative of the work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

Red Hat, Red Hat Network, the Red Hat "Shadow Man" logo, RPM, Maximum RPM, the RPM logo, Linux Library, PowerTools, Linux Undercover, RHmember, RHmember More, Rough Cuts, Rawhide and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Motif and UNIX are registered trademarks of The Open Group.

Compaq and the names of Compaq products referenced herein are either trademarks and/or service marks or registered trademarks and/or service marks of Compaq.

Itanium is a registered trademark of Intel Corporation.

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

Windows is a registered trademark of Microsoft Corporation.

SSH and Secure Shell are trademarks of SSH Communications Security, Inc.

FireWire is a trademark of Apple Computer Corporation.

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

Printed in Canada, Ireland, and Japan

Contents

Red Hat Linux 7.1

| | | |
|------------------|--|----|
| Chapter 1 | Introduction | 7 |
| 1.1 | Documentation Conventions | 7 |
| 1.2 | How to Use This Manual..... | 9 |
| Chapter 2 | Steps to Get You Started | 11 |
| 2.1 | Step 1 - Do You Have the Right Red Hat Linux Components? | 11 |
| 2.2 | Step 2 - Is Your Hardware Compatible?..... | 12 |
| Chapter 3 | Before You Begin the Installation Process | 13 |
| 3.1 | Installation Overview | 13 |
| 3.2 | Checking for Errata | 14 |
| 3.3 | The SRM Firmware Console | 14 |
| Chapter 4 | Installing Red Hat Linux | 23 |
| 4.1 | The Installation Program User Interface | 23 |
| 4.2 | Selecting an Installation Method | 24 |
| 4.3 | Language Selection | 25 |
| 4.4 | Keyboard Configuration | 25 |
| 4.5 | Mouse Configuration | 26 |
| 4.6 | Welcome to Red Hat Linux | 28 |
| 4.7 | Install Options..... | 28 |
| 4.8 | Disk Partitioning | 29 |
| 4.9 | Partitioning Your System..... | 30 |
| 4.10 | Partitioning with fdisk..... | 35 |
| 4.11 | Formatting Partitions | 38 |
| 4.12 | Network Configuration | 40 |
| 4.13 | Firewall Configuration | 41 |
| 4.14 | Language Support Selection | 44 |
| 4.15 | Time Zone Configuration | 45 |

| | | |
|--|---|-----------|
| 4.16 | Account Configuration | 47 |
| 4.17 | Authentication Configuration | 49 |
| 4.18 | Package Group Selection..... | 50 |
| 4.19 | GUI X Configuration Tool | 54 |
| 4.20 | Preparing to Install | 58 |
| 4.21 | Installing Packages..... | 58 |
| 4.22 | Note the Kernel Filename and Partition Number..... | 59 |
| 4.23 | Installation Complete..... | 60 |
| 4.24 | Booting Your Machine and Post-Installation Setup..... | 60 |
| Chapter 5 Installing Red Hat Linux via Text Mode..... | | 63 |
| 5.1 | Things You Should Know | 63 |
| 5.2 | The Installation Program User Interface | 65 |
| 5.3 | Starting the Installation Program..... | 67 |
| 5.4 | Installation Cross-reference Table | 68 |
| 5.5 | Identify Disk Partition to Install From | 69 |
| 5.6 | Installing over a Network..... | 70 |
| 5.7 | Configuring Your Video Adapter | 74 |
| 5.8 | Configuring the X Window System | 75 |
| Appendix A Upgrading Your Current System..... | | 83 |
| A.1 | What it Means to Upgrade..... | 83 |
| A.2 | Upgrading Your System..... | 83 |
| A.3 | Customizing Your Upgrade | 84 |
| A.4 | Selecting Packages to Upgrade..... | 85 |
| A.5 | Upgrading Packages | 87 |
| A.6 | Upgrade Complete | 88 |
| Appendix B MILO..... | | 89 |
| B.1 | MILO Installation Overview..... | 89 |
| B.2 | MILO and Alpha Hardware | 89 |
| B.3 | Choosing Disk Images | 90 |
| B.4 | Disk Partitioning | 92 |
| B.5 | Installations Using MILO..... | 92 |

| | | |
|--|---|------------|
| B.6 | Post-Installation MILO Setup | 98 |
| Appendix C Additional Resources about Alphas and Linux . | | 101 |
| Appendix D Writing Image Files to Diskettes | | 103 |
| D.1 | Using the dd Command..... | 103 |
| D.2 | Using the rawrite Utility | 103 |
| Appendix E Getting Technical Support | | 105 |
| E.1 | Remember to Sign Up | 105 |
| E.2 | An Overview of Red Hat Support | 105 |
| E.3 | Scope of Red Hat Support | 106 |
| E.4 | How to Get Technical Support | 107 |
| E.5 | Questions for Technical Support | 108 |
| E.6 | Support Frequently Asked Questions (FAQ)..... | 109 |
| Appendix F Troubleshooting Your Installation of Red Hat Linux | | 111 |
| F.1 | You are Unable to Boot Red Hat Linux | 111 |
| F.2 | Trouble Beginning the Installation..... | 111 |
| F.3 | Trouble During the Installation | 112 |
| F.4 | Problems After Installation..... | 115 |

1 Introduction

Welcome to the *Official Red Hat Linux Alpha Installation Guide*.

The *Official Red Hat Linux Alpha Installation Guide* contains useful information to assist you during the installation of Red Hat Linux. From fundamental concepts such as installation preparation (gathering information about your current system) 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 installation using the Red Hat Linux CD-ROMs. Once you have completed the installation as outlined in this manual, you will have a fully functioning Red Hat Linux desktop system.

1.1 Documentation 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 in 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 (like 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.

The `/usr/share/doc` directory contains documentation for various programs.

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.

Combinations of keystrokes

A combination of keystrokes is represented like this: [key]-[key]. For example:

The [Ctrl]-[Alt]-[Backspace] key combination will restart the X Window System.

Sequences of commands in Graphical User Interface (GUI) menus

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

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

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 a note, a 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.



CAUTION

Do not do routine tasks as root — use a regular user account unless you need to use the root account to administer your system.

WARNING

If you choose not to partition manually, a server-class 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.

1.2 How to Use This Manual

This manual focuses on a CD-ROM based installation, so it 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 3.0.3 (or greater), you can perform an upgrade. Skim Chapter 2, *Steps to Get You Started* to review the basics, then read Chapter 4, *Installing Red Hat Linux*, following the directions as you go. Once you have chosen to perform an upgrade in the installation program, refer to Appendix A, *Upgrading Your Current System*.

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 4, *Installing Red Hat Linux* 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/FAQ.html

1.2.1 We Need Feedback!

If you spot a typo in the *Official Red Hat Linux Alpha 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 in Bugzilla (<http://www.redhat.com/bugzilla/>) against the component *rhl-ig-alpha*.

Be sure to mention the manual's identifier:

```
rhl-ig-alpha(EN)-7.1-Print-RHI (2001-04-20T16:12-0400)
```

That way we will know exactly which version of the guide you have.

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]), please register your product at <http://www.redhat.com/apps/activate/> and use the online support system for those type of requests.

2 Steps to Get You Started

Before you install Red Hat Linux, you will need to perform the following steps:

2.1 Step 1 - Do You Have the Right Red Hat Linux 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 boxed set.

A black, red, and white Registration Information card is included in your boxed set. 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 its 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 boxed set may not have been actually produced by Red Hat.

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

2.1.1 Where to Find Other Red Hat Linux Manuals

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

To find the manuals in both HTML and PDF formats online, visit <http://www.redhat.com/support/manuals>.

2.1.2 Why You Should Register Your Red Hat Linux Boxed Set

If you have purchased an official Red Hat Linux 7.1 boxed set, you should register your product. Registration offers many useful services, such as installation support, priority FTP access, 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 E, *Getting Technical Support*.

2.1.3 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 disk. It may also be possible 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 Appendix D, *Writing Image Files to Diskettes*.

2.2 Step 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 7.1 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.

3 Before You Begin the Installation Process

3.1 Installation Overview

Installing Red Hat Linux on an Alpha system is slightly more complex than installing Red Hat Linux on an x86 machine because a variety of Alpha machine architectures exist. In general, the sequence of steps to a successful installation are the following:

1. Use SRM commands to identify system information that you will need to boot the installation program and Red Hat Linux.
2. If you can not boot from the CD-ROM, make diskettes from the boot image and RAM disk image files provided with Red Hat Linux.
3. Using the SRM console and the `aboot` bootloader, load and run the kernel, and boot into the Red Hat Linux installation program.

MILO Is Not Supported

As a resource for the Linux and Alpha communities, information that might be useful for MILO users is provided in Appendix B, *MILO*. Before you use MILO to boot your Alpha, however, you should realize that MILO is not supported by Red Hat, Inc.

3.1.1 Disk Space Requirements

Before you begin the installation, be aware that you will need certain amounts of free disk space to install Red Hat Linux on an Alpha.

- At minimum, you will need 64 MB of RAM and 700 MB of free hard disk space to install Red Hat Linux.
- You should have 128 MB of RAM and 4.5 GB of free hard disk space on your machine for a full installation of Red Hat Linux.

You may need a 3.5-inch diskette drive — referred to by Red Hat Linux as `/dev/fd0` (known as the A: drive in certain circles). Note that many modern Alphas can boot the installation program from the CD-ROM drive.

3.2 Checking for Errata

Occasionally, revised diskette images are needed for the installation to work properly. In these cases, we make special images available via the Red Hat Linux errata listing. Since this is relatively rare, you will save time if you try to use the standard diskette images first. Review the errata only if you experience problems completing the installation.

Although most of the time it is not necessary to check for errata before the installation, it is also not a bad idea, either.

At the Errata page on Red Hat's website, located at <http://www.redhat.com/support/errata/>, you can read errata online, and you can download diskette images.

If you experience problems, focus on entries that include new diskette images (the filenames always end in `.img`). If you find an entry that applies to your problem, download a copy of the diskette images, and create them using the instructions in Appendix D, *Writing Image Files to Diskettes*.

3.3 The SRM Firmware Console

Before you start to install Red Hat Linux on an Alpha, you will need a basic understanding of the SRM console, what it does, and the information it can provide.

SRM is a UNIX-like firmware console used by many Alpha systems to boot into an operating system. When the system is booted, the SRM console is displayed as a blue screen with a prompt. The prompt looks like `P00>>>` (the 00 indicates the primary processor). On multi-processor machines, the prompt could be `P01>>>`, `P02>>>`, etc.

We recommend that you upgrade your SRM firmware to the most recent version available. Use the `show version` command at the SRM console prompt to see what version of SRM you are running. For more information about SRM firmware upgrades, see the following URL:

```
http://ftp.digital.com/pub/DEC/Alpha/firmware/
```

Additional information on the SRM console can be found at the *SRM Firmware HOWTO* at the following URL:

```
http://www.alphalinux.org/faq/srm.html
```

The SRM console provides important system information which you will need to use when installing Red Hat Linux and when administering your Alpha.

3.3.1 SRM Device Names

To see information about system devices that SRM recognizes, use the following command:

```
show device
```

The `show device` command will display a list of devices. You can alternatively use the shorthand `show dev` or a command such as `show dev dk`, which will only display SRM-recognized devices that begin with `dk` (in other words, all SCSI disk drives).

SRM does not identify devices the same way as Red Hat Linux. As an example of SRM's device-naming syntax, the output provided after a `show device` on a Compaq AlphaServer ES40 will look similar to the following:

```
P00>>> show device
resetting all I/O buses
dka0.0.0.4.1          DKA0          RZZDD-LS 0306
dqa0.0.0.15.0         DQA0          TOSHIBA CD-ROM XM-6302B 1017
dva0.0.0.1000.0       DVA0
ewa0.0.0.6.1          EWA0          00-00-F8-1B-42-D5
pka0.7.0.4.1          PKA0          SCSI Bus ID 7
```

The following list provides explanations for the most common SRM device names:

dka0

The `dk` is SRM's designation for a SCSI device. The `a` after the `dk` signifies that it is the first SCSI controller; the `0` indicates the SCSI ID number. So, `dka0`, `dka100`, `dka200` and `dka300` are all SCSI devices on the first SCSI controller. Their SCSI ID numbers are, respectively, 0, 1, 2 and 3. (Note that the `a` and numeric portions of the device identifier work the same way for the other devices; they identify the ordering of the same type of device and provide an ID number.)

If these four devices are all SCSI hard drives, Red Hat Linux would refer to the same four devices as `/dev/sda`, `/dev/sdb`, `/dev/sdc` and `/dev/sdd`.

If your Alpha has a SCSI CD-ROM (instead of an IDE CD-ROM), usually the CD-ROM drive device will be on the second SCSI controller. The SRM device name will then be something like `dkb0` or `dkb500` (`b` for the second SCSI controller, followed by the SCSI ID number). If your machine is set up like this, then Red Hat Linux would refer to the CD-ROM drive as `/dev/scd0`.

Identification of CD-ROM Drives

If the device is a CD-ROM drive, the entry in the third column will identify it as a CD-ROM or be something like `CD-224E` (older Alphas may identify a CD-ROM drive as `RRD46 1337` or something similar).

dqa0

The `dq` is SRM's designation for an IDE device (for example, an IDE hard drive or an IDE CD-ROM drive). As in the previous example, the `a` indicates that it is the first IDE device, and the `0` is the device's ID number.

Red Hat Linux would refer to the same IDE device as `/dev/hda`.

If your Alpha has an IDE hard drive and an IDE CD-ROM, then SRM will refer to the hard drive as something like `dqa0` and to the CD-ROM drive as something like `dqb0` or `dqb500`. Red Hat Linux would refer to the same devices as `/dev/hda` and `/dev/hdb`.

dva0

The `dva0` device name is the floppy disk drive; Red Hat Linux would refer to the floppy disk drive as `/dev/fd0`.

ewa0

The `ewa0` device name is the network card; if it is an Ethernet network card, Red Hat Linux would refer to it as `eth0`.

Examples of other SRM device names include `mka0`, the first SCSI tape device, and `pka0`, a SCSI host adapter for the first SCSI bus.

Before you begin the installation process, look at the list provided in response to `show device` on your Alpha's SRM console. Write down how SRM refers to the devices in Table 3-1, *SRM Device Names for your System*.

Table 3-1 SRM Device Names for your System

| Device | SRM Device Name |
|--|-----------------|
| 3.5-inch diskette drive (usually <code>dva0</code>) | |
| Hard drive | |
| CD-ROM drive | |

In this document, we have used the SRM device name `dka0` to refer to the hard drive and `dqa0` to refer to the CD-ROM drive.

You must choose one of these devices from which to boot Red Hat Linux, during the installation process and every time you wish to boot your machine, unless you set up your SRM boot parameters to boot automatically.

Devices that SRM Can Not Detect

You may have a SCSI adapter, IDE drives, or networking hardware that is compatible with Red Hat Linux but is neither OEM Digital or Compaq hardware or otherwise detectable by the SRM console. These types of hardware will not be displayed after the `show device` command on the SRM console. Using unrecognized hardware may not stop you from running Red Hat Linux on your Alpha machine, but you must boot from a device recognized by the SRM console.

3.3.2 SRM Boot Parameters

Certain SRM system parameters designate information needed to boot an operating system. These parameters include the device from which SRM will boot, which kernel to use, and what filesystem it will use as the root filesystem.

To see the boot parameters recognized by SRM, use the following command:

```
show boot*
```

For example, the output provided for a `show boot*` on the SRM console might look like the following:

```
P00>>> show boot*
boot_dev           dka0.0.0.4.1
boot_file          3/boot/vmlinuz
boot_osflags       root=/dev/sda3
boot_reset         ON
bootdef_dev        dka0.0.0.4.1
booted_dev         dka0.0.0.4.1
booted_file        3/boot/vmlinuz
booted_osflags     0
```

The `bootdef_dev` parameter specifies the device from which SRM will try to boot. In this example, SRM will boot from `dka0`, which is the first SCSI hard drive.

The `boot_file` parameter sets the partition, location, and name of the kernel file from which SRM will boot. In this example, `boot_file` is set to boot from the `vmlinuz` kernel, which is located in the `/boot` directory on the third partition of the hard drive. The `boot_file` variable is equivalent to the `-file` value for `boot` on the SRM command line.

The `boot_osflags` parameter can specify which filesystem will be used as the root filesystem. Note that if this parameter is used, it requires Red Hat Linux (instead of SRM) device naming conventions. The `boot_osflags` variable is equivalent to the `-flags` value for `boot` on the SRM command line.

You can use the `set` command to change the boot parameters. After you have finished the installation, you will need to set these variables to the right values, if you want your Alpha to automatically boot Red Hat Linux. Setting the boot parameters to automatically boot Red Hat Linux is covered in *Setting SRM Boot Variables* in Section 4.24.1.

To set a boot parameter, use the following command:

```
set variable "value"
```

When you are setting a value for an SRM console parameter, integer values and single word string values do not have to be in quotes; however, since multiple word string values need to be quoted, it is recommended that you enclose all string values in quotes.

To clear an SRM console variable, use the following command:

```
set variable ""
```

In the above command, *variable* is the name of the SRM console parameter, and the empty quotes ("") denote a blank value. Some integer values cannot be cleared, but can be set to 0. Also note that some SRM console variables are integer in value, but Boolean in nature.

3.3.3 Displaying System Information Using SRM

The SRM console can provide more information about your Alpha. Use the following command:

```
show config | more
```

Note the video card on your system. The video card information will be helpful to you if the Red Hat Linux installation program has any trouble probing your video card.

3.3.4 Booting the Installation Program

Your Alpha system will probably be able to boot the Red Hat Linux installation program directly from the Red Hat Linux CD 1. If your Alpha can not boot the installation program from the CD-ROM (or if you want to perform a hard drive, NFS image, FTP, or HTTP installation) you will need to boot from a diskette. See *Booting the Installation Program from a Diskette* in Section 3.3.4 for more information on booting from a diskette.

