Red Hat Linux 7.0

The Official Red Hat Linux Reference Guide

ISBN: 1-58569-020-1

Red Hat, Inc.
2600 Meridian Parkway Durham NC 27713 US 919-547-0012 1-888-733-4281 919-547-0024
docs@redhat.com 13588 Research Triangle Park NC 27709

RefGuide(EN)-7.0-Print-RHI (2000-07-31T12:19-0400)

Printed in the United States, Ireland, and Japan

Contents

Red Hat Linux 7.0

Introduction

Welcome

Welcome to the *Official Red Hat Linux Reference Guide*.

The *Official Red Hat Linux Reference Guide* contains useful information about your Red Hat Linux system. In fact, much of the information you'll find within can be extended to just about any Linux distribution. From fundamental concepts such as using RPM and Gnome-RPM to the finer points of using disk partitioning, we hope you'll find this book to be a valuable resource.

This guide is for you if you want to learn a bit more about how your Red Hat Linux system works. Among the featured entries, you'll learn about:

- Partitioning concepts — Both an introduction to disk partitions and the strategies behind "finding a home" for more than one operating system on hard drives.
- Text mode installation — Despite Red Hat Linux's GUI installation, you may want the control of a text mode install. Here's what you'll find, and what to expect.
- RPM — From both the Gnome-RPM front-end to using RPM at the console.
- RAID concepts — Take one disk drive, add another, and another… Make them appear as a single logical unit, and you've got power and performance.
- Post-installation configuration — Want to do a little tweaking after the installation? Here's where you can get a start.

Getting the Documentation That's Right for You

While the *Official Red Hat Linux Reference Guide* digs into more of the nuts and bolts of your Red Hat Linux system, it is critical to make sure you have documentation that is appropriate to your level of Linux expertise. Regardless of your experience with Linux, it can be easy to feel overwhelmed without the right documentation.

Let’s take a look at three categories of people using Red Hat Linux, and try to be more explicit in terms of the documentation you’ll need. Let’s start by figuring out your experience level. Here are the three basic categories:

New to Linux

Has never used any Linux (or Linux-like) operating system before, or has had only limited exposure to Linux. May or may not have experience using other operating systems (such as Windows). Is this you? If so, please turn to *Documentation For First-Time Linux Users*.

Some Linux Experience

Has installed and successfully used Linux (but not Red Hat Linux) before. Or, may have equivalent experience with other Linux-like operating systems. Does this describe you? If so, please turn to *For the More Experienced*.

Old-Timer

Has installed and successfully used Red Hat Linux before. Are you an old-timer? If so, please turn to *Documentation for Linux Gurus*.

Documentation For First-Time Linux Users

"A journey of a thousand miles begins with a single step." This old saying can be applied to learning about your Red Hat Linux system. Learning to use a Linux system effectively can be a long, rewarding journey, in which you find that you can easily do things about which people with other operating systems can only dream. But like all journeys, you’ve got to start somewhere, and take that first step.

First, get yourself some documentation! This cannot be stressed enough; without documentation you will only become frustrated at your inability to get your Red Hat Linux system working the way you want.

Here’s the sort of Linux documentation you should get your hands on:

- A brief history of Linux — Many aspects of Linux are the way they are because of historical precedent. There is also a Linux culture that, again, is based to a great deal on past history. A bit of knowledge about the history of Linux will

serve you well, particularly as you interact with more experienced Linux users on the Internet.

- An explanation of how Linux works — While it's not necessary to delve into the most arcane aspects of the Linux kernel, it's a good idea to know something about how Linux is put together. This is particularly important if you've been working with other operating systems; some of the assumptions you hold about how computers work may not transfer from that operating system to Linux. A few paragraphs that discuss how Linux works (and particularly how it differs from the operating system you're used to) can be invaluable in getting off to a good start with your Red Hat Linux system.
- An introductory command overview (with examples) — This is probably the most important thing to look for in Linux documentation. The design philosophy behind Linux is that it's better to use many small commands connected together in different ways than it is to have a few large (and complex) commands that do the whole job themselves. Without some examples that illustrate the Linux approach to doing things, you will find yourself intimidated by the sheer number of commands available on your Red Hat Linux system.

Here is some additional direction that may help to match all of your requirements:

- Books — *Linux for Dummies*, by John "maddog" Hall, published by IDG; *Using Linux*, by William H. Ball, published by Que; *Running Linux*, by Matt Welsh and Lar Kaufman, published by O'Reilly & Associates; *Red Hat Linux Secrets*, by Naba Barkakati, published by IDG.
- Red Hat's website — At our very own website (http://www.redhat.com), you'll find links to the Linux Documentation Project (LDP), the *Official Red Hat Linux Installation Guide*, the *Official Red Hat Linux Getting Started Guide*, FAQs (Frequently Asked Questions), a database which can help you search for a Linux Users Group near you, a knowledgebase of information, and more. In short, you'll find a wealth of information to help you get started.
- Newsgroups — Linux users are second to none when it comes to helping new users understand Linux. You can find dozens of Linux-related newsgroups on the Usenet, but a quick search through Deja.com (http://www.deja.com) shows:

  - linux.help
  - linux.redhat
  - linux.redhat.digest
  - linux.redhat.misc
  - linux.redhat.rpm

  Also, from the Deja.com website, you can frequently search for specific information from Linux newsgroups.

Continue reading the next section to find out more about the kinds of documentation that will help you at that point.

For the More Experienced

If you've used other Linux distributions, you probably already have a basic grasp of the most frequently used commands. You may have installed your own Linux system, and maybe you've even downloaded and built software you found on the Internet. What sorts of information will you need?

- Task-oriented items — Many times, you will find that you would like to configure your Red Hat Linux system in a certain way, but you're not sure where to begin. In this case, it's often a big help to see what others in similar circumstances have done. This is where the Linux Documentation Project (also known as the LDP) can come in handy. Each of their HOWTOs document a particular aspect of Linux, from low-level kernel esoterica, to using Linux for amateur radio station work.

  If you selected one of the various `HOWTO` packages when you installed Red Hat Linux, you'll find the HOWTOs on your system in `/usr/share/doc/HOWTO`.

Documentation for Linux Gurus

If you're a long-time Red Hat Linux user, you probably already know that the following pretty much says it all when it comes to documentation:

Use the Force — Read the source!

There are times when you'll just have to sit there and look at the sources to understand things. Fortunately, because of the freely available nature of Linux, it's easy to get the sources. Now if it were only that easy to understand them…

More to Come

The *Official Red Hat Linux Reference Guide* is part of the Red Hat's growing commitment to provide useful and timely support to Red Hat Linux users. Future editions will feature expanded information on system administration, console tools and other resources to help you extend the power of your Red Hat Linux system — and yourself.

That's also where you come in.

Send in Your Feedback

If you'd like to make suggestions about the *Official Red Hat Linux Reference Guide*, please mention this guide's identifier:

RefGuide(EN)-7.0-Print-RHI (2000-07-31T12:19-0400)

You can send mail to:

`docs@redhat.com`

Sign Up for Support

If you have an official edition of Red Hat Linux 7.0, please remember to sign up for the benefits you're entitled to as a Red Hat customer.

You'll 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.
- Priority FTP access — No more late-night visits to congested mirror sites. Owners of Red Hat Linux 7.0 receive free access to priority.redhat.com, Red Hat's preferred customer FTP service, offering high bandwidth connections day and night.

- Red Hat Update Agent — Receive e-mail directly from Red Hat as soon as updated RPMs are available. Use Update Agent filters to receive notification and quickly download updated packages about those subjects that interest you. Also receive — automatically — kernel updates, security updates and other packages.
- 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/now. You'll find your **Personal Product ID** on a red and white card in your Official Red Hat Linux box.

To read more about technical support for Official Red Hat Linux, refer to the Appendix in the Official Red Hat Linux Installation Guide.

Good luck, and thank you for choosing Red Hat Linux!!

The Red Hat Documentation Team

Part I System-Related Reference

1 Red Hat Linux 7.0 New Features

This chapter describes features that are new to Red Hat Linux 7.0.

1.1 Installation-related Enhancements

The Red Hat Linux 7.0 installation program includes a number of new features. For more information, please refer to the *Official Red Hat Linux Installation Guide*.

1.2 System-Related New Features

There are many features new to Red Hat Linux 7.0 that are not part of the installation process. Some new features are tools or applications that you can use, others are new versions of the kernel or desktop environments. This list provides a little more information about what to expect from Red Hat Linux 7.0 once you are actually using the OS itself.

Linux 2.2.*x* kernel:

Red Hat Linux 7.0 includes the latest stable version of the 2.2.*x* Linux kernel.

Kickstart Enhancements

Kickstart has been improved with new commands, as well as partitioning improvements.

XFree86 version 4.0.1:

Red Hat Linux 7.0 contains the latest version of XFree86 (version 4.0.1 which supports many new drivers).

Improved Update Agent

Red Hat now offers a customizable way of receiving updates of Linux technology. Utilizing Update Agent and product registration, Red Hat will now help you keep up with the latest in hardware drivers and security fixes, enable automatic notification of updates, and much more. For more information please refer to http://www.redhat.com/now.

GNOME 1.2

GNOME 1.2 is now included in Red Hat Linux 7.0.

Sawfish window manager now included:

The sawfish window manager is now included in Red Hat Linux 7.0 as the default window manager for GNOME. Based on a Lisp-like language, sawfish is extensible, and GNOME-aware.

KDE 2.0

KDE 2.0 is now included in Red Hat Linux 7.0.

GCC Compiler 2.9.6

GCC Compiler 2.9.6 allows for faster optimized code and more complete C++ support.

Encryption-related changes:

Due to relaxation of U.S. encryption laws, encryption-related changes have been made to the following packages:

- Kerberos authentication has been added to the installation program.
- LDAP authentication has been added to the installation program.
- OpenSSH encryption tools are included in Red Hat Linux 7.0, allowing remote logins to your system for most anything you need.
- The OpenSSL cryptography library is included in Red Hat Linux 7.0 enabling secure transactions for mail, Web and FTP communications.

2 System Administration

This chapter provides an overview of the Red Hat Linux system. This overview is intended to provide guidance on certain aspects of Red Hat Linux that you may not know. Additionally, this chapter will point out some of the differences between Red Hat Linux and other UNIX systems.

2.1 Filesystem Structure

Red Hat is committed to the **Filesystem Hierarchy Standard** (**FHS**), a collaborative document that defines the names and locations of many files and directories. We will continue to track and follow the standard to keep Red Hat Linux compliant.

The current FHS document is the authoritative reference to any FHS compliant filesystem, but the standard leaves many areas undefined or extensible. In this section we provide an overview of the standard and a description of the parts of the filesystem not covered by the standard.

The complete standard can be viewed at:

http://www.pathname.com/fhs/

Compliance with the standard means many things, but the two most important are compatibility with other compliant systems, and the ability to mount the `/usr` partition as read-only (because it contains common executables and is not meant to be changed by users). Since `/usr` can be mounted read-only, `/usr` can be mounted from the CD-ROM or from another machine via read-only NFS.

2.1.1 Overview of the FHS

The directories and files noted here are a small subset of those specified by the FHS document. Check the latest FHS document for the most complete information.

The `/dev` Directory

The `/dev` directory contains filesystem entries which represent devices that are attached to the system. These files are essential for the system to function properly.

The /etc Directory

The /etc directory is reserved for configuration files that are local to your machine. No binaries are to be put in /etc. Any binaries that were formerly put in /etc should now go into /sbin or possibly /bin.

The X11 and skel directories should be subdirectories of /etc:

```
/etc
  |- X11
  +- skel
```

The X11 directory is for X11 configuration files such as XF86Config. The skel directory is for "skeleton" user files, which are used to populate a home directory when a user is first created.

The /lib Directory

The /lib directory should contain only those libraries that are needed to execute the binaries in /bin and /sbin.

The /proc Directory

The /proc directory contains special files that either extract information or send information to the kernel. /proc provides an easy method of accessing information about the operating system using the cat command.

The /sbin Directory

The /sbin directory is for executables used only by the root user. The executables in /sbin are only used to boot and mount /usr and perform system recovery operations. The FHS says:

> "/sbin typically contains files essential for booting the system in addition to the binaries in /bin. Anything executed after /usr is known to be mounted (when there are no problems) should be placed in /usr/sbin. Local-only system administration binaries should be placed into /usr/local/sbin."

At a minimum, the following programs should be in /sbin:

```
arp, clock, getty, halt, init, fdisk,
fsck.*, ifconfig, lilo, mkfs.*, mkswap, reboot,
route, shutdown, swapoff, swapon, update
```

The /usr Directory

The /usr directory is for files that can be shared across a whole site. The /usr directory usually has its own partition, and it should be mountable read-only. The following directories should be subdirectories of /usr:

```
/usr
  |- X11R6
  |- bin
  |- doc
  |- etc
  |- games
  |- include
  |- lib
  |- libexec
  |- local
  |- sbin
  |- share
  +- src
```

The X11R6 directory is for the X Window System (XFree86 on Red Hat Linux), bin contains executables, doc contains non-manpage documentation, etc contains site-wide configuration files, games is for (you guessed it!) games, include contains C header files, lib contains libraries, libexec contains small helper programs called by other programs, sbin is for system administration binaries (those that do not belong in /sbin), share contains files that aren't architecture-specific, and src is for source code.

The /usr/local Directory

The FHS says:

> "The /usr/local hierarchy is for use by the system administrator when installing software locally. It needs to be safe from being overwritten when the system software is updated. It may be used for programs and data that are shareable amongst a group of machines, but not found in /usr."

The /usr/local directory is similar in structure to the /usr directory. It has the following subdirectories, which are similar in purpose to those in the /usr directory:

```
/usr/local
```

```
        |- bin
        |- doc
        |- etc
        |- games
        |- info
        |- lib
        |- man
        |- sbin
        +- src
```

The /var Directory

Since the FHS requires that you be able to mount /usr read-only, any programs that write log files or need spool or lock directories probably should write them to the /var directory. The FHS states /var is for:

> "…variable data files. This includes spool directories and files, administrative and logging data, and transient and temporary files."

The following directories should be subdirectories of /var:

```
/var
  |- cache
  |- db
  |- ftp
  |- gdm
  |- lib
  |- local
  |- lock
  |- log
  |- named
  |- nis
  |- opt
  |- preserve
  |- run
  +- spool
       |- anacron
       |- at
       |- cron
       |- fax
       |- lpd
       |- mail
       |- mqueue
       +- news
       |- rwho
```

```
        |- samba
        |- slrnpull
        |- squid
        |- up2date
        |- uucp
        |- uucppublic
        |- vbox
        |- voice
|- tmp
|- yp
```

System log files such as `wtmp` and `lastlog` go in `/var/log`. The `/var/lib` directory also contains the RPM system databases. Lock files go in `/var/lock`. The `/var/spool` directory has subdirectories for various systems that need to store data files.

2.1.2 `/usr/local` in Red Hat Linux

In Red Hat Linux, the intended use for `/usr/local` is slightly different from that specified by the FHS. The FHS says that `/usr/local` should be where you store software that is to remain safe from system software upgrades. Since system upgrades from Red Hat are done safely with the RPM system and Gnome-RPM, you don't need to protect files by putting them in `/usr/local`. Instead, we recommend you use `/usr/local` for software that is local to your machine.

For instance, let's say you have mounted `/usr` via read-only NFS from *beavis*. If there is a package or program you would like to install, but you are not allowed to write to *beavis*, you should install it under `/usr/local`. Later perhaps, if you've managed to convince the system administrator of *beavis* to install the program on `/usr`, you can uninstall it from `/usr/local`.

2.2 Special Red Hat File Locations

In addition to the files pertaining to the RPM system that reside in `/var/lib/rpm` (see Chapter 5, *Package Management with RPM* for more information on RPM), there are two other special locations that are reserved for Red Hat Linux configuration and operation.

The control-panel and related tools puts many scripts, bitmaps and text files in `/usr/lib/rhs`. There is probably nothing here that you would want to edit.

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

2.3 Users, Groups and User-Private Groups

Managing users and groups has traditionally been tedious, but Red Hat Linux has a few tools and conventions that make users and groups easier to manage.

While you can use `useradd` to create a new user from the shell prompt, the easiest way to manage users and groups is through Linuxconf (see Chapter 3, *System Configuration*).

Next, we'll discuss the basic structure behind managing users and groups.

2.3.1 Standard Users

In Table 2–1, *Standard Users*, you'll find the standard users set up by the installation process (this is essentially the `/etc/passwd` file). The **Group ID** (GID) in this table is the *primary group* for the user. See Section 2.3.3, *User Private Groups* for details on how groups are used.

Table 2–1 Standard Users

User	UID	GID	Home Directory	Shell
root	0	0	`/root`	`/bin/bash`
bin	1	1	`/bin`	
daemon	2	2	`/sbin`	
adm	3	4	`/var/adm`	
lp	4	7	`/var/spool/lpd`	
sync	5	0	`/sbin`	`/bin/sync`
shutdown	6	0	`/sbin`	`/sbin/shutdown`

User	UID	GID	Home Directory	Shell
halt	7	0	/sbin	/sbin/halt
mail	8	12	/var/spool/mail	
news	9	13	/var/spool/news	
uucp	10	14	/var/spool/uucp	
operator	11	0	/root	
games	12	100	/usr/games	
gopher	13	30	/usr/lib/gopher-data	
ftp	14	50	/var/ftp	
nobody	99	99	/	

2.3.2 Standard Groups

In Table 2–2, *Standard Groups*, you'll find the standard groups as set up by the installation process (this is essentially the `/etc/group` file).

Table 2–2 Standard Groups

Group	GID	Members
root	0	root
bin	1	root, bin, daemon
daemon	2	root, bin, daemon
sys	3	root, bin, adm
adm	4	root, adm, daemon
tty	5	
disk	6	root
lp	7	daemon, lp

Group	GID	Members
mem	8	
kmem	9	
wheel	10	root
mail	12	mail
news	13	news
uucp	14	uucp
man	15	
games	20	
gopher	30	
dip	40	
ftp	50	
nobody	99	
users	100	

2.3.3 User Private Groups

Red Hat Linux uses a **user private group** (UPG) scheme, which makes UNIX groups much easier to use. The UPG scheme does not add or change anything in the standard UNIX way of handling groups. It simply offers a new convention for handling groups. Whenever you create a new user, by default, he or she has a unique group. The scheme works as follows:

User Private Group

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

umask = 002

The traditional UNIX umask is 022, which prevents other users *and other members of a user's primary group* from modifying a user's files. Since every user

has their own private group in the UPG scheme, this "group protection" is not needed. A umask of 002 will prevent users from modifying other users' private files. The umask is set in `/etc/profile`.

setgid bit on Directories

If you set the setgid bit on a directory (with `chmod g+s` *`directory`*), files created in that directory will have their group set to the directory's group.

Most IT organizations like to create a group for each major project and assign people to the groups they should be in. Managing files traditionally has been difficult, though, because when someone creates a file it is owned by the primary group he or she belongs to. When a single person works on multiple projects, it becomes hard to associate the right files to the right ownership group. In the UPG scheme, groups are automatically assigned to files on a project-by-project basis, which makes managing group projects very simple.

Let's say you have a big project called *devel*, with many people editing the devel files in a `devel` directory. Make a group called `devel`, `chgrp` the `devel` directory to `devel`, and add the all the devel users to the `devel` group. Now, all devel users will be able to edit the devel files and create new files in the `devel` directory, and these files will always retain their `devel` group. Thus, they will always be edit-able by other devel users.

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

Since each user's home directory is owned by the user and their private group, it is safe to set the setgid bit on the home directory. However, by default, files are created with the primary group of the user, so the setgid bit would be redundant.

User Private Group Rationale

Although UPG is not new to Red Hat Linux 7.0, many people still have questions about it, such as why UPG is necessary. The following is the rationale for the scheme.

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

```
chown -R root.emacs /usr/lib/emacs/site-lisp
```

  and you add the proper users to the group.

- To allow the users to actually create files in the directory you enter:

```
chmod 775 /usr/lib/emacs/site-lisp
```

- But when a user creates a new file it is assigned the group of the user's default group (usually `users`). To prevent this you enter:

```
chmod 2775 /usr/lib/emacs/site-lisp
```

  which causes everything in the directory to be created with the "emacs" group.

- But the new file needs to be mode 664 for another user in the emacs group to be able to edit it. To do this you make the default umask 002.
- Well, this all works fine, except that if your default group is "users," every file you create in your home directory will be writable by everybody in "users" (usually everyone).
- To fix this, you make each user have a "private group" as their default group.

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

2.4 Configuring Console Access

When normal (non-root) users log in to a computer locally, they are given two types of special permission: they can run certain programs that they would not otherwise

be able to run, and they can access certain files (normally special device files used to access diskettes, CD-ROMs, and so on) that they would not otherwise be able to access.

Since there are multiple consoles on a single computer, and multiple users can be logged into the computer locally at the same time, one of the users has to "win" the fight to access the files. The first user to log in at the console owns those files. Once the first user logs out, the next user who logs in will own the files.

In contrast, *every* user who logs in at the console will be allowed to run programs normally restricted to the root user. By default, those programs will ask for the user's password. This will be done graphically if X is running which makes it possible to include these actions as menu items in a graphical user interface. As shipped, the console-accessible programs are `shutdown`, `halt`, and `reboot`.

2.4.1 Disabling Console Program Access

In environments where the console is otherwise secured (BIOS and LILO passwords are set, [Ctrl]-[Alt]-[Delete] is disabled, the power and reset switches are disabled, etc.), it may not be desirable to allow arbitrary users at the console to run `shutdown`, `halt`, and `reboot`.

In order to disable all access by console users to console programs, you should run the command:

```
rm -f /etc/security/console.apps/*
```

2.4.2 Disabling All Console Access

In order to disable all console access, including program and file access, in the `/etc/pam.d/` directory, comment out all lines that refer to `pam_console.so`. The following script will do the trick:

```
cd /etc/pam.d
for i in * ; do
sed '/[^#].*pam_console.so/s/^/#/' < $i > foo && mv foo $i
done
```

2.4.3 Defining the Console

The `/etc/security/console.perms` file defines the console group. The syntax of that file is very flexible; you can edit the file so that these instructions no longer apply. However, the default file has a line that looks like this:

```
<console>=tty[0-9][0-9]* :[0-9]\.[0-9] :[0-9]
```

When users log in, they are attached to some sort of named terminal, either an X server with a name like `:0` or `mymachine.example.com:1.0`; or a device like `/dev/ttyS0` or `/dev/pts/2`. The default is to define that local virtual consoles and local X servers are considered local, but if you want to consider the serial terminal next to you on port `/dev/ttyS1` to also be local, you can change that line to read:

```
<console>=tty[0-9][0-9]* :[0-9]\.[0-9] :[0-9] /dev/ttyS1
```

2.4.4 Making Files Console-Accessible

In `/etc/security/console.perms`, there is a section with lines like:

```
<floppy>=/dev/fd[0-1]*
<cdrom>=/dev/cdrom
<jaz>=/dev/zip
```

You can also add your own lines:

```
<scanner>=/dev/sga
```

(Of course, make sure that `/dev/sga` is really your scanner and not, say, your hard drive.)

That's the first part. The second part is to define what is done with those files. Look in the last section of `/etc/security/console.perms` for lines similar to:

```
<console> 0660 <floppy> 0660 root.floppy
<console> 0600 <cdrom>  0600 root.disk
<console> 0600 <jaz>    0660 root.disk
```

and add a line like:

```
<console> 0600 <scanner> 0600 root
```

Then, when you log in at the console, you will be given ownership of the `/dev/sga` device and the permissions will be 0600 (readable and writable by you only). When you log out, the device will be owned by root and still have 0600 (now: readable and writable by root only) permissions.

2.4.5 Enabling Console Access for Other Applications

If you wish to make other applications besides `shutdown`, `reboot`, and `halt` accessible to console users, you will have to do just a little bit more work.

First of all, console access *only* works for applications which reside in `/sbin` or `/usr/sbin`, so the application that you wish to run must be there.

Create a link from the name of your application to the `/usr/bin/consolehelper` application:

```
cd /usr/bin
ln -s consolehelper foo
```

Create the file `/etc/security/console.apps/foo`:

```
touch /etc/security/console.apps/foo
```

Create a PAM configuration file for the *foo* service in `/etc/pam.d/`. We suggest that you start with a copy of the shutdown service, then change it if you want to change the behavior:

```
cp /etc/pam.d/shutdown /etc/pam.d/foo
```

Now, when you run `/usr/bin/foo`, it will call `consolehelper`, which, with the help of `/usr/sbin/userhelper` will authenticate the user (asking for the user's password if `/etc/pam.d/foo` is a copy of `/etc/pam.d/shutdown`; otherwise, it will do precisely what is specified in `/etc/pam.d/foo`) and then run `/usr/sbin/foo` with root permissions.

2.5 The floppy Group

If, for whatever reason, console access is not appropriate for you, and you need to give non-root users access to your system's diskette drive, this can be done using the `floppy` group. Simply add the user(s) to the `floppy` group using the tool of your choice. Here's an example showing how `gpasswd` can be used to add user fred to the `floppy` group:

```
[root@bigdog root]# gpasswd -a fred floppy
Adding user fred to group floppy
[root@bigdog root]#
```

User fred will now be able to access the system's diskette drive.

2.6 User Authentication with PAM

Programs which give users access to privileges of any sort need to be able to authenticate the users. When you log into a system, you provide your name and password, and the login process uses those to authenticate the login — to verify that you are who you say you are. Forms of authentication other than passwords are possible, and it is possible for the passwords to be stored in different ways.

PAM, which stands for **Pluggable Authentication Modules**, is a way of allowing the system administrator to set authentication policy without having to recompile authentication programs. With PAM, you control how the modules are plugged into the programs by editing a configuration file.

Most Red Hat Linux users will never need to touch this configuration file. When you use RPM to install programs that require authentication, they automatically make the changes that are needed to do normal password authentication. However, you may want to customize your configuration, in which case you must understand the configuration file.

2.6.1 PAM Modules

There are four types of modules defined by the PAM standard.

- `auth` modules provide the actual authentication, perhaps asking for and checking a password, and they set "credentials" such as group membership or Kerberos "tickets."
- `account` modules check to make sure that the authentication is allowed (the account has not expired, the user is allowed to log in at this time of day, and so on).
- `password` modules are used to set passwords.
- `session` modules are used once a user has been authenticated to allow them to use their account, perhaps mounting the user's home directory or making their mailbox available.

These modules may be *stacked*, so that multiple modules are used. For instance, `rlogin` normally makes use of at least two authentication methods: if `rhosts` authentication succeeds, it is sufficient to allow the connection; if it fails, then standard password authentication is done.

New modules can be added at any time, and PAM-aware applications can then be made to use them. For instance, if you have a one-time-password calculator system, and you can write a module to support it (documentation on writing modules is included with the system in `/usr/share/doc/pam*`), PAM-aware programs can use the new module and work with the new one-time-password calculators without being recompiled or otherwise modified in any way.

2.6.2 Services

Each program using PAM defines its own "service" name. The login program defines the service type `login`, ftpd defines the service type `ftp`, and so on. In general, the service type is the name of the program used to *access* the service, not (if there is a difference) the program used to *provide* the service.

2.6.3 The Configuration Files

The directory `/etc/pam.d` is used to configure all PAM applications. (This used to be `/etc/pam.conf` in earlier PAM versions; while the `pam.conf` file is still

read if no /etc/pam.d/ entry is found, its use is deprecated.) Each application (really, each *service*) has its own file. A file looks like this:

```
#%PAM-1.0
auth       required  /lib/security/pam_securetty.so
auth       required  /lib/security/pam_unix.so shadow nullok
auth       required  /lib/security/pam_nologin.so
account    required  /lib/security/pam_unix.so
password   required  /lib/security/pam_cracklib.so
password   required  /lib/security/pam_unix.so shadow nullok use_authtok
session    required  /lib/security/pam_unix.so
```

The first line is a comment. (Any line that starts with a # character is a comment.) Lines two through four stack up three modules to use for login authorization. Line two makes sure that *if* the user is trying to log in as root, the tty on which they are logging in is listed in the /etc/securetty file *if* that file exists. Line three causes the user to be asked for a password and the password checked. Line four checks to see if the file /etc/nologin exists, and if it does, displays the contents of the file, and if the user is not root, does not let him or her log in.

Note that all three modules are checked, *even if the first module fails*. This is a security decision — it is designed to prevent the user from knowing why their authentication was disallowed, because knowing why it was disallowed might allow them to break the authentication more easily. You can change this behavior by changing required to requisite; if any requisite module returns failure, PAM fails immediately without calling any other modules.

The fifth line causes any necessary accounting to be done. For example, if shadow passwords have been enabled, the pam_unix.so module will check to see if the account has expired, or if the user has not changed his or her password and the grace period for changing the password has expired.

The sixth line subjects a newly changed password to a series of tests to ensure that it cannot, for example, be easily determined by a dictionary-based password cracking program.

The seventh line (which may be wrapped) specifies that if the login program changes the user's password, it should use the pam_unix.so module to do so. (It will do

so only if an `auth` module has determined that the password needs to be changed — for example, if a shadow password has expired.)

The eighth and final line specifies that the `pam_unix.so` module should be used to manage the session. Currently, that module doesn't do anything; it could be replaced (or supplemented by stacking) by any necessary module.

Note that the order of the lines within each file matters. While it doesn't really matter much in which order required modules are called, there are other *control flags* available. While optional is rarely used, and never used by default on a Red Hat Linux system, sufficient and requisite cause order to become important.

Let's look at the `auth` configuration for `rlogin`:

```
auth       sufficient   /lib/security/pam_rhosts_auth.so
auth       required     /lib/security/pam_securetty.so
auth       required     /lib/security/pam_stack.so service=system-auth
auth       required     /lib/security/pam_nologin.so
```

First, if `pam_rhosts_auth.so` authenticates the user, PAM immediately returns success to `rlogin` without any password checking. If `pam_rhosts_auth.so` fails to authenticate the user, that failed authentication is ignored.

Second, `pam_securetty.so` keeps root logins from happening on insecure terminals. This effectively disallows all root `rlogin` attempts. If you wish to allow them (in which case we recommend that you not be Internet-connected or be behind a good firewall), you can simply remove that line.

Third, if `pam_rhosts_auth.so` has failed to authenticate the user, the `pam_stack.so` module performs normal password authentication.

Finally `pam_nologin.so` checks `/etc/nologin`, as specified above.

Note that if you do not want to prompt for a password if the securetty check fails, you can change the `pam_securetty.so` module from required to requisite.

2.6.4 Shadow Passwords

The `pam_unix.so` module will automatically detect that you are using shadow passwords and make all necessary adjustments. Please refer to Section 2.7, *Shadow Utilities* for more information.

2.6.5 Rexec and PAM

For security reasons, rexec is not enabled in Red Hat Linux 7.0. Should you wish to enable it, you will need to comment out one line in the file `/etc/pam.d/rexec`. Here is a sample of the file (note that your file may differ slightly):

```
#%PAM-1.0
auth       required     /lib/security/pam_securetty.so
auth       required     /lib/security/pam_stack.so service=system-auth
auth       required     /lib/security/pam_nologin.so
account    required     /lib/security/pam_stack.so service=system-auth
```

To enable rexec, the line referring to the `pam_nologin.so` module must be commented out:

```
#%PAM-1.0
auth       required     /lib/security/pam_securetty.so
auth       required     /lib/security/pam_stack.so service=system-auth
#auth       required     /lib/security/pam_nologin.so
account    required     /lib/security/pam_stack.so service=system-auth
```

After this file is modified, rexec will be enabled.

Please Note

If your `/etc/pam.d/rexec` file contains a line referring to the `pam_securetty.so` module, you will not be able to `rexec` as root. To do so, you must also comment out the line referring to the `pam_securetty.so` module.

Please Note

Most configuration files have been rewritten to simplify system-wide changes, so that when a setup needs to be changed, it only needs to be changed in one place. This change occurs because of the `pam_stack` file which lets you "call" from inside of the stack for a particular service, the stack defined for any other service. See the man page for `pam_stack` for more information.

More Information

This is just an introduction to PAM. More information is included in the `/usr/share/doc/pam*` directory, including a *System Administrators' Guide*, a *Module Writers' Manual*, an *Application Developers' Manual*, and the PAM standard, DCE-RFC 86.0.