Booting the Installation Program from the CD-ROM

Red Hat Linux CD 1 has a valid `aboot` block, so many modern Alphas which use SRM and `aboot` can boot directly from the Red Hat Linux CD 1 shipped in the boxed set. The command to boot from an IDE CD-ROM drive looks like the following:

```
boot dqa0 -flags 0
```

If the previous command is problematic for your system, the format for the full command is as follows:

```
boot dqa0 -file /kernels/vmlinux.gz -flags "root=/dev/hda initrd=images/ramdisk.img"
```

Please note that in the above command, we have used `dqa0` (SRM) and `/dev/hda` (Linux) to refer to an IDE CD-ROM device. Your system may be set up differently, so the command may need to be changed. For example, if your Alpha uses a SCSI CD-ROM device, you might use a command like the following:

```
boot dkb0 -file /kernels/vmlinux.gz -flags "root=/dev/scd0 initrd=images/ramdisk.img"
```

If you would like to perform a text mode installation instead of a GUI installation, you will need to add the `text` flag to the flags at the end of the full boot command. For example, if you wanted to use the previous boot command to boot a text mode installation, you would use the following command:

```
boot dkb0 -file /kernels/vmlinux.gz -flags "root=/dev/scd0 initrd=images/ramdisk.img text"
```

If you want to run the installation program at 640x480 resolution instead of the default 800x600 resolution, you can boot in `lowres` mode with the following command:

```
boot dkb0 -file /kernels/vmlinux.gz -flags "root=/dev/scd0 initrd=images/ramdisk.img lowres"
```

If you want to run the installation program without using frame buffer, you can boot in `nofb` mode with the following command:

```
boot dkb0 -file /kernels/vmlinux.gz -flags "root=/dev/scd0 initrd=images/ramdisk.img nofb"
```

If you are having problems detecting your hardware in the installation program, try using the `noprobe` mode. Instead of trying to autodetect your hardware, it will ask you a series of questions about your hardware. You can boot in `noprobe` mode with the following command:

```
boot dkb0 -file /kernels/vmlinux.gz -flags "root=/dev/scd0 initrd=images/ramdisk.img noprobe"
```

See Section 3.3.1, *SRM Device Names* for more information on how SRM refers to system devices.

Returning to SRM from the `aboot` Prompt

If your boot command does not work for some reason, you will end up at the `aboot` prompt. To return to the SRM console, use the `q` command.

Once the installation program has booted, follow the instructions contained in Chapter 4, *Installing Red Hat Linux* to install Red Hat Linux in graphical mode. Refer to Chapter 5, *Installing Red Hat Linux via Text Mode* to install Red Hat Linux in text mode.

Booting the Installation Program from a Diskette

If your Alpha will not boot the Red Hat Linux installation program directly from Red Hat Linux CD 1, you will need to boot from a diskette. If you want to perform a hard drive, NFS image, FTP, or HTTP installation, you will need to boot from a diskette.

You must create a boot image file diskette from the boot image file on CD 1: `images/generic.img`. You will also need a RAM disk image made from `images/ramdisk.img`, also on CD 1. Instructions for creating disk images are located in Appendix D, *Writing Image Files to Diskettes*. Remember to label the diskettes appropriately.

Put the boot image diskette into your floppy drive. If you are going to load the installation program from the CD-ROM drive, you will need to put CD 1 into the CD-ROM drive. Use the following command to boot the Red Hat Linux installation program:

```
boot dva0 -flags 0
```

This command assumes that the `aboot` bootloader file is on the diskette and that the boot parameters on the diskette are all correct. If you are booting from the diskette provided with Red Hat Linux, this boot command will work.

If for some reason you need the long version of the command to boot from a diskette, it looks like the following:

```
boot dva0 -file vmlinux.gz -flags "root=/dev/fd0 load_ramdisk=1 prompt_ramdisk=1"
```

In the previous command, you will need to modify the parameters to match the location of the kernel file on your diskette and the location of the root filesystem.

After you have booted from the floppy, you will be prompted to take out the boot image diskette and insert the RAM disk diskette with the following message:

```
VFS: Insert root floppy disk to be loaded into RAM disk and press ENTER
```

Insert the RAM disk diskette and press [Enter].

If you are not using the Red Hat Linux CD, the installation program will start in text mode. You will need to choose a few basic options for your system and then you will need to select which installation method you would like to use. See Section 4.2, *Selecting an Installation Method* for more information on different installation methods.

If you are using the CD-ROM to load the installation program, follow the instructions contained in Chapter 4, *Installing Red Hat Linux* to install Red Hat Linux.

The `text`, `lowres`, `nofb`, and `noprobe` modes described in *Booting the Installation Program from the CD-ROM* in Section 3.3.4 also work when booting from a diskette.

4 Installing Red Hat Linux

This chapter explains how to install Red Hat Linux from the CD-ROM using the graphical, mouse-based installation program. The following topics will be 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

4.1 The Installation Program User Interface

If you have used a **graphical user interface (GUI)** before, you will be familiar with this process. If not, 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.

4.1.1 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 4–1, *Console, Keystrokes, and Contents* for a listing of the virtual consoles, keystrokes to switch to them, and their contents.

Table 4–1 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 |

| Console | Keystrokes | Contents |
|---------|-------------------|---------------------|
| 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.

4.2 Selecting an Installation Method

Next, you will be asked what type of installation method you wish to use. If you want to install from a Red Hat Linux CD-ROM and follow the instructions contained in this chapter, choose **CD-ROM**. Alternatively, you can install Red Hat Linux using any of the following basic methods:

CD-ROM

If you have a CD-ROM drive and the Red Hat Linux CD-ROM, use this method. You will need a boot disk or a bootable CD-ROM.

Hard drive

If you copied the Red Hat Linux files to a local hard drive, you can use this method. Refer to Chapter 5, *Installing Red Hat Linux via Text Mode* for hard drive installation instructions. You will need a boot disk.

NFS image

If you are installing from an NFS image server which is exporting the Red Hat Linux CD-ROM or a mirror image of Red Hat Linux, you can use this method. You will need a network boot disk. Refer to Chapter 5, *Installing Red Hat Linux via Text Mode* 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 network boot disk. Refer to Chapter 5, *Installing Red Hat Linux via Text Mode* for FTP installation instructions.

HTTP

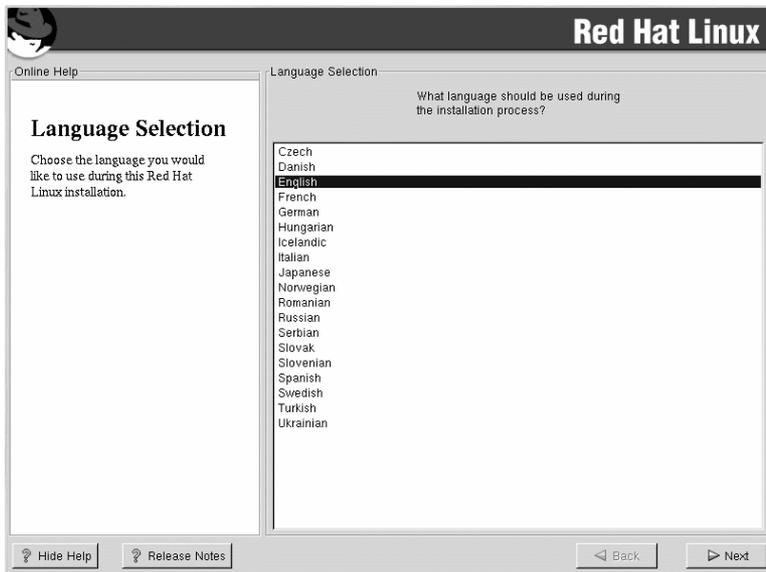
If you are installing directly from a Web (HTTP) server, use this method. You will need a network boot disk. Refer to Chapter 5, *Installing Red Hat Linux via Text Mode* for HTTP installation instructions.

4.3 Language Selection

Using your mouse, select the language you would prefer to use for the installation and as the system default (see Figure 4–1, *Language Selection*).

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.

Figure 4–1 Language Selection



4.4 Keyboard Configuration

Choose the model of keyboard that best fits your system (see Figure 4–2, *Keyboard Configuration*). 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**.

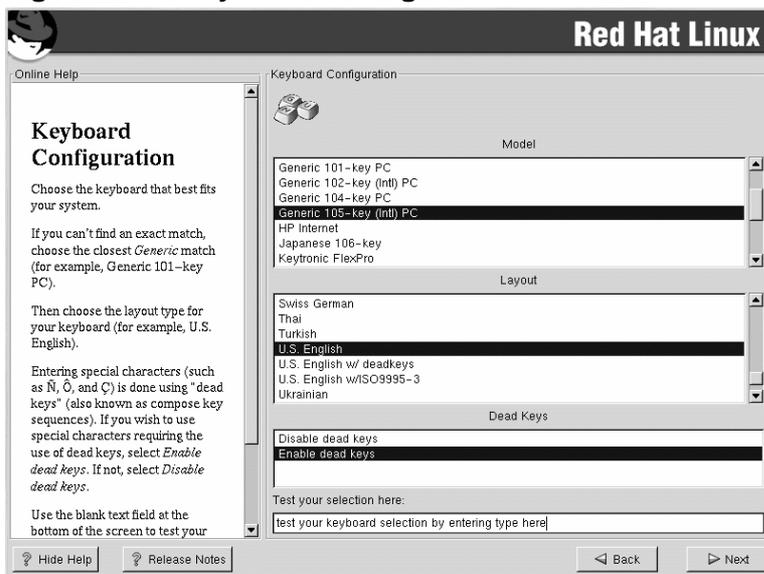
To test your keyboard configuration, use the blank text field at the bottom of the screen to enter text.

Tip

To change your keyboard type after you have installed Red Hat Linux, as root 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 and press [Enter].

Figure 4–2 Keyboard Configuration



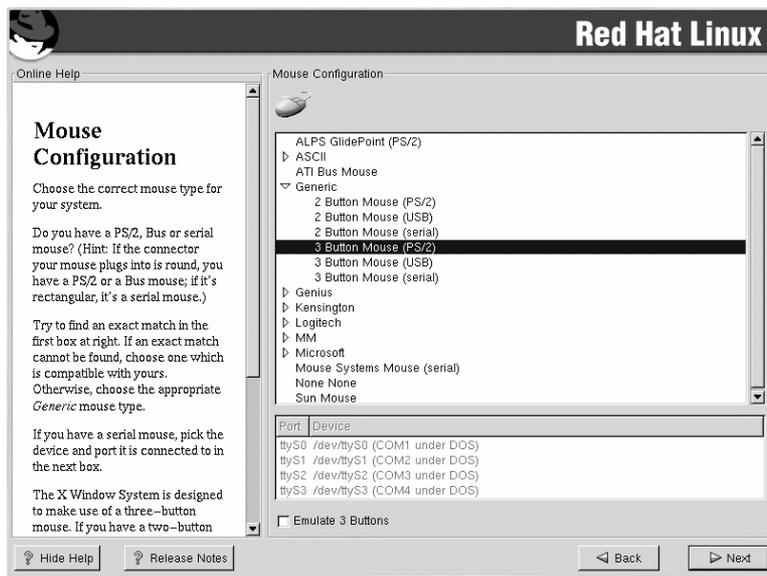
4.5 Mouse Configuration

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 4–3, *Mouse Configuration*).

To determine your mouse's interface, follow the mouse cable back to where it plugs into your system. If the connector at the end of the mouse cable plugs into a rectangular connector, you have a serial mouse; if the connector is round, you have a PS/2 mouse.

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.

Figure 4–3 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.

4.6 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.

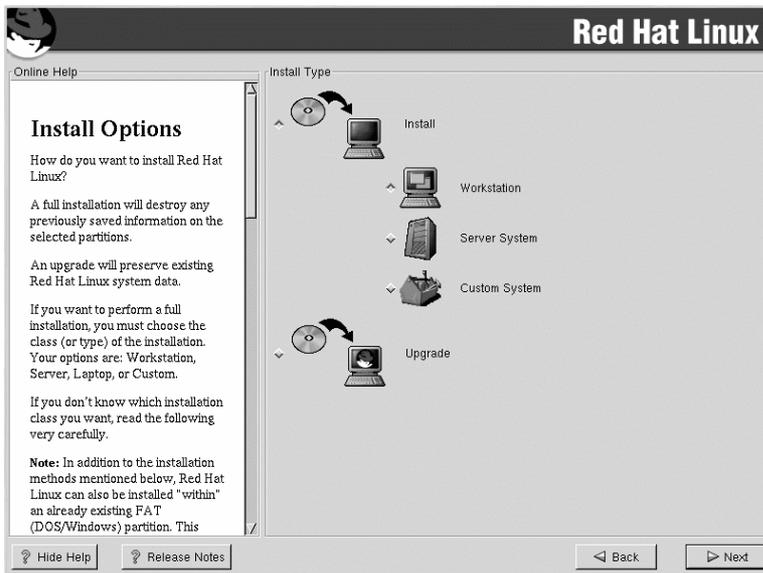
Click on the **Next** button to continue.

4.7 Install Options

Choose whether you would like to perform a full installation or an upgrade (see Figure 4–4, *Choosing Install or Upgrade*).

Red Hat Linux allows you to choose the installation type that best fits your needs. Your options are **Workstation**, **Server**, **Custom**, or **Upgrade**.

Figure 4–4 Choosing Install or Upgrade



To perform an upgrade, please refer to Appendix A, *Upgrading Your Current System*.

4.8 Disk Partitioning

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 *do not* choose to partition manually and instead let the installation program partition for you.

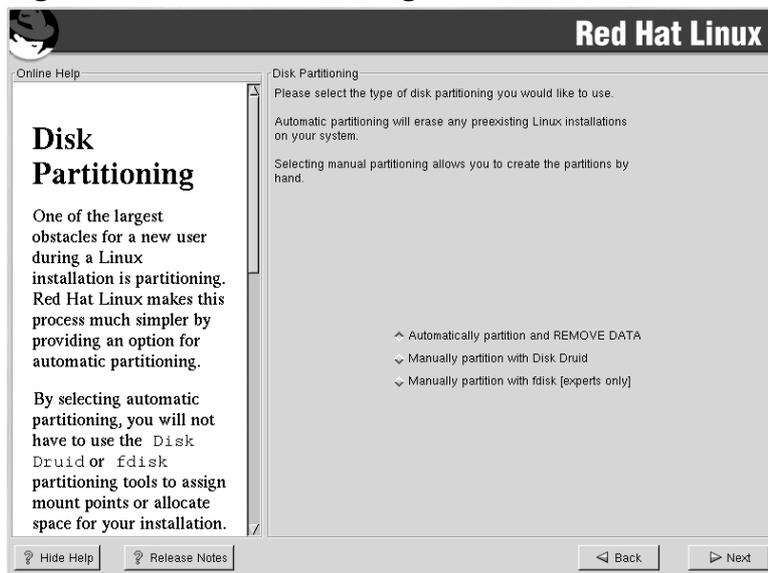
Note

If the **Automatic Partitioning** option does not appear, then you do not have enough free space available for this installation option. You can either exit out of the installation program and make space available, or you can partition manually.

In this screen, you can choose to perform automatic partitioning, to partition manually using Disk Druid, to partition manually using `fdisk`, or use the **Back** button to choose a different installation method (see Figure 4–5, *Disk Partitioning*).

If you do *not* want to lose some or all of your data, choose to partition manually or choose a different installation class.

Figure 4–5 Disk Partitioning





A workstation- or custom-class installation using automatic partitioning will remove all data on all currently existing Linux partitions.



A server-class installation using automatic partitioning will remove all data on all partitions of all hard drives.

If you are not sure how you want your system to be partitioned, please read *An Introduction to Disk Partitions* in the *Official Red Hat Linux Reference Guide*.

4.9 Partitioning Your System

If you chose **Automatic Partitioning**, please skip ahead to Section 4.12, *Network Configuration*.

If you chose **Manually partition with fdisk**, please skip ahead to Section 4.10, *Partitioning with fdisk*.

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 4–6, *Partitioning with Disk Druid*).

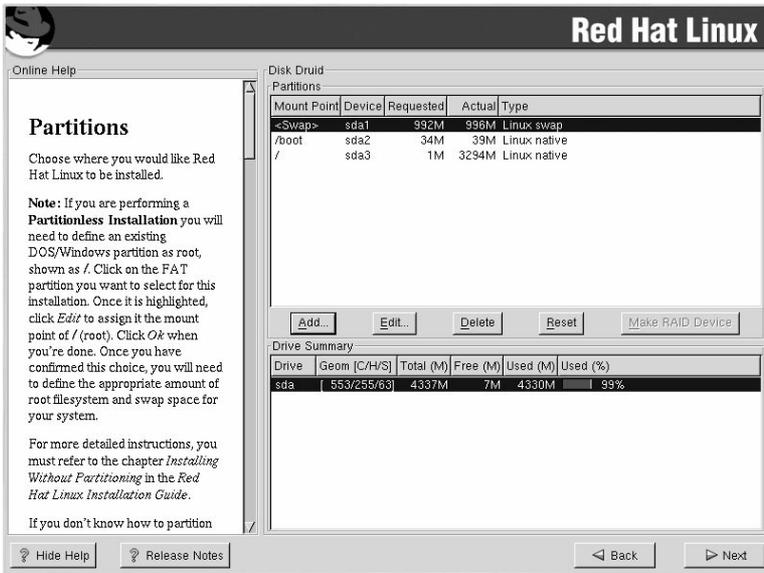
Note

If you have not yet planned how you will set up your partitions, refer to the partitioning appendix in the *Official Red Hat Linux Reference Guide*. At a bare minimum, you will need a `/boot` partition, a swap partition, and an appropriately-sized root (`/`) partition.

Partitionless Installations are Not Officially Supported

Partitionless installations on Alpha machines are not officially supported by Red Hat.

Figure 4–6 Partitioning with Disk Druid



The partitioning tool used in Red Hat Linux 7.1 is Disk Druid. With the exception of certain esoteric situations, Disk Druid can handle the partitioning requirements for a typical Red Hat Linux installation.

SRM has two special partitioning requirements: the use of BSD disk labels and that unallocated space is left at the beginning of the hard drive. Disk Druid is aware of these requirements and will partition your system accordingly — unless your hard drive has a valid MS-DOS partition table. Disk Druid will not convert a hard drive with a valid MS-DOS partition table to use BSD disk labels. To remove the MS-DOS partition table, you will need to use `fdisk` to delete all the partitions and write the partition table to the hard drive. After removing the MS-DOS partition table, you can then partition your hard drives using BSD disk labels.