2.7 Shadow Utilities

Shadow passwords are a method of improving system security by moving the encrypted passwords (normally found in `/etc/passwd`) to `/etc/shadow` which is readable only by root. During the installation of Red Hat Linux, you were given the option of setting up shadow password protection on your system.

The `shadow-utils` package contains a number of utilities that support:

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

Please Note

There are some additional points of interest concerning these utilities:

- The utilities will work properly whether shadowing is enabled or not.
- The utilities have been slightly modified to support Red Hat's user private group scheme. For a description of the modifications, please see the `useradd` man page. For more information on user private groups, please turn to Section 2.3.3, *User Private Groups*.
- The `adduser` script has been replaced with a symlink to `/usr/sbin/useradd`.
- The tools in the `shadow-utils` package are not Kerberos or LDAP enabled. New users will be local only.

2.8 Building a Custom Kernel

Many people new to Linux often ask, "why should I build my own kernel?" Given the advances that have been made in the use of kernel modules, the most accurate response to that question is, "unless you know why you need to build your own kernel, you probably don't." So unless you have a specific reason to build a customized kernel (or you're just the curious sort), you may skip ahead to Section 2.9, *Sendmail*.

In the past, you would need to recompile the kernel if you added new hardware on your system. The kernel was, in other words, **static**. Improvements in the Linux 2.0.*x* kernels allowed for much of the drivers for hardware to be **modularized** into components that could only be inserted on demand. However, there were major problems with having multiple kernels on your system that had been compiled for different advancements (a good case being SMP versus UP kernels). Further advancements with the modularization of the Linux 2.2.x kernel have allowed for multiple kernels to more easily co-exist (though *not* share modules).

For information on handling kernel modules see Section 3.2.2, *Loading Kernel Modules*. Most of the changes are hidden except when recompiling a customized kernel for your system.

2.8.1 Building a modularized kernel

These instructions enable you to take advantage of the power and flexibility available through kernel modularization. If you do not wish to take advantage of modularization, please see Section 2.8.3, *Building a monolithic kernel* for an explanation of the different aspects of building and installing a monolithic kernel. It's assumed that you've already installed the `kernel-headers` and `kernel-source` packages and that you issue all commands from the `/usr/src/linux` directory.

The most important step is to make sure that you have a working emergency boot disk in case you make a mistake below. If you didn't make a boot disk during the installation, use the `mkbootdisk` command to make one. The standard command is similar to `mkbootdisk --device /dev/fd0 2.2.x`, where 2.2.x is the full version of your kernel (such as 2.2.14-5.0). Once done, test the boot disk to make sure that it will boot the system.

It is important to begin a kernel build with the source tree in a known condition. Therefore, it is recommended that you begin with the command `make mrproper`. This will remove any configuration files along with the remains of any previous builds that may be scattered around the source tree. Now you must create a configuration file that will determine which components to include in your new kernel. Available methods for kernel configuration are listed below:

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

- `make oldconfig` — This is a non-interactive script that will set up your Makefile to be the default settings. If you're using the Red Hat patched kernel, it will set up the configuration to be that of the kernel that is shipped for your box. This is useful in setting up your kernel to known working defaults and then turning off features that you don't want.

Please Note

In order to use `kmod` (see Section 3.2.2, *Loading Kernel Modules* for details) and kernel modules you must answer **`Yes`** to `kmod support` and `module version (CONFIG_MODVERSIONS) support` in the configuration.

If you wish to build a kernel with a configuration file (`/usr/src/linux/.config` — this file is created once one of the above methods has been performed) that you have already created with one of the above methods, you can omit the `make mrproper` and `make config` commands and use the command `make dep` followed by `make clean` to prepare the source tree for the build.

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

- Edit the `Makefile` and change the line: `EXTRAVERSION =` to match a "unique" name (such as adding your initials to the end of the string, as in `EXTRAVERSION = -2.5.0sjs`). This will allow you to have the old working kernel and the new kernel on your system at the same time.
- Build the kernel with `make bzImage`.
- Build any modules you configured with `make modules`.

- Install the new modules (even if you didn't build any) with `make modules_install`. This will install the kernel modules into the directory path `/lib/modules/` using the path name that was specified in the `Makefile`. Our example would be `/lib/modules/2.2.15-2.5.0sjs/`.

If you have a SCSI adapter and made your SCSI driver modular, build a new `initrd` image (see Section 2.8.2, *Making an initrd image*; note that there are few practical reasons to make the SCSI driver modular in a custom kernel). Unless you have a specific reason to create an `initrd` image, do not create one and do not add it to `lilo.conf`.

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

```
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
message=/boot/message
linear
default=linux

image=/boot/vmlinuz-2.2.16-12
        label=linux
        initrd=/boot/initrd-2.2.16-12.img
        read-only
        root=/dev/hda8

other=/dev/hda1
        label=dos
```

Now you must update `/etc/lilo.conf`. If you built a new `initrd` image you must tell LILO to use it. In this example of `/etc/lilo.conf` we have added four lines in the middle of the file to indicate another kernel to boot from. We have

renamed `/boot/vmlinuz` to `/boot/vmlinuz.old` and changed its label to `old`. We have also added an `initrd` line for the new kernel:

```
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
message=/boot/message
linear
default=linux

image=/boot/vmlinuz-2.2.16-12
     label=linux
     initrd=/boot/initrd-2.2.16-12.img
     read-only
     root=/dev/hda8

image=/boot/vmlinuz-2.2.16-12.sjs
        label=test
        initrd=/boot/initrd-2.2.16-12sjs.img
        read-only
        root=/dev/hda8

other=/dev/hda1
        label=dos
```

Now when the system boots and you press [Tab] at the LILO `boot:` prompt, available choices will be shown;

```
LILO boot:
linux  test  dos
```

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

Here is a summary of the steps:

- Copy the resulting compiled kernel into your `/boot` directory using the name that resulted from your earlier changes to the `Makefile`. Here is an example:

```
cp -p
/usr/src/linux/arch/i386/boot/bzImage
```

```
/boot/vmlinuz-2.2.15-2.5.0sjs
/usr/src/linux/System.map  /boot/System.map-2.2.15-2.5.0sjs
```

- Edit `/etc/lilo.conf`.
- Make a new initial ramdisk, `initrd` image (see Section 2.8.2, *Making an initrd image*) if needed.
- Run `/sbin/lilo` (you may want to use `/sbin/lilo -t` first — this command will test your `lilo.conf` without actually writing a new boot sector or map file). You can add a `-v` flag to `lilo` to get more verbose reporting if you think there might be a problem.

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

2.8.2 Making an initrd image

An `initrd` image is needed for loading your SCSI module at boot time. If you do not need an `initrd` image, do not make one and do not edit `lilo.conf` to include this image.

The shell script `/sbin/mkinitrd` can build a proper `initrd` image for your machine if the following conditions are met:

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

```
alias scsi_hostadapter BusLogic
```

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

```
/sbin/mkinitrd /boot/newinitrd-image 2.2.15-2.5.0sjs
```

Where `/boot/newinitrd-image` is the file to use for your new image, and `2.2.15` is the kernel whose modules (from `/lib/modules`) should be used in

the `initrd` image (not necessarily the same as the version number of the currently running kernel).

2.8.3 Building a monolithic kernel

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

- When configuring the kernel only answer **`Yes`** and **`No`** to the questions (don't make anything modular). Also, you should answer **`No`** to `kmod support` and `module version (CONFIG_MODVERSIONS) support` in the configuration.
- Omit the steps:

```
make modules
make modules_install
```

- Edit `lilo.conf` and add the line `append=nomodules`.

2.9 Sendmail

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

Please Note

Although SMTP servers are supported automatically, **IMAP** (Internet Message Access Protocol) servers are not. If your ISP uses an IMAP server rather than an SMTP sever, you must install the IMAP package. Without it, your system won't know how to pass information to the IMAP server or retrieve your mail.

To generate a new `sendmail.cf`, you will need to install `m4` and the `sendmail` source package. Read the README file in the sendmail sources for more details on creating sendmail configuration files. Also, O'Reilly & Associates publishes a good sendmail reference entitled *sendmail*, by Bryan Costales. Lastly, www.sendmail.net offers a thorough breakdown of sendmail features and configuration examples.

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

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

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

```
# who I send unqualified names to
# (null means deliver locally)
DRmail.redhat.com

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

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

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

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

2.10 Controlling Access to Services

Maintaining security on your Red Hat Linux system is extremely important. One way to manage security on your system is to carefully manage access to system services. Your system may need to provide open access to particular services (for example, `httpd` if you're running a Web server). However, if you don't need to provide a service, you should turn it off — this will minimize your exposure to any possible bug exploits.

There are several different methods for managing access to system services. You'll need to decide which of them you'd like to use, based on the service, your system's configuration and your level of Linux expertise.

The easiest way to deny access to a service is to simply turn it off. Both the services managed by `xinetd` (which we'll talk about more later in this section) and the services in the `/etc/rc.d` hierarchy can be configured to start or stop using either the ntsysv utility or using `chkconfig`. You may find that these tools are easier to use than the alternatives — editing the numerous symbolic links located in the directories below `/etc/rc.d` by hand or editing the `xinetd` configuration files in `/etc/xinetd.d`.

ntsysv provides a simple interface for activating or deactivating services. You can use ntsysv to turn an `xinetd`-managed service on or off. You can also use ntsysv to start or stop a service in the `/etc/rc.d` hierarchy; in which case, the `ntsysv` command without options configures your current runlevel. If you want to configure a different runlevel, use something like `ntsysv --levels 016`. (In this example, you'd be setting the services for runlevels 0, 1 and 6.)

The ntsysv interface works like the text-mode installation program. Use the up and down arrows to navigate up and down the list. The space bar selects/un-selects services and is also used to "press" the **Ok** and **Cancel** buttons. To move between the

list of services and the **Ok** and **Cancel** buttons, use the [Tab] key. An * signifies that a service is set to on. The [F1] key will pop up a short description of each service.

`chkconfig` can also be used to activate and deactivate services. If you use the `chkconfig --list` command, you'll see a list of system services and whether they are started (`on`) or stopped (`off`) in runlevels 0-6 (at the end of the list, you'll see a section for the services managed by `xinetd`, which we'll discuss later in this section).

You can also use `chkconfig` to find out whether a specific service is running. For example, the following command checks for the `finger` daemon:

```
$ chkconfig --list finger
finger          on
```

As shown above, `finger` is on in the current runlevel.

If you use `chkconfig --list` to query a service in `/etc/rc.d`, you'll see the service's settings for each runlevel, like the following:

```
$ /sbin/chkconfig --list anacron
anacron         0:off   1:off   2:on    3:on
4:on    5:on    6:off
```

More importantly, `chkconfig` can be used to set a service to be started (or not) in a specific runlevel. For example, if we wanted to turn `nscd` off in runlevels 3, 4, and 5, we'd use a command like this:

```
chkconfig --level 345 nscd off
```

See the `chkconfig` man page for more information on how to use it.

Another way of controlling access to Internet services is using `xinetd`, a secure replacement for `inetd`. `xinetd` conserves system resources, provides access control and logging, and can be used to start special-purpose servers. `xinetd` can be used to provide access only to particular hosts, to deny access to particular hosts, to only provide access to a service at certain times, to limit the rate of incoming connections and/or the load created by connections, etc.

`xinetd` runs constantly, and listens on all of the ports for the services it manages. When a connection request arrives for one of its managed services, `xinetd` starts up the appropriate server for that service.

The configuration file for `xinetd` is `/etc/xinetd.conf`, but you'll notice upon inspection of the file that it just contains a few defaults and an instruction to include the `/etc/xinetd.d` directory. The files within the `/etc/xinetd.d` directory contain the configuration options for each service managed by `xinetd`, so you'll need to edit these files to configure `xinetd`.

See the `xinetd` (8) man page and the `xinetd.conf` (8) man page for specific instructions on how to set up the files in `/etc/xinetd.d`. The `xinetd` Web page, located at http://www.xinetd.org/, is another good source of information.

Many UNIX system administrators are accustomed to using TCP wrappers to manage access to certain network services. Any network services managed by `xinetd` (as well as any program with built-in support for libwrap) can use TCP wrappers to manage access. `xinetd` can use the `/etc/hosts.allow` and `/etc/hosts.deny` files to configure access to system services. If you'd like to use TCP wrappers, see the hosts_access (5) man pages for more detailed information.

Another way to manage access to system services is by using `ipchains` to configure an IP firewall. If you're a new Linux user, please realize that `ipchains` may not be the best solution for you. Setting up `ipchains` can be complicated and is best tackled by experienced UNIX/Linux system administrators.

On the other hand, the benefit of using `ipchains` is flexibility. For example, if you need a customized solution which provides access to certain services to certain hosts, `ipchains` can provide it for you. See the *Linux IPCHAINS-HOWTO* at http://www.linuxdoc.org/HOWTO/IPCHAINS-HOWTO.html for more information about `ipchains`. The *Linux IPCHAINS-HOWTO* is also available on the Documentation CD.

Alternatively, if you're looking for a utility which will set general access rules for your home machine, and/or if you are new to Linux, you should try the gnome-lokkit utility. gnome-lokkit is a GUI utility which will ask you questions about how you want to use your machine. Based on your answers, gnome-lokkit will then configure a simple firewall for you.

2.11 Anonymous FTP

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

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

`/etc/ftpaccess`

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

`/etc/ftphosts`

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

`/etc/ftpusers`

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

2.12 NFS Configuration

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

2.12.1 Mounting NFS Filesystems

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

```
mkdir /mnt/local # Only required if /mnt/local doesn't exist
mount bigdog:/mnt/export /mnt/local
```

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

2.12.2 Exporting NFS Filesystems

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

```
directory       hostname(options)
```

the (*options*) are discretionary. For example:

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

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

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

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

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

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

Or you can restart the daemons with this command:

```
/etc/rc.d/init.d/nfs restart
```

The following will also work:

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

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

2.13 The Boot Process, Init, and Shutdown

This section contains information on what happens when you boot or shut down your Red Hat Linux system.

2.13.1 Behind the Scenes of the i386 Boot Process

When a computer is booted, the processor looks at the end of the system memory for the **BIOS** (Basic Input/Output System) and runs it. The BIOS program is written into read-only permanent memory, and is always ready to go. The BIOS provides the lowest level interface to peripheral devices and controls the first step of the boot process.

The BIOS tests the system, looks for and checks peripherals and then looks for a drive to boot from. Usually, it checks the floppy drive (or CD-ROM drive on many newer systems), if present, and then it looks on the hard drive. On the hard drive, the BIOS looks for a **Master Boot Record** (MBR) starting at the first sector on the first hard drive and starts the MBR running.

The MBR looks for the first active partition and reads the partition's boot record. The boot record contains instructions on how to load the boot loader, LILO (*LI*nux *LO*ader). The MBR then loads LILO and LILO takes over the process.

LILO reads the file `/etc/lilo.conf`, which spells out which operating system(s) to configure or which kernel to start and where to install itself (for example, `/dev/hda` for your hard drive). LILO displays a `LILO:` prompt on the screen and waits for a preset period of time (also set in the `lilo.conf` file) for input from the user. If your `lilo.conf` is set to give LILO a choice of operating systems, at this time you could type in the label for whichever OS you wanted to boot.

After waiting for a set period of time (five seconds is common), LILO proceeds to boot whichever operating system appears first in the `lilo.conf` file.

If LILO is booting Linux, it first boots the kernel, which is a `vmlinuz` file (plus a version number, for example, `vmlinuz-2.2.15-xx`) located in the `/boot` directory. Then the kernel takes over.

The kernel looks in several different places for init (`/sbin` is a common location) and runs the first one it finds. Init takes over.

Init starts (and becomes the parent or grandparent of) all of the processes which make up your Linux system. First, it runs `/etc/rc.d/rc.sysinit`, which sets your path, starts swapping, checks the filesystems, and so on. Basically, `rc.sysinit` is taking care of everything that your system needs to have done at system initialization. For example, on a networked system `rc.sysinit` uses the information in the `/etc/sysconfig/network` and `/etc/sysconfig/clock` files to initialize network processes and the clock. It may also run `rc.serial`, if you have serial port processes that need to be initialized.

Init looks at and implements the `/etc/inittab` file. The `/etc/inittab` file describes how the system should be set up in each runlevel and sets the default runlevel. This file states that `/etc/rc.d/rc` and `/sbin/update` should be run whenever a runlevel starts.

The `/sbin/update` file flushes dirty buffers back to disk.

Whenever the runlevel changes, `/etc/rc.d/rc` starts and stops services. First, rc sets the source function library for the system (commonly `/etc/rc.d/init.d/functions`), which spells out how to start/kill a program and how to find out the PID of a program. The rc file then finds out the current and the previous runlevel and tells linuxconf the appropriate runlevel.

The rc file starts all of the background processes necessary for the system to run, and looks for an rc directory for that runlevel (`/etc/rc.d/rc<x>.d`, where the *<x>* is numbered 0-6). rc kills all of the kill scripts (their file name starts with a K). Then it initializes all of the start scripts (their file names start with an S) in the appropriate runlevel directory (so that all services and applications are started correctly).

For example, for runlevel 5, rc looks into the `/etc/rc.d/rc5.d` directory and finds that it needs to kill rusersd, rwalld, rwhod, mcserv, mars-nwe, apmd, and pcmcia. In the bloody aftermath, rc looks into the same directory and finds start

scripts for kmod, network, nfsfs, randomc, syslog, atd, crond, portmap, snmpd, inet, xntpd, lpd, dhcpd, ypbind, autofs, keytable, sendmail, gpm, and sound. And life begins anew.

The `/etc/inittab` file forks a getty process for each virtual console (login prompts) for each runlevel (runlevels 2-5 get all six; runlevel 1, which is single user mode, only gets one console; runlevels 0 and 6 get no virtual consoles).

In runlevel 5, `/etc/inittab` also runs a script called `/etc/X11/prefdm`. The `prefdm` script runs the preferred X display manager (`gdm` if you're running GNOME, `kdm` if you're running KDE, or `xdm` if you're running AnotherLevel) based on the contents of the `/etc/sysconfig/desktop` directory.

Also, /etc/inittab describes how the system should handle translating [Ctrl]-[Alt]-[Delete] into something like the command `/sbin/shutdown -t3 -r now`. And finally, `/etc/inittab` states what the system should do in case of power failures.

At this point, you should be looking at a login prompt. All that, and it only took a few seconds.

Next, we'll discuss information on the files in `/etc/sysconfig`.

2.13.2 Sysconfig Information

The following information outlines the various files in `/etc/sysconfig`, their function, and their contents.

Files in `/etc/sysconfig`

The following files are normally found in `/etc/sysconfig`:

- `/etc/sysconfig/apmd`
- `/etc/sysconfig/clock`
- `/etc/sysconfig/harddisks`
- `/etc/sysconfig/hwconf` (this should be ignored for editing)
- `/etc/sysconfig/init`
- `/etc/sysconfig/keyboard`

- /etc/sysconfig/mouse
- /etc/sysconfig/network
- /etc/sysconfig/pcmcia
- /etc/sysconfig/sendmail
- /etc/sysconfig/soundcard (which is written by sndconfig)

Let's take a look at each one.

/etc/sysconfig/apmd

The /etc/sysconfig/apmd is used by apmd, as a configuration for what things to start/stop/change on suspend or resume. It is set up to turn on or off apmd during startup, depending on whether your hardware supports Advanced Power Management (apm), or if you choose not to use it.

/etc/sysconfig/clock

The /etc/sysconfig/clock file controls the interpretation of values read from the system clock. Earlier releases of Red Hat Linux used the following values (which are deprecated):

- CLOCKMODE=*mode*, where *mode* is one of the following:
  - GMT — indicates that the clock is set to UTC.
  - ARC — on Alpha only indicates the ARC console's 42-year time offset is in effect.

Currently, the correct values are:

- UTC=*boolean*, where *boolean* is the following:
  - true — indicates that the clock is set to UTC. Any other value indicates that it is set to local time.
- ARC=*boolean*, where *boolean* is the following:
  - true — (for Alpha-based systems only) Indicates the ARC console's 42-year time offset is in effect. Any other value indicates that the normal UNIX epoch is assumed.

- `ZONE="filename"` — indicates the zonefile under `/usr/share/zoneinfo` that `/etc/localtime` is a copy of, for example:

```
ZONE="America/New York"
```

`/etc/sysconfig/desktop`

The `/etc/sysconfig/desktop` file specifies the desktop manager to run.

`/etc/sysconfig/harddisks`

The `/etc/sysconfig/harddisks` file allows you to tune your hard drive(s).

It may contain the following:

- `USE_DMA=1`, where setting this to 1 enables DMA. However, with some chipsets and hard drive combinations, this could cause some data corruption.
- `Multiple_IO=16`, where setting to 16 allows for multiple sectors per I/O interrupt. When enabled, this feature reduces operating system overhead by 30-50%. *Use with caution.*
- `EIDE_32BIT=3`, enables (E)IDE 32-bit I/O support to an interface card.
- `LOOKAHEAD=1`, enables drive read-lookahead.
- `EXTRA_PARAMS=`, where extra parameters can be added.

`/etc/sysconfig/hwconf`

The `/etc/sysconfig/hwconf` file lists all the hardware that kudzu detected on your system, as well as the drivers used, vendor ID and device ID information. It is not meant to be edited. If you do edit it, devices could suddenly show up as being *added* or *removed*.

`/etc/sysconfig/init`

The `/etc/sysconfig/init` file controls how the system will look during bootup.

The following values may be used:

- BOOTUP=*<some bootup mode>*, where <*some bootup mode*> is one of the following:
  - BOOTUP=color means a new (as of Red Hat Linux 6.0) boot display.
  - BOOTUP=verbose means an old style display.
  - Anything else means a new display, but without ANSI-formatting.

- LOGLEVEL=*<a number>*, where <*a number*> sets the initial console logging level for the kernel. The default is 7; 8 means everything (including debugging); 1 means nothing except kernel panics. syslogd will override this once it starts.
- RES_COL=*<a number>*, where <*a number*> is a column of the screen to start status labels at. Defaults to 60.
- MOVE_TO_COL=*<a command>*, where <*a command*> moves the cursor to $RES_COL. Defaults to ANSI sequences output by echo -e.
- SETCOLOR_SUCCESS=*<a command>*, where <*a command*> sets the color to a color indicating success. Defaults to ANSI sequences output by echo -e, setting the color to green.
- SETCOLOR_FAILURE=*<a command>*, where <*a command*> sets the color to a color indicating failure. Defaults to ANSI sequences output by echo -e, setting the color to red.
- SETCOLOR_WARNING=*<a command>*, where <*a command*> sets the color to a color indicating warning. Defaults to ANSI sequences output by echo -e, setting the color to yellow.
- SETCOLOR_NORMAL=*<a command>*, where <*a command*> sets the color to 'normal'. Defaults to ANSI sequences output by echo -e.
- PROMPT=*an answer*, where *an answer* is one of the following:
  - yes — Enables the key check for interactive mode.
  - no — Disables the key check for interactive mode.

/etc/sysconfig/i18n

The /etc/sysconfig/i18n file sets the default language, etc.

/etc/sysconfig/keyboard

The /etc/sysconfig/keyboard file controls the behavior of the keyboard. The following values may be used:

- KEYTABLE=*file*, where *file* is the name of a keytable file. For example: KEYTABLE="/usr/lib/kbd/keytables/us.map"
- KEYBOARDTYPE=sun|pc, which is used on SPARCs only. sun means a Sun keyboard is attached on /dev/kbd, pc means a PS/2 keyboard is on a PS/2 port.

/etc/sysconfig/mouse

The /etc/sysconfig/mouse file is used to specify information about the available mouse. The following values may be used:

- MOUSETYPE=*type*, where type is one of the following:
  - microsoft — A Microsoft mouse.
  - mouseman — A MouseMan mouse.
  - mousesystems — A Mouse Systems mouse.
  - ps/2 — A PS/2 mouse.
  - msbm — A Microsoft bus mouse.
  - logibm — A Logitech bus mouse.
  - atibm — An ATI bus mouse.
  - logitech — A Logitech mouse.
  - mmseries — An older MouseMan mouse.
  - mmhittab — An mmhittab mouse.

- XEMU3=*emulation*, where *emulation* is one of the following:

  - `yes` — Three mouse buttons should be emulated.
  - `no` — The mouse already has three buttons.

In addition, `/dev/mouse` is a symlink that points to the actual mouse device.

`/etc/sysconfig/network`

The `/etc/sysconfig/network` file is used to specify information about the desired network configuration. The following values may be used:

- `NETWORKING=`*`answer`*, where *answer* is one of the following:
  - `yes` — Networking should be configured.
  - `no` — Networking should not be configured.

- `HOSTNAME=`*`hostname`*, where *hostname* should be the FQDN (Fully Qualified Domain Name), but can be whatever hostname you want.

> **Please Note**
>
> For compatibility with older software that people might install (such as `trn`), the `/etc/HOSTNAME` file should contain the same value as here.

- `GATEWAY=`*`gw-ip`*, where *gw-ip* is the IP address of the network's gateway.
- `GATEWAYDEV=`*`gw-dev`*, where *gw-dev* is the gateway device (e.g. eth0).
- `NISDOMAIN=`*`dom-name`*, where *dom-name* is the NIS domain name.

`/etc/sysconfig/pcmcia`

The `/etc/sysconfig/pcmcia` file is used to specify PCMCIA configuration information. The following values may be used:

- `PCMCIA=`*`answer`*, where *answer* is one of the following:

  - `yes` — PCMCIA support should be enabled.
  - `no` — PCMCIA support should not be enabled.

- `PCIC=`*`pcic-type`*, where *pcic-type* is one of the following:
  - `i82365` — The computer has an i82365-style PCMCIA socket chipset.
  - `tcic` — The computer has a tcic-style PCMCIA socket chipset.

- `PCIC_OPTS=`*`option`*, where *option* is the socket driver (i82365 or tcic) timing parameters.
- `CORE_OPTS=`*`option`*, where *option* is the list of `pcmcia_core` options.
- `CARDMGR_OPTS=`*`option`*, where *option* is the list of options for the PCMCIA cardmgr (such as `-q`, quiet mode; `-m`, looks for loadable kernel modules in the specified director; and so on, read the cardmgr man page for more information).

`/etc/sysconfig/sendmail`

The `/etc/sysconfig/sendmail` allows messages to be sent to one or more recipients, routing the message over whatever networks are necessary. The file sets the default values for the sendmail program to run. Its default values are to run as a background daemon, and to check its queue once an hour in case something has backed up.

The following values may be used:

- `DAEMON=`*`answer`*, where *answer* is one of the following:
  - `yes` — Sendmail should be configured to listen to port 25 for incoming mail. `yes` implies `-bd`.
  - `no` — Sendmail should not be configured to listen to port 25 for incoming mail.

- QUEUE=1h which is given to sendmail as -q$QUEUE. The -q option is not given to sendmail if /etc/sysconfig/sendmail exists and QUEUE is empty or undefined.

/etc/sysconfig/soundcard

The /etc/sysconfig/soundcard file is generated by sndconfig and should not be modified. The sole use of this is to determine what card entry in the menu to pop up by default the next time sndconfig is run.

It may contain the following:

- CARDTYPE=*<a card>*, where <*a card*> is seen as, for example, CARDTYPE=SB16.

Files in /etc/sysconfig/network-scripts/

The following files are normally found in /etc/sysconfig/network-scripts:

- /etc/sysconfig/network-scripts/ifup
- /etc/sysconfig/network-scripts/ifdown
- /etc/sysconfig/network-scripts/network-functions
- /etc/sysconfig/network-scripts/ifcfg-*<interface-name>*
- /etc/sysconfig/network-scripts/ifcfg-*<interface-name>*-*<clone-name>*
- /etc/sysconfig/network-scripts/chat-*<interface-name>*
- /etc/sysconfig/network-scripts/dip-*<interface-name>*
- /etc/sysconfig/network-scripts/ifup-post

Let’s take a look at each one.

/etc/sysconfig/network-scripts/ifup,
/etc/sysconfig/network-scripts/ifdown

These are symbolic links to /sbin/ifup and /sbin/ifdown, respectively. These are the only two scripts in this directory that should be called directly; these two scripts call all the other scripts as needed. These symlinks are here for legacy purposes only — they will probably be removed in future versions, so only /sbin/ifup and /sbin/ifdown should currently be used.

These scripts take one argument normally: the name of the device (e.g. "eth0"). They are called with a second argument of "boot" during the boot sequence so that devices that are not meant to be brought up on boot (ONBOOT=no, [see below]) can be ignored at that time.

/etc/sysconfig/network-scripts/network-functions

Not really a public file. Contains functions which the scripts use for bringing interfaces up and down. In particular, it contains most of the code for handling alternative interface configurations and interface change notification through netreport.

/etc/sysconfig/network-scripts/ifcfg-*<interface-name>*,
/etc/sysconfig/network-scripts/ifcfg-*<interface-name>*:*<clone-name>*

The first file defines an interface, while the second file contains only the parts of the definition that are different in a "alias" (or alternative) interface. For example, the network numbers might be different, but everything else might be the same, so only the network numbers would be in the clone file, while all the device information would be in the base ifcfg file.

The items that can be defined in an ifcfg file depend on the interface type.

The following values are common:

- DEVICE=*name*, where *name* is the name of the physical device (except dynamically-allocated PPP devices where it is the "logical name").
- IPADDR=*addr*, where *addr* is the IP address.
- NETMASK=*mask*, where *mask* is the netmask value.
- NETWORK=*addr*, where *addr* is the network address.

- BROADCAST=*addr*, where *addr* is the broadcast address.
- GATEWAY=*addr*, where *addr* is the gateway address.
- ONBOOT=*answer*, where *answer* is one of the following:
  - yes — This device should be activated at boot-time.
  - no — This device should not be activated at boot-time.
- USERCTL=*answer*, where *answer* is one of the following:
  - yes — Non-root users are allowed to control this device.
  - no — Non-root users are not allowed to control this device.

- BOOTPROTO=*proto*, where *proto* is one of the following:
  - none — No boot-time protocol should be used.
  - bootp — The BOOTP protocol should be used.
  - dhcp — The DHCP protocol should be used.

The following values are common to all SLIP files:

- PERSIST=*answer*, where *answer* is one of the following:
  - yes — This device should be kept active at all times, even if deactivated after a modem hang up.
  - no — This device should not be kept active at all times.

- MODEMPORT=*port*, where *port* is the modem port's device name (for example, "/dev/modem").
- LINESPEED=*baud*, where *baud* is the modem's linespeed (for example, "115200").
- DEFABORT=*answer*, where *answer* is one of the following:

- yes — Insert default abort strings when creating/editing the script for this interface.
- no — Do not insert default abort strings when creating/editing the script for this interface.

/etc/sysconfig/network-scripts/chat-*<interface-name>*

This file is a chat script for or SLIP connections, and is intended to establish the connection. For SLIP devices, a DIP script is written from the chat script .

/etc/sysconfig/network-scripts/dip-*<interface-name>*

This write-only script is created from the chat script by netcfg. Do not modify this file. In the future, this file may disappear and instead will be created on-the-fly from the chat script.

/etc/sysconfig/network-scripts/ifup-post

This file is called when any network device (except a SLIP device) comes up. It calls /etc/sysconfig/network-scripts/ifup-routes to bring up static routes that depend on that device, brings up aliases for that device, and sets the hostname if it is not already set — and a hostname can be found for the IP for that device. ifup-post sends SIGIO to any programs that have requested notification of network events.

This file could be extended to fix up name service configuration, call arbitrary scripts, and more, as needed.

2.13.3 System V Init

This section is a brief description of the internals of the boot process. It discusses how the machine boots using SysV init, as well as the differences between the init used in older Linux releases, and SysV init.

The Init program is run by the kernel at boot time. It is in charge of starting all the normal processes that need to run at boot time. These include the getty processes that

allow you to log in, NFS daemons, FTP daemons, and anything else you want to run when your machine boots.

SysV init is quickly becoming the standard in the Linux world to control the startup of software at boot time, because it is easier to use and more powerful and flexible than the traditional BSD init.

SysV init also differs from BSD init in that the configuration files are in a subdirectory of `/etc` instead of residing directly in `/etc`. In `/etc/rc.d`, you will find `rc.sysinit` and the following directories:

```
init.d
rc0.d
rc1.d
rc2.d
rc3.d
rc4.d
rc5.d
rc6.d
```

The `init.d` directory contains a variety of scripts. Basically, you must have one script for each service you may need to start at boot time or when entering another runlevel. Services include things like networking, `nfs`, `sendmail`, `httpd`, and so on. Services do not include things like `setserial` that must only be run once and then exited. Things like that should go in `rc.local` or `rc.serial`.

If you want `rc.local`, it should be in `/etc/rc.d`. Most systems include one even though it doesn't do much. You can also include an `rc.serial` file in `/etc/rc.d` if you need to perform serial port specific tasks at boot time.

The chain of events is as follows:

- The kernel looks in several places for init and runs the first one it finds
- init runs `/etc/rc.d/rc.sysinit`
- `rc.sysinit` handles most of the boot loader's processes and then runs `rc.serial` (if it exists)
- init runs all the scripts for the default runlevel.
- init runs `rc.local`

The default runlevel is decided in `/etc/inittab`. You should have a line close to the top like:

```
id:3:initdefault:
```

From this, you'd look in the second column and see that the default runlevel is 3. If you want to change it, you can edit `/etc/inittab` by hand. Be very careful when you are editing the `inittab` file. If you do mess up, you can fix it by rebooting and typing:

```
LILO boot:  linux single
```

This *should* allow you to boot into single-user mode so you can re-edit `inittab` to its previous value.

Now, how does it run all the right scripts? If you enter `ls -l` on `rc3.d`, you might see something like:

```
lrwxrwxrwx 1 root root 17 3:11 S10network -> ../init.d/network
lrwxrwxrwx 1 root root 16 3:11 S30syslog -> ../init.d/syslog
lrwxrwxrwx 1 root root 14 3:32 S40cron -> ../init.d/cron
lrwxrwxrwx 1 root root 14 3:11 S50inet -> ../init.d/inet
lrwxrwxrwx 1 root root 13 3:11 S60nfs -> ../init.d/nfs
lrwxrwxrwx 1 root root 15 3:11 S70nfsfs -> ../init.d/nfsfs
lrwxrwxrwx 1 root root 18 3:11 S90lpd -> ../init.d/lpd.init
lrwxrwxrwx 1 root root 11 3:11 S99local -> ../rc.local
```

What you'll notice is that there are no "real" files in the directory. Everything there is a link to one of the scripts in the `init.d` directory. The links also have an "S" and a number at the beginning. The "S" means to start this particular script and a "K" would mean to stop it. The number in the file name is for ordering purposes. Init will start all the services based on the order in which they appear. You can duplicate numbers, but it will only confuse you somewhat. You only need to use a two digit number, along with an upper case "S" or "K" to start or stop the services you require.

How does init start and stop services? Simple. Each of the scripts is written to accept an argument which can be "`start`" and "`stop`". You can execute those scripts by hand, in fact, with a command like:

```
/etc/rc.d/init.d/httpd stop
```

This will stop the `httpd` server. init reads the name and if it has a "K," it calls the script with the "`stop`" argument. If it has an "S" it calls the script with a "`start`" argument.

Why all these runlevels? Some people want an easy way to set up machines to be multi-purpose. You could have a "server" runlevel that just runs `httpd`, `sendmail`, networking, etc. Then you could have a "user" runlevel that runs `gdm`, networking, and so on.

2.13.4 Init Runlevels

Generally, Red Hat Linux operates in runlevel 3 — full multi-user mode. The following runlevels are defined in Red Hat Linux:

- `0` — Halt
- `1` — Single-user mode
- `2` — Multi-user mode, without networking
- `3` — Full multi-user mode
- `4` — Not used
- `5` — Full multi-user mode (with an X-based login screen)
- `6` — Reboot

If your machine gets into a state where it will not boot due to a bad `/etc/inittab`, or will not let you log in because you have a corrupted `/etc/passwd` or have simply forgotten your password, boot into single-user mode by typing `linux single` at the LILO boot prompt. A very bare system will boot and you will have a shell from which you can fix things.

2.13.5 Initscript Utilities

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

In addition, there is the `ntsysv` utility, that provides a screen-oriented interface, versus `chkconfig`'s command-line interface.

Please see Section 2.10, *Controlling Access to Services* or the `chkconfig` and `ntsysv` man pages for more information.

2.13.6 Running Programs at Boot Time

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

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

2.13.7 Shutting Down

To shut down Red Hat Linux, issue the `shutdown` command. You can read the `shutdown` man page for complete details, but the two most common usages are:

```
shutdown -h now
shutdown -r now
```

Each will cleanly shutdown the system. After shutting everything down, the `-h` option will halt the machine, and the `-r` option will reboot.

Although the `reboot` and `halt` commands are now "smart" enough to invoke `shutdown` if run while the system is in runlevels 1-5, it is a bad habit to get into, as not all Linux-like operating systems have this feature.

2.14 Rescue Mode

When things go wrong, there are several ways to work on fixing them. However, they require that you understand the system well. We will present the ways that you can boot into rescue modes where you can use your own knowledge to rescue the system.

2.14.1 What is Rescue Mode?

Rescue mode is a term used to describe a method of booting a small Linux environment completely from a diskette, CD or other method.

What follows in this section may help you recover from a problem at some point.

As the name implies, rescue mode is there to rescue you from something. In normal operation, your Red Hat Linux system uses files located on your system's hard drive to do everything — run programs, store your files, and more.

However, there may be times when you are unable to get Linux running completely enough to access its files on your system's hard drive. Using rescue mode, it's possible to access the files stored on your system's hard drive, even if you can't actually run Linux from that hard drive.

Normally, you'll need to get into rescue mode for one of two reasons:

- You are unable to boot Linux, and you'd like to fix it.
- You are having hardware or software problems, and you want to get a few important files off your system's hard drive.

Let's take a closer look at each of these scenarios.

Unable to Boot Linux

Many times this is caused by the installation of another operating system after you've installed Red Hat Linux. Some other operating systems assume that you have no other operating systems on your computer, and overwrite the Master Boot Record (or MBR) that originally contained the LILO bootloader. If LILO is overwritten in this manner, you're out of luck — unless you can get into rescue mode.

Hardware/Software Problems

There can be as many different situations under this category as there are systems running Linux. Things like failing hard drives and forgetting to run LILO after building a new kernel are just two instances that can keep you from booting Red Hat Linux. If you can get into rescue mode, you might be able to resolve the problem — or at least get copies of your most important files.

To boot your system in rescue mode, enter the following parameter at the installation boot prompt:

```
boot: linux rescue
```

You can get to the installation boot prompt in one of these ways:

- By booting your system from the diskette or CD-ROM that came with your Red Hat Linux boxed set.
- By booting from a network or PCMCIA boot diskette. These methods assume your network connection is working and require you to identify the network host and transfer type. For an explanation of how to specify this information, see "Installing over the Network" in Chapter 14, *Installing Red Hat Linux via Text Mode*.

Once you have your system in rescue mode, a prompt appears on VC (virtual console) 2 (use the [Ctrl]-[Alt]-[F2] key combination to access VC 2 :

```
bash#
```

From this prompt, you can run the commands listed below:

```
anaconda          gzip        mkfs.ext2   ps
badblocks         head        mknod       python
bash              hwclock     mkraid      python1.5
cat               ifconfig    mkswap      raidstart
chatter           init        mlabel      raidstop
chmod             insmod      mmd         rcp
chroot            less        mmount      rlogin
clock             ln          mmove       rm
collage           loader      modprobe    rmmod
cp                ls          mount       route
cpio              lsattr      mpartition  rpm
dd                lsmod       mrd         rsh
ddcprobe          mattrib     mread       sed
depmode           mbadblocks  mren        sh
df                mcd         mshowfat    sync
e2fsck            mcopy       mt          tac
fdisk             mdel        mtools      tail
fsck              mdeltree    mtype       tar
fsck.ext2         mdir        mv          touch
ftp               mdu         mzip        traceroute
genhdlist         mformat     open        umount
gnome-pty-helper  minfo       pico        uncpio
grep              mkdir       ping        uniq
gunzip            mke2fs      probe       zcat
```

However, if your root filesystem is undamaged, you can mount it and then run any standard Linux utility. For example, suppose your root filesystem is in `/dev/hda5`. Here's how to mount this partition:

```
mount -t ext2 /dev/hda5 /foo
```

Where */foo* is a directory that you have created.

Now you can run `chroot`, `fsck`, `man`, and other utilities. At this point, you are running Linux in single-user mode.

If you don't know the names of your Linux partitions, you can guess what they are; mounting non-existent partitions will do no harm.

Booting Single-User Mode Directly

You may be able to boot single-user mode directly. If your system boots, but does not allow you to log in when it has completed booting, try rebooting and specifying one of these options at the LILO boot prompt:

```
LILO boot: linux single
LILO boot: linux emergency
```

In single-user mode, you computer boots to runlevel 1. Your local filesystems will be mounted but your network will not be activated. You get a usable system maintenance shell.

In emergency mode, you are booted into the most minimal environment possible. The root filesystem will be mounted read-only and almost nothing will be set up. The main advantage of this over `linux single` is that your init files are not loaded. If init is corrupted or not working, you can still mount filesystems to recover data that could be lost during a re-installation.

A Handy Trick

Have you ever rebuilt a kernel and, eager to try out your new handiwork, rebooted before running LILO? And you didn't have an entry for an older kernel in `lilo.conf`? Read on…

In many cases, it's possible to boot your Red Hat Linux/Intel system from the Red Hat Linux boot disk with your root filesystem mounted and ready to go. Here's how:

Enter the following command at the boot disk's `boot:` prompt:

```
linux single root=/dev/hdXX initrd=
```

(Replace the *XX* in `/dev/hd`*XX* with the appropriate letter and number for your root partition.)

What does this do? First, it starts the boot in single-user mode, with the root partition set to your root partition. The empty `initrd` specification bypasses the installation-related image on the boot disk, which will cause you to enter single-user mode immediately.

Is there a downside? Unfortunately, yes. Because the kernel on the Red Hat Linux boot disk only has support for IDE built-in, those of you with SCSI-based systems won't be able to use this trick. In that case, you'll have to use the boot/rescue disk combination mentioned above.

3 System Configuration

One of the main strengths of Red Hat Linux is that the operating system can be configured to do just about anything. In the past, this meant editing sometimes cryptic configuration files by hand, then restarting system services and crossing your fingers in the hope that your changes were correctly made. Additionally, the variety of configuration options can be bewildering to new users, who may not know where to look for a particular configuration file.

Red Hat Linux provides two system configuration utilities: linuxconf and the control panel. The control panel provides a launcher for various Red Hat system configuration tools, including linuxconf. Instructions on how to use the control panel and the tools it includes can be found in Section 3.2, *System Configuration with the Control Panel*.

The linuxconf configuration tool can ease some of the burdens of system configuration. Linuxconf is both:

- A configuration interface — You type the values needed to configure your system into a user interface.
- A configuration activator — When you're satisfied with your edits, you tell linuxconf to apply the changes you've made.

While linuxconf can do nearly everything the control panel tools can, there are two areas in which the control panel still holds the upper hand:

- Printer configuration
- Loading kernel modules to support new hardware

Let's take a look at linuxconf.

3.1 System Configuration with linuxconf

Linuxconf allows you to configure and control various aspects of your system, and is capable of handling a wide range of programs and tasks. Complete documentation of linuxconf could be a separate book in its own right and is certainly more than we

can cover in this chapter. Instead, we'll focus on common tasks such as adding new users and getting connected to a network.

More information on linuxconf, including its most recent release can be found at the linuxconf website:

http://www.solucorp.qc.ca/linuxconf/

The linuxconf website includes fairly extensive information, including a description, rationale, history, list of contacts, and a lot of other information in addition to the software. The website is maintained by linuxconf's creator and developer, Jacques Gelinas, so it includes the latest news about linuxconf.

This chapter will go into detail on just a few of linuxconf's capabilities. If you just need a quick reference to show you where to go in linuxconf for the tasks covered by this chapter, see Section 3.1.15, *Finding Your Way Through linuxconf.*

If you need more help with linuxconf, please try these sources:

- The linuxconf FAQ, which is available at http://www.xc.org/jonathan/linuxconf-faq.html.
- The archives of the linuxconf mailing list, which are available at http://hub.xc.org/scripts/lyris.pl?visit=linuxconf.
- After you've checked the linuxconf FAQ and the archives of the linuxconf list, you might try posting your question to the linuxconf list. Subscription information for the linuxconf list is available at the linuxconf website (http://www.solucorp.qc.ca/linuxconf/); click on the "Mailing lists" link.

  Please note that this list is for questions pertaining to linuxconf, and is not intended for general Linux questions.

3.1.1 Running linuxconf

You'll need to be root to run linuxconf, so if you're in your user account, `su` to become root.

Now, type `linuxconf` at the shell prompt to begin the program.

3.1.2 Linuxconf User Interfaces

Linuxconf has four user interfaces:

- Text-based — Using the same user interface style as the Red Hat Linux text-mode installation program, the text-based interface makes it easy to navigate your way through linuxconf if you aren't running X. If you are running X, you can switch to a virtual console, log in as root, and type `linuxconf` to bring up text-mode linuxconf.

  Use the [Tab] and [arrow] keys to navigate the text-mode screens. A **down arrow** on a line indicates that a pull-down menu exists on that line. The [Ctrl]-[X] key combination will make pull-down menus appear.

- Graphical user interface (GUI) — linuxconf can take advantage of the X Window System. Red Hat Linux includes a GUI interface for linuxconf called gnome-linuxconf.

  This document will display linuxconf screens using the gnome-linuxconf interface, but you shouldn't have any trouble using the other interfaces with the instructions provided here.

- Web-based — A Web-based interface makes remote system administration a breeze; it can also be displayed with the Lynx text-mode browser.

  To use the linuxconf Web interface, use your browser to connect to port 98 on the machine running linuxconf (i.e., http://your_machine:98).

  Before you use the Web-based interface, you'll need to configure linuxconf to allow connections from the machine running the browser. See Section 3.1.4, *Enabling Web-Based linuxconf Access* for instructions on enabling Web access to linuxconf.

- Command line — linuxconf's command-line mode is handy for manipulating your system's configuration in scripts.

Linuxconf will start in either character-cell or X mode, depending on your **DISPLAY** environment variable. The first time you run linuxconf, an introductory message will be shown; although it is only displayed once, accessing help from the main screen will give you the same basic information.

Linuxconf includes some context-specific help. For information on any specific aspect of linuxconf, select **Help** from the screen you'd like help with. Note that not all help screens are complete at this time; as help screens are updated, they will be included in subsequent versions of linuxconf.

3.1.3 gnome-linuxconf Interface

gnome-linuxconf makes it easy to navigate the hierarchical structure of linuxconf.

Figure 3–1 Linuxconf Menu View

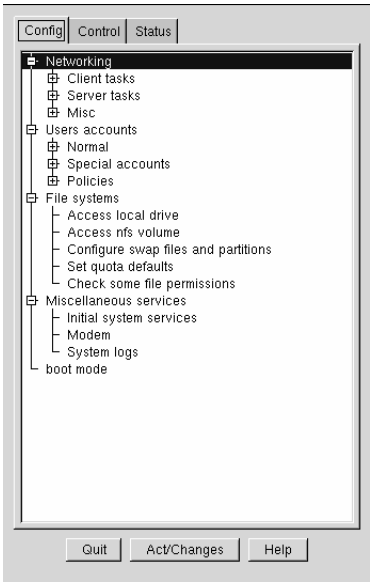

Please Note

If you don't see the tree menu interface shown above, follow these instructions:

1. Open **Control** => **Control files and systems** => **Configure linuxconf modules**
2. Select the treemenu check box

3. Click **Accept**
4. Click **Quit**
5. Restart linuxconf

When you use the tree menu view, finding the appropriate panel should be simple and fast. Collapse and expand sections by clicking on the + or - next to the menu item.

Selected entries will appear as tabs in the right-hand panel and will remain there until closed. If you end up with more tabs open than you like, just select **Cancel** on the bottom of each tab to close it without making any changes, or **Accept** to implement them.

3.1.4 Enabling Web-Based linuxconf Access

For security reasons, Web-based access to linuxconf is disabled by default. Before attempting to access linuxconf with a Web browser, you'll need to enable access. Here's how to do it:

1. Open **Config** => **Networking** => **Misc** => **Linuxconf network access**
2. In the **Linuxconf html access control** dialog box, enter the hostname of any computers that should be allowed access to Linuxconf. This includes your own system, if you wish to use the Web-based interface locally. Web accesses related to linuxconf may be logged to your system's `htmlaccess.log` file by selecting the check box.
3. Select the **Accept** button.

Web-based access should be enabled. To test it out, go to a system that you added to the access control list. Then, launch your Web browser, and enter the following URL:

```
http://<host>:98/
```

(Replace *<host>* with your system's hostname, of course. Also, remove the "disable=yes" line from the `/etc/xinetd.d/linuxconf` file and then run the command `/sbin/service xinetd reload` from a shell prompt.) You should see the main linuxconf page. Note that you will need to enter your system's root password to gain access beyond the first page.

You can also enable network-wide access to linuxconf by following the same steps, by entering a network name instead of a hostname.

3.1.5 Adding a User Account

Adding a user is one of the most basic tasks you will encounter in administering your system. To add a user:

- Open **Config** => **Users accounts** => **Normal** => **User accounts**. Linuxconf may show you a filter screen (see Figure 3–2, *Filter Control Screen*).

Figure 3–2 Filter Control Screen

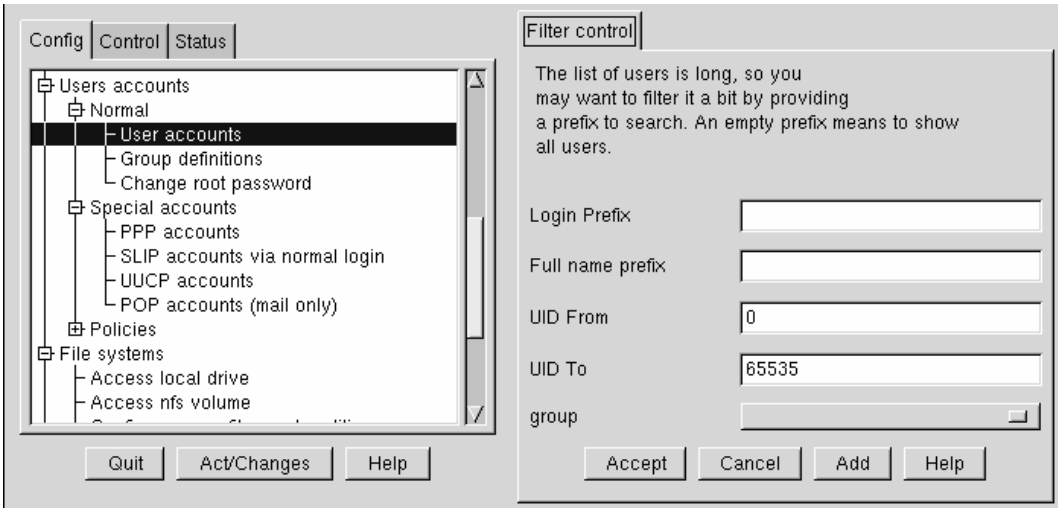

You can use the filter screen to select a smaller range of accounts than the full list. To get the full list, select **Accept** without changing any of the parameters. For detailed information on the various filters, select the **Help** button on the **Filter control** screen. Once you've applied or bypassed the filter, you'll see the **Users accounts** tab (see Figure 3–4, *Users Accounts Screen*).

(You can control the filter using **Control** => **Features**. You'll see the **Features** tab, which allows you to set the **Trigger for filter** parameter, as shown in Figure 3–3, *Setting the Trigger for Filter*.

Figure 3–3 Setting the Trigger for Filter

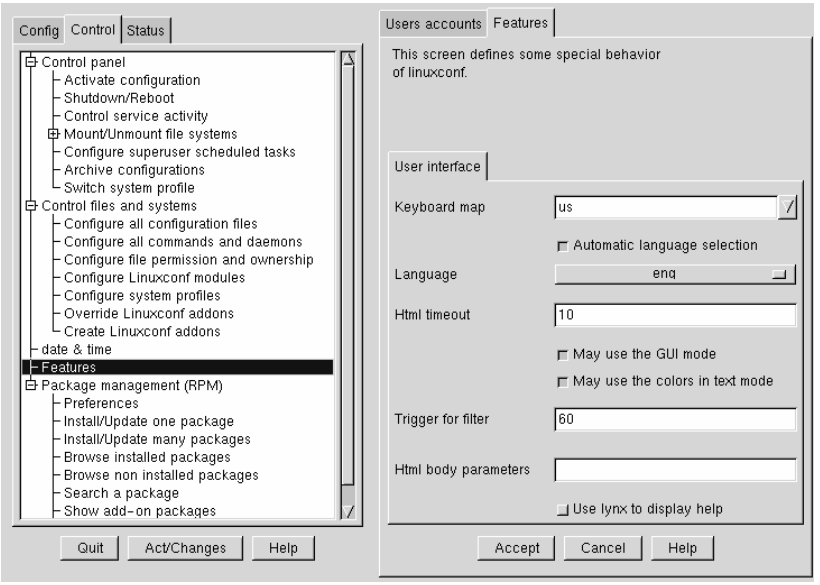

The **Trigger for filter** field sets the number of entries that will pop up a filter screen.)

Figure 3–4 Users Accounts Screen

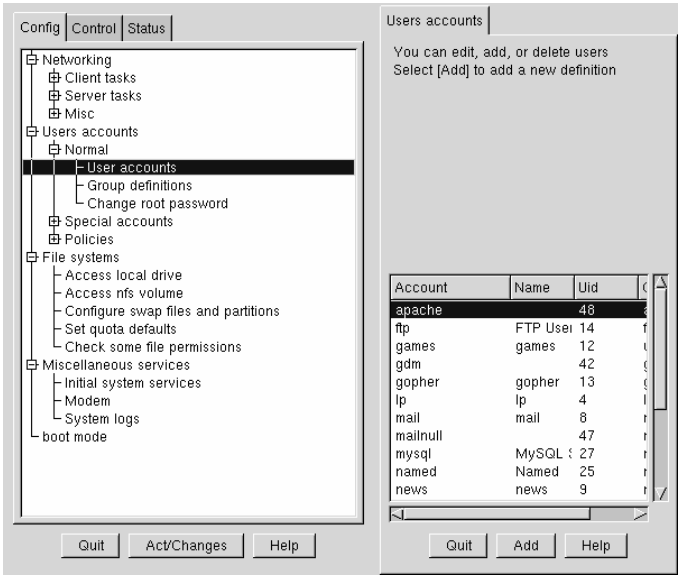

- Select **Add**. This will open the **User account creation** tab (see Figure 3–5, *User Account Creation*).

The **User account creation** screen includes the **Base info, Params** and **Privileges** sections. Only the **Login name** is required, but you should be aware of the other fields, which you may or may not want to fill in.

Figure 3–5 User Account Creation

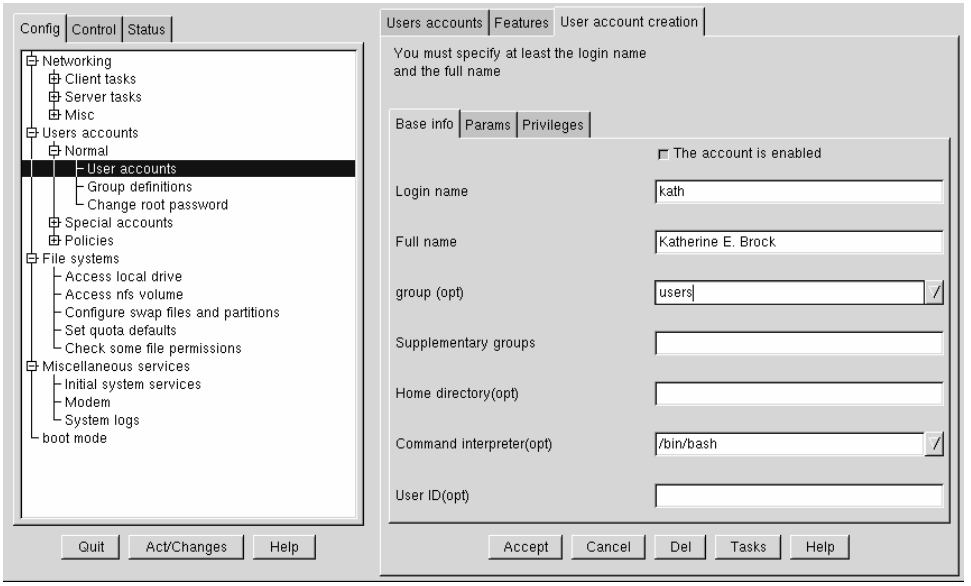

Base info for User Accounts

The **Login name** is the name of the account and is usually all lowercase letters. First or last names, initials or some combination thereof are fairly common login names. For a user named John T. Smith, **smith**, **john**, **jts**, or **jsmith** would be common user names. Of course **spike** or something else works just fine, too. You can also use numbers, so **jts2** would be fine if you had a second person with the same initials. There is no default for this field.

The **Full name** is the name of the user or the account. For an individual, it would be their name, **John T. Smith** for example. If the account represents a position rather than a person, the full name might be the title. So an account called **webmaster** might have a full name of **Red Hat Webmaster** or just **Webmaster**. There is no default for this field.

Since Red Hat Linux uses the User Private Group scheme, each user will be assigned to a default **group** consisting only of the user. For more information on User Private Groups, see Section 2.3.3, *User Private Groups*.

In the **Supplementary groups** field, you can specify additional groups. Group names should be separated by spaces. The default for this field is blank, meaning no supplementary groups are specified.

The **Home directory** specifies the home or login directory for the account. The default is `/home/`*`login`*, where *`login`* is replaced by the login name. A home directory is your starting point in the directory structure when you log in, or if in X, for each Xterm window opened. This is also where account specific preference files are stored.

The **Command interpreter** is the default shell for the account. The `bash` shell is the default shell for Red Hat Linux.

The **User ID** (UID) is the number associated with each user account. This is automatically generated by the system when the account is created, so just leave this field blank. The system uses the UID to identify an account.

Params for User Accounts

The **Params** are used for password and account management. By default, all of the settings are **Ignored**, so they are unused. **Must keep # days** sets a minimum number of days for a user's password.

The **Must change after # days** field can be set to make a user's password expire after a certain number of days. If you want to warn them that the password is going to expire (a good idea), the **Warn # days before expiration** field should be used.

If you'd like to have their account set to expire after a certain number of days, use the **Account expire after # days** field. You could alternatively set a hard **Expiration date**.

Privileges for User Accounts

In the **Privileges** section, you can grant access and/or control over various aspects of system configuration. As a default, regular users are denied all privileges on this screen. You may instead choose to grant or to silently grant them specific privileges. The difference between **Granted** and **Granted/silent** is that if the privilege is granted,

linuxconf will ask for the user's password before allowing them the privilege. If the privilege is granted silently, linuxconf will not prompt for their password.

Generally, careful system administrators won't grant users any system configuration privileges unless it is absolutely necessary. If you do grant privileges, be careful when granting them silently. If a user with silently granted privileges logs in to his/her machine and walks away, their privileges are wide open for the next person who sits down at their desk. Silently granted privileges are less risky if used on machines in a physically restricted area.

May use linuxconf: the user is allowed to access all of linuxconf's capabilities, and they can set up or change **linuxconf** parameters. Note that use of **linuxconf** is separate from the privilege of activating configuration changes. System administrators might want to grant the use of linuxconf, but deny the activation privilege, so that the sysadmin has a final "yes/no" on whether to activate any configuration changes.

May activate config changes: After you change a parameter in linuxconf, at some point you'll have to indicate to linuxconf that the changes you made should be applied. Depending upon the flavor of linuxconf that you're using, you might do this by clicking on an **Activate the changes** button in GUI linuxconf, or clicking on an **Accept** button in Web-based linuxconf, or selecting an **Accept** button in text-mode linuxconf, etc.

You can grant the privilege of activating changes to a user. In that case, the user will be able to activate any changed system configuration parameters in linuxconf.

May shutdown: A user can be granted the right to shutdown the system. Note that Red Hat Linux is set in `/etc/inittab` to cleanly shutdown following the [Ctrl]-[Alt]-[Del] keystroke combination.

You can also grant the user the privileges to switch network modes, to view system logs, and even give someone superuser equivalence.

Once you have entered the login name and any other desired information select the **Accept** button at the bottom of the screen. If you decide against creating a new user, select **Cancel** instead.

When you click on **Accept** linuxconf will prompt you to enter the password, as in Figure 3–6, *Change Password Screen*. You'll have to re-type the password, to prevent unusable passwords caused by typos. Passwords must be at least six characters in

length, but you can increase the required length and set other parameters for users' passwords at the **Users Accounts** => **Policies** => **Password & Account Policies** screen.

Good passwords contain a combination of letters, numbers, and special characters. It should use both upper case and lower case letters. Don't use your username, your anniversary, your social security number, your dog's name, your middle name or the word root. Don't use any variation of a word associated with your account or with yourself. Don't use a word that can be found in a dictionary; dictionary words are easy to crack.

A simple technique for creating a password is to use the first letters from each word of a phrase that is familiar to you (a line from a favorite song might be appropriate). Make a few letters uppercase, and insert a few numbers and/or special characters in place of letters and you'll have a decent password.

Press the **Accept** button again when finished. The system will let you know if it thinks the password is easy to crack; if you get a warning message, don't use the password.

3.1.6 Modifying a User Account

- Go to **Config** => **Users accounts** => **Normal** => **User accounts**, use the filter if necessary, and then select the account that you wish to modify.
- See Section 3.1.5, *Adding a User Account* if you need guidance for how to fill in the user accounts fields.

To implement the changes select **Accept**. If you decide against making any changes select **Cancel**. This guarantees that no changes are made.

3.1.7 Changing a User's Password

- Open **Config** => **Users accounts** => **Normal** => **User accounts**. This will open the **Users accounts** tab (see Figure 3–4, *Users Accounts Screen*).
- You may see a filter screen, depending upon the settings you've provided on the **Control** => **Features** screen. If you want the full list, select `Accept` without changing any of the parameters. For detailed information on the various filters, select the `Help` button on the **Filter control** screen.

- Select the account whose password you wish to change. This will open the **User information** screen.
- Select **Passwd** from the options at the bottom of the screen.

Linuxconf will prompt you to enter the new password. There is also a field called `Confirmation` where you will need to type the password again. This is to prevent you from mistyping the password. See Section 3.1.5, *Adding a User Account* for guidance on choosing a password. If you decide against changing the password, select `Cancel`. Once you have entered the new password select `Accept`.

Figure 3–6 Change Password Screen

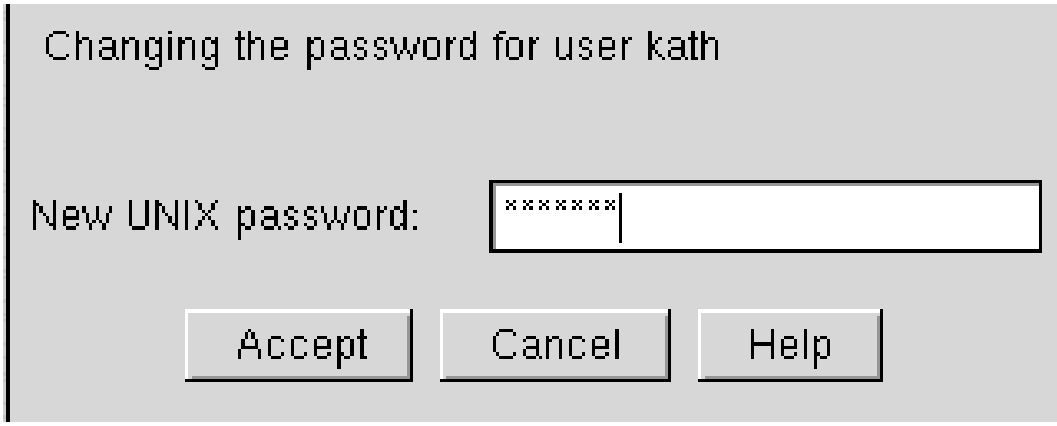

3.1.8 Changing the Root Password

Because of the security implications of root access, linuxconf requires you to verify that you currently have access to the root account.

- Open **Config** => **Users accounts** => **Normal** => **Change root password**.

You'll first need to enter the current root password to verify access to the root account.

Once you have entered root's current password, it will prompt you for a new password. In the **Confirmation** field, type the password again. This is to prevent you from mistyping the password. See Section 3.1.5, *Adding a User Account* if you need guidance on choosing a password. Be sure to choose a good password! If you decide

against changing the root password, just select **Cancel**. Once you have entered the new password select **Accept**.

3.1.9 Disabling a User Account

Disabling a user's account is preferable to deleting a user's account, unless you need the storage space or you're certain that his/her data will not be needed in the future. If a user's account is disabled, they will not be allowed to log in.

- Open **Config** => **Users accounts** => **Normal** => **User accounts**.
- Select an account.
- De-select the check-box that states that `The account is enabled`. Select the `Accept` button at the bottom of the window and you're all set.

The account is disabled and can be enabled later using a similar method.

3.1.10 Enabling a User Account

By default, all newly created user accounts are enabled. If you need to enable an account, you can use linuxconf to do it.

Open **Config** => **Users accounts** => **Normal** => **User accounts**. Select an account. Select the **The account is enabled** checkbox.

3.1.11 Deleting a User Account

Please Note

While there are options for retaining files associated with an account, any files that are deleted are gone and effectively unrecoverable. Take care when using this option!

To delete an account:

- Open **Config** => **Users accounts** => **Normal** => **User accounts**.
- On the **User accounts** screen (see Figure 3–4, *Users Accounts Screen*) select the account you wish to delete.

- At the bottom of the **User information** screen, select **Del** to delete the account.

Linuxconf will then prompt you with a list of options.

Figure 3–7 Deleting Account Screen

The default option is to archive the account's data. The archive option has the following effects:

1. Removes the user from the user accounts list;
2. Takes everything contained in the user's home directory and archives it (using tar and gzip compression), storing the resulting file in the `/default_home_directory/oldaccounts` directory. For an account named *`useraccount`* the filename would be similar to:

   `useraccount``-2000-01-10-497.tar.gz`

   The date indicates when the account was deleted, and the number following it is the ID of the process that actually performed the deletion. The `oldaccounts`

directory is created in the same place as all of your user directories, and is created automatically the first time you remove a user account using this option.

3. Files not contained in the user's home directory, but owned by that user remain. The file is owned by the deleted account's user ID (UID). If you create a new account and specifically assign it the UID of a deleted account, it will then become the owner of any remaining files.

Selecting **Delete the account's data** on the **Deleting account *<accountname>*** screen (see Figure 3–7, *Deleting Account Screen*) will:

1. Remove the user from the user accounts list;
2. Remove the user's home directory and all its contents.

Please Note

Files not contained in the user's home directory, but owned by that user will remain on the system. The file will still be owned by the deleted account's user ID (UID). If you create a new account and specifically assign it the UID of a deleted account, it will then become the owner of any such "orphaned" files.

Selecting **Leave the account's data in place** on the **Deleting account *<accountname>*** screen (see Figure 3–7, *Deleting Account Screen*) will:

1. Remove the user from the user accounts list;
2. Leave the user's home directory (with all its files) in place.

Please Note

Files and directories owned by the deleted account's user ID (UID) will remain on the system. If you create a new account and specifically assign it the UID of a deleted account, it will then become the owner of these "orphaned" files.

3.1.12 Groups

All users belong to one or more groups. Just as each file has a specific owner, each file belongs to a particular group as well. The group might be specific to the owner of the file, or may be a group shared by all users. The ability to read, write or execute a file can be assigned to a group; this is separate from the owner's rights. For example, the owner of a file will be able to write to a document, while other group members may only be able to read it.

Creating a Group

To create a new group:

- Open **Config** => **Users accounts** => **Normal** => **Group definition.**

You may see a filter screen, depending upon the settings in **Control** => **Filters.** Either provide a filter, or select **Accept** to bypass the filter.

Figure 3–8 User Groups Screen

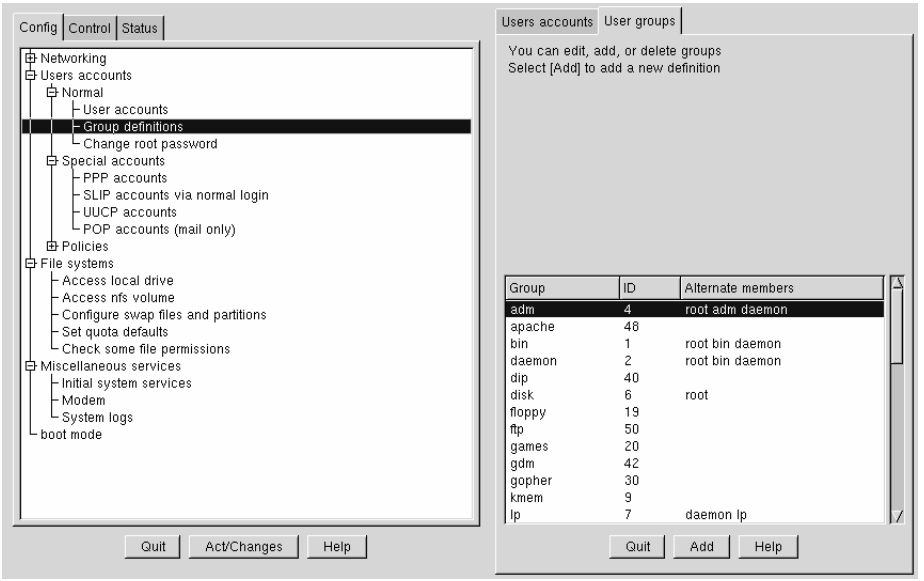

Select Add at the bottom of the **User groups** screen.

Figure 3–9 Group Specification Screen

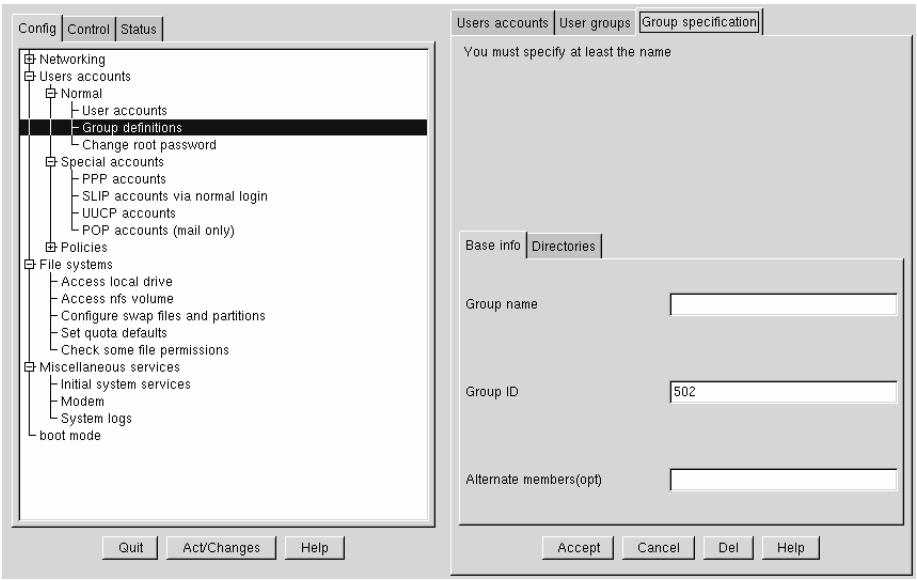

Enter a group name. You may also wish to specify members of the group and can do so in the **Alternate members** field. The list of users should be space delimited, meaning that each username must have a space between it and the next one. Leave the **Group name** field blank, so that the system will assign a **Group ID** (GID) to your new group. When you're finished, select `Accept` and the group will be created.

Deleting a Group

To delete a group:

- Open **Config** => **Users accounts** => **Normal** => **Group definitions**.

  You may see a filter screen, depending upon the filter setting in **Control** => **Features**. You can use the filter to narrow your choice of groups by specifying a prefix.

- With or without a prefix select `Accept` at the bottom of the screen.

- On the **User groups** screen (see Figure 3–8, *User Groups Screen*) select the group you wish to delete.

- You'll be presented with the **Group specification** screen (see Figure 3–9, *Group Specification Screen*).
- Select `Del` to delete the group. Linuxconf will then prompt you to confirm the deletion. Choose **yes** to delete the group.

The group's files will still remain and their respective owners will still have sole control over them. The group name will be replaced with the deleted group's ID. The files may be assigned to a new group by using the `chgrp` command. More information on `chgrp` can be found by typing the command `info chgrp` or `man chgrp` at the shell prompt. If a new group is created and the deleted group's ID is specified then the new group will have access to the deleted group's files. Don't worry, linuxconf doesn't recycle old group numbers any more than it does old user IDs, so it won't happen by accident.

Modifying Group Membership

There are two ways to modify the list of users that belong to a group. You can either update each user account itself, or you can update the group definitions. In general, the fastest way is to update each of the group definitions. If you're planning on changing more information for each user than just the group information, then updating each user account may prove easier.

We'll start by detailing the group definitions method.

- Start linuxconf by typing `linuxconf` at the shell prompt.
- Open **Config** => **Users accounts** => **Normal** => **Group definitions.**

  Depending on the filter settings in **Control** => **Features**, you may see a filter screen. Use the filter to narrow the list, or just select **Accept** to bypass the filter.
- Select the group you wish to modify. This will open the `Group specification` screen (see Figure 3–9, *Group Specification Screen*).
- Add or remove each user from the **Alternate members** field. Make sure that all of the user names are separated by a space character.
- Select `Accept`, which can be found at the bottom of the screen.

This will automatically update each user account with the group showing up in the **Supplementary groups** field if added or absent if removed.

Adding and removing groups can also be done by modifying each individual user account.

- Start linuxconf by typing `linuxconf` at the shell prompt.
- Open **Config** => **Users accounts** => **Normal** => **User accounts.**

  You may see a filter screen, depending on the settings in **Control** => **Features.** Use the filter to narrow the list or select **Accept** to bypass the filter.
- On the **User accounts** screen (see Figure 3–4, *Users Accounts Screen*), select a user that you wish to update. You will be presented with the **User information** screen.
- Add or remove the desired groups from the **Supplementary groups** field. Each group should be separated by a space character.
- Once you've made all the changes you'd like, select `Accept` at the bottom of the screen.

This will automatically update the group definitions. Repeat the process for each user.

3.1.13 Filesystems

A filesystem is composed of files and directories, all starting from a single root directory. The root directory may contain any number of files and other directories, with each directory in turn following suit. The average filesystem often looks like an inverted tree with the directories as branches and the files as leaves. Filesystems reside on mass storage devices such as diskette drives, hard drives, and CD-ROMs.

For example, a diskette drive on DOS and Windows machines is typically referenced by `A:\`. This describes both the device (`A:`), and the root directory on that device (\). The primary hard drive on the same systems is typically referred to as the "C" drive because the device specification for the first hard drive is `C:`. To specify the root directory on the C drive, you would use `C:\`.

Under this arrangement, there are two filesystems — the one on `A:`, and the one on `C:`. In order to specify *any* file on a DOS/Windows filesystem, you must either

explicitly specify the device on which the file resides, or it must be on the system's default drive (which is where DOS' C prompt comes from — that's the default drive in a system with a single hard drive).

Under Linux, it is possible to link the filesystems on several mass storage devices together into a single, larger, filesystem. This is done by placing one device's filesystem "under" a directory on another device's filesystem. So while the root directory of a diskette drive on a DOS machine may be referred to as `A:\`, the same drive on a Linux system may be accessible as `/mnt/floppy`.

The process of merging filesystems in this way is known as **mounting**. When a device is mounted, it is then accessible to the system's users. The directory "under" which a mounted device's filesystem becomes accessible is known as the **mount point**. In the previous paragraph's example, `/mnt/floppy` was the diskette drive's mount point. Note that there are no restrictions (other than common conventions) as to the naming of mount points. We could have just as easily mounted the floppy to `/long/path/to/the/floppy/drive`.

One thing to keep in mind is that all of a device's files and directories are relative to its mount point. Consider the following example:

- A Linux System:
  - `/` — system root directory
  - `/foo` — mount point for the CD-ROM

- A CD-ROM:
  - `/` — CD-ROM's root directory
  - `/images` — A directory of images on the CD-ROM
  - `/images/old` — A directory of old images

So, if the above describes the individual filesystems, and you mount the CD-ROM at `/foo`, the new operating system directory structure would be:

- A Linux System (with the CD-ROM mounted):
  - `/` — System root directory
  - `/foo` — CD-ROM root directory
  - `/foo/images` — A directory of images on the CD-ROM
  - `/foo/images/old` — A directory of old images

To mount a filesystem make sure to be logged in as root, or become root using the `su` command. For the latter, type `su` at the shell prompt and then enter the root password. Once you are root, type `mount` followed by the device and then the mount point. For example, to mount the first diskette drive on `/mnt/floppy`, you would type the command `mount /dev/fd0 /mnt/floppy`.

At installation, Red Hat Linux will create `/etc/fstab`. This file contains information on devices and associated mount points. The advantage to this file is that it allows you to shorten your mount commands and it controls which filesystems are automatically mounted when the system is booted.

Using the information in `/etc/fstab`, you can type `mount` and then either the mount point or the device. The `mount` command will look for the rest of the information in `/etc/fstab`. It's possible to modify this file by hand, or by using linuxconf.

Reviewing Your Current Filesystem

We'll start by looking at your current directory structure.

- Open **Config** => **File systems** => **Access local drive.**

Figure 3–10 Local Volume Screen

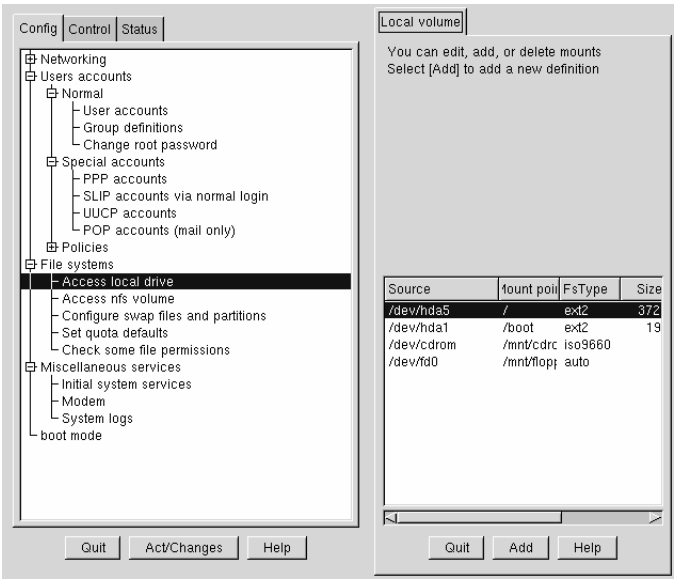

The fields, as shown in Figure 3–10, *Local Volume Screen*, are:

- **Source**: The physical hardware; `hd` indicates an IDE hard drive, `fd` indicates a diskette drive, and `cdrom` typically indicates a CD-ROM drive. If your system has a SCSI drive, you will see an `sd` instead. More than one drive of a type are listed by letters, so `hda` represents the first IDE drive, while `hdb` would be the second. In some cases, you'll see numbers following these letters; on hard drives, the numbers represent the partitions on that drive, while for diskette drives, this number refers to the actual unit.
- **Mount point**: This is where in the system the drive is accessible from when mounted.
- **FsType**: The type of filesystem. A standard Linux partition uses the ext2 filesystem type. A filesystem type of vfat indicates a DOS filesystem with long filename support, while a fat filesystem type is for DOS filesystems supporting traditional 8.3 filenames. The iso9660 filesystem type indicates a CD-ROM drive.

Please Note

Red Hat Linux 7.0 can access FAT32 filesystems using the `vfat` filesystem type.

- **Size**: Size may indicate the size of the filesystem in megabytes (M), or it may not be filled in.
- **Partition type**: A description of the filesystem used on that partition (may not be filled in).
- **Status**: Whether the device is mounted or not.

Filesystems from other machines on a network may also be available. These can range from single small directories to entire volumes. No information on **Size** or **Partition type** is available for these partitions, either. Additional information on NFS filesystems (should you have any available) will be contained under:

Config => **File systems** => **Access nfs volume**

The screen is similar to the **Local Volume** screen (see Figure 3–10, *Local Volume Screen*, with some notable differences in the information provided for each entry:

- **Source**: This will be the name of the machine serving the filesystem, followed by the remote directory. For example: `foo:/var/spool/mail` where `foo` is the machine serving the directory, and `/var/spool/mail` is the directory being served.
- FsType — This will always be "nfs."

Adding NFS Mounts

NFS (Network File System) is a way for computers to share sections of their local filesystem across a network. These sections may be as small as a single directory, or include thousands of files in a vast hierarchy of directories. For example, many companies will have a single mail server with individuals' mail files served as an NFS mount to each users' local systems.

To add an NFS mount:

- Open **Config** => **File systems** => **Access nfs volume.**
- On the **NFS volume** screen, select Add.

Figure 3–11 Volume Specification Screen

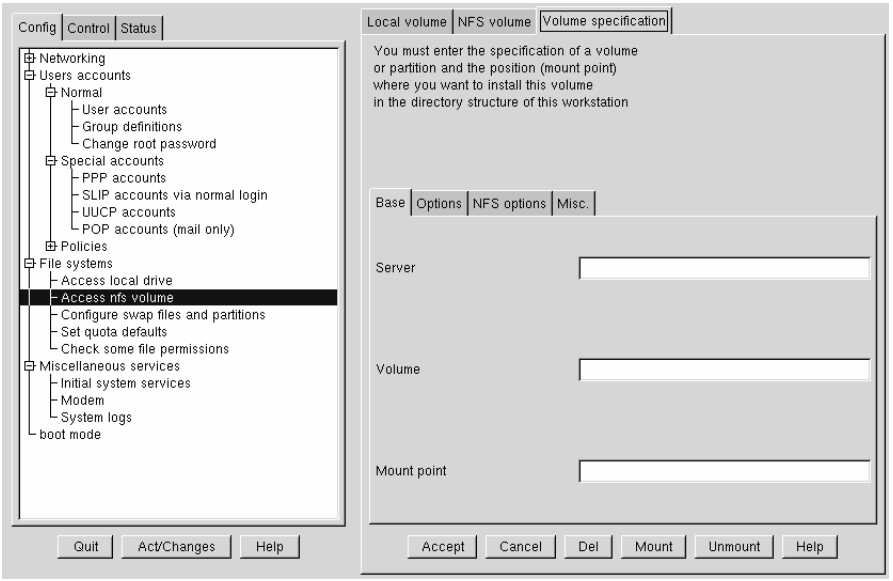

The three fields on the **Base** tab are what you'll need to concern yourself with next (see Figure 3–11, *Volume Specification Screen*).

- **Server**: The host name of the machine the desired filesystem resides on. For example, foo.bar.com.
- **Volume**: The filesystem you wish to add. For example, `/var/spool/mail`.
- **Mount point**: Where in your system you want the remote file system accessible from. For example, `/mnt/mail`.

This is all you need to get the mount created. Linuxconf will update your `/etc/fstab` file accordingly. If you are aware of additional requirements, please

read the help file on the **Volume specification** screen and see the `mount` man page for more information.

Once you have entered the information, select `Accept`.

3.1.14 Getting Connected with Linuxconf (Network Configuration)

The first thing to determine when getting hooked up is whether you're connecting to a local area network, such as a group of computers in an office, or a wide area network, such as the Internet. Before continuing, it's important to know what hardware you have and how you intend to connect. If you're going to dial into another computer, then make sure your modem is installed and that the cables are arranged properly. If you're using a network card, make sure it is installed properly and that the cables are correctly connected. Regardless of what network configuration you specify, if every phone line or cable is not in place, you'll never get connected.

Network Connections

Setting up a network connection over Ethernet requires an entirely different type of setup. Network connections to Token Ring or ARCnet networks follow a similar procedure, but will not be discussed here.

- First you will need to have an Ethernet card installed.
- Start linuxconf by typing `linuxconf` at the shell prompt.
- Open **Config** => **Networking** => **Client tasks** => **Basic host information**. The `Host name` tab will request a host name, which should be specified by default unless you did not setup your networking during the installation process. If it is not already specified, please take the time now to configure it. It should be specified as `localhost.localdomain`. Skip this tab. Select the tab for `Adaptor 1`.

Figure 3–12 Adaptor 1

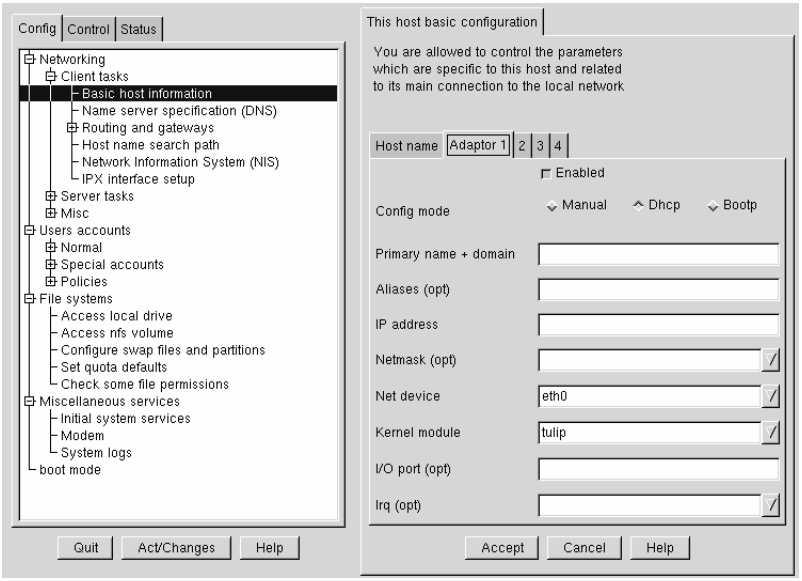

The first item on this screen is a check box to indicate whether this adaptor is enabled or not. It should be checked if this is the one you intend to use. Below that is a choice of **Config modes**. **Manual** means that you will be providing all the information and entering it yourself. **Dhcp** and **Bootp** mean that your machine will be getting its network configuration information from a remote DHCP or BOOTP server. If you're not sure what option to choose, talk to your network administrator.

Required fields for DHCP or BOOTP:

- **Net device** — The type of network card you are using; for example, eth0 would be the appropriate entry to use the first Ethernet card.
- **Kernel module** — The correct module based on your network card; for further information see the list below.

For DHCP and bootp configurations you only need to specify the **Net device** and the **Kernel module**. For the **Net device**, you will choose from a list where the **eth** prefix represents Ethernet cards, **arc** specifies an ARCnet card and **tr** specifies Token Ring

cards. A complete list of network cards and their respective modules can be found in Appendix A, *General Parameters and Modules*. For the most up-to-date list, please see our website at:

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

The netmask information may be set by default, although depending on what kind of network you are setting up, or becoming a part of, you may need to specify this. If you are connecting to an ISP, ask them for the information. Most likely it will be 255.255.255.0.

Required fields for manual configuration:

- **Primary name + domain** — The primary name is the name of your computer, while the domain is how your network is specified. For example, `foo.bar.com`; `foo` is the primary name and `bar.com` is the domain.
- **IP address** — The address of the machine and will follow the pattern of *x.x .x.x*. For example, 192.168.0.13.
- **Net device** — The type of network card you are using; eth0 would be the appropriate entry to use the first Ethernet card.
- **Kernel module** — The correct module based on your network card.

Information on net devices and kernel modules is described above. The appropriate primary name + domain and IP address will depend on whether you are adding the computer to an existing network or creating a new network. For connecting to an existing network, contact your network administrator for the information. Getting a network connected to the Internet is beyond the scope of this book, and we recommend the following starting point:

TCP/IP Network Administration, 2nd Edition, by Craig Hunt (O'Reilly and Associates).

If you're setting up a private network that won't *ever* be connected to the Internet, then you can choose any primary name + domain name you would like and have several choices for IP addresses (See Table 3–1, *Addresses and Examples*).

Table 3–1 Addresses and Examples

Addresses available	Examples
10.0.0.0 - 10.255.255.255	10.5.12.14
172.16.0.0 - 172.31.255.255	172.16.9.1, 172.28.2.5
192.168.0.0 - 192.168.255.25	192.168.0.13

The three sets of numbers above correspond to class a, b, and c networks respectively. The classes are used to describe the number of IP addresses available as well as the range of numbers. The numbers above have been set aside for private networks.

Please Note

You should not use these IP addresses if you connect to the Internet since 192.168.0.* and 192.168.255.* are not reliably considered private. If you want your network to be connected to the Internet, or think you might want to at some point in the future, do yourself a favor and get yourself non-private addresses now.

Name Server Specification

A name server and default domain are also needed to establish a network connection. The name server is used to translate host names such as `private.network.com` to their corresponding IP address such as `192.168.7.3`.

The default domain tells the computer where to look if a fully qualified hostname isn't specified. Fully qualified means that the full address is given, so `foo.redhat.com` is the fully qualified hostname, while the hostname is simply `foo`. If you specified your default domain as `redhat.com`, then you could use just the hostname to connect successfully. For example `ftp foo` would be sufficient if your search domain is `redhat.com`, while `ftp foo.redhat.com` would be required if it wasn't.

To specify the nameserver, open **Config** => **Networking** => **Client tasks** => **Name server specification (DNS).**

Figure 3–13 Resolver Configuration Screen

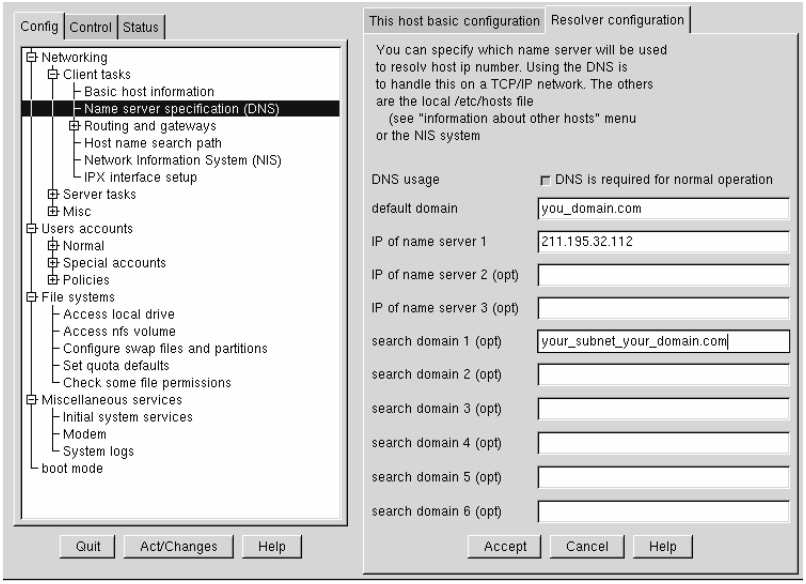

Nameservers are ranked according to the order in which they are accessed, so it's not unusual to see nameservers referred to as primary, secondary, tertiary and so on down the list if more than one is specified. Each of these must be an IP address and not a name, since the computer has no way to resolve the name until it connects to a nameserver.

In addition to a default domain, you can also specify search domains. Search domains work differently; they progress from one to six in a similar manner to the nameserver. However, they all take precedence over the default domain! Keep this in mind when specifying search domains. Search domains are not commonly used.

The one item not yet covered is the check box for DNS usage. If you are running a small private network with no Internet connection, then using `/etc/hosts` files and keeping them all synchronized will work. As you add more and more machines, the complexity increases until it is easier to have a single machine run a DNS than to continue to sync `/etc/hosts` files.

Another reason for not using DNS is if your network is going to use NIS instead. Note that NIS can be used in conjunction with DNS. So to sum it all up, unless you know why using `/etc/hosts` or NIS would be best for your situation, DNS is probably going to be your best choice.

You can add, modify, or delete entries from the `/etc/hosts` file using linuxconf. Open **Config** => **Networking** => **Misc** => **Information about other hosts.**

Figure 3–14 `/etc/hosts` Screen

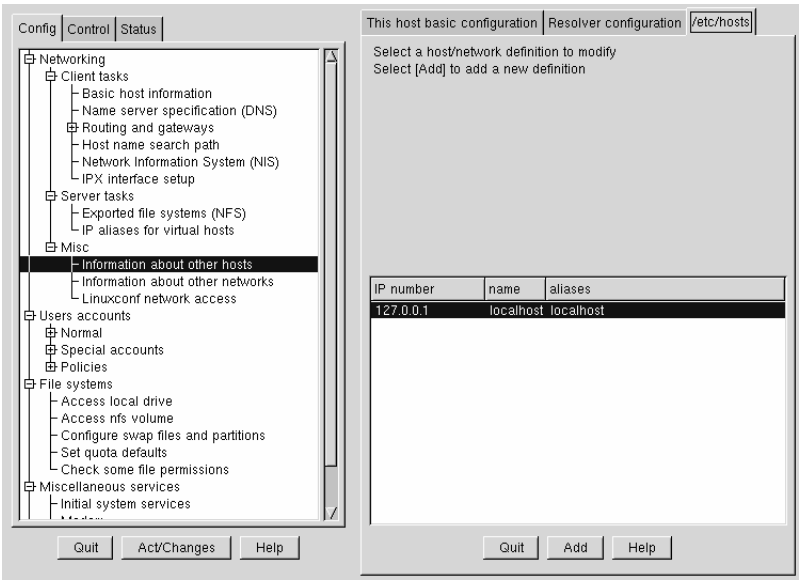

To modify or delete an entry select it. To delete the entry, select **Del** at the bottom of the **host/network definition** screen.

Figure 3–15 Host/Network Definition Screen

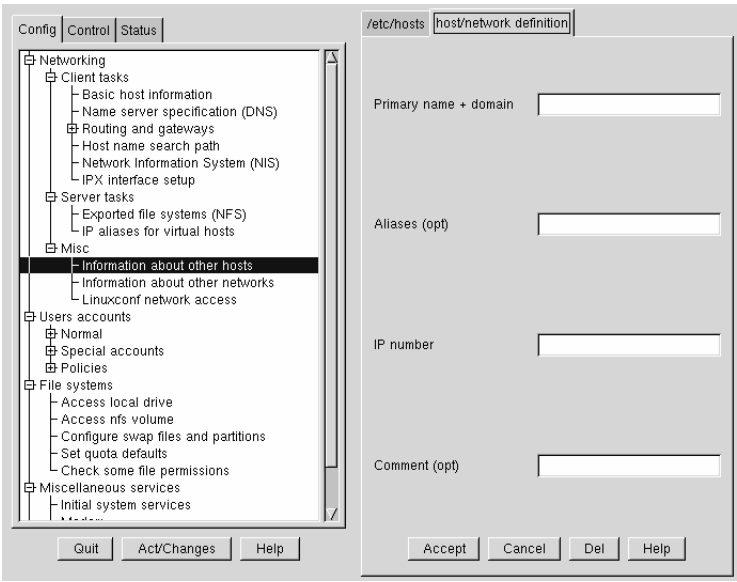

To modify it, change the information as necessary. To add a new entry, select Add at the bottom of the **/etc/hosts** screen. This will also open the **host/network definition** screen.

Required Fields:

- **Primary name + domain** — The primary name is the name of the computer, while the domain is how the network it is attached to is specified. For example, given `foo.bar.com`, `foo` is the primary name and `bar.com` is the domain.
- **IP number** — Also referred to as IP address; this is the address of the machine and will follow the pattern of *x.x. x.x.* For example, 192.168.0.13.

Optional Fields:

- **Aliases** — A shorthand for the fully qualified domain name. This is often the same as the primary name. So, for example, if the fully qualified domain name is foo.bar.com, you could select foo as the alias.

- **Comment** — A comment on the machine. For example, "The remote nameserver."

Once finished, select `Accept`.

3.1.15 Finding Your Way Through linuxconf

This table provides a quick reference for this chapter. Unfortunately, it doesn't provide a complete quick reference for linuxconf, which has many more capabilities than this documentation provides.

Table 3–2 Linuxconf Quick Reference

What do you want to do?	Where to find it in linuxconf
Add/modify/disable/delete a user account	**Config** => **Users accounts** => **Normal** => **User accounts**
Change a user's password	**Config** => **Users accounts** => **Normal** => **User accounts**
Change the root password	**Config** => **Users accounts** => **Normal** => **Change root password**
Configure networking	**Config** => **Networking** => **Client tasks** => **Basic host information**
Create/delete a group	**Config** => **Users accounts** => **Normal** => **Group definitions**
Edit parameters for passwords	**Users Accounts** => **Password & Account Policies**
Disable tree menu	**Control** => **Control files and systems** => **Configure linuxconf modules**
Enable Web-based access to linuxconf	**Config** => **Networking** => **Misc** => **Linuxconf network access**
Modify `/etc/hosts`	**Config** => **Networking** => **Misc** => **Information about other hosts**

What do you want to do?	Where to find it in linuxconf
Modify group membership	**Config** => **Users accounts** => **Normal** => **Group definitions** or **Config** => **Users accounts** => **Normal** => **User accounts**
Set filter parameters	**Control** => **Features**
Specify a nameserver (DNS)	**Config** => **Networking** => **Client tasks** => **Name server specification (DNS)**
View filesystem	**Config** => **File systems** => **Access local drive** or **Config** => **File systems** => **Access nfs volume**

3.2 System Configuration with the Control Panel

Please Note

Most of what can be done with the control panel applications can also be done using linuxconf. In addition, linuxconf supports both character-cell *and* graphical user interfaces. Please refer to Section 3.1, *System Configuration with linuxconf* for an introduction to linuxconf.

The control panel is a launching pad for a number of different system administration tools (see Figure 3–16, *The Control Panel*). These tools make your life easier by letting you configure things without remembering configuration file formats and awkward command line options.

Figure 3–16 The Control Panel

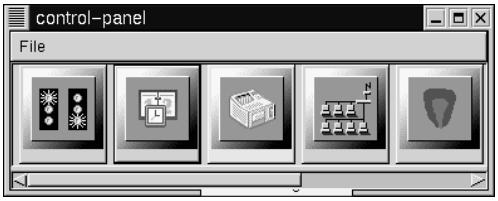

To start the control-panel, start the X Window System as root with `startx` and type `control-panel` in an Xterm. You will need to be root to run the control-panel tools successfully. You can do this as well if you already have X running as a normal user. Just type `su -c control-panel` and then type the root password when prompted. If you plan to do other tasks as root, you could type `su` followed by the root password when prompted.

Please Note

If you are not running X as root, you may need to give root access to your system's X server. To do this, enter the following command on a *non-root* terminal window:

```
xhost +localhost
```

After starting the control panel, simply clicking on an icon starts up a tool. Please note that you are not prevented from starting two instances of any tool, but doing so is a very bad idea because you may try to edit the same files in two places and end up overwriting your own changes.

Please Note

If you do accidentally start a second copy of a tool, you should quit it immediately. Also, do not manually edit any files managed by the control-panel tools while the tools are running. Similarly, do not run any other programs (such as linuxconf) that may change those files while the tools are running.

3.2.1 Printer Configuration

Please note that the *Official Red Hat Linux Getting Started Guide* contains more up-to-date documentation on printtool, so be sure to check that document before using printtool.

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

- plain text (ASCII) files
- PostScript files
- TeX `.dvi` files
- GIF, JPEG, TIFF, and other graphics formats
- RPMs

In other words, simply printing a GIF or RPM file using the `lpr` command will result in the printer doing "the right thing."

Figure 3–17 Print Tool

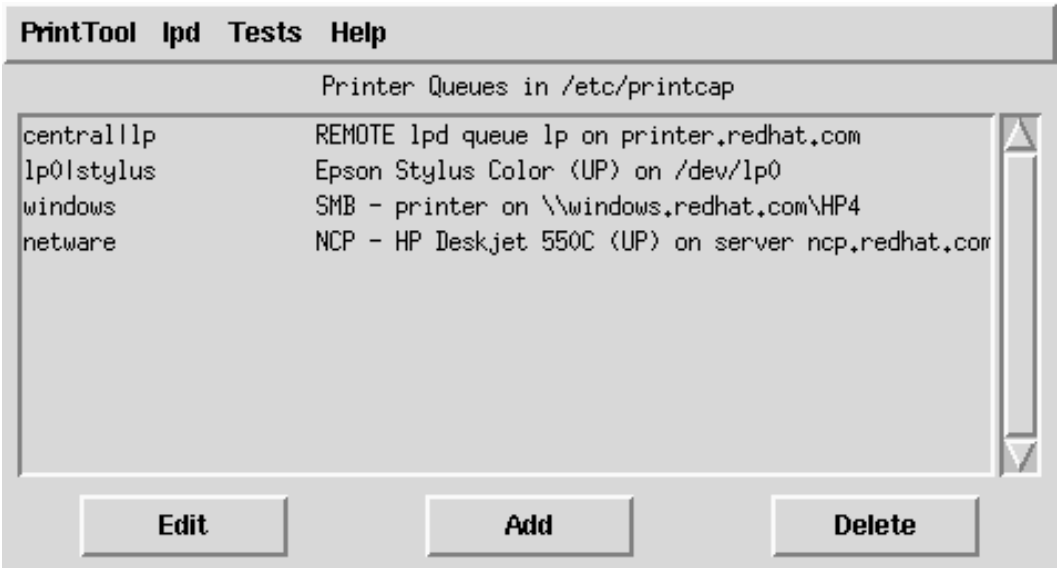

In order to create a new **print queue**, choose Add. Then, select what type of printer is being added. There are four types of print queues which can be configured with printtool:

- Local print queues are for printers attached to a printer or serial port on your Red Hat Linux system
- Remote print queues are attached to a different system which you can access over a TCP/IP network
- SMB print queues are attached to a different system which uses LAN-Manager-type (SMB) networking
- NCP print queues are attached to a different system which uses Novell's NetWare network technology

Figure 3–18 Selecting a Printer Type

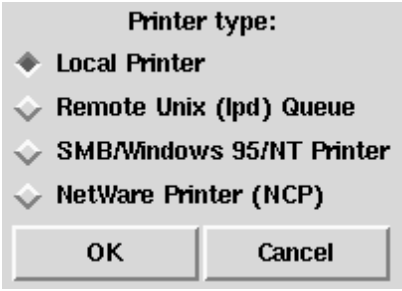

After choosing the printer type, a dialog box requests further information about the print queue (see Figure 3–19, *Adding a Local Printer*). All types of print queues require the following information:

- Queue Name — What the queue will be called. Multiple names can be specifed with the | (pipe) character separating entries.
- Spool Directory — This is the directory on the local machine where files are stored before printing occurs. Be careful to not have more than one printer queue use a given spool directory.
- File Limit — Maximum size print job accepted, in kilobytes (1 kb = 1024 bytes). A size of 0 indicates no limit should be imposed.
- Input Filter — Filters convert printed files into a format the printer can handle. Press `Select` to choose the filter which best matches your printer (see Figure 3–20, *Configuring a Print Filter*).

  In addition to configuring print queues able to print graphical and PostScript output, you can configure a **text-only** printer, which will only print plain ASCII text. Most printer drivers are also able to print ASCII text without converting it to PostScript first; simply choose `Fast text printing` when you configure the filter.

Please Note

This only works for non-PostScript printers.

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

For **local** printers, the following information is also required:

- Printer Device — Usually `/dev/lp1`; the name of the port which the printer is attached to. Serial printers are usually on `/dev/ttyS?` ports. Note that you will need to manually configure serial parameters.

Figure 3–19 Adding a Local Printer

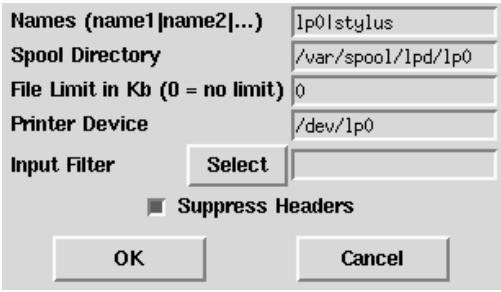

Figure 3–20 Configuring a Print Filter

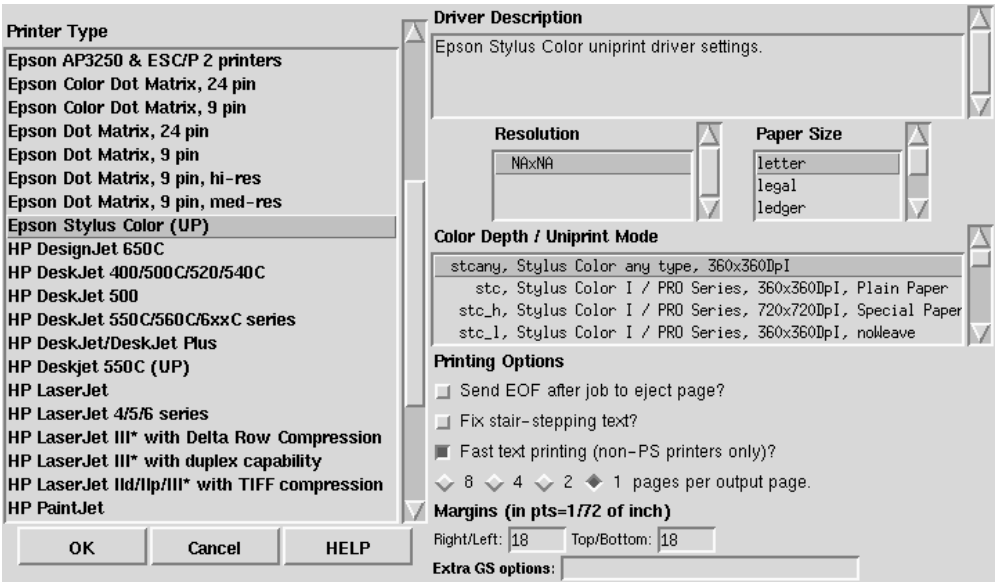

For **remote** printers, the dialog box contains additional fields; fill in the following information:

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

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

Figure 3–21 Adding a Remote Printer

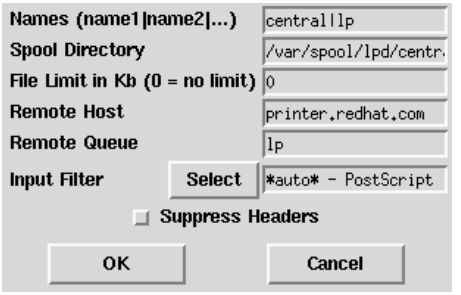

Figure 3–22 Adding an NCP Printer

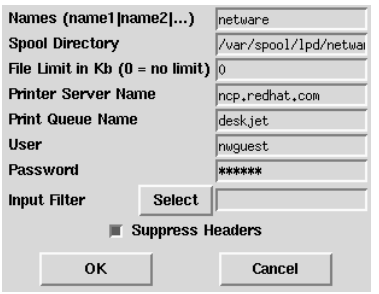

Figure 3–23 Adding an SMB Printer

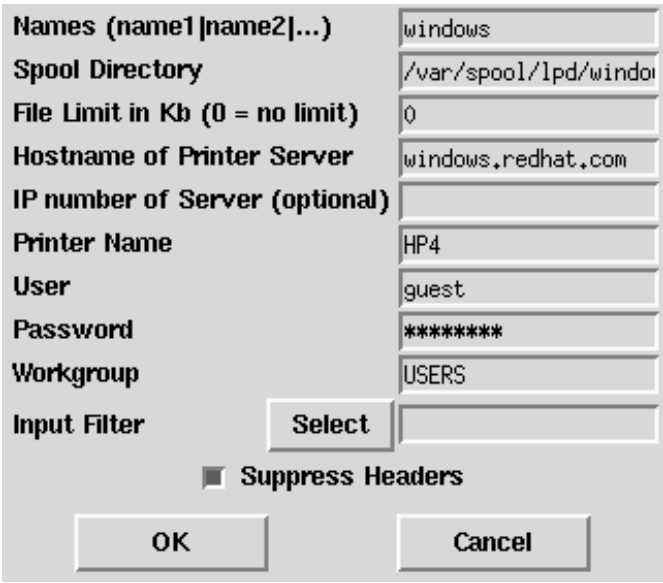

For SMB and NCP printers, fill in the following information:

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

Please Note

If you require a username and password for an SMB (LAN Manager) or NCP (NetWare) print queue, they are stored unencrypted in a local script. Thus, it is possible for another person to learn the username and password. It is therefore recommended that the username and password for use of the printer to be different than that for a user account on the local Red Hat Linux system, so that the only possible security compromise would be unauthorized use of the printer. If there are file shares from the SMB server, it is recommended that they also use a different password than the one for the print queue.

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

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

Figure 3–24 Printing a Test Page

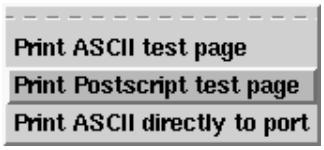

3.2.2 Loading Kernel Modules

The Linux kernel has a modular design. At boot time, only a minimal **resident** kernel is loaded into memory. Thereafter, whenever a user requests a feature that is not present in the resident kernel, a kernel **module** is dynamically loaded into memory. After a specified period of inactivity, the module may be removed from memory. This design promotes leanness and efficiency.

The mechanism that supports dynamic loading of modules is a kernel thread called `kmod`. When the kernel requests a module, kmod wakes up and calls `modprobe`(8) to get it.

When you install Red Hat Linux, the hardware on your system is probed and you provide information about how the system will be typically used and which programs should be loaded. Based on this probing and the provided usage information, the installation program decides which features to compile into the resident kernel and which to put in loadable modules, and sets up the dynamic loading mechanism to work transparently. But this is a highly configurable procedure. If you build your own custom kernel, you can make all of these decisions for yourself.

If you add new hardware after installation requiring support provided in a kernel module, you need to set up the dynamic loading mechanism. You do this by editing the module configuration file, `/etc/conf.modules`.

For example, if at the time you installed Red Hat Linux your system included a model SMC EtherPower 10 PCI network adapter, the module configuration file will contain this line: `alias eth0 tulip`. If, after installation, you install a second identical network adapter to your system, add this line to `/etc/conf.modules`:`alias eth1 tulip`.

See Appendix A, *General Parameters and Modules* for an alphabetical list of kernel modules and the hardware the modules support.

3.2.3 Network Configuration

Please Note

Documentation on network configuration using linuxconf can be found in Section 3.1.14, *Getting Connected with Linuxconf (Network Configuration)*.

The network configuration tool (`netcfg`) shown in Figure 3–25, *Network Configuration Panel* is designed to allow easy manipulation of parameters such as IP address, gateway address, and network address, as well as name servers and `/etc/hosts`.

Figure 3–25 Network Configuration Panel

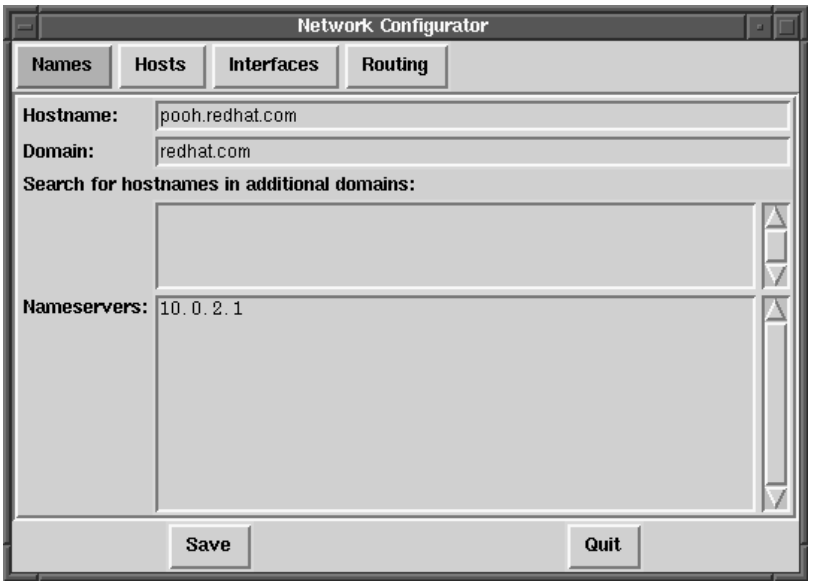

Network devices can be added, removed, configured, activated, deactivated and aliased. Ethernet, arcnet, token ring, pocket (ATP), SLIP, PLIP and loopback devices are supported. SLIP/PLIP support works well on most hardware, but some hardware setups may exhibit unpredictable behavior. When using the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

Managing Names

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

Figure 3–26 Adding/Editing Hosts

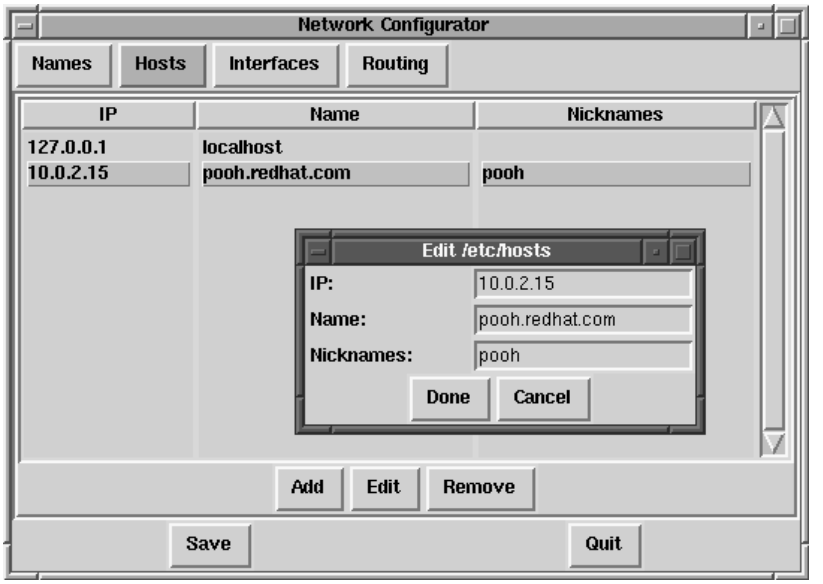

Managing Hosts

In the `Hosts` management panel you have the ability to add, edit, or remove hosts from the `/etc/hosts` file. Adding or editing an entry involves identical actions. An edit dialog box will appear, simply type the new information and click `Done` when you are finished. See Figure 3–26, *Adding/Editing Hosts* for an example.

Adding a Networking Interface

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

Please Note

You may need to configure `kerneld` to load a driver for the network interface you are adding (e.g., `eth0`); see Section 3.2.2, *Loading Kernel Modules* for more information.

Begin adding an interface by clicking on `Interfaces` in the main panel. This will bring up a window of configured devices with a row of available options, see Figure 3–27, *Configured Interfaces*.

Figure 3–27 Configured Interfaces

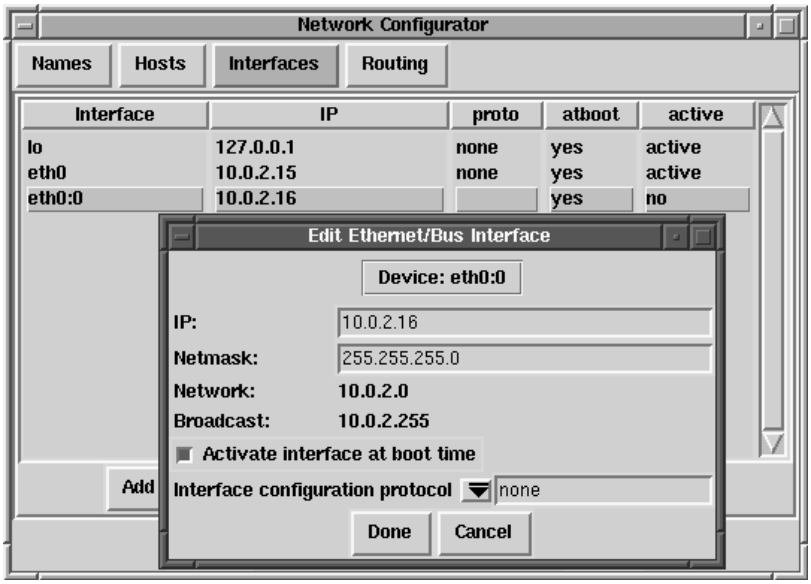

To add a device, first click the `Add` button then select the type of interface you want to configure from the box that appears.

Please Note

There is now a `clone` button available in `netcfg`. This button can be used to create a "clone" of an already-existing interface. By using clone interfaces, it is possible for a laptop to have one Ethernet interface defined for a work LAN, and a clone Ethernet device defined for a home LAN.

SLIP Interface

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

PLIP Interface

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

Ethernet, Arcnet, Token Ring and Pocket Adaptor Interfaces

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

- Device — This is determined by netconfig based on the devices already configured.
- IP Address — Enter an IP address for your network device.
- Netmask — Enter the network mask for your network device.

  The network and broadcast addresses are calculated automatically based on the IP address and netmask you enter.
- Activate interface at boot time:

If you want the device to be configured automatically when your machine boots select this by clicking on the box.

- Allow any user to (de)activate interface:

  Check this if you want any user to be able to activate or deactivate the interface.

- Interface configuration protocol:

  If you have a BOOTP or DHCP server on your network and would like to use it to configure the interface, choose the appropriate option; otherwise, choose `none`.

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

Managing Routes

In the Routes management screen you have the ability to add, edit, or remove static networking routes. Adding or editing an entry involves identical actions, just like the Hosts panel. An edit dialog box will appear; simply type the new information and click `Done` when you are finished. See Figure 3–28, *Adding/Editing Routes* for an example.

Figure 3–28 Adding/Editing Routes

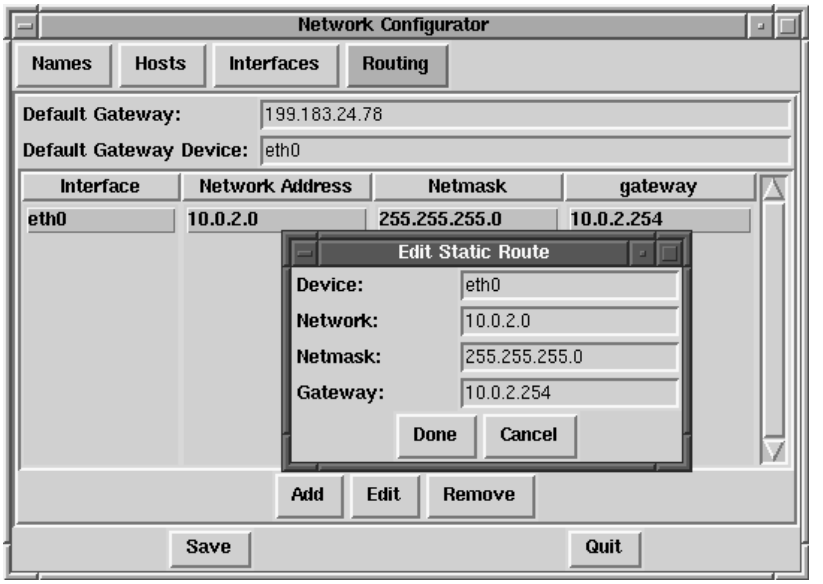

3.2.4 Time and Date

The "time machine" allows you to change the time and date by clicking on the appropriate part of the time and date display and clicking on the arrows to change the value.

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

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

Please Note

Changing the time can seriously confuse programs that depend on the normal progression of time, and could possibly cause problems. Try to quit as many applications and processes as possible before changing the time or date.

---

4 PowerTools

4.1 PowerTools Packages

Red Hat PowerTools is a collection of software packages built for the Red Hat Linux 7.0 operating system. PowerTools includes the latest versions (as of this product's release date) of hundreds of programs -- so finding an interesting application should be easy.

Among the many applications are audio programs, chat clients, development tools, editors, file managers, emulators, games, graphics programs, productivity applications, math/statistics packages, systems administration and network management tools, and window managers.

Now that you know what you can find on PowerTools, you may want to know how to install them. See Section 4.1.2, *Installing PowerTools Packages* for installation information.

4.1.1 Reading the Contents of the CD-ROM

You can read the contents of the PowerTools CD-ROM from a shell prompt (either in a terminal window or in console mode). Then mount the CD-ROM (`mount /mnt/cdrom`). Once this is done, change directories by issuing the `cd /mnt/cdrom` command. Finally, type `less CONTENTS` to view the available applications.

4.1.2 Installing PowerTools Packages

Installing PowerTools in a GUI Environment

If you're using GNOME or KDE, place the CD-ROM in your CD-ROM drive. You'll be prompted for the root password (you must be root in order to install packages). After you type in the root password, either the Gnome-RPM or the Kpackage package management program will start automatically (depending on you GUI environment) and can be used to install PowerTools.

See Chapter 6, *Gnome-RPM* for specific instructions on how to use Gnome-RPM. See http://www.general.uwa.edu.au/u/toivo/kpackage/ for more information on how to use Kpackage.

If you're not using GNOME or KDE, you'll need to use the shell prompt to install PowerTools. See *Installing PowerTools from the Shell Prompt* in Section 4.1.2 for more information.

Installing PowerTools from the Shell Prompt

First, mount the PowerTools CD-ROM on your CD-ROM drive.

Place the PowerTools CD in your CD-ROM drive. As root, type the following:

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

Please note

On your system, you or the system administrator may already allow users (instead of only root) to mount the CD-ROM drive. Users have this privilege if the `user` option is included in the `/dev/cdrom` line in the `/etc/fstab` file. However, keep in mind that you must be logged in as root to install any PowerTools RPMs.

After you've mounted the drive, `cd` to the mounted CD-ROM directory with the following command:

```
# cd /mnt/cdrom
```

When you list the contents of the CD with `ls`, you'll see the following directories: `SRPMS` and `RedHat/` The `SRPMS` directory contains the PowerTools source RPMs. The `RedHat/RPMS` directory contains the RPMs for the three specified operating system architectures.

The `RedHat/RPMS` path is used as a general example. You should substitute the correct directory for `RedHat/RPMS`, depending upon your architecture and which package you're installing.

`cd` to the `RedHat/RPMS` directory:

```
# cd RedHat/RPMS
#
```

List the RPM files in the directory with `ls` to see the complete list of RPM packages included for Intel-compatible systems.

You will probably want more information about a specific package before you can decide whether you want to install it. You can use RPM's querying capability to find out more information about the packages, such as the packages' functions and origination. See Section 5.3, *Impressing Your Friends with RPM* for instructions on how to query packages using RPM.

You can install your selected packages with RPM. RPM is a powerful command line-driven package management system. See Chapter 5, *Package Management with RPM* for more information on how to use RPM to install and manage PowerTools packages.

Once you have finished installing your packages, you'll need to unmount your CD-ROM. First, change directories so that you will be one level above the `/mnt/cdrom/` directory by issuing the command `cd ..`. Then, type `umount /mnt/cdrom` to unmount the CD-ROM. Type `eject /mnt/cdrom` and the CD-ROM drive will open so that you can remove the CD.

5 Package Management with RPM

The *R*PM *P*ackage *M*anager (RPM), is an open packaging system available for anyone to use, and works on Red Hat Linux as well as other Linux and UNIX systems. Red Hat, Inc. encourages other vendors to take the time to look at RPM and use it for their own products. RPM is distributable under the terms of the GPL.

For the end user, RPM provides many features that make maintaining a system far easier than it has ever been. Installing, uninstalling, and upgrading RPM packages are all one line commands, and all the messy details have been taken care of for you. RPM maintains a database of installed packages and their files, which allows you to perform powerful queries and verification of your system.

During upgrades, RPM handles configuration files specially, so that you never lose your customizations -- a feature that is impossible with straight `.tar.gz` files.

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

Please Note

Although it can be important to understand the concepts behind RPM, for those who prefer a graphical interface to the command line, we suggest you use Gnome-RPM. Please see Chapter 6, *Gnome-RPM* for more information.

5.1 RPM Design Goals

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

Upgradability

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

Powerful Querying

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

System Verification

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

Pristine Sources

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

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

5.2 Using RPM

RPM has five basic modes of operation (not counting package building): installing, uninstalling, upgrading, querying, and verifying. This section contains an overview of each mode. For complete details and options try `rpm --help`, or turn to Section 5.4, *Other RPM Resources* for more information on RPM.

5.2.1 Installing

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

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

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

Please Note

Although using the command **`rpm -ivh foo-1.0-1.i386.rpm`** to install is common, you may want to consider using **`rpm -Uvh foo-1.0-1.i386.rpm`** instead. `-U` is commonly used for upgrading a package, but it will also install new packages. See Section 5.2.3, *Upgrading* for more information about using the `-U` RPM command.

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

Package Already Installed

If the package is already installed, you will see:

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

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

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

Conflicting Files

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

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

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

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

Unresolved Dependency

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

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

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

5.2.2 Uninstalling

Uninstalling a package is just as simple as installing:

```
# rpm -e foo
#
```

Please Note

Notice that we used the package *name* "`foo`," not the name of the original package *file* "`foo-1.0-1.i386.rpm`". To uninstall a package, you will need to replace `foo` with the actual filename of the original package.

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

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

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

5.2.3 Upgrading

Upgrading a package is similar to installing.

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

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

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

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

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

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

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

To cause RPM to "upgrade" anyway, use `--oldpackage` on the command line:

```
# rpm -Uvh --oldpackage foo-1.0-1.i386.rpm
foo                         ##################################
#
```

5.2.4 Freshening

Freshening a package is similar to upgrading:

```
# rpm -Fvh foo-1.2-1.i386.rpm
foo                         ##################################
#
```

RPM's freshen option checks the versions of the packages specified on the command line against the versions of packages that have already been installed on your system. When a newer version of an already-installed package is processed by RPM's freshen option, it will be upgraded to the newer version. However, RPM's freshen option will not install a package if no previously-installed package of the same name exists. This differs from RPM's upgrade option, as an upgrade *will* install packages, whether or not an older version of the package was already installed.

RPM's freshen option works well with single packages or with a group of packages. It's especially handy if you've just downloaded a large number of different packages, and you only want to upgrade those packages that are already installed on your system. Using the freshen option means that you won't have to pick through the downloaded packages, deleting any unwanted ones before using RPM.

In this case, you can simply issue the following command:

```
# rpm -Fvh *.rpm
```

RPM will automatically upgrade only those packages that have already been installed.

5.2.5 Querying

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

```
# rpm -q foo
foo-2.0-1
#
```

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

- `-a` queries all currently installed packages.
- `-f <file>` will query the package owning *<file>*.
- `-p <packagefile>` queries the package *<packagefile>*.

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

- `-i` displays package information including name, description, release, size, build date, install date, vendor, and other miscellaneous information.
- `-l` displays the list of files that the package contains.
- `-s` displays the state of all the files in the package.
- `-d` displays a list of files marked as documentation (man pages, info pages, README's, etc).
- `-c` displays a list of files marked as configuration files. These are the files you change after installation to adapt the package to your system (`sendmail.cf`, `passwd`, `inittab`, etc).

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

5.2.6 Verifying

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

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

- To verify a package containing particular file:

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

- To verify ALL installed packages:

```
rpm -Va
```

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

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

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

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

- 5 -- MD5 checksum
- S -- File size
- L -- Symbolic link
- T -- File modification time
- D -- Device
- U -- User
- G -- Group
- M -- Mode (includes permissions and file type)
- ? -- Unreadable file

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

5.3 Impressing Your Friends with RPM

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

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

```
rpm -Va
```

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

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

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

  The output would look like:

```
gv-3.5.8-10
```

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

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

  and the appropriate package will be verified.

- Want to find more about a particular program? You can find out by entering the following to locate the documentation which came with the package that "owns" that program (in this case ispell):

```
rpm -qdf /usr/bin/md5sum
```

  The output would be:

```
/usr/share/doc/textutils-2.0a/NEWS
/usr/share/doc/textutils-2.0a/README
/usr/info/textutils.info.gz
/usr/man/man1/cat.1.gz
/usr/man/man1/cksum.1.gz
/usr/man/man1/comm.1.gz
```