4.9.1 Partition Fields

Each line in the **Partitions** section represents a disk partition. Each line in this section has five different fields:

- **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.

- **Device:** This field displays the partition's device name.
- **Requested:** This field shows the partition's original size. To redefine the size of a partition, you must delete the current partition and recreate it using the **Add** button.
- **Actual:** This field shows the space currently allocated to the partition.
- **Type:** This field shows the partition's type (for example, Linux Native or DOS).

4.9.2 Recommended Partitioning Scheme

You need to create at least three partitions (or more if you chose to put `/usr` or `/var` in a separate partition) We recommend the following partitions as a minimum:

- A `/boot` partition — the partition mounted on `/boot` will be the SRM boot partition and will contain `about`, the boot loader which will load Red Hat Linux. It is strongly recommended that the `/boot` partition be on the first 2 gigabytes of the hard drive to make sure `about` can read it. In Disk Druid, the partition field for `/boot` should look similar to:

```
/boot sda2 34M 39M Linux native
```

- A swap partition — 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. If you are unsure about what size swap partition to create, make it twice the amount of RAM on your machine. In Disk Druid, the partition field for `swap` should look similar to the following:

```
<Swap> sda1 992M 996M Linux swap
```

- A `root` Linux partition — where `/` (the root directory) is located. In this setup, all files (except those stored in `/boot`) are on the root partition. You may wish to have more than one Linux partition. For example, if you would like to have `/usr` on its own disk partition, you will need to create the `/usr` partition now. In Disk Druid, the partition field for `/` should look similar to:

```
/ sda3 1M 3294M Linux native
```

4.9.3 Problems When Adding a Partition

If you attempt to add a partition and Disk Druid cannot handle your request, you will see a dialog box listing partitions that are currently unallocated, along with the reason they could not be allocated. Unallocated partition(s) are also displayed on Disk Druid's main screen (though you may have to scroll through the **Partitions** section to see them).

As you scroll through the **Partitions** section, you might see an **Unallocated Requested Partition** message in red text, followed by one or more partitions. A common reason for this is a lack of sufficient free space for the partition. In any case, the reason the partition remains unallocated will be displayed after the partition's requested mount point.

To fix an unallocated requested partition, you must move the partition to another drive which has the available space, resize the partition to fit on the current drive, or delete the partition entirely. Make changes using the **Edit** button or by double-clicking on the partition.

4.9.4 Drive Summaries

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

- **Drive:** This field shows the hard disk's device name.
- **Geom [C/H/S]:** This field shows the hard disk's geometry. The geometry consists of three numbers representing the number of cylinders, heads, and sectors as reported by the hard disk.
- **Total:** This field shows the total available space on the hard disk in megabytes.
- **Free:** This field shows how much of the hard disk's space is still unallocated, in megabytes.
- **Used:** These fields show how much of the hard disk's space is currently allocated to partitions, in megabytes and percentage.

The **Drive Summaries** section is displayed only to indicate your computer's disk configuration. It is not meant to be used as a means of specifying the target hard drive for a given partition; to do that, use the **Allowable Drives** field in Section 4.9.6, *Adding Partitions*.

4.9.5 Disk Druid's Buttons

These buttons control Disk Druid's actions. They are used to add and delete partitions, and to change partition attributes. 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.

- **Add:** Used to request a new partition. When selected, a dialog box will appear 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** will open up a dialog box. Some or all of the fields can be edited, depending on whether the partition information has already been written to disk.
 - **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 Device:** **Make RAID Device** 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 Reference Guide*.

4.9.6 Adding Partitions

To add a new partition, select the **Add** button. A dialog box will appear.

Note

You will need to dedicate at least one partition to Red Hat Linux, and optionally more. This is discussed more completely in *An Introduction to Disk Partitions* in the *Official Red Hat Linux Reference Guide*.

- **Mount Point:** Highlight and 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 pulldown menu to choose the correct mount point for your partition.
 - **Size (Megs):** Enter the size (in megabytes) of the partition. Note this field starts with a "1" in it; unless changed you will end up with a 1 MB partition.
 - **Use remaining space:** This checkbox indicates whether the size you entered in the previous field is to be considered the partition's exact size, or its minimum size. When selected, the partition will grow to fill all available space on the hard disk. The partition's size will increase and decrease as other partitions are modified. You can make multiple partitions growable; if you do, the additional free space will be shared among all growable partitions.
 - **Partition Type:** This field contains a list of different partition types (such as Linux Native or DOS). Select the appropriate partition type by using the mouse.
 - **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.
 - **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.
-

4.9.7 Editing Partitions

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

Note

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

4.9.8 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 4.11, *Formatting Partitions* for further installation instructions.

4.10 Partitioning with fdisk

This section applies only if you chose to use `fdisk` to partition your system. If are not using `fdisk`, please skip to Section 4.8, *Disk Partitioning* for automatic partitioning or Section 4.9, *Partitioning Your System* for partitioning with Disk Druid.



Unless you have previously used `fdisk` and understand how it works, we do not recommend that you use it.

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**.

4.10.1 SRM Partitioning Requirements

When you allow the Red Hat Linux installation program to partition for you (automatic partitioning), the installation program will handle certain special partitioning requirements for SRM. Disk Druid is also aware of SRM's partitioning requirements and will partition your system accordingly. If you use `fdisk` to partition manually, however, you will need to be aware of these issues, since you will need to take care of them yourself. Special partitioning issues for SRM include the following:

1. You will need to use BSD-style disklabels. When you use `fdisk` in the graphical installation program, it will start in BSD disklabel mode.
2. You will need to leave unallocated space at the beginning of your hard drive.

You will need to use BSD-style disklabels because of `about`'s partitioning requirements. When you are using the SRM firmware to boot Linux, you are using the `about` boot loader. The `about` program supports the creation of bootable block devices and contains a program which can load Linux kernels from a filesystem that is bootable by SRM.

The `about` program does not understand DOS-style partition tables. You will need to use BSD-style disklabels, which `about` does understand.

BSD-style disk labeling is different in certain ways from standard `fdisk`-style partitioning (used on x86 installations) in the following ways:

- The partitions are identified by letters (a, b, c...) instead of numbers.
- The partition types are different. For example, instead of using "83" and "82" to correspond to "Linux" and "Linux Swap," respectively, ext2 partitions are identified as "8"; the swap partition is "1".
- You can have a maximum of eight partitions (a to h) per drive.
- Partition types must be explicitly set after creation. The BSD disklabel program uses a default partition type of **unknown** which causes the Disk Setup portion of the installation program not to see the partitions.

Another requirement for partitioning is that you provide initial unallocated space at the beginning of the hard drive. The initial unallocated space is unallocated disk space that is not explicitly allocated to any label other than the "whole disk" label. This space should be approximately 1 MB to 10 MB (2048 to 20480 sectors), depending on exactly what configuration of `about` you are using.

The unallocated disk space provides space at the very beginning of the hard drive for `about` to write the boot block and object file that allows booting to kernels located within the disk labels. For more information on the process `about` uses to boot, please see the *SRM Firmware HOWTO* at the following URL:

<http://www.alphalinux.org/faq/srm.html>

WARNING

To provide the initial unallocated space, you will need to start your partitions at cylinder 2. If you do a workstation- or server-class installation and allow the installation program to partition for you, the installation program will provide the initial unallocated space. If you are using Disk Druid, Disk Druid knows about this requirement, and it will not allow you to start the partitions at cylinder 1 during manual partitioning. If you are using `fdisk`, however, you will need to remember to start your partitions at cylinder 2.

4.10.2 Using fdisk

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.

The following table provides the most common `fdisk` commands.

Table 4–2 `fdisk` commands

| Command | What it Does |
|----------------|--|
| <code>b</code> | enters BSD disklabel command line mode |
| <code>m</code> | displays help |
| <code>p</code> | displays the current partition table |
| <code>d</code> | deletes a partition |
| <code>n</code> | creates a new partition |
| <code>w</code> | writes the partition table to disk |
| <code>t</code> | sets the type of partition |
| <code>L</code> | displays the list of partition types |
| <code>q</code> | quits <code>fdisk</code> |

You will need to be aware of a few issues regarding partitioning with `fdisk`. First, Linux `fdisk` creates partitions of type **unknown** by default. When you create partitions, do not forget to change swap partitions to type **Linux swap**, and Linux partitions to type **ext2** using the `t` command. The value for the **Linux swap** type is 1; **ext2** partitions are 8. For other partition types, use the `L` command to see a list of partition types and values.

Linux allows up to four (4) partitions on one disk. If you wish to create more than that, one (and only one) of the four may be an extended partition, which acts as a container for one or more logical partitions. Since it acts as a container, the extended partition must be at least as large as the total size of all the logical partitions it is to contain.

It is a good idea to write down which partitions (for example, `/dev/sda2`) are meant for which filesystems (for example, `/usr`) as you create each one.

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 choose to partition another drive or continue with your installation.

When you are finished partitioning your disks, press **Next**; you may see a message indicating that the installation program needs to reboot. This is a normal occurrence after changing a disk's partition data; it usually happens if you created, changed or deleted any extended partitions. After you press **OK**, your machine will reboot and the installation will begin again. Repeat the same installation steps you performed earlier until you reach the **Partition Disks** dialog; then simply choose **Next**.

Please 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**. Then 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 will be able to edit mount points for the partitions you have already created.

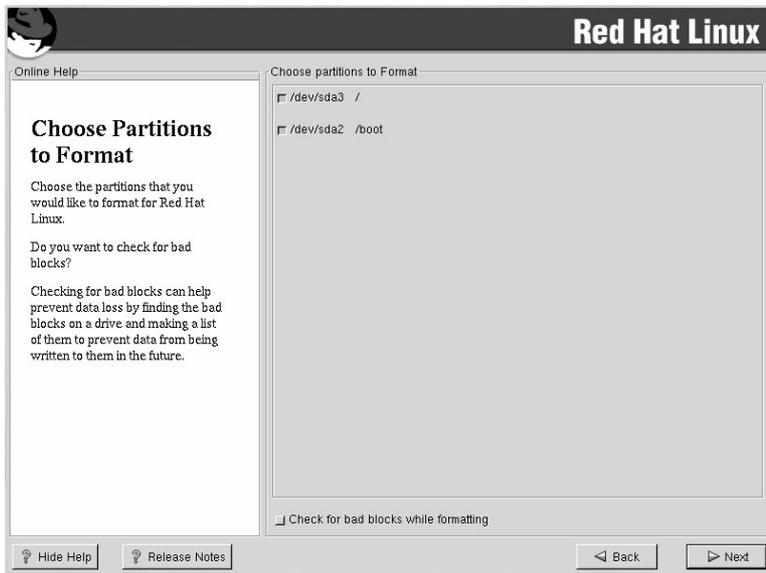
For each partition you created with `fdisk`, click on the **Edit** button, choose the appropriate mount point for that partition from the pulldown menu, and click on **OK**.

4.11 Formatting Partitions

After you have created partitions, you will need to format them. Choose the partitions that you would like to format. All newly created partitions should be formatted. In addition, any existing partitions

that contain data you no longer need should be formatted. However, partitions such as `/home` or `/usr/local` must not be formatted if they contain data you wish to keep (see Figure 4–7, *Choosing Partitions to Format*).

Figure 4–7 Choosing Partitions to Format



If you wish to check for bad blocks while formatting each filesystem, please make sure to select the **Check for bad blocks while formatting** checkbox. 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.

Note

Selecting **Check for bad blocks while formatting** 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.

4.12 Network Configuration

If you do not have a network card, you will not see this screen. Please skip ahead to the next section, Section 4.15, *Time Zone Configuration*.

If you have a network card and you have not already configured your networking, you now have the opportunity (as shown in Figure 4–8, *Network Configuration*).

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.

Choose whether 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.

Figure 4–8 Network Configuration

The screenshot shows the 'Network Configuration' window for the 'eth0' interface. The window title is 'Red Hat Linux'. On the left, there is an 'Online Help' pane titled 'Network Configuration' with the following text: 'Choose your network card and whether you would like to configure using DHCP. If you have multiple Ethernet devices, each device will have its own configuration screen. You can switch between device screens, (for example eth0 and eth1); the information you give will be specific to each screen. If you select *Activate on boot*, your network card will be started when you boot. If you do not have DHCP client access or are unsure as to what this information is, 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'.

The main configuration area for 'eth0' includes:

- Configure using DHCP
- Activate on boot
- IP Address: 192.168.0.1
- Netmask: 255.255.255.0
- Network: 192.168.0.254
- Broadcast: 192.168.0.1
- Hostname: sparky.redhat.com
- Gateway: 192.168.0.1
- Primary DNS: 207.175.42.153
- Secondary DNS: (empty)
- Tertiary DNS: (empty)

At the bottom of the window are buttons for 'Hide Help', 'Release Notes', 'Back', and 'Next'.

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

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 in a name, your system will be known as `localhost`.

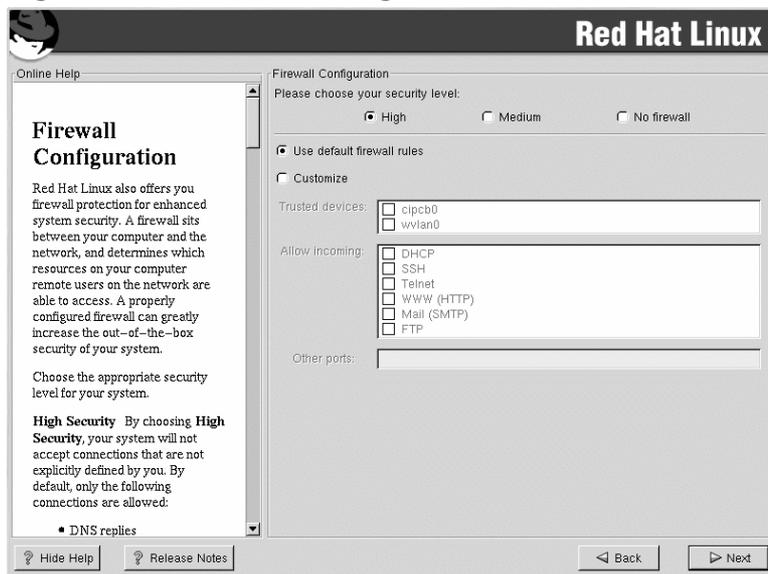
Finally, enter the **Gateway** and **Primary DNS** (and if applicable the **Secondary DNS** and **Tertiary DNS**) addresses.

4.13 Firewall Configuration

Red Hat Linux also offers you 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.

Choose the appropriate security level for your system.

Figure 4–9 Firewall Configuration



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.

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, and HTTP
- The NFS server port (2049)
- 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**. You can select **Customize** to allow specific services through the firewall.

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. It is recommended that this only be selected if you are running on a trusted network (not the Internet), or if you plan to do more detailed 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, but 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-class 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

Secure *SHell* (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 Web pages. 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 Web pages. You will need to install the `apache` package if you want to serve Web pages.

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

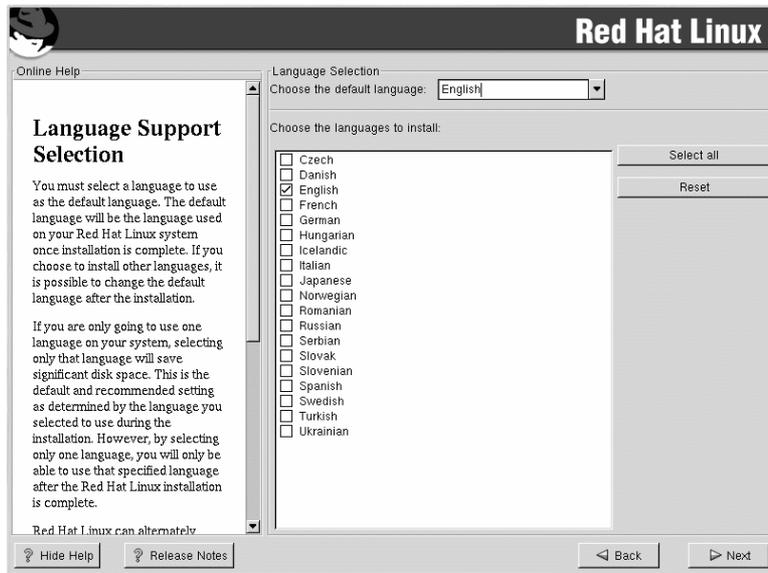
You can allow access to ports which are not listed here, by listing them in **Other ports** field. Use the following format: `port:protocol`. For example, if you want to allow IMAP access through your firewall, you can specify `imap:tcp`. You can also explicitly specify numeric ports; to allow UDP packets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

4.14 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.

Figure 4–10 Language Support Selection



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.

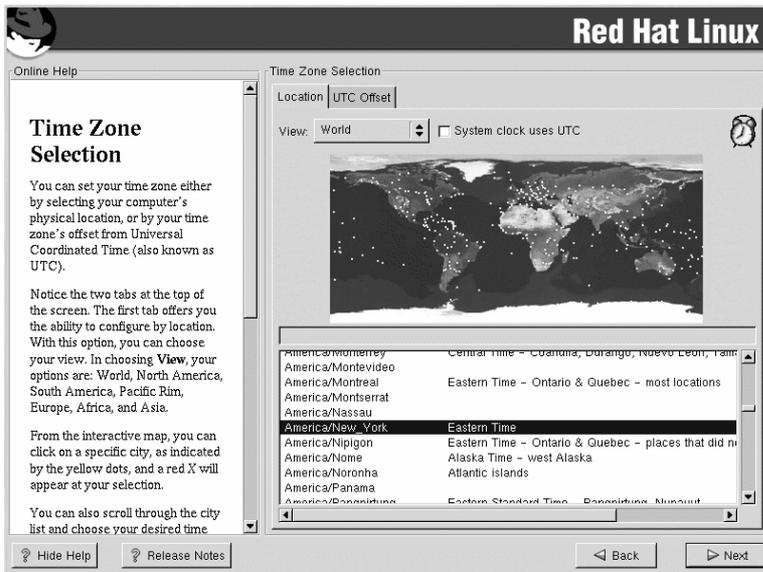
To use more than one language on your system, choose specific languages to be installed or select all languages to have all available languages installed on your Red Hat Linux system.

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.

4.15 Time Zone Configuration

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

Figure 4–11 Configuring the Time Zone



Notice the two tabs at the top of the screen (see Figure 4–11, *Configuring the Time Zone*). 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.

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

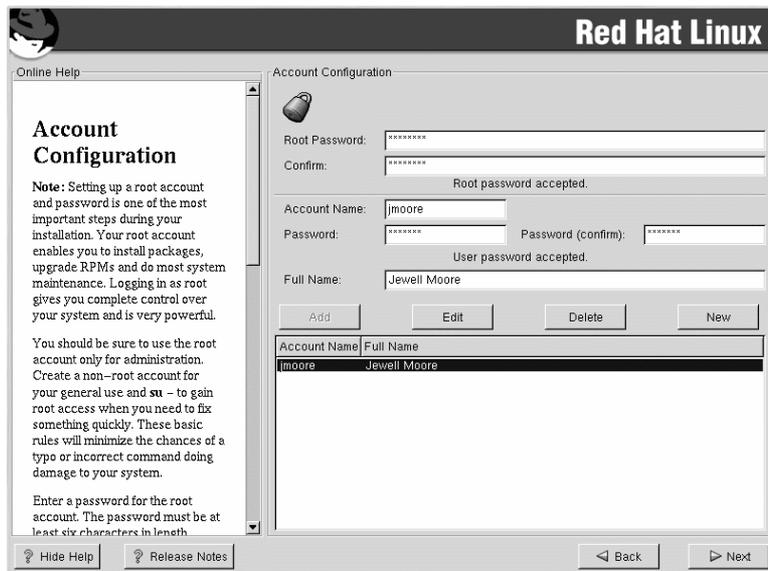
Tip

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/time-config` command.

4.16 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 4–12, *Account Creation*).

Figure 4–12 Account Creation



4.16.1 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.

The installation program will prompt you to set a root password¹ 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 **420BMttNT**, for example. Remember that the password is case-sensitive. If you write down your 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.

4.16.2 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.

¹ 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.

You can also choose **New** to add a new 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.

4.17 Authentication Configuration

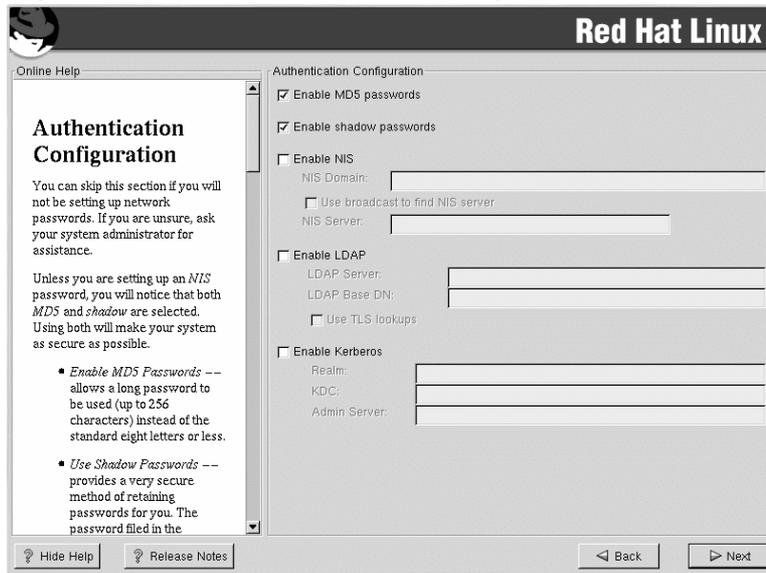
If you are performing a workstation- or server installation, please skip ahead to Section 4.18, *Package Group Selection*.

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 4–13, *Authentication Configuration*). 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 4–13 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 two 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.
- **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 *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`.

4.18 Package Group Selection

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

Note

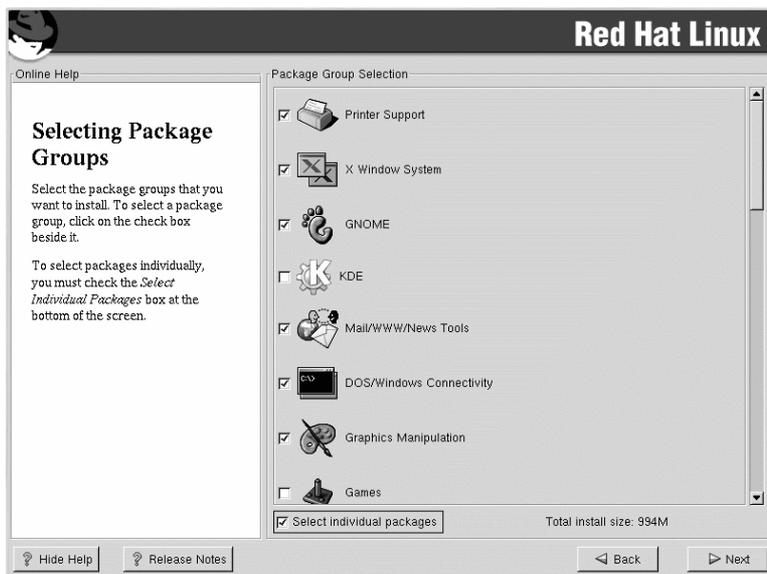
Unless you choose a custom-class installation, the installation program will automatically choose most packages for you. However, you must select either GNOME or KDE (or both) to install a graphical environment.

GNOME and KDE are both graphical desktop environments² 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.

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 4–14, *Package Group Selection*).

Figure 4–14 Package Group Selection



² 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.

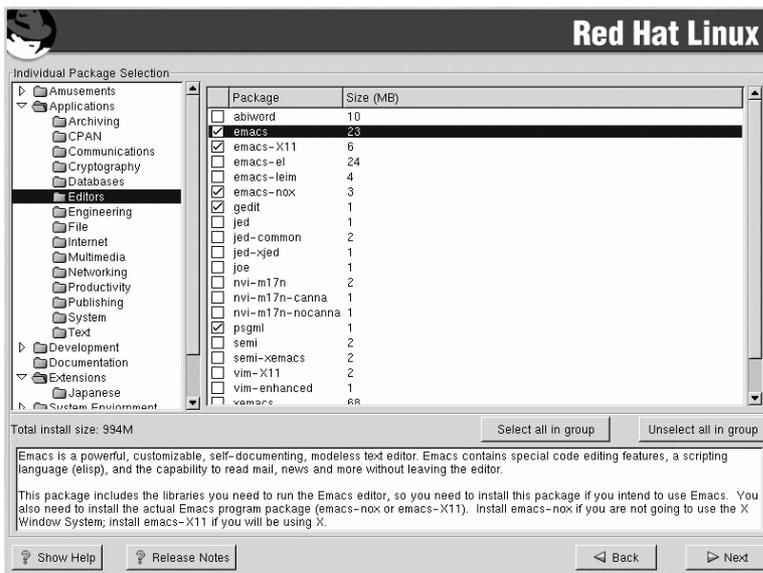
Select each component you wish to install. Selecting **Everything** (at the end of the component list) during a custom-class installation installs all packages included with Red Hat Linux. If you select every package, you will need approximately 1.7 GB of free disk space.

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

4.18.1 Selecting Individual Packages

After selecting the components you wish to install, you can select or deselect individual packages. You can select or deselect packages using your mouse (see Figure 4–15, *Selecting Individual Packages*).

Figure 4–15 Selecting Individual Packages



On the left side of the screen you will see a directory listing of various package groups. When you expand a directory (double-click to select it) and click on a single directory, the list of packages available for installation will appear in the panel on the right.

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.

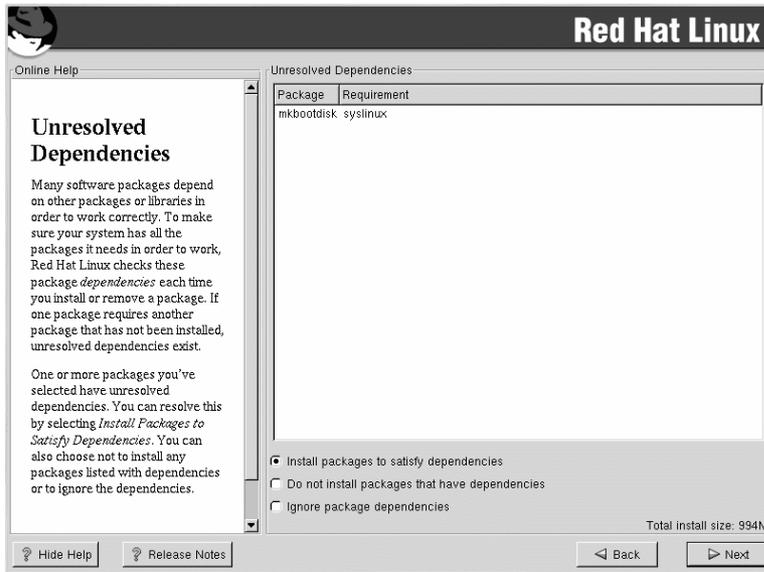
4.18.2 Unresolved Dependencies

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 4–16, *Unresolved Dependencies*).

The **Unresolved Dependencies** screen will only appear 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 package dependencies automatically by adding all required packages to the list of selected packages.

Figure 4–16 Unresolved Dependencies



4.19 GUI X Configuration Tool

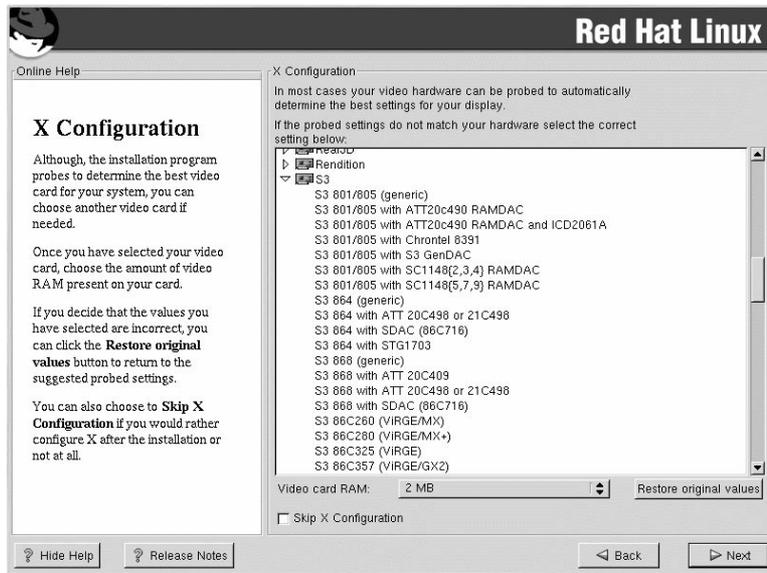
If you decided to install the X Window System packages, you now have the opportunity to configure an X server for your system. If you did not choose to install the X Window System packages, skip ahead to Section 4.21, *Installing Packages*.

4.19.1 Video Card Configuration

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

If your video card does not appear on the list (see Figure 4–17, *Video Card Setup*), 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 4–17 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.

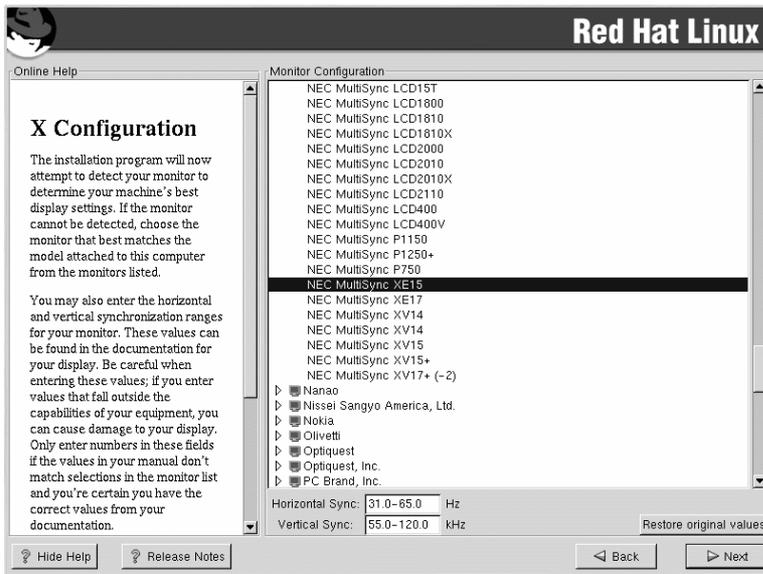
If you decide that the values you have selected are incorrect, you can click the **Restore original values** button to return to the suggested settings.

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

4.19.2 Configuring Your Monitor

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 autodetected for you, or choose another monitor.

Figure 4–18 Monitor Selection



If your monitor does not appear on the list, select the most appropriate **Generic** model available. If you do select 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.



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.

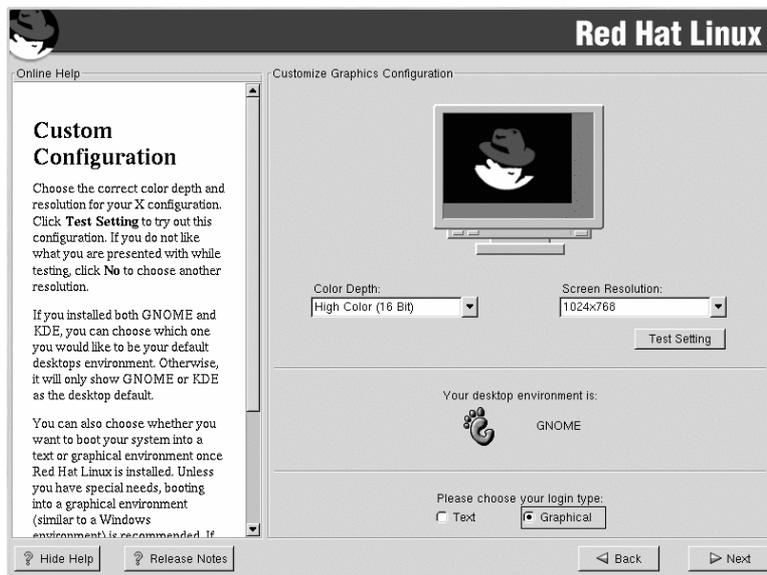
4.19.3 Custom Configuration

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.

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.

Figure 4–19 X Customization



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).

4.20 Preparing to Install

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

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 your hard drive(s) are rewritten.

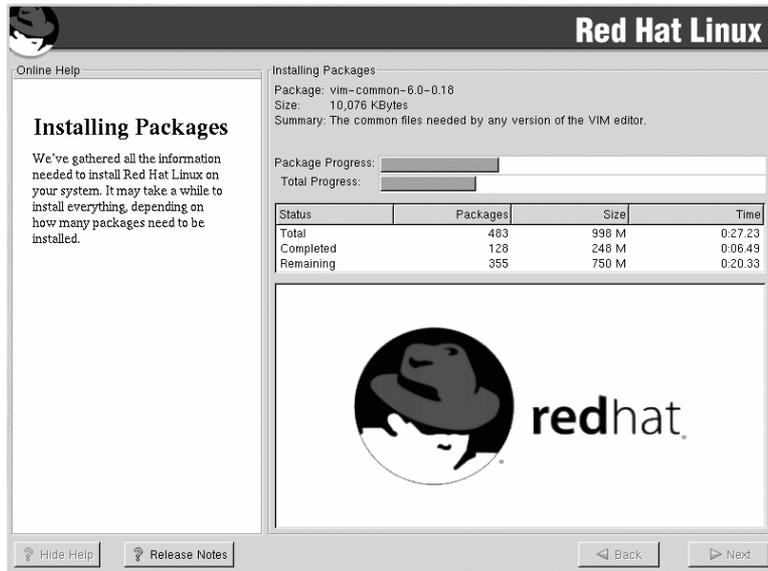
To cancel this installation process, press your computer's Reset button or use the [Control]-[Alt]-[Delete] key combination to restart your machine.

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

4.21 Installing Packages

At this point there is nothing left for you to do until all the packages have been installed (see Figure 4-20, *Installing Packages*). How quickly this happens depends on the number of packages you have selected and your computer's speed.

Figure 4–20 Installing Packages



4.22 Note the Kernel Filename and Partition Number

If you are using SRM to boot your Alpha, you should take a moment to write down the kernel filename and partition number. This is necessary because `about` does not include an `ls` command, and if you encounter problems (for example, if `/etc/about.conf` is not found on the root filesystem on the booted disk or if it contains incorrect information), you would have to guess at the kernel name to boot your machine.

Switch to the second virtual console using `[Alt]-[F2]` in text mode, or using `[Ctrl]-[Alt]-[F2]` in X. Type in the following command:

```
bash# cat /mnt/sysimage/etc/about.conf
```

You should see output similar to the following:

```
# about default configurations
0:2/boot/vmlinuz-2.2.17-4 root=/dev/sda2
```

Use Table 4–3, *Kernel Location* to write down the kernel filename and the location of the root filesystem. Hopefully, you will not need to use it, but you will be prepared. Switch back to the installation program using `[Ctrl]-[Alt]-[F7]`.

Table 4–3 Kernel Location

| Path to and Filename of Kernel | Location of Root Filesystem |
|--------------------------------|-----------------------------|
| | |

4.23 Installation Complete

Congratulations! Your Red Hat Linux 7.1 installation is now complete!

The installation program will prompt you to prepare your system for reboot. Do not forget to remove any diskette in the diskette drive or CD in the CD-ROM drive.

For the commands you will need to boot Red Hat Linux, see Section 4.24, *Booting Your Machine and Post-Installation Setup*.

4.24 Booting Your Machine and Post-Installation Setup

This section describes how to boot your Alpha into Red Hat Linux and how to set your SRM console variables so that Red Hat Linux is automatically booted when the machine is powered on.

4.24.1 Recommended Boot Method

Once the installation has been completed, your machine should be capable of booting in SRM directly from the hard drive, using a command like the following:

```
boot dka0 -f1 0
```

In this example, `dka0` is the hard drive on which the `/boot` partition resides (the SRM device on which the installation program installed the `aboot` secondary boot loader).