```
/usr/man/man1/csplit.1.gz
/usr/man/man1/cut.1.gz
/usr/man/man1/expand.1.gz
/usr/man/man1/fmt.1.gz
/usr/man/man1/fold.1.gz
/usr/man/man1/head.1.gz
/usr/man/man1/join.1.gz
/usr/man/man1/md5sum.1.gz
/usr/man/man1/nl.1.gz
/usr/man/man1/od.1.gz
/usr/man/man1/paste.1.gz
/usr/man/man1/pr.1.gz
/usr/man/man1/ptx.1.gz
/usr/man/man1/sort.1.gz
/usr/man/man1/split.1.gz
/usr/man/man1/sum.1.gz
/usr/man/man1/tac.1.gz
/usr/man/man1/tail.1.gz
/usr/man/man1/tr.1.gz
/usr/man/man1/tsort.1.gz
/usr/man/man1/unexpand.1.gz
/usr/man/man1/uniq.1.gz
/usr/man/man1/wc.1.gz
```

- Let's say you find a new sndconfig RPM, but you don't know what it is. To find out some information on it, enter:

```
rpm -qip sndconfig-0.48-1.i386.rpm
```

The output would be:

```
Name        : sndconfig                Relocations: (not relocateable)
Version     : 0.48                          Vendor: Red Hat
Release     : 1                         Build Date: Mon 10 Jul 2000 02:25:40
Install date: (none)                    Build Host: porky.devel.redhat.com
Group       : Applications/Multimedia   Source RPM: sndconfig-0.48-1.src.rpm
Size        : 461734                       License: GPL
Packager    : Red Hat <http://bugzilla.redhat.com/bugzilla>
Summary     : The Red Hat Linux sound configuration tool.
Description :
Sndconfig is a text based tool which sets up the configuration files
you'll need to use a sound card with a Red Hat Linux system.
Sndconfig can be used to set the proper sound type for programs which
use the /dev/dsp, /dev/audio and /dev/mixer devices.  The sound
```

```
settings are saved by the aumix and sysV runlevel scripts.
```

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

```
rpm -qlp sndconfig-0.48-1.i386.rpm
```

The output is:

```
/usr/sbin/pnpprobe
/usr/sbin/sndconfig
/usr/share/locale/cs/LC_MESSAGES/sndconfig.mo
/usr/share/locale/da/LC_MESSAGES/sndconfig.mo
/usr/share/locale/de/LC_MESSAGES/sndconfig.mo
/usr/share/locale/es/LC_MESSAGES/sndconfig.mo
/usr/share/locale/fr/LC_MESSAGES/sndconfig.mo
/usr/share/locale/hu/LC_MESSAGES/sndconfig.mo
/usr/share/locale/id/LC_MESSAGES/sndconfig.mo
/usr/share/locale/is/LC_MESSAGES/sndconfig.mo
/usr/share/locale/it/LC_MESSAGES/sndconfig.mo
/usr/share/locale/ko/LC_MESSAGES/sndconfig.mo
/usr/share/locale/no/LC_MESSAGES/sndconfig.mo
/usr/share/locale/pt/LC_MESSAGES/sndconfig.mo
/usr/share/locale/pt_BR/LC_MESSAGES/sndconfig.mo
/usr/share/locale/ro/LC_MESSAGES/sndconfig.mo
/usr/share/locale/ru/LC_MESSAGES/sndconfig.mo
/usr/share/locale/sk/LC_MESSAGES/sndconfig.mo
/usr/share/locale/sl/LC_MESSAGES/sndconfig.mo
/usr/share/locale/sr/LC_MESSAGES/sndconfig.mo
/usr/share/locale/sv/LC_MESSAGES/sndconfig.mo
/usr/share/locale/tr/LC_MESSAGES/sndconfig.mo
/usr/share/locale/uk/LC_MESSAGES/sndconfig.mo
/usr/share/man/man8/pnpprobe.8.gz
/usr/share/man/man8/sndconfig.8.gz
/usr/share/sndconfig/sample.au
/usr/share/sndconfig/sample.midi
```

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

5.4 Other RPM Resources

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

http://www.rpm.org/

There is also an RPM book available. It's called *Maximum RPM*, and it is available from Red Hat, at your local bookstore and through online booksellers. It contains a wealth of information about RPM for both the end-user and the package builder.

An on-line version of the book is available at http://www.rpm.org/.

There is also a mailing list for discussion of RPM-related issues, called `rpm-list@redhat.com`.

The list is archived on http://www.redhat.com/support/mailing-lists/. To subscribe, send mail to `rpm-list-request@redhat.com` with the word `subscribe` in the subject line.

6 Gnome-RPM

One of the most convenient package manipulation tools available is Gnome-RPM, a graphical tool which runs under the X Window System. Gnome-RPM was written by James Henstridge (`james@daa.com.au`); RPM 3.0 support was written by Red Hat and additional rpmfind code was written by Daniel Veillard.

Gnome-RPM (which is also referred to as gnorpm) allows the end-user to easily work with RPM technology; it is fast, powerful and features a friendly interface.

To learn more about RPM technology, turn to Chapter 5, *Package Management with RPM*.

Gnome-RPM is "GNOME-compliant," meaning that it seamlessly integrates into GNOME, the X Window System desktop environment.

With Gnome-RPM, you can easily

- install RPM packages
- uninstall RPM packages
- upgrade RPM packages
- find new RPM packages
- query RPM packages
- verify RPM packages

The interface features a menu, a toolbar, a tree and a display window of currently installed packages.

Operations are often performed in Gnome-RPM by finding and selecting packages, then choosing the type of operation to perform via push-button on the toolbar, through the menu or by right-clicking with the mouse.

- Installing a package places all of the components of that package on your system in the correct locations.
- Uninstalling a package removes all traces of the package except for configuration files you have modified.

- Upgrading a package installs the newly available version and uninstalls all other versions that were previously installed. This allows quick upgrading to the latest releases of packages. Refer to Section 6.4, *Configuration* for information about how to alter the default settings for installing and uninstalling packages.

You can also use the **Web find** option to search the Internet for newly released packages. You can direct Gnome-RPM to search for particular distributions when you want to look for new packages. (If you have a slow connection, this option can take some time to fully execute.) See Section 6.4, *Configuration* for more information about this feature.

Please Note

Exercise caution if you choose to use the **Web find** option, since there is no way to verify the integrity of the many packages which are available at numerous repositories. Before installing packages, you should perform a query on that package to help you determine whether it can be trusted. Packages not produced by Red Hat are not supported in any way by Red Hat.

Using Gnome-RPM to perform all of these and many other operations is the same as using `rpm` from the shell prompt. However, the graphical nature of Gnome-RPM often makes these operations easier to perform.

The usual way to work with Gnome-RPM is to display the available packages, select the package(s) you want to operate on, and then select an option from the toolbar or menu which performs the operation. However, Gnome-RPM is flexible enough to display packages in a variety of views, thanks to the use of **filters.** Refer to Section 6.3, *Installing New Packages* for more information on using filters to identify packages.

You can install, upgrade or uninstall several packages with a few button clicks. Similarly, you can query and verify more than one package at a time. Because of Gnome-RPM's integration with GNOME, you can also perform installation, query and verification on packages from within the GNOME File Manager.

6.1 Starting Gnome-RPM

You can start Gnome-RPM from either an Xterm window or from the GNOME desktop Panel (**Main Menu Button** => **System** => **GnoRPM**).

To start Gnome-RPM from an Xterm window, at the shell prompt, simply type

```
gnorpm &
```

That will bring up the main Gnome-RPM window (as shown in Figure 6–1, *Main Gnome-RPM Window*).

Please Note

If you would like to install, upgrade or uninstall packages, you must be in root. The easiest way to do this is to type su to become root, and then type the root password at a shell prompt. However, it isn't necessary to be root in order to query and verify packages.

There are several parts to the Gnome-RPM interface.

- Package Panel - On the left; allows you to browse and select packages on your system.
- Display window - To the right of the package panel; shows you contents from folders in the panel.
- Toolbar - Above the display and panel; a graphical display of package tools.
- Menu - Above the toolbar; contains text-based commands, as well as help info, preferences and other settings.
- Status bar - Beneath the panel and display windows; shows the total number of selected packages.

Figure 6–1 Main Gnome-RPM Window

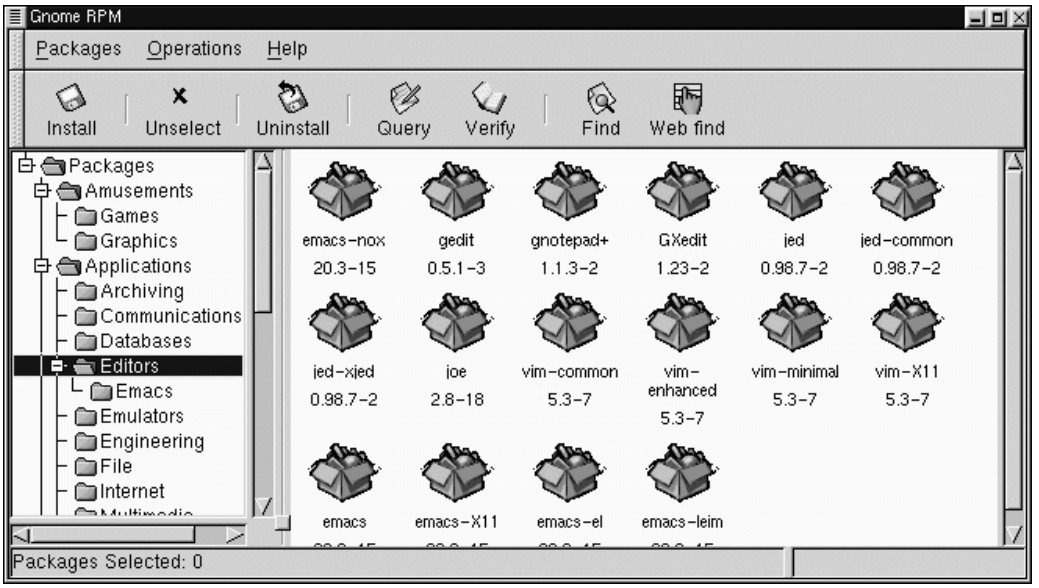

6.2 The Package Display

Each folder icon in the tree view at left represents a group of packages. Each group can contain subgroups. Groups are used to place packages that perform similar functions in similar locations. For example, the folder **Editors** contains text editors such as ed, vim and GXedit. From the tree view on the left, you might find another folder beneath **Editors** called **Emacs**, which would contain both emacs and emacs-X11.

The tree view is also arranged in an expandable and collapsible manner, which helps you to easily navigate through the packages. A folder which appears with a **+** next to it indicates that there are subfolders within that category.

To view the packages and subgroups within a group, click once on a folder or a **+** with your left mouse button. The display window will then show you the contents of that folder. By default, you will be presented with icons to represent the packages. You can change that view to a list view by selecting **View as list** from the **Interface** tab

you'll find under **Operations** => **Preferences**. Refer to Section 6.4, *Configuration* for more information about customizing the settings.

In this manner, you can move about the tree view, opening and expanding folders containing applications, games, tools and more. The contents of each folder will be displayed at the right.

6.2.1 Selecting Packages

To select a single package, click on it with the left mouse button. You'll notice that highlighting will appear around the package title (as shown in Figure 6–2, *Selecting Packages in Gnome-RPM*) which reflects the selection. To unselect it, either click on an empty space in the display panel with the left mouse button, or click on the **Unselect** button on the toolbar. When you unselect a package, the highlighting will disappear.

Figure 6–2 Selecting Packages in Gnome-RPM

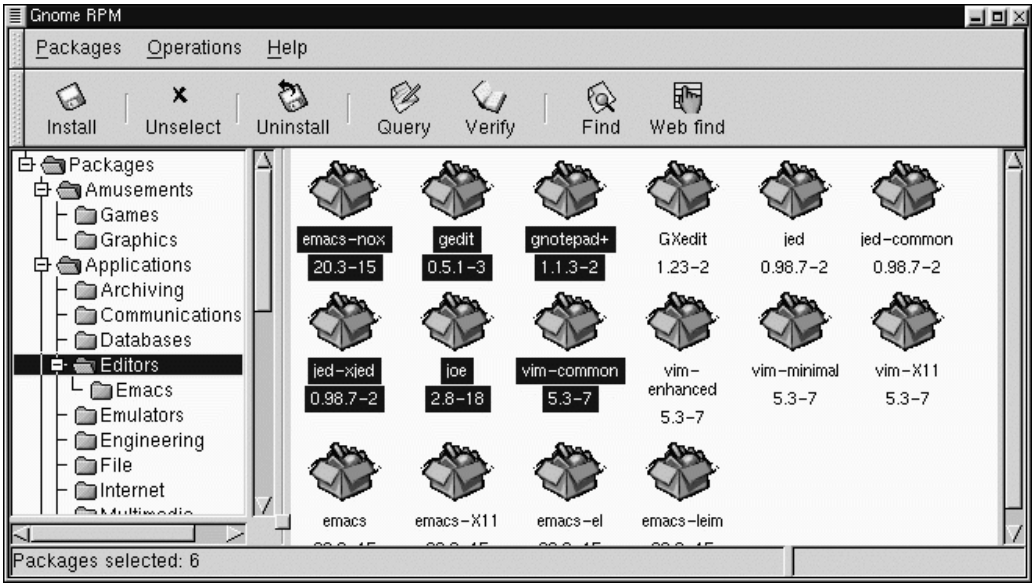

You can select and unselect multiple packages, in more than one folder in the tree panel. To select more than one package incrementally, left-click with your mouse

button, while holding down the [Ctrl] key; you'll see highlighting around each additional selection.

To select more than one package **globally**, that is, make larger selections within a folder, left-click one package, then, while holding down the [Shift] key, left-click on the final package you wish to select. By doing so, you'll notice that individual packages between your starting and ending selections will also be highlighted for selection. Using this option makes selecting groups of packages quicker than selecting each package individually.

The status bar at the bottom of Gnome-RPM will display the total number of packages you have selected.

6.3 Installing New Packages

To install new packages, choose **Install** from the toolbar. The **Install** window will open, revealing packages which are either already installed, newer packages or other options, which can be selected from the **Filter** drop-down bar at the top of the window. See Figure 6–4, *The Add Packages Window* for an example of the **Install** window, using the filter for "**All but installed packages.**"

Using the **Filter** feature allows you to winnow your choices for viewing packages. Available filters for viewing include:

- All packages
- All but installed packages
- Only uninstalled packages
- Only newer packages
- Uninstalled or newer packages

You can switch the display of packages by using the drop-down bar at the top of the window.

Figure 6–3 The Install Window

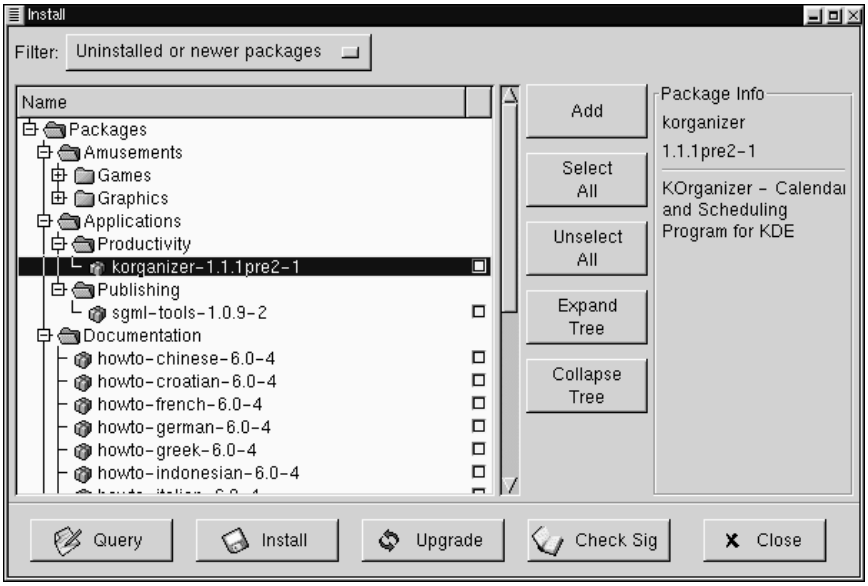

Choose the **Add** button. By default, if your CD-ROM is mounted with a Red Hat Linux CD-ROM, Gnome-RPM will search in `/mnt/cdrom/RedHat/RPMS` for new packages. (You can find this default option in the **Install Window** tab of the **Preferences** dialog. See Section 6.4, *Configuration* for more information on this feature.)

If no packages are available in the default path, you'll be presented with an **Add Packages** window from which you can select the appropriate location of your new package. In this view, you can select the correct path by using the drop-down bar at the top of the window to quickly navigate to pre-set locations; or you can double-click in the left panel of the **Add Packages** window to navigate to the correct path (as in Figure 6–4, *The Add Packages Window*). You can also type the path in the text window at the bottom of the **Add Packages** window.

Figure 6–4 The Add Packages Window

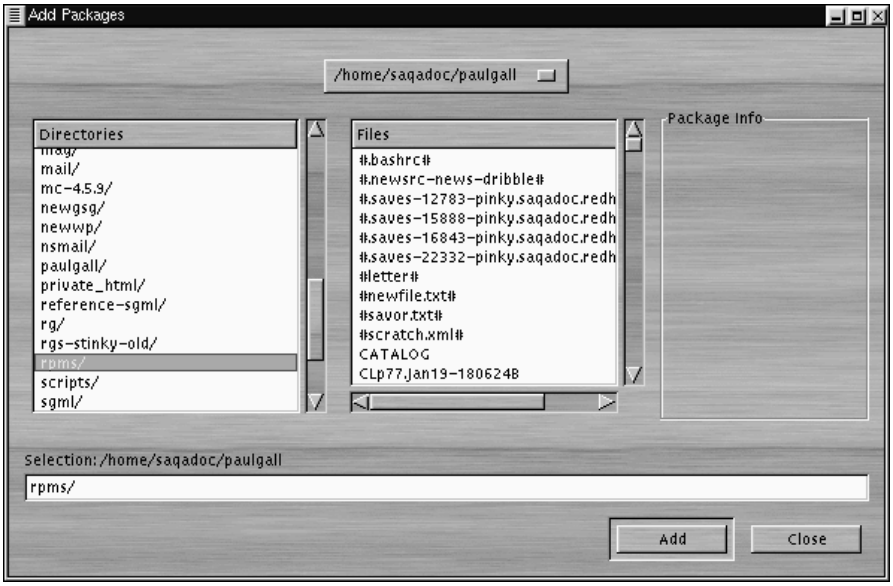

Please Note

If you download RPMs, you'll likely find those downloaded packages in a directory called `rpms` within your home directory. For example, if your home directory is `/home/bill`, your downloaded RPMs will be in `/home/bill/rpms`.

By clicking on the item, you'll find a brief description of the package in the **Package Info** panel of the **Install** window. To perform an installation or a query on the package, click inside the **checkbox** next to the package, then select the **Install** button. You can also query the item once it's checked; from within the **Package Info** window, you can also perform the installation (see Figure 6–10, *Query Window* in Section 6.5.1, *Querying Packages* for an example).

To choose an item, double-click on it with your left mouse button, or click on the **Add** button. The selected package(s) will be added to the **Install** window. You can also install more than one package in the same manner; each selection will be added to the **Install** window.

In addition to choosing to install the packages from within the **Install** window, you can install after performing a query on the selected package. Click on **Query**, which will open the **Package Info** window. Here, you can find a variety of details about the file(s) you've selected to install. Information will include the origination of the package, the date it was built, its size and more.

Within this **Package Info** window, you have the option of installing or upgrading packages.

If the package already exists on your system and you're querying a newer release, the **Package Info** window provides an **Upgrade** button, which will perform an upgrade to newer releases.

You can also "drag and drop" packages from GNOME File Manager. Within the File Manager, left-click on your selected RPM file then, while still holding down the mouse button, "drag" the file to the **Install** window and place it within the **Name** panel.

When dragging files to the **Install** window from the File Manager, you'll notice that the file appears as an icon while it's being dragged toward Gnome-RPM. Once inside the **Name** panel, you'll see that the package is checked for installation by default, and its information appears in the **Package Info** panel to the right.

To install the package now, just select the **Install** button.

You'll see a progress indicator when your package is being installed.

6.4 Configuration

Gnome-RPM offers a wide selection of choices for installing and uninstalling packages, documentation and other features. You can customize Gnome-RPM through the **Preferences** dialog, which you can access from **Operations** => **Preferences** on the menu. To make selections in the many of the **Preferences** dialogs, select the check boxes next to the options.

Under the **Behavior** tab, you'll find a number of options for configuring the way Gnome-RPM installs, uninstalls and upgrades packages. The Behavior tab is split into five sections: Install, Upgrade, Other, Database and Architecture. Note that by default these boxes are not checked (see Figure 6–5, *Behavior Tab in Preferences*).

Figure 6–5 Behavior Tab in Preferences

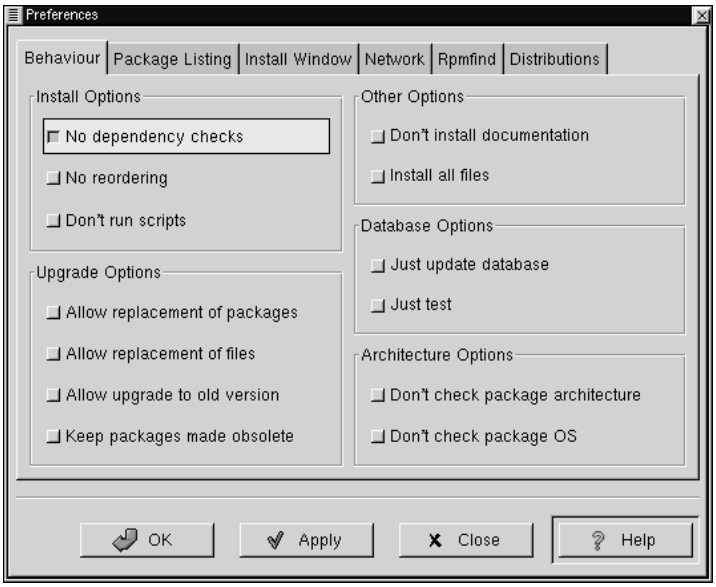

Under **Install Options**, you have the following choices:

- **No dependency checks** - When selected, this will install or upgrade a package without checking for other types of files on which the program may be dependent in order to work. However, unless you know what you're doing we strongly suggest you not use this option as some packages may depend on other packages for files, libraries or programs to function correctly.
- **No reordering** - This option is useful if RPM is unable to change the installation order of some packages to satisfy dependencies.

- **Don't run scripts** - Pre- and post-install scripts are sequences of commands that are sometimes included in packages to assist with installation. This check box is similar to the `--noscripts` option when installing from the shell prompt.

Under **Upgrade Options**, you can select the following:

- **Allow replacement of packages** - Replaces a package with a new copy of itself. Similar to the `--replacepkgs` option from the shell prompt. This option can be useful if an already-installed package has become damaged or may require other repair to function correctly.
- **Allow replacement of files -** Allows the replacement of files which are owned by another package. The shell prompt equivalent for this RPM option is `--replacefiles`. This option can sometimes be useful when there are two packages with the same file name but different contents.
- **Allow upgrade to old version** - Like the shell prompt RPM option equivalent `--oldpackage`, this option allows you to "upgrade" to an earlier package. It can sometimes be useful if the latest version of a package doesn't function correctly for your system.
- **Keep packages made obsolete** - Prevents packages listed in an Obsoletes header from being removed.

In **Other Options**, you can select:

- **Don't install documentation** - Like `--excludedocs`, this option can save on disk space by excluding documentation such as man pages or other information related to the package.
- **Install all files** - Installs all files in the package.

The choices available in **Database Options** and **Architecture Options** allow you to decide, among other things, whether you want to perform a "test" installation (which will check for file conflicts without actually performing an install), or whether you want to exclude packages for other operating systems or system architectures.

In the **Package Listing** tab, you'll find a choice of displays for your packages: either **View as icons**, which will be graphically-based, or **View as list**, which is not graphical but can provide more information about the packages.

In **Install Window,**, you can specify the path through which Gnome-RPM can find new RPMs on your system. Refer to Figure 6–6, *Install Window* for an example of this dialog. If you're using your Red Hat Linux CD-ROM, this path will probably be