Setting SRM Boot Variables

After you have installed Red Hat Linux, at the SRM prompt, use the `show boot*` command to display the list of devices recognized by SRM. The system hard drive should be displayed as the `bootdef_dev`. If it is not, set `bootdef_dev` to the hard drive with a command like the following:

```
set bootdef_dev dka0
```

In the previous example, `dka0` may be a different value, depending upon your system. Refer to Section 3.3.1, *SRM Device Names* if you need more information on how SRM refers to system devices.

Set the `boot_osflags` SRM console variable with the following command:

```
set boot_osflags 0
```

Unset the `boot_file` variable with the following command:

```
set boot_file ""
```

After setting these variables, your Alpha should then boot Red Hat Linux in response to the command `boot` at the SRM console prompt.

4.24.2 Alternate Boot Method

If the `boot dka0 -fl 0` command is problematic for your machine, the complete format of the command for booting from SRM is as follows:

```
boot <device> -file <boot file> -flags "<boot flags>"
```

In the above command, `<device>` is the SRM device on which the `aboot` secondary boot loader is installed. `<boot file>` is the uncompressed kernel file that you installed. You should have noted these values during the installation process, as described in Section 4.22, *Note the Kernel Filename and Partition Number*. You will need to preface the kernel filename with the number of the partition where the kernel resides, and include the full path to the kernel.

Also in the above command, `<boot flags>` are the root device (the device mounted as `/`) and any other kernel flags that you need to pass.

For example, if you installed `aboot` on the first SCSI device on the second SCSI bus (`dkb0`), the root of your Linux filesystem is the third partition of your first SCSI drive (`/dev/sda3`), and you installed version 2.2.17-4 of the Linux kernel, then the boot command to boot your Alpha is:

```
boot dkb0 -file 3/boot/vmlinuz-2.2.17-4 -flags "root=/dev/sda3"
```

Setting SRM Boot Parameters Using Alternate Boot Method

After you have installed Red Hat Linux, at the SRM prompt, use the `show boot*` to display the list of devices recognized by SRM. Set `bootdef_dev` to the hard drive with a command like the following (note that this is an example; you will need to provide the correct values for your machine):

```
set bootdef_dev dkb0
```

In the previous example, `device` should be set to your hard drive, as recognized by SRM. Refer to Section 3.3.1, *SRM Device Names* if you need more information on how SRM refers to system devices.

Set the `boot_osflags` SRM console variable to the boot flags you need with a command like the following:

```
set boot_osflags root=/dev/sda3
```

Set the `boot_file` variable with the following command:

```
set boot_file 3/boot/vmlinuz-2.2.17-4
```

In the above command, you will need to provide the location of the boot file, including the partition and the path to the file.

After setting these variables, your Alpha should then boot Red Hat Linux in response to the command `boot` at the SRM console prompt.

4.24.3 Booting Red Hat Linux Automatically

After you have set your boot parameters according to either Section 4.24.1, *Recommended Boot Method* or Section 4.24.2, *Alternate Boot Method*, and if you have a HALT button on your machine, you can use the `auto_action` variable to automatically boot Red Hat Linux upon power-up.

If you have a HALT button, you can set the SRM console to autoboot using the `auto_action` SRM console variable. The `auto_action` parameter can be set to `HALT` or `BOOT` to set the default power-on function of the Alpha machine when it is booted into the SRM console.

If you set `auto_action` to `BOOT` and your other SRM boot variables are set correctly, your Alpha will automatically boot Red Hat Linux when booted into the SRM console. If you set `auto_action` to `halt`, your Alpha will stop at the SRM console when it is powered on.

Does Your Alpha Have a HALT Button?

When setting the value of the `auto_action` SRM console variable, note that your machine will not be able to return to the SRM console unless you press the HALT button on the machine. Some Alpha machines do not have a HALT button. Do not set the `auto_action` value to `BOOT` unless your Alpha machine has a HALT button (which will allow you to return to SRM).

Other methods for halting your machine may work, depending upon your hardware. You can try pressing `[Ctrl]-[c]` several times in a row. You can also try the `halt` command in Red Hat Linux to halt your machine and return to the SRM console. Neither of these methods are guaranteed to work on your machine.

If you want to set your machine to automatically boot, and you have a HALT button, use the following command:

```
set auto_action BOOT
```

5 Installing Red Hat Linux via Text Mode

This release of Red Hat Linux features a graphical, mouse-based installation program, but you can also install Red Hat Linux using a text mode, keyboard-based installation program. This chapter briefly explains how to use the text mode installation program. Here are some recommendations:

- If you are new to Linux installations, read Chapter 4, *Installing Red Hat Linux* first. The main focus of the chapter is the graphical installation process, but most of the concepts apply to the text mode installation as well. After reading that chapter, Section 5.1, *Things You Should Know* will give you more in-depth information regarding the aspects of installing Red Hat Linux that do not apply to the graphical installation process.

Additionally, the *An Introduction to Disk Partitions* appendix in the *Official Red Hat Linux Reference Guide* may be helpful to you, since it discusses disk partition resizing. If you plan to install Red Hat Linux on a disk where another operating system is currently installed, this knowledge will be crucial.

- If you plan to install over a network (via NFS, FTP, or HTTP), you must make a network boot disk. Appendix D, *Writing Image Files to Diskettes* explains how to do this.
- If you have never used the text mode installation program, or need a refresher on its user interface, read the next sections.
- To begin installation without further delay, turn to Section 5.3, *Starting the Installation Program*.

5.1 Things You Should Know

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

The most recent list of supported hardware can be found at <http://hardware.redhat.com>. You should check your hardware against this list before proceeding.

Note

You can perform a text mode installation of Red Hat Linux 7.1 by following the instructions in this chapter. However, if you are installing from a CD-ROM, you might prefer to use the graphical installation mode, which is easy to use and provides a flexible, custom-class installation mode. For more information on graphical installations, turn to Chapter 4, *Installing Red Hat Linux*.

5.1.1 Basic Hardware Configuration

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

- Hard drive(s) — specifically, the number, size, and type. If you have more than one, it is helpful to know which one is first, second, and so on. It is also good to know if your drives are IDE or SCSI. If you have IDE drives, you should check your computer's BIOS to see if you are accessing them in linear mode. Please refer to your computer's documentation for the proper key sequence to access the BIOS. Note that your computer's BIOS may refer to linear mode by other names, such as "large disk mode." Again, your computer's documentation should be consulted for clarification.
- Memory — the amount of RAM installed in your computer.
- CD-ROM — most importantly, the unit's interface type (IDE, SCSI, or other interface) and, for non-IDE, non-SCSI CD-ROMs, the make and model number. IDE CD-ROMs (also known as AT-API) are the most common type of CD-ROM in recently manufactured, PC-compatible computers.
- SCSI adapter (if one is present) — the adapter's make and model number.
- Network card (if one is present) — the card's make and model number.
- Mouse — the mouse's type (serial, PS/2, or bus mouse), protocol (Microsoft, Logitech, MouseMan, etc.), and number of buttons; also, for serial mice, the serial port it is connected to.

On many newer systems, the installation program is able to automatically identify most hardware. However, it is a good idea to collect this information anyway, just to be sure.

5.1.2 Video Configuration

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

- Your video card — the card's make and model number (or the video chipset it uses) and the amount of video RAM it has. (Most PCI-based cards are auto-detected by the installation program.)
- Your monitor — the unit's make and model number, along with allowable ranges for horizontal and vertical refresh rates. (Newer models may be auto-detected by the installation program.)

5.1.3 Network-related Information

If you are connected to a network, be sure you know the following:

- IP address — usually represented as a set of four numbers separated by dots, such as 10 . 0 . 2 . 15.
 - Netmask — another set of four numbers separated by dots; an example netmask would be 255 . 255 . 248 . 0.
-

- Gateway IP address — another set of four dot-separated numbers; for example, 10.0.2.254.
- One or more name server IP addresses — one or more sets of dot-separated numbers; for example, 10.0.2.1 might be the address of a name server.
- Domain name — the name your organization uses; for example, Red Hat has a domain name of redhat.com.
- Hostname — the name assigned to your individual system; for example, a computer might be named pooh.

Note

The information provided here is as an example only! Do *not* use it when you install Red Hat Linux! If you do not know the proper values for your network, ask your network administrator.

5.2 The 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 5-1, *Installation Program Widgets as seen in Configure TCP/IP* and Figure 5-2, *Installation Program Widgets as seen in Disk Druid* illustrate the screens you will see.

Figure 5-1 Installation Program Widgets as seen in Configure TCP/IP

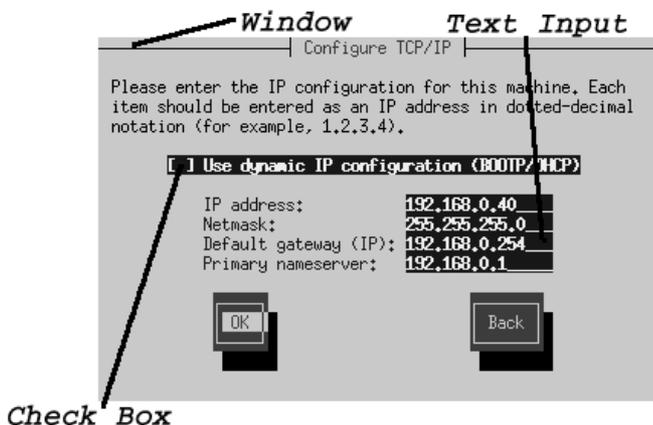
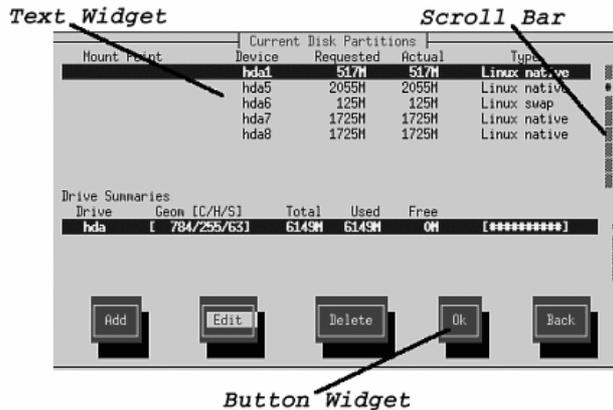


Figure 5–2 Installation Program Widgets as seen in Disk Druid



Here is a list of the most important widgets shown in Figure 5–1, *Installation Program Widgets as seen in Configure TCP/IP* and Figure 5–2, *Installation Program Widgets as seen in Disk Druid*:

- **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 5–1, *Installation Program Widgets as seen in **Configure TCP/IP***, the cursor is positioned on the **OK** button. Figure 5–2, *Installation Program Widgets as seen in **Disk Druid*** shows the cursor on the **Edit** button.

5.2.1 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.



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).

5.3 Starting the Installation Program

The following methods can be used to start the installation:

- Insert a boot diskette that you have created (or network boot diskette that you have created) into the primary diskette drive and reboot your computer.
- Insert the Red Hat Linux CD 1 into the drive and reboot, if your computer can boot from the CD-ROM drive.

Refer to Section 3.3.4, *Booting the Installation Program* for further instructions on how to start the installation program in an Alpha system.

5.3.1 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.

5.4 Installation Cross-reference Table

Most text mode installation screens can be cross-referenced with their GUI screen counterparts. This table lists the installation screens in order and where you can go to get the necessary information for each part of the installation.

For those screens which do not have adequate GUI counterparts, installation-related instructions and screenshots will be provided in this chapter.

Table 5–1 Installation Cross-reference Table

| Text Mode Screen | Text Mode Reference Point |
|--|---|
| Language Screen | Section 4.3, <i>Language Selection</i> |
| Keyboard Screen | Section 4.4, <i>Keyboard Configuration — Layout</i> |
| Installation Method | Section 4.2, <i>Selecting an Installation Method</i> |
| Identify Disk Partition to Install From | Section 5.5, <i>Identify Disk Partition to Install From</i> |
| Installing over a Network | Section 5.6, <i>Installing over a Network</i> |
| Welcome Screen | Section 4.6, <i>Welcome to Red Hat Linux</i> |
| Installation Type | Section 4.7, <i>Install Options</i> |
| Upgrade Only — Customize Packages to Upgrade | Section A.3, <i>Customizing Your Upgrade</i> |
| Automatic Partitioning | Section 4.8, <i>Disk Partitioning</i> |
| Disk Druid | Section 4.9, <i>Partitioning Your System</i> |
| fdisk | Section 4.10, <i>Partitioning with fdisk</i> |
| Hostname Configuration | Section 4.12, <i>Network Configuration</i> |
| Firewall Configuration | Section 4.13, <i>Firewall Configuration</i> |
| Network Configuration | Section 4.12, <i>Network Configuration</i> |

| Text Mode Screen | Text Mode Reference Point |
|--|--|
| Mouse Selection | Section 4.5, <i>Mouse Configuration</i> |
| Language Support and Default Language | Section 4.14, <i>Language Support Selection</i> |
| Time Zone Selection | Section 4.15, <i>Time Zone Configuration</i> |
| Root Password | Section 4.16, <i>Account Configuration</i> |
| Add User/User Account Setup | Section 4.16, <i>Account Configuration</i> |
| Authentication Configuration | Section 4.17, <i>Authentication Configuration</i> |
| Package Installation (Group, Individual, Dependencies) | Section 4.18, <i>Package Group Selection</i> |
| X Configuration — Video Adapter | Section 5.7, <i>Configuring Your Video Adapter</i> |
| Package Installation | Section 4.20, <i>Preparing to Install</i> and Section 4.21, <i>Installing Packages</i> |
| Configuring the X Window System | Section 5.8, <i>Configuring the X Window System</i> |
| Video Card — Manual Selection | Section 4.19.1, <i>Video Card Configuration</i> |
| Installation Complete | Section 4.23, <i>Installation Complete</i> |

5.5 Identify Disk Partition to Install From

Note

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

If you verify that the ISO images are intact before you attempt an installation, you will help to avoid problems that are often encountered during a hard drive installation. To verify the ISO images are intact before 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.

The **Select Partition** screen (Figure 5–3, *Selecting Partition Dialog for Hard Drive Installation*) 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.

Figure 5–3 Selecting Partition Dialog for Hard Drive Installation



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. See Table 5–1, *Installation Cross-reference Table* for more information.

5.6 Installing over a Network

If you are performing a network installation, the **Configure TCP/IP** dialog appears; for an explanation of this dialog, go to Section 4.12, *Network Configuration* and then return here.

Because the Red Hat Linux 7.1 installation program is capable of installing Red Hat Linux 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 7.1 onto a disk drive:

- Insert CD-ROM 1 and execute the following commands:
 - `mount /dev/cdrom /mnt/cdrom`

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

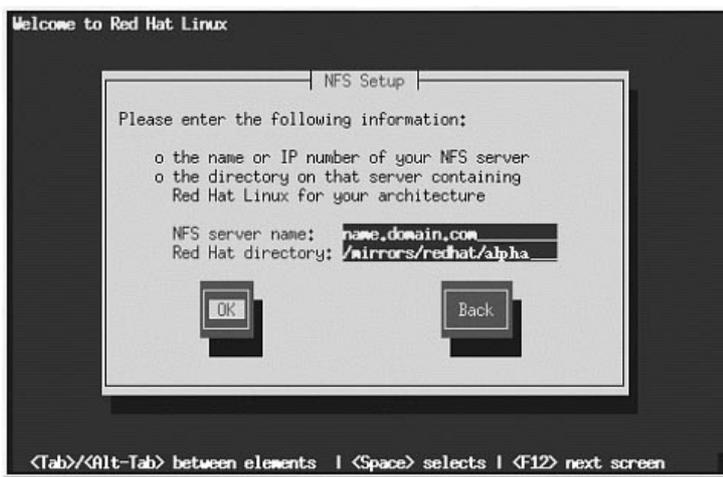
- Insert CD-ROM 2 and execute the following commands:
 - `mount /dev/cdrom /mnt/cdrom`
 - `cp -var /mnt/cdrom/RedHat /location/of/disk/space`
 - `umount /mnt/cdrom`

- You must then make the directory containing the RedHat directory accessible to the installation program (for example, exporting it for NFS installations).

5.6.1 NFS Setup

The NFS setup dialog (Figure 5–4, *NFS Setup Dialog*) applies only if you are installing from an NFS server (that is, if you selected **NFS Image** in the **Installation Method** dialog).

Figure 5–4 NFS Setup Dialog



Enter the fully-qualified domain name or IP address of your NFS server, and the name of the exported directory that contains the Red Hat Linux installation files or CD. 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.

If the NFS server has the Red Hat Linux CD mounted on `/mnt/cdrom`, enter `/mnt/cdrom` in the **Red Hat directory** field. If the NFS server is exporting a mirror of the Red Hat Linux installation tree instead of a CD, enter the directory which contains the RedHat directory. For example, if your NFS server contains the directory `/mirrors/redhat/alpha/RedHat`, enter `/mirrors/redhat/alpha`.

Note

If you are performing an NFS installation and are copying the files from the Red Hat Linux CD-ROM, be sure to check the file permissions to make sure they are set correctly for your installation. If you do not, the files that you copy will not be executable and you will have to change the permissions before you are able to install.

Next you will see the **Welcome** dialog. See Table 5-1, *Installation Cross-reference Table* for more information.

5.6.2 FTP Setup

The **FTP Setup** screen (Figure 5-5, *FTP Setup Dialog*) applies only if you are installing from an FTP server (that is, 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.

Figure 5–5 FTP Setup Dialog



Enter the fully-qualified domain name or IP address of the FTP site you are installing from, and the name of the directory on that FTP site which contains the RedHat installation files for your architecture. For example, if the FTP site contains the directory `/mirrors/redhat/alpha/RedHat`, enter `/mirrors/redhat/alpha`.

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

Next you will see the **Welcome** dialog. See Table 5–1, *Installation Cross-reference Table* for more information.

5.6.3 HTTP Setup

The **HTTP Setup** screen (Figure 5–6, *HTTP Setup Dialog*) applies only if you are installing from an HTTP server (that is, 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.

Figure 5–6 HTTP Setup Dialog

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

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

Next you will see the **Welcome** dialog. See Table 5–1, *Installation Cross-reference Table* for more information.

5.7 Configuring Your Video Adapter

The installation program then probes for your video card. If your card is detected, click **OK**.

Figure 5–7 Video Card Confirmation

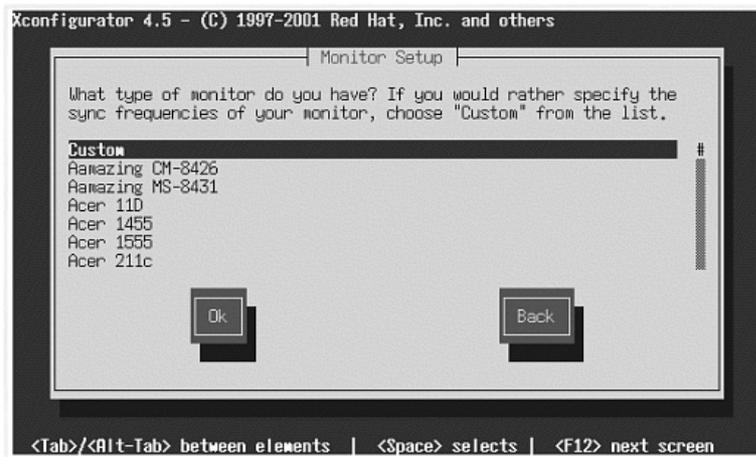
If the installation program cannot detect your video adapter, you will need to select your video card from the available list. See Section 4.19.1, *Video Card Configuration* for more information.

5.8 Configuring the X Window System

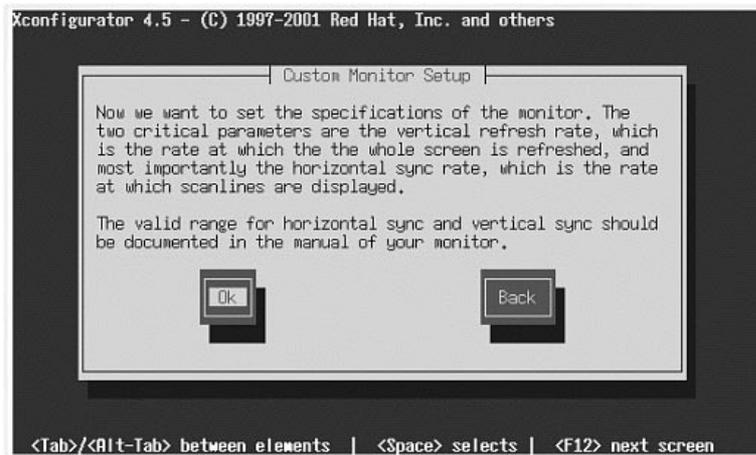
The Xconfigurator utility gives you the opportunity to configure the X server for your system.

First, Xconfigurator displays a list of monitors (see Figure 5–8, *Monitor Setup Dialog*). If your monitor is listed, select it and press [Enter]. Otherwise, select **Custom**.

If you select a monitor from the list, you will need to enter your video memory. See Figure 5–13, *Video Memory*.

Figure 5–8 Monitor Setup Dialog

If you select **Custom**, Xconfigurator prompts you to select the horizontal sync range and vertical sync range of your monitor (these values are usually available in the documentation which accompanies your monitor, or from your monitor's vendor or manufacturer).

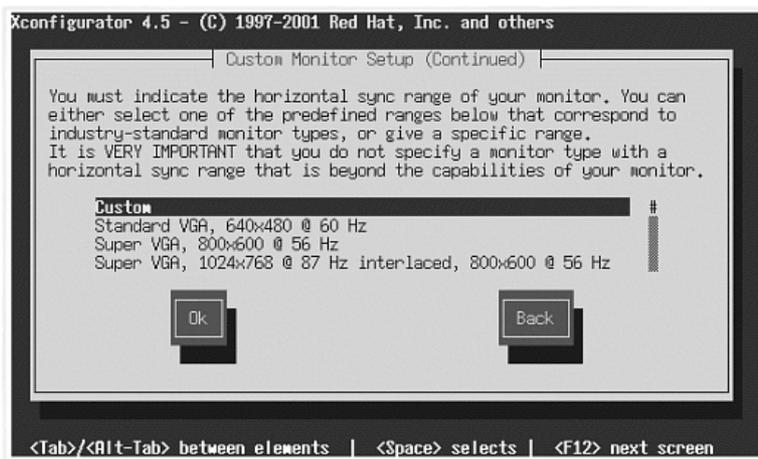
Figure 5–9 Custom Monitor Setup



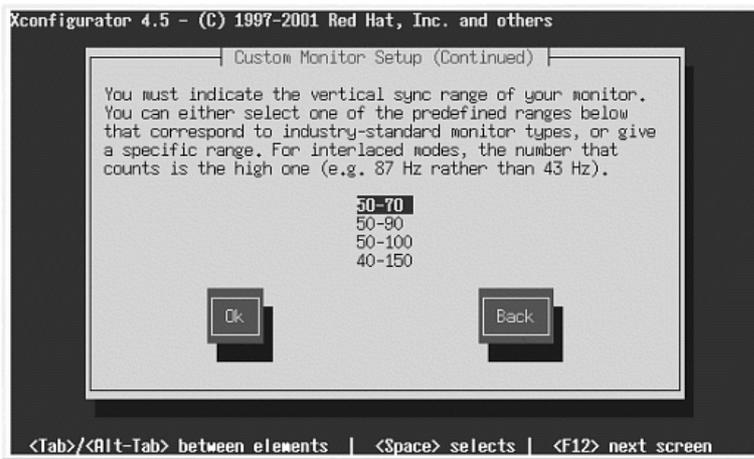
Do not select any range that exceeds the capacity of your monitor. If you do this, it is possible you may overclock your monitor and damage or destroy it.

Choose a horizontal sync range for your monitor (see Figure 5–10, *Custom Monitor Setup - Horizontal Sync*). Be sure *not* to choose a range outside of your monitor's capacity. For custom settings, refer to Figure 5–12, *Custom Monitor Information*.

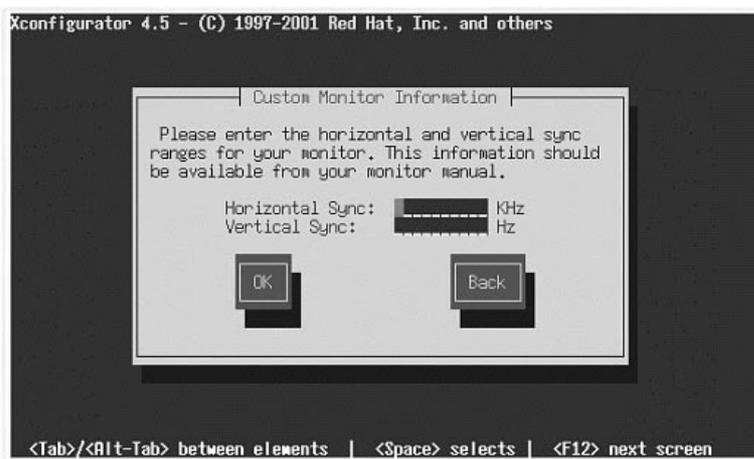
Figure 5–10 Custom Monitor Setup - Horizontal Sync



Next, choose a vertical sync range.

Figure 5–11 Custom Monitor Setup — Vertical Sync

If you chose to enter custom monitor settings, please enter both the horizontal and vertical sync ranges for your monitor (see Figure 5–12, *Custom Monitor Information*).

Figure 5–12 Custom Monitor Information

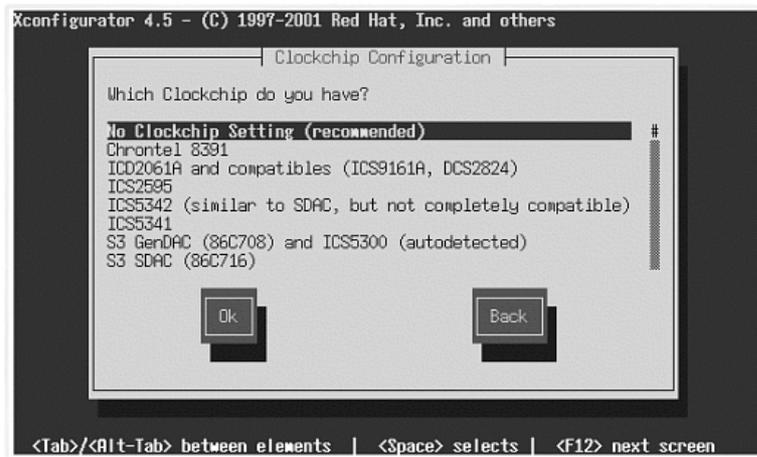
After you have selected your monitor (and its settings), you must choose the amount of memory provided by your video card (see Figure 5–13, *Video Memory*).

Figure 5–13 Video Memory

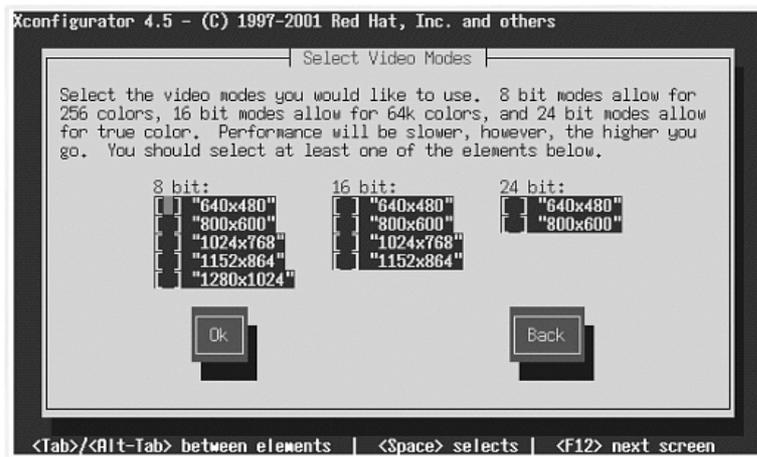


If you are not sure how much memory your card has, consult the documentation accompanying your video card. Choosing more memory than is present in your card will not damage it, but may cause the XFree86 server to start incorrectly.

Next, choose a clockchip setting if your video card supports it (if it has one). The recommended choice is **No Clockchip Setting**, since XFree86 can automatically detect the proper clockchip in most cases.

Figure 5–14 Clockchip Configuration Dialog

To choose your preferred display mode (see Figure 5–15, *Select Video Modes*), select one or more modes by pressing [Space].

Figure 5–15 Select Video Modes

Tip

The minimum recommended settings for video modes are 16 bit for most applications in 1024 x 768. Video modes set higher than 16 bit can run more slowly.

Do not select a mode that exceeds the capacity of your monitor.

Once you have either selected the default video mode or specified a different mode, the **Starting X** dialog lets you test your X configuration.

Figure 5–16 Test Your X Configuration



If you select **OK**, you will have the opportunity to go back and enter different configuration values, if an error occurs. If no error occurs, X will start and a small display window will ask you if you can read its contents clearly; using the mouse, click **Yes** within ten seconds. Then you will be asked whether you want X to start when the system boots.

In most cases, X configuration is complete at this point, and you will see the completion screen.

Xconfigurator then saves all of your choices to the configuration file `/etc/X11/XF86Config` or `/etc/X11/XF86Config-4`. To modify your X configuration after installation, issue the Xconfigurator command as root.

A Upgrading Your Current System

This appendix walks you through a typical Red Hat Linux 7.1 upgrade.

A.1 What it Means to Upgrade

The installation process for Red Hat Linux 7.1 includes the ability to upgrade from prior versions of Red Hat Linux (version 3.0.3 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 `.rpmsave` extension (for example, `sendmail.cf.rpmsave`). The upgrade process also creates a log of its actions in `/tmp/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.

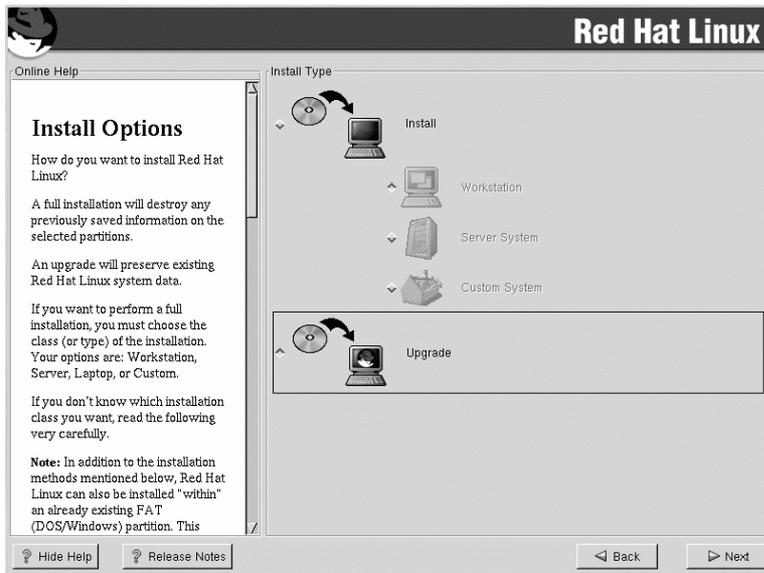
Note

Some upgraded packages may require that other packages are also installed 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.

A.2 Upgrading Your System

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

Figure A-1 Choosing Install or Upgrade



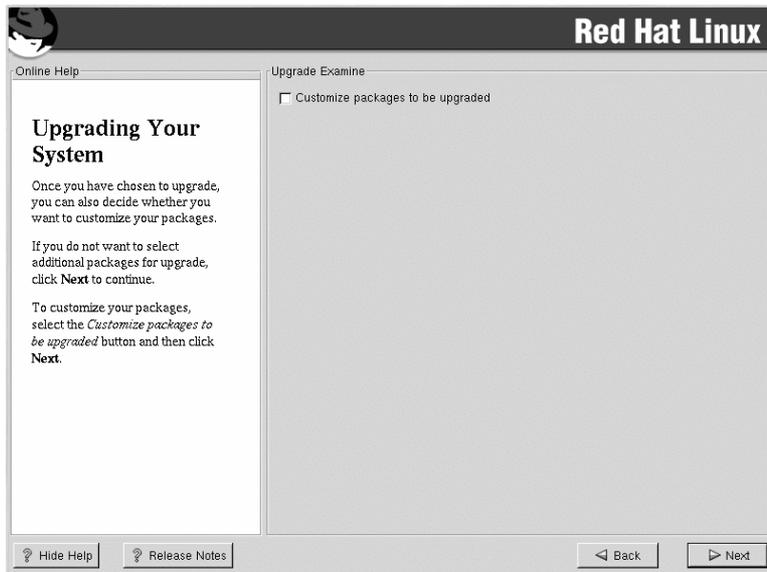
A.3 Customizing Your Upgrade

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

To let the installation program begin the upgrade process for you, make sure **Customize packages to be upgraded** is *not* selected and click **Next** (see Section A.5, *Upgrading Packages*).

To customize your upgrade packages, select this option and then click **Next**.

Figure A–2 Upgrade Customization



A.4 Selecting Packages to Upgrade

On this screen, you can choose which packages you would like to upgrade (see Figure A–3, *Individual Package Selection*).

On the left side of the screen you will see a directory listing of various package groups. When you expand this list (double-click on a package group, such as **Applications**, to select it) and double-click on a single directory, the list of packages available for installation will appear on the right.

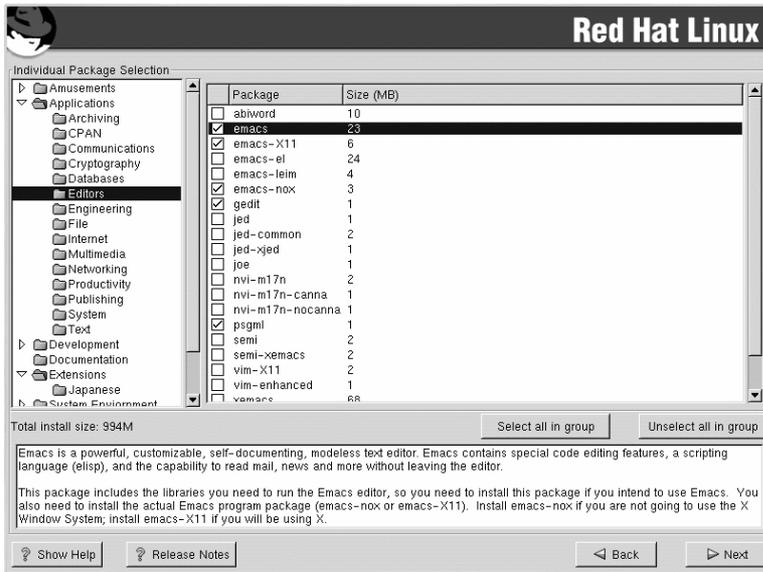
To select an individual package, double-click the checkbox beside the package name. A check mark in the box means that a package has been selected.

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.

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-3 Individual Package Selection

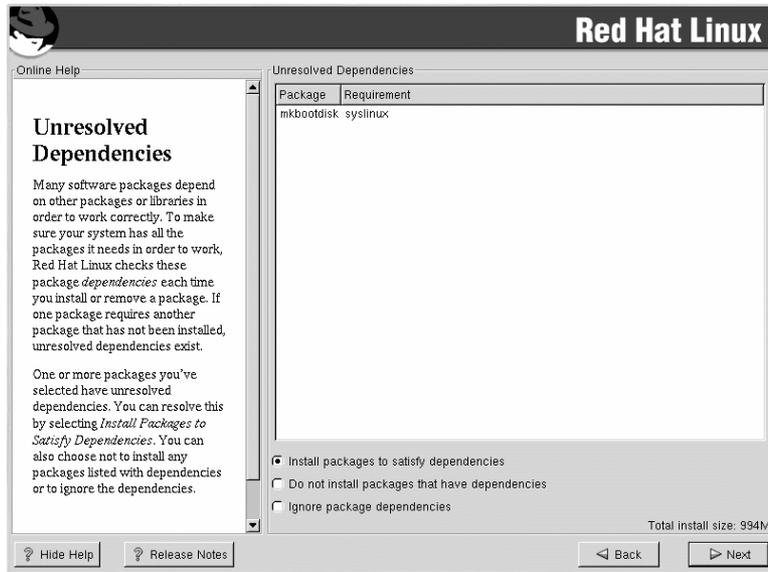


A.4.1 Unresolved Dependencies

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-4, *Unresolved Dependencies*).

The **Unresolved Dependencies** screen will only appear 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** checkbox 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.

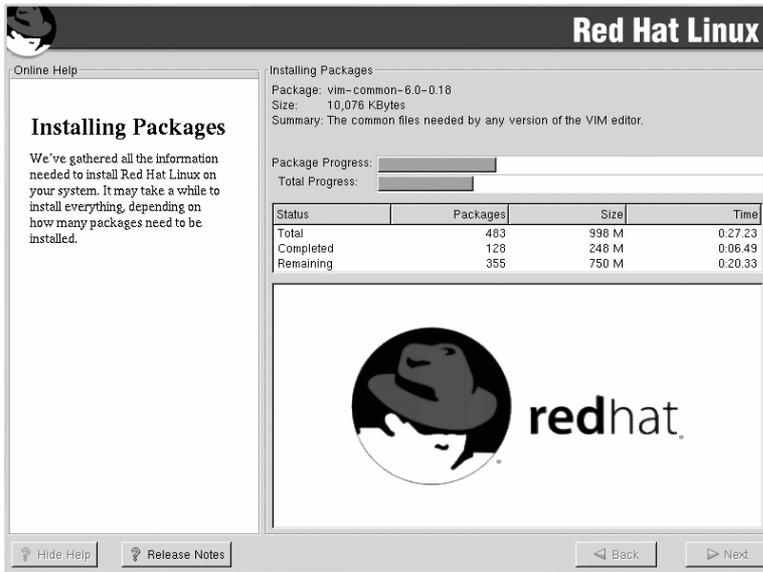
Figure A–4 Unresolved Dependencies



A.5 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–5, *Installing Packages*).

Figure A-5 Installing Packages



A.6 Upgrade Complete

Congratulations! Your Red Hat Linux 7.1 upgrade is now complete!

If you are going to boot your Alpha from the SRM console, take the time to note the kernel filename using the instructions in Section 4.22, *Note the Kernel Filename and Partition Number*, before you reboot your machine.

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. See Section 4.24, *Booting Your Machine and Post-Installation Setup* for information on how to boot Red Hat Linux.

B MILO

This chapter provides information that might be useful for those who want to install Red Hat Linux on an Alpha which boots using MILO, the Alpha Miniloader, instead of the SRM console.

WARNING

Red Hat does not provide support for the MILO miniloader. If you use MILO to boot your Alpha, you will not receive technical support from Red Hat.

You may need more information to help you with MILO. The *Alpha Miniloader HOWTO*, which contains in-depth information about the MILO miniloader, is available from [ftp://gatekeeper.dec.com/](ftp://gatekeeper.dec.com/pub/Digital/Linux-Alpha/Miniloader/docs/HOWTO/milo.howto.html), in the file `pub/Digital/Linux-Alpha/Miniloader/docs/HOWTO/milo.howto.html`.

Please Note

Many people will find the previous information to be very useful. However, this information is provided by sources outside of Red Hat, Inc. and it contains descriptions of procedures that are not supported by Red Hat, Inc.

B.1 MILO Installation Overview

Installing Red Hat Linux on an Alpha system using MILO is slightly more complex than installing Red Hat Linux on an Intel machine. In general, the sequence of steps to a successful installation are:

1. Create a boot image diskette, a MILO image diskette and a RAM disk image diskette.
2. Load and run the Red Hat Linux/Alpha kernel, and boot into the Red Hat Linux installation program using MILO.
3. Install the bootloader on a small partition on your machine.

B.2 MILO and Alpha Hardware

If you are reading this chapter, presumably you are intent upon using MILO to boot your Alpha. As you have already read, you will not receive any technical support from Red Hat if you are using MILO. Therefore, by definition, any hardware that you are using is unsupported.

Disk space requirements for installing Red Hat Linux on an Alpha are listed in Section 3.1.1, *Disk Space Requirements*.

B.3 Choosing Disk Images

To install Red Hat Linux on an Alpha using MILO, you will need to create three diskettes: a boot image diskette, a RAM disk diskette and a MILO image diskette. Create a boot image file diskette from the boot image file on CD 1: `images/generic.img`. You will also need a RAM disk image made from `images/ramdisk.img`, also on CD 1.

The MILO images are located in the `miло/images` directory on CD 1, in compressed `gzip` format. You will need to choose the appropriate image for your system, uncompress it and write it to a diskette. If you do not know which image to choose for your system, see Table B-1, *MILO Images*.

Table B-1 MILO Images

| Machine Class | Model | Alias | MILO Image |
|------------------------------|-----------------|---------|-------------------------|
| Alcor | AS 600 | ALCOR | <code>alcor.img</code> |
| | AS 500 5/5xx | BRET | <code>xlt.img</code> |
| | XL-300 | XLT | |
| | XL-366 | XLT | |
| | XL-433 | XLT | |
| Digital Alpha XL | XL-233 | XL | <code>xl.img</code> |
| | XL-266 | XL | |
| Avanti | AS 200 4/xxx | MUSTANG | <code>avanti.img</code> |
| | AS 205 4/xxx | LX3 | |
| | AS 250 4/xxx | M3 | |
| | AS 255 4/xxx | LX3+ | |
| | AS 300 4/xxx | MELMAC | |
| | AS 400 4/xxx | AVANTI | |
| Digital Personal Workstation | PWS 433a, 433au | MIATA | <code>miata.img</code> |
| | PWS 500a, 500au | MIATA | |
| | PWS 600a, 600au | MIATA | |

| Machine Class | Model | Alias | MILO Image |
|---------------|---------------|-----------|---------------|
| Mikasa | AS 1000 4/xxx | MIKASA | mikasa.img |
| Noname | AXPpci33 | NONAME | noname.img |
| | UDB/Multia | MULTIA | |
| EB64+ | EB64+ | EB64+ | eb64p.img |
| | AlphaPC64 | CABRIOLET | cabriolet.img |
| | AlphaPCI-64 | CABRIOLET | |
| PC164 | AlphaPC164 | PC164 | pcl64.img |
| | AlphaPC164-LX | LX164 | lx164.img |
| | AlphaPC164-SX | SX164 | sx164.img |
| EB164 | EB164 | EB164 | eb164.img |
| EB66+ | EB66+ | EB66+ | eb66p.img |
| EB66 | EB66 | EB66 | eb66.img |
| Takara | Takara | TAKARA | takara.img |
| Ruffian | AlphaPC 164UX | RUFFIAN | ruffian.img |
| AlphaBook1 | | | book1.img |
| Platform2000 | | | p2k.img |

Look through Table B-1, *MILO Images* for your machine or its alias. The appropriate MILO image (if applicable) is listed on the right. All of the MILO image files are located in the `miло/images` directory on Red Hat Linux CD 1.

Try the MILO images included on Red Hat Linux CD 1 first. MILO images are also available from the following URL:

`ftp://gatekeeper.dec.com/pub/Digital/Linux-Alpha/Miniloader/latest-Images/`

The subdirectories in this directory correspond to different Alpha systems. Within the subdirectory, the MILO image file is named `miло.dd`. You will need to write this file to a diskette.

Refer to Appendix D, *Writing Image Files to Diskettes* for instructions on how to make diskettes from disk images. Label all three diskettes appropriately.

B.4 Disk Partitioning

The only special rule for partitioning the hard drive for MILO is that there should be a 2 MB FAT partition at the beginning of the drive where MILO will be installed.

B.5 Installations Using MILO

To boot Red Hat Linux using MILO, you will need to use some version of either the ARC or AlphaBIOS console, plus the `linload.exe` program.

The ARC and AlphaBIOS consoles are designed to load the Windows NT operating system. Like many other things in the Linux world, they're being used to do a job that they were never intended to do. In this case, they're used to load Linux. When you power-on an ARC/AlphaBIOS console machine, you will see the firmware initialize the hardware, and you will see a boot menu. The initialization process varies greatly by machine.

The ARC and AlphaBIOS consoles operate in essentially the same way. Both use MILO and, aside from cosmetic changes, they are identical in functionality. The keystrokes described in each section have been tested against ARC and AlphaBIOS systems, and are known to work with on-site machines. However, small differences may exist for other ARC and AlphaBIOS systems.

The Windows NT ARC/AlphaBIOS firmware is an environment in which programs (for example, the Windows NT `osloader`) can run and make callbacks into the firmware to perform actions.

Linux's `linload.exe` is a simple program which does just enough to load and execute MILO. `linload.exe` loads the appropriate image file into memory at `0x00000000` and then makes a swap-PAL PALcall to the image file. The swap is necessary because MILO, like Linux, uses a different PALcode than Windows NT. MILO relocates itself to `0x200000` and continues on through the PALcode reset entry point as before.

B.5.1 Setting up ARC for Installation

If you are using MILO and ARC to boot the installation program from the CD-ROM, you will only need a MILO image diskette.

If you are using MILO and ARC, but you are not booting the installation program from the CD-ROM, you will need to make three diskettes:

1. The correct MILO image diskette for the class of machine onto which you are installing Red Hat Linux. The MILO images, in compressed format, are located in the `miло/images` directory on CD 1. See Section B.3, *Choosing Disk Images* if you are not sure which MILO image to use.
 2. A generic kernel diskette, made from `images/generic.img` on CD 1.
 3. The RAM disk diskette, made from `images/ramdisk.img` on CD 1.
-

Please refer to Appendix D, *Writing Image Files to Diskettes* for instructions on writing image files to diskettes.

Once you have created these diskettes, boot your Alpha machine into the ARC boot menu. The first step is to set up a boot menu entry to enable MILO to be loaded from the diskette. At the boot menu, select the option **Supplementary menu**.

Next, select the command: **Set up the system**, which takes you to the **Setup Menu**. From here, select the command **Manage boot selection menu**, which takes you to the **Boot Selections Menu**.

You need to add a boot command to load MILO from the diskette you created. Choose the command **Add a boot selection**.

You should see several devices listed. To select the A: drive, choose **Floppy Disk 0**.

Please Note

If you do not see an entry for the diskette drive, consult your AlphaBIOS/ARC installation manual to determine the correct method for configuring your hardware in the console.

Next, enter the name of the OS loader to use. Enter `\linload.exe` as the OS loader directory and name. The MILO disk you created has the `linload.exe` file on it, and the correct MILO for your Alpha machine class. The MILO command is always named `miilo` on the diskette, and both files should reside in the root directory of the diskette.

Select the **yes** option when you are asked if the operating system is on the same partition as the OS loader (it is, both are in the root directory), and enter `\` as the operating system root directory. As the name of the boot selection, enter something descriptive like **MILO Diskette** to indicate that this entry boots MILO from the diskette drive.

If prompted, select **no** to indicate that you do not want to start the debugger at boot time.

You should now be back in the boot selections menu. Press [Esc] to get back to the **Boot Selections Menu**. Choose the option **Supplementary menu, and save changes** to commit the changes to NVRAM. Once you have saved your boot entry, [Esc] will get you back to the boot menu and you can now attempt to boot MILO.

You should have a boot selection that looks something like this:

```
LOADIDENTIFIER=Linux
SYSTEMPARTITION=multi(0)disk(0)fdisk(0)
OSLOADER=multi(0)disk(0)fdisk(0)\linload.exe
OSLOADPARTITION=multi(0)disk(0)fdisk(0)\
OSLOADOPTIONS=
```

This boot selection enables ARC to boot into MILO from the diskette using the `linload.exe` OS loader.

B.5.2 Setting up AlphaBIOS for Installation

First, you will need to install the latest version of AlphaBIOS, which can be obtained from Compaq's Alpha Firmware update site at

`http://ftp.digital.com/pub/DEC/Alpha/firmware/`

As with ARC, if you can boot from the CD-ROM, you will only need to make a MILO disk. If your machine can not boot from the CD-ROM, you will need three diskettes for the MILO installation process:

1. The correct MILO diskette for the class of machine onto which you are installing Red Hat Linux. The MILO images, in compressed format, are located in the `miло/images` directory on CD 1. See Section B.3, *Choosing Disk Images* if you are not sure which MILO image you need.
2. A generic installation kernel diskette, from `images/generic.img` on CD 1.
3. A RAM disk diskette, made from `images/ramdisk.img` on CD 1.

Please refer to Appendix D, *Writing Image Files to Diskettes* for instructions on writing image files to diskettes.

Once your AlphaBIOS is at the latest revision level, you can start the installation process. Turn on the system and insert the MILO diskette. At the opening screen, press [F2] to enter setup mode. You are going to add a boot selection that allows you to boot MILO from a diskette. Select the following command from the menu: **Utilities -> OS Selection Setup.**

Press [Insert] to add a new operating system selection. For the **Boot Name** parameter, enter something like: **MILO Floppy** to indicate that you are loading MILO from a diskette. Press [Tab] to get to the next field. Use the [DOWN-ARROW] key to move through the options until the selection for **Boot File** is **A:**. Then, [Tab] to the next field and enter: **linload.exe** as the name of the OS loader.

[Tab] past the **OS Path Load Device** option, since it's irrelevant and `linload.exe` ignores it. [Tab] to the **OS Path Load File** option. Enter: \ for the OS Path load file. Finally, press [Enter] to add the selection to the boot menu.

At this point, AlphaBIOS may display an error message such as `Warning: Operating System Selection not valid!`

Ignore this error (AlphaBIOS considers it an error whenever your OS is not Windows NT). Press [Enter] to continue. Press [F10] to save the changes you just made, and press [Enter] to confirm the changes.

Press [Esc] twice to get back to the opening screen. Use the [UP-ARROW] and [DOWN-ARROW] keys to select the boot selection you just added, and press [Enter] to boot it.

AlphaBIOS will load `linload.exe`, which will in turn load MILO.

B.5.3 The MILO User Interface

Once the Windows NT firmware is running and you have the correct MILO image for your system loaded, the rest of the installation details are completely generic.

Once you boot into MILO, you get a prompt that is very familiar to the standard `bash#` prompt in Linux. As mentioned earlier, MILO uses native Linux device drivers; therefore, all devices in MILO are referred to in the same manner as a running Linux system (`/dev/scd0` for a SCSI CD-ROM 0, `/dev/hda` for the first IDE device, etc.).

MILO has a very simple interface that is designed to allow you to boot a Linux kernel image, and perform some basic diagnostic functions. Typing **help** is a good idea, since it provides a useful summary of the commands, as follows:

```
MILO> help
MILO command summary: ls [-t fs] [dev:[dir]]
                        - List files in directory on device
boot [-t fs] [dev:file] [boot string]
-
Boot Linux from the specified device and file
run [-t fs] dev:file
    - Run the standalone program dev:file
show
-
Display all known devices and file systems
set VAR VALUE
-
Set the variable VAR to the specified VALUE
unset VAR
- Delete the specified variable
reset
- Delete all variables
print
- Display current variable settings
help [var]
- Print this help text
Devices are specified as: fd0, hda1, hda2, sda1...
Use the '-t filesystem-name' option if you want
to use anything but the default filesystem ('ext2').
Use the 'show' command to show known devices and file-
systems.
Type 'help var' for a list of variables.
```

Please Note

The `bootopt` command only appears on AlphaPC64 (and similar) systems. Refer to the board's documentation to find out just what it means.

Devices

Until you use a command that needs to make use of a device, no device initialization will take place. The first time you issue the `show`, `ls`, or `run` command, the devices within MILO will be initialized. Devices are named exactly the same way that Linux names devices. So, the first IDE drive will be called `hda` and its first partition will be `hda1`. Use the `show` command to show what devices are available.

Filesystems

MILO supports three filesystems: MSDOS, ext2 and ISO9660. If a device is available to it, MILO can `ls`, `boot` or `run` an image stored on one of these filesystems. MILO will assume ext2 as the default filesystem, and so you have to explicitly tell MILO if the filesystem is something other than ext2. All of the commands that use filenames allow you to pass the file system using the `-t [filesystem]` option. So, if you wanted to list the contents of a SCSI CD-ROM, you might type the following:

```
MILO> ls -t iso9660 scd0:
```

Variables

MILO provides certain variables to help the boot process. If you are loading via Windows NT ARC or AlphaBIOS firmwares, then MILO makes use of the boot option environment variables set up by that firmware. For some systems (for example, the AlphaPC64), MILO maintains its own set of environment variables that do not change from boot to boot. These variables can be displayed with the `help var` command:

```
MILO> help var
Variables that MILO cares about:
MEMORY_SIZE - System memory size in megabytes
BOOT_DEV    - Specifies the default boot device
BOOT_FILE   - Specifies the default boot file
BOOT_STRING -
Specifies the boot string to pass to the kernel
SCSIn_HOSTID - Specifies the host id of the
n-th SCSI controller
PCI_LATENCY - Specifies the PCI master device latency
AUTOBOOT    - If set, MILO attempts to boot on powerup
and enters command loop only on failure.
AUTOBOOT_TIMEOUT - Seconds to wait before auto-
```

booting on powerup.

Please Note

If you set AUTOBOOT, you need to set the timeout. Setting AUTOBOOT without setting the timeout can result in your machine automatically booting after a timeout of 0 seconds, which will not allow for any user intervention in the case of errors.

B.5.4 Booting with MILO

The primary function of MILO is to boot the Linux kernel. To boot the installation program to use the CD-ROM (assuming the CD-ROM drive is SCSI) the command would look similar to this:

```
MILO> boot scd0:/kernels/generic.gz root=/dev/scd0
```

Note that the above command may be different for your system, because it depends upon how your system and its CD-ROM drive are set up.

To boot from the diskette, put the `generic.img` diskette into the diskette drive and type in the following command:

```
MILO> boot fd0:
```

The installation program should boot, and you should be able to proceed as normal. IDE CD-ROM devices use the standard IDE device references. You will see output like the following:

```
available configurations: linux
(or 'halt' to return back to MILO prompt)
boot:
```

Type in **linux** and press [Enter].

If needed, you will be prompted to place the RAM disk diskette in the diskette drive, and then the installation should begin.

B.5.5 Installation

Refer to Chapter 4, *Installing Red Hat Linux* for instructions on the graphical installation process. The next MILO-specific step you will need to take is the installation of MILO to your hard drive, which should be done after the packages have been installed (see Section 4.21, *Installing Packages*). Instructions for installing MILO to the hard drive are as follows:

Stop when the packages you've selected are finished being transferred to the hard drive. At this point, the packages have been installed and you will need to press the **Next** button to finish the installation.

Do *NOT* press the **Next** button yet. Instead, use the following instructions to install MILO to your hard drive.

The installation process puts all the necessary components of the Red Hat Linux operating system on your machine, but MILO must be installed manually. Switch to the command prompt on the second virtual console by pressing [Ctrl]-[Alt]-[F2].

Move to the `/mnt/sysimage` directory and temporarily make `/mnt/sysimage` the root directory with these commands:

```
# cd /mnt/sysimage
# /usr/sbin/chroot .
```

Now, you must install MILO. First, insert your MILO diskette into the diskette drive. Then you will need to issue a command which will vary slightly in syntax, depending on your boot device. For example, if you plan to boot your Alpha from the first SCSI hard disk, issue the following command:

```
bash# dd if=/dev/fd0 of=/dev/sda1
```

If you plan to boot your Alpha from the second IDE device, the `dd` command should be executed like this:

```
bash# dd if=/dev/fd0 of=/dev/hdb1
```

The command will copy MILO (along with `linload.exe`) to the small MILO partition you created. Once MILO has been installed, type the following command to exit the temporary root device set using the `chroot` command:

```
bash# exit
```

Now switch back to the GUI using [Ctrl]-[Alt]-[F1] and select the **Next** button to finish the installation process.

B.6 Post-Installation MILO Setup

If you dedicated the first partition of your first hard drive to a small FAT partition for booting, you can modify the **Boot Menu** to have a selection to load MILO from this partition.

Add another boot selection, as described in Section B.5.1, *Setting up ARC for Installation* for ARC and in Section B.5.2, *Setting up AlphaBIOS for Installation* for AlphaBIOS. However, this time you should select the appropriate hard drive device. If you installed Linux to the first hard drive, the device should be:

```
scsi(0)disk(0)rdisk(0)partition(1)
```

If you are using an ARC console or the AlphaBIOS console, select the following device:

```
Disk 0 Partition 1
```

If you wish to autoboot into Linux, set the `OSLOADOPTIONS` variable. In the ARC and AlphaBIOS consoles, the contents of the `OSLOADOPTIONS` parameter are passed to MILO as a command. In order to boot Linux automatically in MILO, enter a value for `OSLOADOPTIONS` similar to this one:

```
boot sda2:/boot/vmlinux root=/dev/sda2
```

Once you have done this, booting and running Linux on an Alpha system should be very similar to doing so on an x86 system.

C Additional Resources about Alphas and Linux

Other reference materials, related to running Red Hat Linux on an Alpha system, are available on the Web. A few of the available resources are as follows:

The AlphaLinux website provides a wealth of information regarding running Linux on an Alpha:

<http://www.alphalinux.org>

The *SRM Firmware HOWTO* describes how SRM firmware can be used to boot Linux:

<http://www.alphalinux.org/faq/srm.html>

The *Brief Introduction to Alpha Systems and Processors* document contains an overview of the Alpha architecture:

<http://www.alphalinux.org/faq/alpha-sys.html>

The Compaq website provides many additional resources for Alpha users. See the following URLs:

http://www5.compaq.com/alphaserver/support_doc.html

<http://www.tru64unix.compaq.com/linux/documentation.htm>

Please Note

Many people will find the previous information to be very useful. However, this information is provided by sources outside of Red Hat, Inc. and it may contain descriptions of procedures that are not supported by Red Hat, Inc. For a description of Red Hat, Inc.'s installation support, please refer to Red Hat's website at <http://www.redhat.com/apps/support/programs.html>.

Another source for information about Alpha systems is the Alpha mailing list. Information on how to join the Alpha mailing list, alp-list@redhat.com, is available from <http://www.redhat.com/mailling-lists/>.

D Writing Image Files to Diskettes

You must write each of the image files you need to a floppy diskette. The easiest methods of writing an image file to a diskette are using `dd` or `rawrite`.

D.1 Using the `dd` Command

The `dd` command in a Linux/UNIX OS is used to do direct dumps from one device (or file) to another. It is an excellent tool to create a disk image or to make a diskette from an image. In Linux, the command is invoked as follows:

```
bash# dd if=imagefile of=/dev/fd0 bs=72k
```

You should see something like the following to indicate that the image transfer was successful:

```
20+0 records in
20+0 records out
```

If you see a smaller block count, your image did not transfer correctly. If this is the case, it will usually be accompanied by a disk error.

After you make a disk, make sure to label it according to its contents.

D.2 Using the `rawrite` Utility

Under MS-DOS, you can use the `rawrite` utility to write an image file to a diskette.

If you need to obtain `rawrite`, it is available on the Red Hat FTP site at:

```
ftp://ftp.redhat.com/pub/redhat/redhat-7.0/i386/en/dosutils/
```

The Red Hat FTP site is usually very busy, so you may want to try a mirror site. The list of mirror sites is available at the following URL:

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

The `rawrite` command looks like the following (assuming you have copied `rawrite` to the `dosutils` folder on your C: drive):

```
C:\> dosutils> rawrite
```

`Rawrite` first asks you for the filename of a disk image; enter the directory and name of the image you wish to write.

```
Enter disk image source file name:D:\images\<image file>
```

Then `rawrite` asks for a disk drive to write the image to:

```
Enter target diskette drive:a:
```

Finally, `rawrite` asks for confirmation that a formatted disk is in the drive you have selected. After pressing `[Enter]` to confirm, `rawrite` copies the image file onto the disk.

```
Please insert a formatted diskette into drive A: and press  
--ENTER-- : [Enter]  
C:\ dosutils>
```

The `rawrite` utility will return a message indicating whether your disk was created successfully, or it will inform you if your image failed to transfer correctly.

Make sure that you label the disk according to its contents.

E Getting Technical Support

E.1 Remember to Sign Up

If you have an official edition of Red Hat Linux 7.1, 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, Inc.'s support team.
- Red Hat Network — Easily update your packages and receive security notices that are customized for your system. Go to <https://www.redhat.com/network> for details.
- Priority FTP access — No more late-night visits to congested mirror sites. Owners of Red Hat Linux 7.1 receive free access to priority.redhat.com, Red Hat's preferred customer FTP service, offering high bandwidth connections day and night.
- 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.

E.2 An Overview of Red Hat Support

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.

Red Hat Support will attempt to answer any questions you may have before the installation process is initiated. This includes the following:

- Hardware compatibility questions
- Basic hard drive partitioning strategies

Red Hat, Inc. Support can also provide assistance during the installation process:

- Getting supported hardware recognized by the Red Hat Linux operating system
- Assistance with drive partitioning

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 on-line 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 available. If you are new to Linux, a book that covers Linux basics will be invaluable. We can recommend several titles: *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. The Red Hat technical support website is located at the following URL:

<http://www.redhat.com/support/>

E.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. If you have obtained Linux from any other company, you must contact that company for support. Examples of such companies are as follows:

- 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

E.4 How to Get Technical Support

In order to receive technical support for your Official Red Hat product, you must first sign up.

Every Official Red Hat product comes with a Product Identification code: a 16-character alphanumeric string. The Product ID for Red Hat Linux 7.1 is located on the Registration Information Card that can be found inside the box. Your Product ID is on a perforated card that you can punch out and keep in a safe place. You need this code, so do not lose the card!

Note

Do not throw away the card with your Product ID. You need 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 your technical support and any other benefits or services that you purchased from Red Hat, depending upon which Red Hat product you purchased. The Product ID may also enable priority FTP access, depending on the product that you purchased, for a limited amount of time.

E.4.1 Signing up for Technical Support

To sign up for technical support, you will need to:

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. If you do not already have a customer profile on the Red Hat website, please create a new one.
 2. Using your login name and password you created during the customer profile, please login at the Red Hat Support website at <http://www.redhat.com/support>.
 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
 - A description of the hardware on which the Red Hat Linux product will be installed
 - 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). Please do not lose your Product ID, or you might not be able to receive support.

E.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 NIC do you have?)

E.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>.

E.6 Support Frequently Asked Questions (FAQ)

E.6.1 Q: Email Messages to support@redhat.com Bounce

I send email to support@redhat.com but my messages bounce back to me. What is the problem?

E.6.2 A: support@redhat.com Is Not Used at This Time

To better serve our customers, Red Hat is re-engineering our email support process. At this time, the support@redhat.com address is not functional. In the meantime, please use support via the Web or by telephone.

E.6.3 Q: System Will Not Allow Login

I know that I have already signed up, but the system will not let me log in.

E.6.4 A: Old Logins and Passwords Will Not Work

You could be trying to use an old login and password, or simply mistyping your login or password. Refer to Section F.4.2, *Problems When You Try to Log In* for more information.

F Troubleshooting Your Installation of Red Hat Linux

This appendix discusses some common installation problems you could encounter and provides solutions to these problems.

F.1 You are Unable to Boot Red Hat Linux

F.1.1 Is Your System Displaying Signal 11 Errors?

If you receive a fatal signal 11 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/>.

F.1.2 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.

F.2 Trouble Beginning the Installation

F.2.1 Is Your Mouse Not Detected?

If the **Mouse Not Detected** screen (see Figure F-1, *Mouse Not Detected*) 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 Figure 4–3, *Mouse Configuration*).

Figure F–1 Mouse Not Detected



For an overview of text mode installation instructions, please refer to Chapter 5, *Installing Red Hat Linux via Text Mode*.

F.2.2 Problems with Booting into the Graphical Installation

The installation program will first try to run at 800 x 600 resolution. If that fails, it will try to run at 640 x 480. To force the installation program to run at 640 x 480, boot into `lowres` mode. Refer to *Booting the Installation Program from the CD-ROM* in Section 3.3.4 for details.

F.3 Trouble During the Installation

F.3.1 Do You Not See the Automatic Partitioning Screen?

If the **Automatic Partitioning** screen does not appear during your installation, you probably do not have enough free space in which to install Red Hat Linux.

If you do not want to partition manually, you should exit out of the installation (remove any diskettes and/or CD-ROMs from their respective drives and reboot your system). Then you can re-allocate space for your installation.

The best way to re-allocate space is to use the DOS-based `fips` utility. Review *An Introduction to Disk Partitions* in the *Official Red Hat Linux Reference Guide* for more information.

F.3.2 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.

F.3.3 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.

F.3.4 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.

ALPHA Specific

If you are using `fdisk`, be sure to use BSD-style disklabels and leave allocated space at the beginning of your hard drive. Refer to Section 4.10.1, *SRM Partitioning Requirements* for more information.

F.3.5 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>.

F.4 Problems After Installation

F.4.1 Problems with Server Installations and X

If you performed a server installation and you are having trouble getting X to start, that is because the X Window System is not installed during a server-class installation.

If you want the X Window System, but you did a server-class installation, you can perform an upgrade to install X. During the upgrade, select the X Window System packages, and choose GNOME, KDE, or both.

Alternatively, you can install the XFree86 RPMs. For more information, refer to <http://www.redhat.com/support/docs/howto/XFree86-upgrade/XFree86-upgrade.html>.

F.4.2 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 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-class 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>.

F.4.3 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.

F.4.4 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 `printtool` program. To access `printtool`, log in as root, open a terminal window, and type `printtool`.

F.4.5 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.

If the `sndconfig` application did not help, you may need to uncheck the mute buttons in the Audio Mixer. Click on **Panel=>Programs=>Applications=>Audio Mixer** to launch the application.

Index

A

- about
 - returning to SRM from 19
- adding partitions 34
- Alpha
 - information about 101
 - installation overview 13
- AlphaBIOS
 - setting up for installation 94
- ARC
 - setting up for installation 92
- authentication
 - configuration 49
 - Kerberos 49
 - LDAP 49
 - MD5 passwords 49
 - NIS 49
 - shadow passwords 49
- autoboot 67
 - SRM 62
- automatic partitioning 29
 - server 29
 - workstation 29

B

- /boot partition 32
- bootable CD-ROM 67
- booting
 - after installation of Red Hat Linux 60
 - boot parameters set in SRM 17
 - FTP installation 20
 - hard drive installation 20
 - HTTP installation 20
 - NFS image installation 20
 - the installation program with MILO 97
- BSD
 - style disklabeling 36

C

- CD-ROM
 - bootable 67
- class
 - installation 28
- clock 45
- configuration
 - clock 45
 - hardware 64
 - network 40
 - time 45
 - time zone 45
 - video 64
 - X Window System 75
 - XFree86 75
- consoles, virtual 23

D

- dd 103
- deleting partitions 35
- dependencies
 - installing packages 53
 - upgrading packages 86
- devices
 - MILO 96
 - recognized by SRM 14
- Disk Druid
 - adding partitions 34
 - buttons 33
 - deleting partitions 35
 - drive summaries 33
 - editing partitions 35
 - partitions 30
 - problems adding partitions 32
- disk labeling
 - BSD style 36
- disk partitioning 29
- disk space requirements 13
- documentation

- other manuals..... 11
 - documentation conventions 7
- E**
-
- editing partitions 35
 - errata 14
- F**
-
- fdisk..... 35
 - commands 37
 - feedback
 - contact information for this manual 10
 - filesystems
 - MILO support for 96
 - formatting partitions 38
 - fsck 38
 - FTP
 - installation 72
- H**
-
- hardware
 - compatibility 12
 - configuration 64
 - support for specific 89
 - hostname configuration 40
 - how to use this manual 9
 - HTTP
 - installation 74
- I**
-
- image files
 - writing to diskettes 103
 - information
 - network 64
 - pre-installation 63
 - installation
 - Alpha overview 13
 - booting without diskette 67
 - class 28
 - FTP 72
 - getting Red Hat Linux 11
 - GUI
 - CD-ROM 23
 - HTTP..... 74
 - keyboard navigation 67
 - lowres 19
 - method
 - CD-ROM 24
 - FTP 24
 - hard drive 24
 - HTTP..... 24
 - NFS image 24
 - selecting..... 24
 - NFS server information..... 71
 - no boxed set 12
 - nofb 19
 - noprobe 19
 - of MILO to hard drive 97
 - partitioning 30
 - program
 - text mode user interface 65
 - user interface 23
 - virtual consoles..... 23
 - registering your product 12
 - text mode 19, 63, 67
 - cross-reference table 68
 - online help 68
 - user interface 65
 - using MILO 89, 92
 - installing packages..... 50
 - introduction..... 7
- K**
-
- keyboard
 - configuration 25
 - navigating the installation program
 - using 67
 - keymap
 - selecting type of keyboard 25

L

- language
 - selecting 25
 - support for multiple languages..... 44

M

- manuals 11
- MILO 89
 - autoboot..... 99
 - boot images needed..... 90
 - booting the installation program..... 97
 - devices 96
 - HOWTO 89
 - installation 92
 - installation overview 89
 - installing to hard drive..... 97
 - partitioning considerations 92
 - support is unavailable..... 89
 - supported filesystems 96
 - updated images..... 91
 - user interface 95
 - variables 96
- mouse
 - configuring 26
 - not detected..... 111
 - selecting 26

N

- network
 - configuration 40
 - information 64
 - installations
 - FTP..... 72
 - HTTP..... 74

O

- online help
 - text mode installation 68

P

- packages
 - groups 50
 - selecting..... 50
 - individual 52
 - installing 50
 - selecting 50
- partitioning 30
 - automatic 29
 - changing partition table..... 38
 - for MILO 92
 - initial unallocated space 36
 - manual
 - SRM requirements 35
 - MS-DOS partition tables 31
 - problems 32
 - recommended 32
 - with fdisk 35
- partitionless installations
 - not supported..... 30
- password
 - setting root..... 47
 - user accounts 48
- pre-installation information 63

R

- RAM requirements 13
- rawrite..... 103
- recursion
 - (See recursion)
- Red Hat FAQ 9
- registering your product 12
- root partition ("/")..... 32
- root password 47

S

- selecting
 - packages 50
- SRM 14

- autoboot..... 62
- boot commands 60
- boot parameters 17
- booting from a diskette 20
- booting installation program from
 - CD-ROM..... 19
- disk images 18
- displaying system information with 18
- post-installation boot setup..... 60
- show device command..... 14
- SRM Firmware HOWTO*..... 14
- upgrading to latest version 14
- starting
 - installation 67
- steps
 - hardware compatibility 12
 - Red Hat Linux components 11
- support, technical
 - (See technical support)
- swap partition 32
- system information
 - using SRM to display 18

T

- tables
 - text mode cross-reference 68
- technical support 105
 - FAQ 109
 - how to send questions for 109
 - how to state problems for..... 108
 - not provided for other companies'
 - products 106
 - policy overview 105
 - registering online..... 107
 - signing up for..... 107
- text mode installation
 - cross-reference table 68
- time zone
 - configuration 45
- troubleshooting..... 111

- after the installation..... 115
 - logging in 115
 - Netscape Navigator 115
 - printers and X 116
 - server-class installations and X 115
 - sound configuration 116
- beginning the installation..... 111
 - GUI installation method unavailable 112
 - mouse not detected..... 111
- booting 111
 - network boot disk 111
 - signal 11 error 111
- during the installation..... 112
 - automatic partitioning 112
 - completing partitions 113
 - creating partitions 113
 - Python errors 113
 - using remaining hard drive space 113

U

- unallocated partition(s)..... 32
- unresolved dependencies
 - full installation 53
 - upgrade 86
- upgrade 83
 - customizing..... 84-85
 - description of..... 83
 - package selection 85
 - packages 84
 - starting 83
 - unresolved dependencies 86
- user accounts
 - creation 48
 - setting up 48
- user interface
 - installation program..... 23
 - text mode installation 65

V

| | |
|---------------------------|----|
| variables | |
| MILO | 96 |
| video configuration | 64 |
| virtual consoles | 23 |

X

| | |
|-----------------------|----|
| X | |
| configuration | 54 |
| X Window System | 54 |
| configuration | 75 |
| Xconfigurator..... | 75 |
| monitor setup..... | 55 |
| video card setup..... | 54 |
| XFree86 | |
| configuration | 75 |