```
/mnt/cdrom/RedHat/RPMS
```

or a similar path which had been set as the default path for Gnome-RPM. If you download new RPMs from the Internet or want to install RPMs via a NFS-mounted CD-ROM this path will be different for you.

Figure 6–6 Install Window

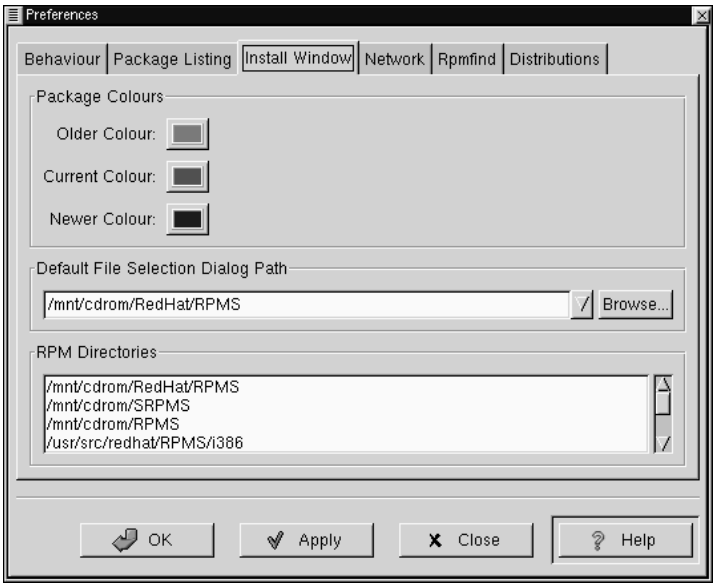

To change this path, type the full path to the RPMs you'd like to work with. Choosing the **Apply** or **OK** buttons will save this path, making it the default path for future sessions. You can also determine the default path by selecting the **Browse...** button, and visually navigating through the **RPMPath** window.

After changing install path and closing the dialog box, you can use the **Install** button to view the packages available in your new location.

(If the path for your RPMs doesn't match the default path in your preferences, you'll be presented with a browser window, which will allow you to select the correct path for your new RPMs.)

Under **Package Colors**, you'll find color coding for packages. The default setting for older packages is gray; for current packages, the color is green; for newer packages than those installed, the color is blue. These color values can be customized to suit your needs.

The **RPM Directories**, contains a listing of default locations where Gnome-RPM will search for packages.

In **Network**, you have the ability to specify proxies for use with HTTP and FTP transfers, as well as user and password names (see Figure 6–7, *Network Settings*). Note, however, that the password will not be stored securely.

In the **Cache expire** window, you can set the length of time before data from the rpmfind database is considered out of date.

Figure 6–7 Network Settings

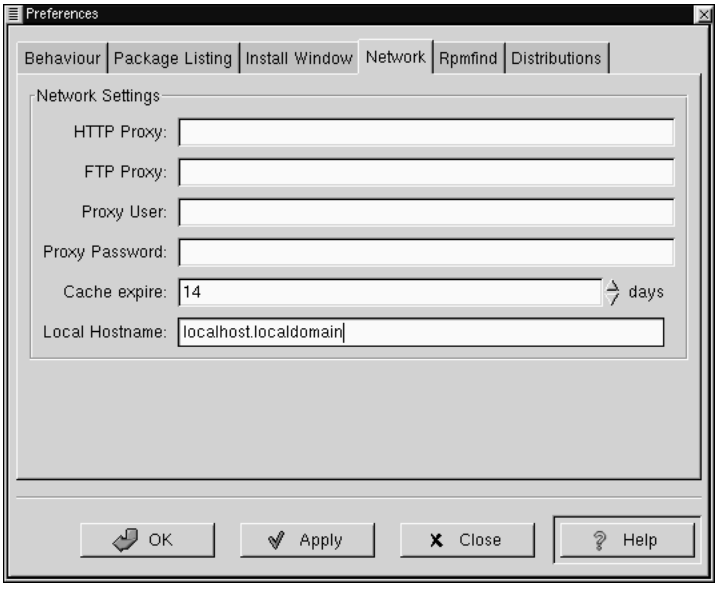

In **Rpmfind** and **Distributions**, you'll find settings and options which correspond to the **Web find** feature.

The Rpmfind system was devised by Daniel Veillard, and allows the user to search the Internet for packages by name, summary, architecture and more (see Figure 6–8, *The Rpmfind Window*). The user is then given the option of downloading and installing the most appropriate packages for their system. To learn more about Rpmfind, go to http://rpmfind.net/.

Figure 6–8 The Rpmfind Window

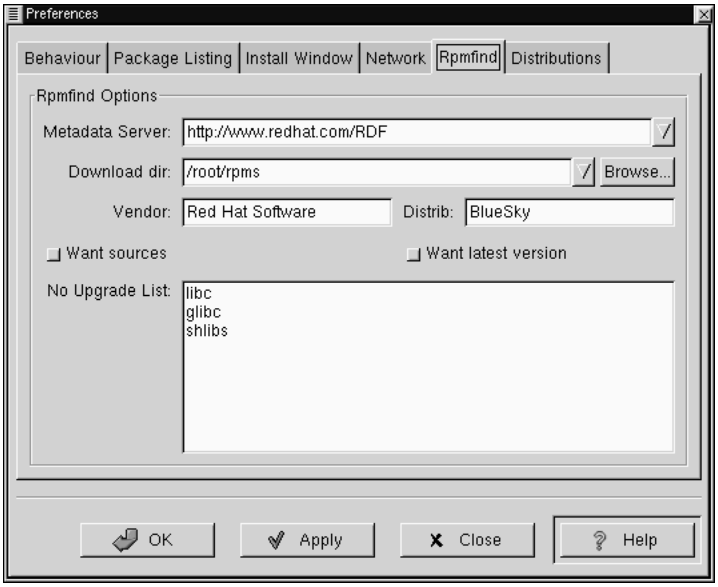

The **Metadata server** sets the server to be used for searches. The **Download dir:** entry allows you to specify where you want the files to be placed.

You can also specify the vendor, distribution name and whether to find sources and/or the latest files.

Figure 6–9 Distribution Settings in Preferences

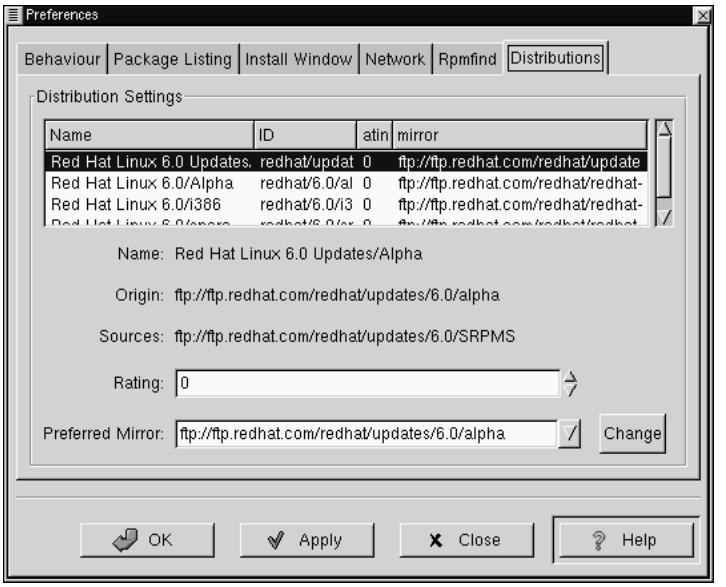

In **Distribution Settings**, you can set the options for choosing the most appropriate package out of the selections Rpmfind returns, as well as which mirror you would like to use. The higher the rating you indicate for your selection (as shown in Figure 6–9, *Distribution Settings in Preferences*) the higher the priority it will receive; the lower rating, such as "-1," will specify that packages not be recommended.

6.5 Package Manipulation

6.5.1 Querying Packages

The easiest way to query packages is to use the **Query** option from the menu at the top. If you want to query more than one package, make all your selections then press the **Query** button on the menu.

You'll be presented with a window like the one shown in Figure 6–10, *Query Window*. The more packages you've queried, the more "tabs" you'll find within the **Query** box, each tab representing a **Query** window for a package.

Figure 6–10 Query Window

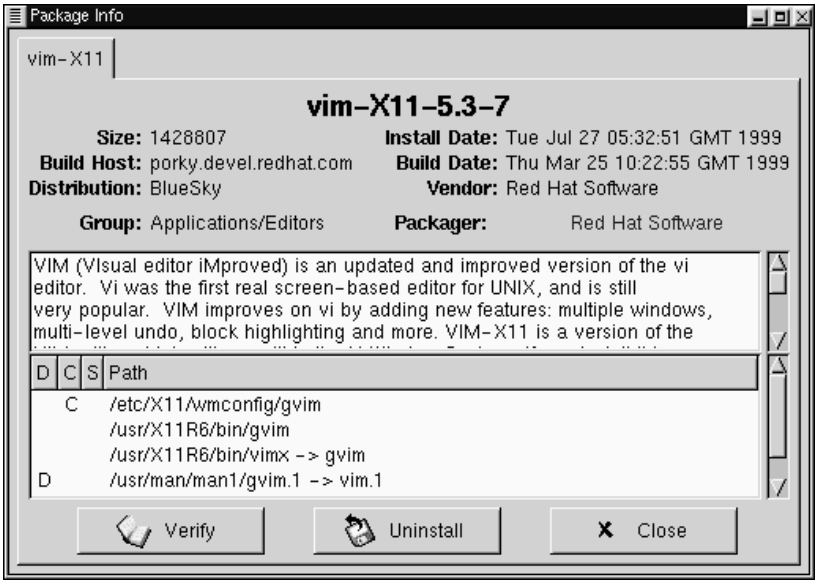

The name of the package is centered at the top of the box. Below, the box is divided into two columns of listed information; below this information, you'll see a display area showing package files.

In the left column in the information list, you'll find the size of the file, the machine on which the file is found, the name of the package distribution and the group to which its function belongs.

In the right column, you'll find the date of the package's installation on your machine, the date the package was built, the name of the vendor and the name of the group who packaged the software. If the package has not been installed on your machine, that space will simply read, "not installed." Clicking on the name following **Packager** will cause your browser's e-mail application to open in compse, so that you can write to the packager.

Centered at the bottom of the information list is the URL of the application's developer (see Figure 6–11, *URL in the Query Window*). Similar to the e-mail function of

the **Packager** entry, clicking on the URL will cause your browser to open to the corresponding website.

Figure 6–11 URL in the Query Window

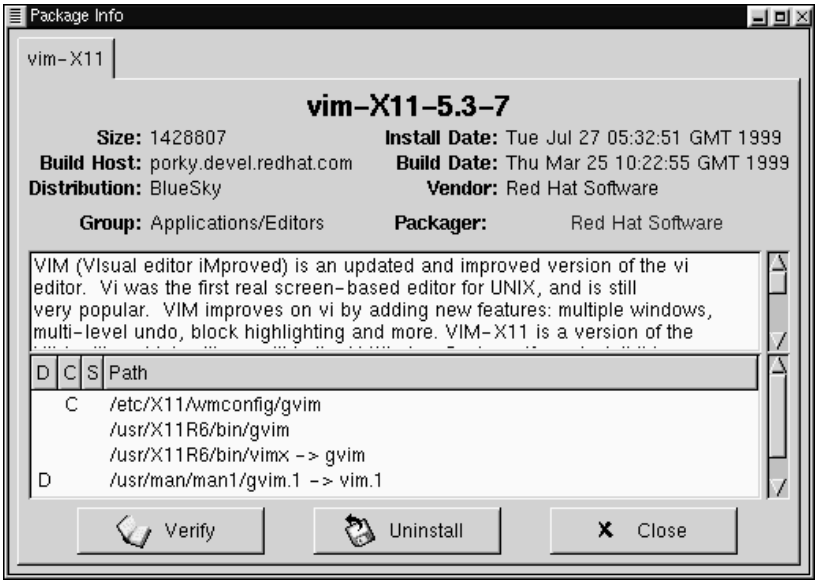

Below the description is a list of the files contained in the package. If a D appears in its related column to the left of the path, that file is a documentation file and would be a good thing to read for help on using the application. If a C appears in its respective column, the file is a configuration file. Under the S column, you can view the "state" of the package; here, you'll receive information if any files are reported as "missing" from the package (and therefore probably mean there's a problem with the package).

If you're querying a package that's already installed, you'll also find two additional buttons beneath at the bottom of this window: **Verify** and **Uninstall**. If you're performing a query on a package that hasn't been installed yet, the buttons on the bottom will be labeled **Install**, **Upgrade** and **Check Sig**.

To close the query window without performing any action, left-click on the **X** at the top right of the window bar.

6.5.2 Verifying Packages

Verifying a package checks all of the files in the package to ensure they match the ones present on your system. The checksum, file size, permissions, and owner attributes are all checked against the database. This check can be used when you suspect that one of the program's files has become corrupted for some reason.

Choosing the packages to verify is like choosing the packages to query. Select the packages in the display window and use the **Verify** button on the toolbar or from **Packages** => **Verify** on the menu. A window opens like the one in Figure 6–12, *Verify Window*.

Figure 6–12 Verify Window

As the package is being checked, you'll see the progress in the **Verify** window. If there are any problems discovered during the verify process, they'll be described in the main display area.

6.5.3 Uninstalling Packages

Uninstalling a package removes the application and associated files from your machine. When a package is uninstalled, any files it uses that are not needed by other packages on your system are also removed. Changed configuration files are copied to *<filename>*.rpmsave so you can reuse them later.

Please Note

Remember that you must be root to uninstall packages.

If uninstalling a package would break "dependencies" (which could hobble other applications that require one or more of the removed files in the package), a dialog will pop up, asking you to confirm the deletion. This will occur if you haven't selected the **"No dependency checks"** box from the **Preferences** menu (as shown in Figure 6–13, *The Behavior Tab in Preferences*).

Figure 6–13 The Behavior Tab in Preferences

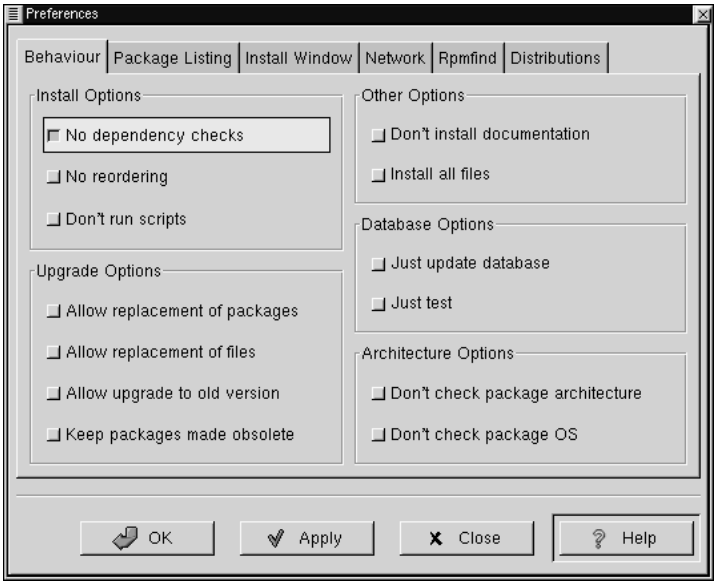

There are a variety of methods through which you can remove a selected package: from the menu, under **Packages**; from the toolbar and from the **Query** function. If you decide to remove more than one package at a time, you can choose either an incremental or global selection in the same way as you would when installing, querying or verifying. The total of your selections will be reflected in the status bar on the bottom of the main window. Because you can remove more than one package at a time, use caution to select only those which you wish to remove.

Figure 6–14 Uninstall Window

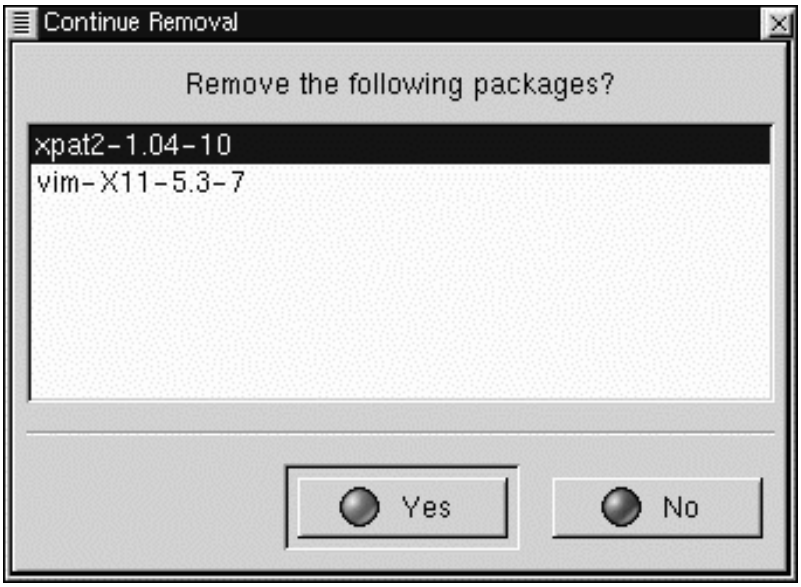

Once you've begun the uninstall, Gnome-RPM asks for confirmation, showing a window like the one in Figure 6–14, *Uninstall Window*. All of the packages that are about to be uninstalled are listed. You should look at them all to ensure you're not about to remove something you want to keep. Clicking the **Yes** button will start the uninstallation process. After it completes, the packages and groups that have been removed will disappear from any windows they were in.

Upgrading Packages

When a new version of a package has been released, it is easy to install it on your system. Select the packages from the window of available packages in the same way you select packages for installation. Both the **Upgrade** button on the toolbar and, from the menu, under **Operations** => **Upgrade** will begin the process. You simply **Add** packages in the same manner as you would a new package installation.

During the upgrade, you'll see a progress indicator like the one for installing packages. When it's finished, any old versions of the packages will be removed, unless you specify otherwise (refer to Section 6.4, *Configuration* for more information).

It is much better to use the upgrade option than to uninstall the old versions of a package and then install the new one. Using upgrade ensures that any changes you made to package configuration files are preserved properly, while uninstalling and then reinstalling a new package could cause those changes to be lost.

If you run out of disk space during an installation, the install will fail. However, the package which was being installed when the error occurred may leave some files around. To clean up after this error, reinstall the package after you've made more disk space available.

7 Lightweight Directory Access Protocol (LDAP)

7.1 What is LDAP?

LDAP (Lightweight Directory Access Protocol) is a proposed open standard for global or local directory services over a network and/or the Internet. A directory, in this sense, is very much like a phone book. LDAP can handle other information, but at present it is typically used to associate names with phone numbers and e-mail addresses. Directories are designed to support a high volume of queries, but the data in the directory doesn't change all that often.

LDAP is much more useful than a paper phone book, because LDAP's design is intended to support propagation over LDAP servers throughout the Internet, much like the Domain Name Service (DNS). The DNS system acts as the address book of the Internet by keeping track of domain name/IP address pairs. DNS servers tell networked machines where packets need to go. In the future, LDAP could provide the same type of global access to many types of directory information: at present, LDAP is more commonly used within a single large organization, like a college or a company, for directory services.

LDAP is a client-server system. An LDAP client connects to an LDAP server and either queries for information or provides information that needs to be entered into the directory. The server either answers the query, refers the query to another LDAP server, or accepts the information for incorporation into the directory.

LDAP is sometimes known as **X.500 Lite**. X.500 is an international standard for directories. X.500 is full-featured, but it is complex and requires lots of computing resources and the full OSI stack. LDAP, in contrast, can run easily on a PC and over TCP/IP. LDAP can access X.500 directories, but it does not support every capability of X.500.

This chapter will refer to the configuration and use of OpenLDAP, an open source implementation of LDAP. OpenLDAP includes `slapd`, a stand-alone LDAP server;

slurpd, a stand-alone LDAP replication server; libraries implementing the LDAP protocol; utilities; tools; and sample clients.

7.2 Pros and Cons of LDAP

The main benefit of using LDAP is the consolidation of 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. This directory can be queried by any LDAP-enabled applications that need this information. The directory can also be used by users who need directory information.

Other LDAP benefits include its ease of implementation (compared to X.500), and its well-defined Application Programming Interface (API), which means that the number of LDAP-enabled applications and LDAP gateways should increase in the future.

On the negative side, if you want to use LDAP, you'll need LDAP-enabled applications or you'll need to use LDAP gateways. As mentioned previously, LDAP will only increase in usage, but at present, there aren't a plethora of LDAP-enabled applications available for Linux. Also, while LDAP does support some access control, it does not support as many security features as X.500.

7.3 Uses for LDAP

Several Netscape applications, including Netscape Roaming Access are LDAP-enabled. Sendmail can use LDAP to look up addresses. Your organization can use LDAP as an organization-wide directory and/or name service (in place of NIS or flat files). You can even use a personal LDAP server to keep track of your own e-mail address book (see Section 7.10, *LDAP Resources on the Web*).

LDAP can be used as an authentication service via the pam_ldap module. LDAP is commonly used as a central authentication server so that users have a unified login that covers console logins, POP servers, IMAP servers, machines connected to the network using Samba, and even Windows NT machines. All of these login situations can rely on the same user ID and password combination, using LDAP. The pam_ldap module is provided in the nss_ldap package.

7.4 LDAP Terminology

An **entry** is one unit in an LDAP directory. An entry is identified or referenced by its unique **Distinguished Name** (DN).

An entry has attributes; attributes are pieces of information which are directly associated with the entry. For example, an organization could be an LDAP entry. Attributes associated with the organization might be its fax number, its address, and so on. People can also be entries in the LDAP directory. Common attributes for people include their telephone numbers and their e-mail addresses.

Certain attributes are required, while other attributes are optional. An **objectclass** sets which attributes are required and which are optional. Objectclass definitions are found in the `slapd.oc.conf` file.

The **LDAP Data Interchange Format** (LDIF) is an ASCII text format for LDAP entries. Files that import or export data to and from LDAP servers must be in LDIF format. An LDIF entry looks like this:

```
[<id>]
dn: <distinguished name>
<attrtype>: <attrvalue>
<attrtype>: <attrvalue>
<attrtype>: <attrvalue>
```

An entry can contain as many <attrtype>: <attrvalue> pairs as needed. A blank line indicates that the entry is finished and that another entry is about to begin.

Everything enclosed within < > is variable, and can be set by you when you add an LDAP entry, with the exception of the <id>. The <id> is a number normally set by the LDAP tools when you add an entry, and you'll probably never need to manually set one.

7.5 OpenLDAP Files

OpenLDAP configuration files are installed into the `/etc/openldap` directory. If you do an `ls` on `/etc/openldap`, you'll see something like:

```
ldap.conf        ldapsearchprefs.conf  slapd.at.conf  slapd.oc.conf
ldapfilter.conf  ldaptemplates.conf    slapd.conf
```

The files you should know about are `slapd.conf`, `slapd.at.conf` and `slapd.oc.conf`. The `slapd.conf` file contains configuration information for the `slapd` daemon and for all of the database back-ends. You will need to make some changes to `slapd.conf` before you start the `slapd` daemon.

The `slapd.at.conf` file holds **attribute syntax definitions** for the LDAP directory. Attribute syntax definitions describe the type of information that the attribute provides. You have five choices:

- bin -- binary information
- ces -- case exact string (a string whose case matters)
- cis -- case ignore string (a string whose case doesn't matter)
- tel -- a telephone number (blank spaces and hyphens are ignored)
- dn -- the distinguished name

For example, here are attribute syntax definitions from a typical `slapd.at.conf`:

```
attribute       photo                           bin
attribute       personalsignature               bin
attribute       jpegphoto                       bin
attribute       audio                           bin
attribute       labeledurl                      ces
attribute       ref                             ces
attribute       userpassword                    ces
attribute       telephonenumber                 tel
```

Certain LDAP-enabled applications may require that you edit the `slapd.at.conf` file, usually to add particular attribute syntax definitions.

The `slapd.oc.conf` file includes the **objectclass definitions** for an LDAP directory. The objectclass definitions state which attributes are required and which attributes are optional for particular objectclasses.

The following excerpt from a typical `slapd.oc.conf` file contains the objectclass definitions for the top, alias, and referral objectclasses:

```
objectclass top
        requires
                objectClass
```

```
objectclass alias
        requires
                aliasedObjectName,
                objectClass

objectclass referral
        requires
                ref,
                objectClass
```

You may need to edit the objectclass definitions in your `slapd.oc.conf`, depending upon how you're going to use your LDAP directory. For example, if you're creating an LDAP directory of employees for use in your organization, you'll probably have specific required attributes for certain objectclasses that might not be used outside the organization (e.g., an employee ID number internal to your organization might be a required attribute for an objectclass of "person").

7.6 OpenLDAP Daemons and Utilities

The OpenLDAP package includes two daemons: `slapd` and `slurpd`. The `slapd` daemon is the stand-alone LDAP daemon, which you'll need to run to support LDAP.

The `slurpd` daemon controls the replication of LDAP directories over a network. `Slurpd` sends changes from the master LDAP directory to slave LDAP directories. You won't need to run `slurpd` unless you have more than one LDAP server on your network. If you have two or more LDAP servers, you'll need to run `slurpd` to keep the LDAP directories in sync.

OpenLDAP also includes some utilities for adding, modifying and deleting entries in an LDAP directory. The `ldapmodify` tool is used to modify entries in an LDAP database. The `ldapadd` utility is used to add entries to your directory (`ldapadd` is actually a hard link to `ldapmodify -a`). `Ldapsearch` is used to search for entries and `ldapdelete` is used to delete entries. The `ldif2ldbm` tool converts an LDIF file into an LDBM back-end database.

See their man pages for more information on all of these utilities.

7.7 Modules for Adding Extra Functionality to LDAP

Red Hat Linux includes the following packages which add functionality to LDAP:

The nss_ldap module is an LDAP module for the **Solaris Nameservice Switch** (NSS). NSS is a set of C library extensions necessary for accessing LDAP directory information, instead of or in addition to the **Network Information Service** (NIS) name service and/or flat files. The nss_ldap module is needed to use LDAP as a native name service.

The pam_ldap module is needed to integrate LDAP authentication into the Pluggable Authentication Modules (PAM) API. If you use pam_ldap, users can authenticate and change their password using LDAP directories. The nss_ldap and pam_ldap modules are provided in the nss_ldap package.

Red Hat Linux also includes LDAP modules for the Apache Web server. The auth_ldap module is for authenticating HTTP clients against the user entries in an LDAP directory. The php-ldap module adds LDAP support to the PHP4 HTML-embedded scripting language. The auth_ldap and php-ldap modules will need to be compiled into Apache as **Dynamic Shared Objects** (DSOs).

7.8 LDAP How To: A Quick Overview

This section provides a quick overview of the steps you'll need to take to get an LDAP directory working.

1. Make sure the openldap RPM, and any other LDAP-related RPMS that you need, are installed.
2. See either the Quick Start Guide at the OpenLDAP site (http://www.openldap.org/faq/data/cache/172.html; start at "Create configuration file for slapd," since the LDAP files are already installed), or see the Linux-LDAP HOWTO (http://www.linuxdoc.org/HOWTO/LDAP-HOWTO.html) for instructions on using LDAP on your system. Both cover the rest of these steps.
3. Edit the slapd.conf file to get it right for your system.
4. Start slapd.
5. Create your LDAP directory (examples of LDAP entries are provided at the PADL Software website at http://www.padl.com/ldap_examples.html).
6. Add entries to your LDAP directory with ldapadd or with a script.

7. Use `ldapsearch` to see if `slapd` is working.
8. At this point, your LDAP directory should exist. The next step is to configure your LDAP-enabled applications so that they can use the LDAP directory.

7.9 Configuring Your System to Authenticate Using OpenLDAP

This section provides a brief overview of how to configure your Red Hat Linux system to authenticate using OpenLDAP. Unless you're an OpenLDAP expert, you will probably need more documentation than is provided here. Please refer to the references provided in Section 7.10, *LDAP Resources on the Web* for more information.

7.9.1 Install the Necessary LDAP Packages

First, you'll need to make sure that the appropriate packages are installed on both the LDAP server and the LDAP client machines. The LDAP server needs the `openldap` package.

The LDAP client machines need the following packages installed: `openldap`, `auth_ldap`, and `nss_ldap`.

7.9.2 Edit Configuration Files

Edit `/etc/openldap/slapd.conf`

The `slapd.conf` file, located in `/etc/openldap`, contains the configuration information needed by your `slapd` LDAP server. You'll need to edit this file to make it specific to your domain and your server.

The suffix line names the domain for which the LDAP server will provide information. The suffix line should be changed from:

```
suffix   "dc=your-domain, dc=com"
```

so that it reflects your domain name. For example:

```
suffix  "dc=acmewidgets, dc=com"
```

or

```
suffix   "dc=acmeuniversity, dc=org"
```

The rootdn entry is the DN for a user who is unrestricted by the access control or administrative limit parameters set for operations on the LDAP directory. The rootdn user can be thought of as the root user for the LDAP directory. The rootdn line needs to be changed from:

```
rootdn  "cn=root, dc=your-domain, dc=com"
```

to something like:

```
rootdn  "cn=root, dc=redhat, dc=com"
```

or

```
rootdn  "cn=ldapmanager, dc=my_organization, dc=org"
```

Change the rootpw line from:

```
rootpw   secret
```

to something like

```
rootpw  {crypt}s4L9sOIJo4kBM
```

In the above example, you're using an encrypted root password, which is a much better idea than leaving a plain text root password in the `slapd.conf` file. To make this crypt string, you should either copy it out of a `passwd` file, or use Perl:

```
perl -e "print crypt('passwd','a_salt_string');"
```

In the previous Perl line, `salt_string` is a two character salt, and `passwd` is the plain text version of the password.

You could also copy a `passwd` entry out of `/etc/passwd`, but this won't work if the `passwd` entry is an MD5 password (the default in Red Hat Linux 7.0).

Edit `ldap.conf`

Edit the `ldap.conf` files in `/etc` and in `/etc/openldap` on the LDAP server and clients.

Edit `/etc/ldap.conf`, the configuration file for `nss_ldap` and `pam_ldap`, to reflect your organization and search base. The file `/etc/openldap/ldap.conf` is the configuration file for the command line tools like `ldapsearch`, `ldapadd`,

etc., and it will also need to be edited for your LDAP setup. Client machines will need to have both of these files modified for your system.

Edit `/etc/nsswitch.conf`

To use `nss_ldap`, you'll need to add `ldap` to the appropriate fields in `/etc/nsswitch.conf`. (Be very careful when editing this file; be sure that you know what you're doing.) For example:

```
passwd: files ldap
shadow: files ldap
group: files ldap
```

PAM and LDAP

To have standard PAM-enabled applications use LDAP for authentication, run `authconfig` and select **Use LDAP**. (PAM is beyond the scope of this LDAP overview, so if you need help, consult Section 2.6, *User Authentication with PAM* and/or PAM man pages.)

7.9.3 Migrate Your Old Authentication Information to LDAP Format

The `/usr/share/openldap/migration` directory contains a set of shell and Perl scripts for migrating your old authentication information into LDAP format. (Yes, you'll need to have Perl on your system to use these scripts.)

First, you'll need to modify the `migrate_common.ph` file so that it reflects your domain. The default DNS domain should be changed from:

```
$DEFAULT_MAIL_DOMAIN = "padl.com";
```

to something like:

```
$DEFAULT_MAIL_DOMAIN = "your_company.com";
```

The default base should also be changed, from:

```
$DEFAULT_BASE = "dc=padl,dc=com";
```

to something like:

```
$DEFAULT_BASE = "dc=your_company,dc=com";
```

Next, you'll need to decide which script to use. The following table should tell you:

Table 7–1 LDAP Migration Scripts

Existing name service	Is LDAP running?	Use this script:
`/etc` flat files	yes	`migrate_all_online.sh`
`/etc` flat files	no	`migrate_all_offline.sh`
NetInfo	yes	`migrate_all_netinfo_online.sh`
NetInfo	no	`migrate_all_netinfo_offline.sh`
NIS (YP)	yes	`migrate_all_nis_online.sh`
NIS (YP)	No	`migrate_all_nis_offline.sh`

Run the appropriate script based on your existing name service.

The `README` and the `migration-tools.txt` files in `/usr/share/openldap/migration` provide more details.

7.10 LDAP Resources on the Web

Lots of useful information about LDAP can be found on the Web. Please review these sources, especially the OpenLDAP website and the LDAP HOWTO, before you start to set up LDAP on your system.

OpenLDAP

http://www.openldap.org

University of Michigan

http://www.umich.edu/~dirsvcs/ldap/

The SLAPD and SLURPD Administrator's Guide

http://www.umich.edu/~dirsvcs/ldap/doc/guides/slapd

Innosoft/Critical Angle

http://www.innosoft.com/ldapworld

Jeff Hodges' LDAP Road Map and FAQ

http://www.kingsmountain.com/ldapRoadmap.shtml

PADL (nss_ldap, pam_ldap and ypldapd)

http://www.padl.com/

auth_ldap

http://www.rudedog.org/auth_ldap/1.4/auth_ldap.html

The LDAP HOWTO

http://www.linuxdoc.org/HOWTO/LDAP-HOWTO.html

Sendmail — using LDAP to do address lookups

http://www.stanford.edu/~bbense/Inst.html

Using a personal LDAP server, the Apache Web server and PHP to manage your personal e-mail addressbook

http://www.webtechniques.com/archives/1999/05/junk/junk.shtml

Also, remember that man pages exist for the various LDAP daemons and utilities. Please check the man pages if you need more information.

8 Using Kerberos 5 on Red Hat Linux

Kerberos is a secure system for providing network authentication services. Authentication means:

- the identities of entities on the network are verified
- traffic on the network is from the source who claims to have sent it

Kerberos uses users' passwords to verify the identity of users, but passwords are never sent unencrypted over the network.

8.1 Why Use Kerberos?

Most conventional network systems use password-based authentication schemes. When a user needs to authenticate to a service running on a network server, they type in their password for each service that requires authentication. Their password is sent over the network, and the server verifies their identity using the password.

Transmission of passwords in plaintext in this way, while commonly done, is a tremendous security risk. Any system cracker with access to the network and a packet analyzer (commonly called a packet sniffer) can intercept any passwords sent this way.

The primary design goal of Kerberos is to ensure that passwords are *never* sent across a network unencrypted, and preferably never sent over the network at all. The proper use of Kerberos will eradicate the threat of packet sniffers intercepting passwords on your network.

8.2 Why Not Use Kerberos?

Kerberos removes a common security threat, so why isn't it in use on every network? For several reasons, Kerberos may be difficult to implement:

- No quick "script-o-matic" solution exists for migrating user passwords from a standard UNIX password database to a Kerberos password database. Migration is technically feasible, but conversion scripts aren't provided with Kerberos. See the Kerberos FAQ Question 2.23 for more detailed information on this issue.

- Kerberos is only partially-compatible with the Pluggable Authentication Modules (PAM) system used by most servers on Red Hat Linux. For more information on this issue, see Section 8.7, *Kerberos and Pluggable Authentication Modules (PAM)*.
- For an application to use Kerberos, its sources must be modified to make the appropriate calls into the Kerberos libraries. For some applications, this may require too much programming effort. For other applications, changes must be made to the protocol used between network servers and their clients; again, this may require too much effort. Furthermore, it may be impossible to make certain closed-source applications work with Kerberos.
- Finally, if you decide to use Kerberos on your network, you must realize that it is an all-or-nothing proposition. If *any* services that transmit plaintext passwords remain in use, passwords can still be compromised, and your network gains no net benefit from the use of Kerberos. To secure your network with Kerberos, you must either **kerberize** (i.e., make it work with Kerberos) *all* applications that send plaintext passwords or stop using those applications on your network.

8.3 Kerberos Terminology

Like any other system, Kerberos has its own terminology. Before we talk about how it works, here is a list of terms that you'll need to be familiar with:

- ciphertext — encrypted data
- client — an entity on the network (a user, a host or an application) that can get a ticket from Kerberos
- credential cache or ticket file — a file which contains the keys for encrypting communications between a user and various network services. Kerberos 5 provides a framework for using other cache types (such as shared memory), but files are more well-supported
- key — a chunk of data, used when encrypting or decrypting data. Encrypted data can't be decrypted without the proper key or really good guessing
- Key Distribution Center (KDC) — a machine that issues Kerberos tickets (sometimes called a Ticket Granting Server or TGS)

- keytab — short for **key table**, a file that includes an unencrypted list of principals and their keys. Servers retrieve the keys they need from keytab files instead of using `kinit`. The default keytab file is `/etc/krb5.keytab`, with `kadmind` the only known service that uses any other file (it uses `/var/kerberos/krb5kdc/kadm5.keytab`)
- plaintext — unencrypted data
- principal — a user or service that can authenticate using Kerberos. A principal's name is in the form "*root*[/*instance*]*@REALM*". For a typical user, the *root* is the same as their login ID. The *instance* is optional. If the principal has an instance, it is separated from the root with a forward slash ("/"). The empty string ("") is actually a valid instance (which differs from the default, *NULL* instance), but using it can be confusing. All principals in a realm have their own *key*, which is derived from their password (for users) or randomly set (for services)
- realm — a network that uses Kerberos, composed of one or a few servers (also known as KDCs) and a (potentially very large) number of clients
- service — a program or computer accessed over the network
- ticket — a temporary set of electronic credentials that verify the identity of a client for a particular service
- Ticket Granting Ticket (TGT) — a special ticket which allows the client to obtain additional tickets without applying for them from the KDC

8.4 How Kerberos Works

Now that you've heard a few of the terms that Kerberos uses, here is a simplified explanation of how a Kerberos authentication system works:

On a "normal" network which uses passwords to authenticate users, when a user requests a network service that requires authentication, the user is prompted to type in their password. Their password is transmitted in plaintext over the network, and access to the network service is granted.

As mentioned previously, the central problem solved by Kerberos is how to use passwords for authentication without sending them over the network. On a kerberized network, the Kerberos database contains principals and their keys (for users, their

keys are derived from their passwords). The Kerberos database also contains keys for all of the network services.

When a user on a kerberized network logs in to their workstation, their principal is sent to the KDC as a request for a TGT. This request can be sent by the login program (so that it is transparent to the user) or can be sent by the `kinit` program after the user logs in.

The KDC checks for the principal in its database. If the principal is found, the KDC creates a TGT, encrypts them using the user's key, and sends it back to the user.

The login program or `kinit` decrypts the TGT using the user's key (which it computes from the user's password). The TGT, which is set to expire after a certain period of time, is stored in your credentials cache. An expiration time is set so that a compromised TGT can only be used for a certain period of time, usually eight hours (unlike a compromised password, which could be used until changed). The user won't have to re-enter their password until the TGT expires or they logout and login again.

When the user needs access to a network service, the TGT requests a ticket for the service from the Ticket Granting Service (TGS), which runs on the KDC. The TGS issues a ticket for the desired service, which is used to authenticate the user.

As you might have guessed, the preceding explanation was vastly oversimplified. If you need a more in-depth explanation of how Kerberos works, see Section 8.8, *Sources of Information about Kerberos*.

Please note

Kerberos depends on certain network services to work correctly. First, Kerberos needs (loose) clock synchronization between the machines on your network. If you haven't set up a clock syncing program for your network, you'll need to do so. And since certain aspects of Kerberos rely on the Domain Name Service (DNS), be sure that the DNS entries and hosts on your network are all correctly set up. See the *Kerberos V5 System Administrator's Guide*, provided in PostScript and HTML formats, in `/usr/share/doc/krb5-server-`*`versionnumber/`*, for more information on these issues.

8.5 Setting Up a Kerberos 5 Server on Red Hat Linux 7.0

When you're setting up Kerberos, install the server(s) first. If you need to set up slave servers, the details of setting up relationships between master and slave servers are covered in the *Kerberos 5 Installation Guide* (in `/usr/share/doc/krb5-server-`*`versionnumber/`*).

To install a Kerberos server:

1. Install the `krb5-libs`, `krb5-server`, and `krb5-workstation` packages on the dedicated machine which will run your KDC. This machine needs to be secure — if possible, it shouldn't run any services besides the KDC.

   If you'd like to use a Graphical User Interface (GUI) utility to administrate Kerberos, you should also install the `gnome-kerberos` package. `gnome-kerberos` contains `krb5`, a GUI tool for managing tickets, and `gkadmin`, a GUI tool for managing Kerberos realms.

2. Edit the `/etc/krb5.conf` and `/var/kerberos/krb5kdc/kdc.conf` configuration files to reflect your realm name and domain-to-realm mappings.

A simple realm can be constructed by replacing instances of *EXAMPLE.COM* and *example.com* with your domain name (keeping uppercase names uppercase and lowercase names lowercase) and by changing the KDC from *kerberos.example.com* to the name of your Kerberos server. By convention, all realm names are uppercase and all DNS hostnames and domain names are lowercase. For full details on the formats of these files, see their respective man pages.

3. Create the database using the `kdb5_util` utility from a shell prompt:

```
/usr/kerberos/sbin/kdb5_util create -s
```

The `create` command creates the database that will be used to store keys for your Kerberos realm. The `-s` switch forces creation of a **stash** file in which the master server key is stored. If no stash file is present to read the key from, the Kerberos server (`krb5kdc`) will prompt the user for the master server password (which can be used to regenerate the key) every time it is started.

4. Edit the `/var/kerberos/krb5kdc/kadm5.acl` file. `kadmind` uses this file to determine which principals have access to the Kerberos database, and what kind of access they have. Most organizations will be able to get by with a single line:

```
*/admin@EXAMPLE.COM  *
```

Most users will be represented in the database by a single principal (with a *NULL* instance, i.e., *joe@EXAMPLE.COM*). With this configuration, users with a second principal with an instance of *admin* (for example, *joe/admin@EXAMPLE.COM*) will be able to wield full power over the realm's Kerberos database.

Once `kadmind` is started on the server, any user will be able to access its services by running `kadmin` or `gkadmin` on any of the clients or servers in the realm. However, only users listed in the `kadm5.acl` file will be able to modify the database in any way except for changing their own passwords.

Please note

The `kadmin` and `gkadmin` utilities communicate with the `kadmind` server over the network. Of course, you need to create a principal before you can connect to the server over the network to administer it, so do that with the `kadmin.local` command:

```
/usr/kerberos/sbin/kadmin.local -q addprinc joeuser/admin
```

5. Start Kerberos using the following commands:

```
krb5kdc start
kadmin start
krb524 start
```

6. Add principals for your users using `kadmin`'s `addprinc` command, or using `gkadmin`'s **Principal** => **Add** menu option.

7. Verify that your server will issue tickets. First, run `kinit` to obtain a ticket and store it in a credential cache file. Then use `klist` to view the list of credentials in your cache and use `kdestroy` to destroy the cache and the credentials it contains.

Please note

By default, `kinit` attempts to authenticate you using the login name of the user you're currently logged in as. If that user doesn't correspond to a principal in your Kerberos database, you will get an error message. If that happens, just give `kinit` the name of your principal as an argument on the command line.

Once you've completed the steps listed above, your Kerberos server should be up and running. Next, you'll need to set up your Kerberos clients.

8.6 Setting Up a Kerberos 5 Client on Red Hat Linux 7.0

Setting up a Kerberos 5 client is less involved than setting up a server. At minimum, you'll need to install the client packages and provide your clients with a valid `krb5.conf` configuration file. Kerberized versions of `rsh` and `rlogin` will also require some configuration changes.

1. Install the `krb5-libs` and `krb5-workstation` packages on all of the clients in your realm. You will need to supply your own version of `/etc/krb5.conf` for your client workstations; usually this can be the same `krb5.conf` used by the KDC.

2. Before a particular workstation in your realm can allow users to connect using kerberized `rsh` and `rlogin`, that workstation will need to have the `xinetd` package installed and have its own host principal in the Kerberos database. The `kshd` and `klogind` server programs will also need access to the keys for their service's principal.

   Using `kadmin`, add a host principal for the workstation. The instance in this case will be the hostname of the workstation. Because you'll never need to type the password for this principal again, and you probably don't want to bother with coming up with a good password, you can use the `-randkey` option to kadmin's `addprinc` command to create the principal and assign it a random key:

   ```
   addprinc -randkey host/blah.example.com
   ```

   Now that you have created the principal, you can extract the keys for the workstation by running `kadmin` *on the workstation itself*, and using `kadmin`'s `ktadd` command:

   ```
   ktadd -k /etc/krb5.keytab host/blah.example.com
   ```

   In order to use the kerberized versions of `rsh` and `rlogin`, you'll need to use either ntsysv or `chkconfig` to enable `klogin`, `eklogin`, and `kshell`.

3. Other kerberized network services will need to be started. To use kerberized `telnet`, you'll need to use ntsysv or `chkconfig` to enable `ktelnet`.

If you want to provide FTP access as well, you'll need to create and extract a key for a principal with a root of ftp, and the instance set to the hostname of the FTP server. Then use ntsysv or `chkconfig` to enable `gssftp`.

The IMAP server included in the `imap` package will use GSS-API authentication using Kerberos 5 if it finds the proper key in `/etc/krb5.keytab`. The root for the principal should be imap. The CVS gserver uses a principal with a root of (surprise!) cvs, and is otherwise identical to a pserver.

That should be all you need to do to set up a simple Kerberos realm.

8.7 Kerberos and Pluggable Authentication Modules (PAM)

Currently, kerberized services do not make use of PAM at all — a kerberized server bypasses PAM completely. Applications that use PAM can make use of Kerberos for password-checking if the pam_krb5 module (provided in the `pam_krb5` package) is installed. The `pam_krb5` package contains sample configuration files that will allow services like `login` and `gdm` to authenticate users and obtain initial credentials using their passwords. Provided that access to network servers is always done using kerberized services (or services that use GSS-API, like IMAP), the network can be considered reasonably safe.

Careful system administrators will not add Kerberos password-checking to network services, because most of the protocols used by these services do not encrypt the password before sending it over the network — obviously something you'd want to avoid.

8.8 Sources of Information about Kerberos

If you need more information on Kerberos, there are good sources on the Web:

http://www.nrl.navy.mil/CCS/people/kenh/kerberos-faq.html

The Kerberos Frequently Asked Questions (FAQ)

http://web.mit.edu/kerberos/www/

The Kerberos home page on MIT's website

ftp://athena-dist.mit.edu/pub/kerberos/doc/usenix.PS

Kerberos: An Authentication Service for Open Network Systems by Jennifer G. Steiner, Clifford Neuman, and Jeffrey I. Schiller — the original paper describing Kerberos, in PostScript format

http://web.mit.edu/kerberos/www/dialogue.html

Designing an Authentication System: a Dialogue in Four Scenes originally by Bill Bryant in 1988, modified by Theodore Ts'o in 1997. This document is a conversation between two developers who are thinking through the creation of a Kerberos-style authentication system. The conversational style and "from the ground up" nature of the discussion make this a good starting place for Kerberos newbies.

http://www.ornl.gov/~jar/HowToKerb.html

Practical advice on kerberizing your network

When you install the Kerberos RPM packages, the *Kerberos V5 Installation Guide* and the *Kerberos V5 System Administrator's Guide* are installed in PostScript and HTML formats, in `/usr/share/doc/krb5-server-`*`versionnumber`*`/`. The *Kerberos V5 UNIX User's Guide* is installed in PostScript and HTML formats in `/usr/share/doc/krb5-workstation-`*`versionnumber`*`/`.

9 Credit Card Verification System (CCVS) Basics

The Credit Card Verification System (CCVS) uses your computer and modem to simulate a credit card swipe box (also known as a Point of Sale [POS] terminal). A stand-alone product, CCVS includes several Application Programming Interfaces (APIs) that facilitate customization and integration with third party software applications or database products.

CCVS is safe, secure, and easy to use. Written in ANSI C and conforming to POSIX standards, CCVS is portable and designed to be easily integrated with modern operating systems, programming languages, and the Internet. Designed for easy scripting and programming, CCVS can be used to automate batch processing or enhance any application that requires credit card processing.

CCVS can be used in countries other than the US if your bank or merchant services representative can support one of the protocols supported by CCVS. If you're in Canada, CCVS supports the NDC protocol, which can be used by any bank in Canada to configure your merchant account. If you're in a country other than the US or Canada, you'll need to check with your merchant services representative. The protocol supported by CCVS that has the best chance of being supported by a financial institution outside the US is the the Visa 2nd Generation "K Format" protocol (VITAL).

A demonstration version of CCVS is included with Red Hat Linux. The demo version is fully functional and can be used for testing CCVS and your system; the demo version will do everything except contact your financial institution. If you choose to purchase CCVS to process credit cards, you'll need to contact Red Hat to purchase a license key. See http://www.redhat.com/products/ccvs/ for more information on how to activate CCVS.

Examples of how CCVS can be used (depending upon the protocol you're using — see www.redhat.com/products/ccvs/support/CCVS3.3docs/protocol-specific.html for more information on what different protocols will support):

- CCVS can support a system for telephone operators taking catalog orders over the phone. CCVS' Tcl extensions can be used to create a Tcl/Tk Graphical User Interface (GUI) that presents a simple interface for telephone operators. The operators can then use simple X terminals; all the software will run on the central server. CCVS only needs to be installed on one computer, and the operators don't have to wait for an available phone line — all of their transactions will go out over the same phone call.
- CCVS can be used to help automate billing. For example, an Internet Service Provider (ISP) might have a customer database on a database server. The ISP's database administrator could write a Perl script, combining the CCVS Perl module with a module for the ISP's database system. The script would then be run every month. The script will read the customer data, process monthly billing, and update the records in the database to indicate payment has taken place.

These are only two examples of CCVS capabilities. CCVS can be used to enhance any aspect of your operations that require credit card processing. CCVS' many features include the following:

- A C library with a documented API empowers users to integrate CCVS seamlessly with existing applications.
- A Tcl extension enables use of CCVS with server-side Tcl such as NeoWebScript.
- A Perl 5.0 module allows CCVS to work with the most popular CGI programming language in use today.
- The ability to quickly construct custom GUIs using Tcl/Tk — typical development time is less than a day.
- Python, PHP3 and Java modules enable CCVS to work with other common programming languages.
- Command Line Interface (CLI) programs for interactive use. Call programs from any UNIX shell and program in the UNIX language you like best.

- AVS fraud protection, which allows merchants to check for stolen credit cards. Many clearinghouses offer a better rate to merchants who use AVS, even on orders taken over the phone.
- Support for multiple merchant accounts, allowing users to open their very own virtual malls with unlimited store fronts. A "merchant account" is a special type of bank account which allows a business to accept credit card payments from customers; the merchant account holds the proceeds from credit card transactions.
- The ability to conduct multiple transactions in a single session, approaching leased line performance (two seconds per transaction!) with no extra cost or complexity.
- The reassurance of being able to test and do development programming on the product without charging real credit cards.

9.1 The Credit Card Verification Process

How does that little piece of plastic tell the nice people that you really can afford the big screen TV?

First, a consumer presents credit card information to the merchant. The merchant transmits this data, along with their merchant ID code, to a clearinghouse (also referred to as a processor or acquirer). The clearinghouse might be the bank that has issued the merchant their credit card account, but it is more likely a firm that has contracted with the merchant's bank to clear charges in exchange for a flat fee and a percentage of every charge processed.

The data is transmitted by reading the card and merchant numbers over the phone, by using a credit card POS terminal, or by using CCVS or some other piece of software to transmit the information from a computer.

The clearinghouse contacts the bank that issued the consumer's credit card and verifies that the charge is acceptable. If it is accepted, the clearinghouse then sends a confirmation message to the merchant. At the same time, the available credit from the customer's credit card is frozen by the amount of the transaction.

At the end of a business day, the merchant (actually, the merchant's computer or credit card terminal) calls the clearinghouse and verifies all transactions for that day

to ensure that the merchant's system and the clearinghouse agree on the transactions that have occurred during that day. Once the merchant and the clearinghouse agree on the day's transactions, the clearinghouse starts the process of transferring the money from the credit card bank to the merchant's bank account.

9.2 What You'll Need to Run CCVS

To run CCVS, you'll need a modem and a merchant account. You'll also need to follow a few guidelines so that CCVS will run correctly.

9.2.1 Modems

You need at least one modem dedicated to CCVS use. Credit card protocols do not support compression or error correction during modem connects, so compression and error correction cannot be used. We can provide you with information about how to turn off such features on the following modems:

- Hayes Optima
- US Robotics Courier
- US Robotics Sportster
- Chase Research PCI-RAS

Please Note

Please use a modem or modems from the above list!

If you use a non-supported modem (anything besides the four modems listed above), it may be very difficult to get the unsupported modem to work with CCVS. You should also check the Red Hat Linux Hardware Compatibility Lists at http://www.redhat.com/support/hardware/ to make sure that your modem will work with Red Hat Linux.

If the modem you have to use does not appear on this list, look through your modem manual to find the string which turns off all compression and error correction, and

the string which resets your modem for normal use. You'll need to provide these two strings when you configure CCVS.

9.2.2 Merchant Accounts

If you're just setting up a merchant account or you're modifying an existing merchant account in order to use CCVS, your merchant account provider may want to see proof that CCVS can work with the protocol it uses. Certification letters for specific protocols are available at http://www.redhat.com/certifications.html. Print all pages of the letter corresponding to the protocol you will be using and show it to your merchant account provider.

Your merchant account provider must use one of the protocols supported by CCVS:

- First Data Corporation's ETC PLUS protocol (also known as FDR7, ETC+, ETC7, "Omaha")
- First Data Corporation's South Platform protocol (also known as "Nabanco")
- Global Payment Systems' MAPP protocol (also known as "St. Louis")
- Global Payment Systems' NDC protocol (also known as "Atlanta")
- Visa International's VITAL protocol (also known as VisaNet, Visa 2nd generation, "K format")
- Paymentech's UTF protocol (also known as GENSAR)
- NOVA Information Systems protocol

If your merchant account provider one of these protocols, you will be able to use CCVS.

Once you've identified which protocol you will be using, review the information applicable to that protocol at http://www.redhat.com/CCVS3.3docs/protocol-specific.html before you start the CCVS configuration process. The *CCVS Protocol Guide*, available from the link stated, describes the functionalities supported by different protocols.

9.2.3 Guidelines for Using CCVS On Your System

The following requirements allow CCVS to run smoothly and efficiently. Please make sure you are following all these guidelines before attempting to run CCVS.

Exclusive Use of the Modem(s) While CCVS is Running

Do not run other software applications that need to access the modem while you are running CCVS; they can interfere with CCVS's operations.

Permissions, Privileges and Access to the Modem

Most of the permissions needed for CCVS are set up for you during the installation process through the creation of a special group called "ccvs." However, there are issues involving system permissions of which you'll need to be aware.

All operations for a particular CCVS configuration must be performed from a single user account. One account is required so that all file ownerships and permissions are correctly set and protected. This user account must be added to the ccvs group (by you or by your system administrator) before you run the configuration program.

After the user has been added to the ccvs group, run the CCVS configuration program as that user. After you've run the configuration program, the same user must run the CCVS commands for that configuration.

If you want CCVS to run with a modem, the users in the ccvs group must also be added to the uucp group. Membership in the uucp group may not be sufficient for running the modems; if it isn't on your system, be sure that the ccvs group members also have access to the serial port for the modems that CCVS needs to use.

If you're using PHP with CCVS, you'll need to enable the Web server to run CCVS commands. To accomplish this, you'll must make the Web server user a member of the ccvs group. Usually, the Web server user will also need to be a member of the uucp group.

If you're not using PHP, but you want to make your Web server capable of running CCVS, you have other options (e.g., `suexec`, `setuid`) besides making the Web server user a member of the ccvs group. You can set it up any way you like, unless you're using PHP.

Software Versions

CCVS requires Tcl version 7.6 or greater to run the included GUI or to use the included Tcl/Tk APIs to develop your own graphical front end. Tcl version 8.3 is included in Red Hat Linux 7.0.

CCVS requires Perl version 5.0 or greater to use the included Perl APIs. Perl version 5.6 is included in Red Hat Linux 7.0.

9.3 Installing CCVS

The CCVS RPMs are available on the Linux Applications Library Workstation CD.

You can use RPM, Gnome-RPM or Kpackage to install the CCVS packages:

- `CCVS` — The core CCVS programs
- `CCVS-devel` — The C developer's kit
- `CCVS-perl` — The Perl interface for CCVS
- `CCVS-python` — The Python interface for CCVS
- `CCVS-php3` — The PHP3 interface for CCVS
- `CCVS-tcl` — The Tcl interface for CCVS
- `CCVS-java` — The Java interface for CCVS (included as source code)
- `CCVS-examples` — Sample source code, needed for development

9.4 Before You Configure CCVS

Before configuring CCVS, you need to be able to answer certain questions about your system and about how you want to set up CCVS. To prepare for the configuration process, be sure to follow these steps:

1. Please read through all documentation and errata that came with the program.
2. Fill out `setup.txt`. The `setup.txt` file is a worksheet which explains the different information needed when configuring CCVS to use particular protocols. If you fill out `setup.txt`, you'll have all of the information needed for the configuration process available at your fingertips. You can find

it in the `/usr/share/doc/CCVS-<version>` directory. Alternatively, `setup.txt` is also available at http://www.redhat.com/products/ccvs/support/CCVS3.3docs/setup.txt.

Please Note

On the setup worksheet, you'll be asked for some protocol-specific information. You only need to provide information for the protocol which you are going to use. You don't need to fill in the worksheet information for any of the other protocols.

---

3. The CCVS installation program will ask you several things about your modem, so be prepared with the appropriate information. Currently CCVS only documents the init strings for three modems:

Hayes Optima or ACCURA

```
\r~~~\rAT &D3 X4 E0 &K0 &Q0
```

U.S. Robotics Sportster or Courier

```
\r~~~\rAT E0 L0 M1 V1 X4 &K0 &M0 +FCLASS=0
```

Chase Research PCI-RAS

```
\r~~~\rAT E0 %C0 \\N0
```

If your modem does not appear on this list, look through your modem manual to find the string which turns off all compression and error correction, and the string which resets your modem for normal usage. You'll need to set these two modem strings during the configuration process.

9.5 Configuring CCVS

You must configure CCVS for your system, either in demo mode or for processing real data.

Use `su` to switch to the user account that you created (a member of the ccvs group) for this configuration.

Run the CCVS configuration program with the following command :

```
/usr/sbin/ccvs_configure
```

The rest of this section will walk you through the CCVS configuration program. You should see an entry "splash" screen. Press [Enter] to read the CCVS software license. You can use the standard scrolling and paging commands of `more` (or the paging program set by your $PAGER environment variable) to read the license.

When you have read the license and exited the pager, you will see:

```
Type "accept" to accept this license, or anything else to exit.
```

Type the word **`accept`** to accept the terms of the license and continue configuring CCVS. Any other input will exit the program.

You will then see this screen:

```
This program creates the configuration file for CCVS functions.
To do this, you will require the following information:
 1: The clearing protocol you will be using. This may be MAPP,
ETC+, or any of the other protocols which CCVS supports. There
is also a demo protocol; if you have downloaded the free demo of
CCVS, you will be using the demo protocol.
 2: The unique number which identifies you to the clearing
house. This may be your merchant account number or a terminal id
number, depending on what protocol you will be using. This number
will be supplied when you set up your merchant account.
 3: Your modem type, and the serial port your modem is attached
to. You will also need modem configuration strings. (We can
supply modem configuration strings for many popular modems.)
 4: The location of your data directory. This is where the
configuration file and data directories will be placed.
 5: Other information as needed for particular protocols. This
information will generally be supplied when you set up your
merchant account.

We supply a worksheet which you can use to organize all this
information, including the details for each protocol. See the
file "setup.txt" in /usr/share/doc/CCVS-<version>.

The configuration program is running as user "<username>".
```

```
It is important that this be the same user which the actual CCVS
software will run as. (We recommend creating a special user
account for just this purpose.)

Do you wish to continue configuring CCVS as user "<username>"?

[Enter Y to continue, or N to stop here:]
```

Press [Y] to continue. If you are su'd to root, you will instead get the following error. (If this happens, you should su to the CCVS user and re-run ccvs_configure.)

```
The configuration program may not be run as root. You must run
this as the same user which the actual CCVS software will run as.
(We recommend creating a special user account for just this
purpose.)
```

When you continue, the program will begin prompting you for information. At any time, you can back up to a previous prompt by typing **.** (a period) by itself and pressing [Enter].

```
Do you want to configure CCVS for the free demo, or a working
merchant account? (If you have not purchased a license for CCVS,
only the demo configuration is available.)

[Enter Y to use the demo configuration, N for a real configuration,
or . to exit:]
```

Unless you have purchased a software key and license for CCVS, type [Y]. This installs a demo configuration, which does not dial the modem or use a real merchant account. If you have purchased a license and are ready to install a working configuration, type [N].

```
Where do you want to place the CCVS configuration files and
transaction queues? This should be a directory name which is
writable by the current user.
The default is "/var/ccvs".
Enter directory, or Return for default value, or . by itself to
back up.
>
```

Unless you have specific reasons for moving the CCVS configuration files and transaction queues, leave them in their default locations. If you need to move them, remember that you'll also need to set an environment variable.

```
What do you want to name this configuration? This should be a
short filename.
The default is "ccvs".
Enter name, or Return for default value, or . by itself to back
up.
>
```

For example, you might have a configuration called **`tshirt`** for a merchant who sells T-shirts, and **`music`** for the sheet music retailer. The name entered here is the name used to distinguish between the two configurations.

The demo version of CCVS requires no other information; if you chose it, you will immediately see:

```
Writing "/var/ccvs/ccvs.conf"...

The CCVS system is now configured.
```

You can now begin testing the demo software. The demo acts just like the full CCVS software, except that it does not dial the modem or talk to a real merchant processor.

If you have a license for the full version of CCVS, and you chose to install a real configuration, you will instead see something like this:

```
Which protocol and merchant processor will you be using?

Credit card clearing protocols:
  1: ETC PLUS (FDR7/ETC7/FDR "Omaha"): First Data Corporation
  2: South Platform (FDR "Nabanco"): First Data Corporation
  3: MAPP: Global Payment Systems "St. Louis"
  4: NDC: Global Payment Systems "Atlanta" / NDC
  5: VITAL (Visa 2nd generation, K format): Visa / Total System Services
  6: UTF: Paymentech Inc.
  7: NOVA: NOVA Information Systems

[Enter a number, or . by itself to back up:]
```

Select the protocol for which you have a CCVS license and a valid merchant account.

```
What is the number of your merchant account?
Enter number, or . by itself to back up.
>
```

This number should have been provided with your merchant account.

```
What is your CCVS software customer number?
```

```
Enter number, or . by itself to back up.
>
```

This number will have been provided with your CCVS license.

```
What is your CCVS software license key?
Enter number, or . by itself to back up.
>
```

This number will also have been provided with your CCVS license.

```
What is the phone number of your merchant processor?
Enter number, or . by itself to back up.
>
```

Additional questions may appear, because they are required by particular protocols. If you've filled in the `setup.txt` worksheet section for your protocol, you should be prepared for these questions. For example, VITAL continues with several more prompts about your business's name, address, bank, and so on. You should already have found out this information when you established your VITAL merchant account.

You must now enter information about how to communicate with your modem. The modem configuration information is very important. Be sure that you enter correct information for your system's setup; CCVS will not work if the modem is set up incorrectly.

```
Do you want to configure a pool of several modems? (If you answer
yes, all the modems must be exactly the same make and model. If
you want to use just one modem, answer no.)

[Enter Y or N, or . to back up:]
```

If you have several identical modems, you can configure CCVS to use them all, as a pool. Each CCVS process which needs to use a modem will draw one from the pool, if any are available. Several CCVS configurations can share a set of modems this way. You can also configure a single configuration with two modems, so that authorizations and batch settlement can occur at the same time.

```
What serial port is your modem connected to? (Do not include the
"/dev/" prefix.) The default is ttyS0. The modem should be
connected and ready now, so that the serial port can be tested.

Enter port name, or Return for default value, or . by itself to
```

```
back up.
>
```

The program will test the serial port you enter; if you configure more than one, it will test each of them. Don't include the `/dev/`. This step may take up to thirty seconds if the modem does not respond.

```
What type of modem do you have? This information makes it
possible to suggest modem configuration strings. If your modem
is not listed, you can choose "none of the above"; but you will
then have to create your own configuration strings, which is a
difficult process.

 1: USR Sportster/Courier
 2: Hayes Optima
 3: Chase Research PCI-RAS
 4: None of the above

[Enter a number, or . by itself to back up:]
```

You will be prompted for the modem initialization, dialing, and hang-up strings. (If you configure a pool of modems, they must all be identical, so they will all use the same strings.) If CCVS knows appropriate strings for your modem, they'll be suggested and you can just press [Enter].

```
The modem initialization string should set the modem to do no
protocol
negotiation. What string do you want to use?
A string which works for your modem is:
 \r~~~\rAT E0 L0 M1 V1 X4 &K0 &M0 +FCLASS=0
Enter string, or Return for suggested value.
>

The modem dial string should dial the modem. (Do not include a
phone number.)
What string do you want to use?
A string which works for your modem is:
 ATDT
Enter string, or Return for suggested value.
>

The modem hang-up string should hang the modem up if it's
connected. What string do you want to use?
A string which works for your modem is:
```

```
 ~~~+++~~~~~~~~~~\rATH0\r~~~
Enter string, or Return for suggested value.
>

Initialize: \r~~~\rAT E0 L0 M1 V1 X4 &K0 &M0 +FCLASS=0
Dial: ATDT
Hang up: ~~~+++~~~~~~~~~~\rATH0\r~~~
Are these the values you want?

[Enter Y to accept, N to change, . to back up.]
```

You may not see exactly the same screen as shown above because the suggested defaults will vary depending on the modem you selected.

The next question is baud rate:

```
What baud rate do you want to use? You should use the
default unless you have explicit information that another
value is appropriate.
The default baud rate is 1200.

Enter rate, or Return for default value, or . by itself to
back up.
>
```

When you have finished entering configuration information, you will see:

```
Writing "/var/ccvs/ccvs.conf"...

The CCVS system is now configured.
```

9.6 Multiple Merchant Accounts

If you need to support more merchant accounts, simply follow the configuration procedure again. Use a different configuration name for each merchant account.

Different configurations may share the same serial port, or the same pool of serial ports. The modems will be used first-come-first-served.

9.7 Starting CCVS

To run CCVS for a particular application, you'll need to be su'd to the account which created that configuration.

As the user for the account, to run CCVS, you'll need to start the `ccvsd` daemon for each merchant account and you'll need to run the `cvupload` program on a regular basis (using `cron` to run `cvupload` every day is a good idea).

9.7.1 The ccvsd Daemon

To run CCVS, you must run the `ccvsd` daemon. The `ccvsd` daemon actually makes the phone calls and conducts the transactions. The `ccvsd` command must be followed by the name of the account, which you specified when you configured the account.

For example, if you want to start processing transactions for the sheet music retailer mentioned during the configuration program, and you installed the software in its default location of `/usr/sbin`, you would type in the following command to start `ccvsd`:

```
/usr/sbin/ccvsd music
```

Every time you add a merchant account, you need to start `ccvsd` for that account, if you want to process transactions for that account.

For more information on `ccvsd`, see the `ccvsd` man page.

9.7.2 The cvupload Command

Some transactions (such as authorizations) occur at the time that the credit card is presented. Other transactions (such as sales and returns) are saved up and are not processed immediately. These transactions are batched up and are then processed as a group.

CCVS uses the `cvupload` program to do this batch processing. We recommend invoking `cvupload` as an (at least) daily `cron` job, so that `cvupload` will automatically run every day, without any intervention on your part.

For example, the command to do the periodic processing for the sheet music retailer, we would issue the following command:

```
/usr/sbin/cvupload music
```

For more information on `cvupload`, see the `cvupload` man page.

9.8 Special Language Considerations

- C — The CCVS C library is included in the `CCVS-devel` package. When compiling C programs that use CCVS, add the `-lccvs` flag on the linkage line.
- Java — Please see http://www.redhat.com/CCVS3.3docs/AdminJava.html for more information on building the CCVS Java interface. The source code for the Java interface is provided in the `CCVS-java` package.
- Perl — The Perl interface is provided in the `CCVS-perl` package.
- Python — The Python interface is provided in the `CCVS-python` package.
- PHP — The `CCVS-php3` package provides the PHP3 interface.
- Tcl — The Tcl interface is included in the `CCVS-tcl` package.

9.9 Support for CCVS

Support for CCVS can be purchased from Red Hat. When you purchase your key to activate CCVS, be sure to review the support options available. See http://www.redhat.com/products/ccvs/ for more information about purchasing a key and support options.

If you do need support, be sure to have the following information available before you contact support:

- Your company name
- The version of CCVS you are using
- Your merchant number
- Your CCVS customer number
- Your operating system and version

Red Hat technical support will attempt to address any issues that deal directly with CCVS. We cannot support third party products, except for issues regarding integration with CCVS.

Part II Secure Web Server-Related Reference

10 Installing the Red Hat Linux Secure Web Server

10.1 Introduction

The next few chapters are intended to get you started running the Apache World Wide Web (WWW or Web) server version 1.3.12 with the mod_ssl security module and the OpenSSL library and toolkit. The combination of these three components, provided with Red Hat Linux, will be referred to in this manual as the Red Hat Linux Secure Web Server (or secure server, for short).

Web servers provide Web pages to browsers (e.g., Netscape Navigator, Microsoft Internet Explorer) who request them. In more technical terms, Web servers support the HyperText Transfer Protocol (HTTP), the Internet standard for Web communications. Using HTTP, a Web server sends HyperText Markup Language (HTML) Web pages and CGI and other types of scripts to browsers at the request of the browser. When users click on a link on a Web page, a request is sent to a Web server for the content named by the link. The Web server receives the request and provides the content that was asked for (e.g., an HTML page, an interactive script, a Web page dynamically generated from a database, etc.), or it sends back an error message. Apache, the Web server provided in this product, is the most widely used Web server on the Internet today (see http://www.netcraft.net/survey/).

The mod_ssl module is a security module for the Apache Web server. The mod_ssl module uses the tools provided by the OpenSSL Project to add a very important capability to Apache — the ability to encrypt communications. In contrast, using "regular" HTTP, communications between a browser and a Web server are sent in plain text, which could be intercepted and read along the route between the browser and the server.

The OpenSSL Project includes a toolkit which implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols and a general purpose cryptography library. The SSL protocol is used for secure data transmission over the Internet

today; the TLS protocol is a proposed Internet standard for private (secure) and reliable communications over the Internet. OpenSSL tools are used by the mod_ssl module to provide security for Web communications.

These chapters are not meant to be complete and exclusive documentation for any of these programs. When possible, this guide will point you to appropriate places where you can find more in-depth documentation on particular subjects.

This guide will show you how to install the included programs, as well as the basic options for configuring your Apache Web server. You will also be walked through the steps necessary to get a certificate from a Certificate Authority (CA), how to generate your own self-signed certificate, and how to install a certificate to use with your secure Web server.

10.2 Acknowledgments

The Red Hat Linux Secure Web Server includes the following:

- Software developed by the Apache Group for use in the Apache HTTP server project (http://www.apache.org/httpd.html)
- The mod_ssl security module, developed by Ralf S. Engelschall (http://www.modssl.org/)
- The OpenSSL toolkit, developed by Mark J. Cox, Ralf S. Engelschall, Dr. Stephen Henson, and Ben Laurie (http://www.openssl.org/)
- Software based on the Apache-SSL HTTP server project developed by Ben Laurie (http://www.apache-ssl.org/)
- Software based on SSLeay cryptographic software written by Eric Young and Tim Hudson

Red Hat gratefully acknowledges these contributions to this product.

10.3 Installation Overview

This chapter contains information about the Red Hat Linux Secure Web Server RPM packages and how to install them. Optional packages are also included with the Red

Hat Linux Secure Web Server; you can choose to install them or not, depending upon whether you want the functionality they offer.

You can install the Red Hat Linux Secure Web Server in three different ways, depending upon the configuration of your system. None of these procedures is difficult, but you should choose the correct one for your situation, depending upon how you're installing the Red Hat Linux operating system:

1. Installing Red Hat Linux using the installation program — since the Red Hat Linux Secure Web Server is included with the Red Hat Linux operating system, the easiest method is during the installation of Red Hat Linux. If you're about to do a new or "fresh" installation of Red Hat Linux, this is how you should install your secure server. See Section 10.5, *Installing the Red Hat Linux Secure Web Server During the Installation of Red Hat Linux* for more information on installing the Red Hat Linux Secure Web Server with a new installation of Red Hat Linux.

2. Upgrading Red Hat Linux using the installation program — if you already have a previous version of Red Hat Linux running on your system and you're upgrading to Red Hat Linux 7.0, you'll need to install the secure server packages during the upgrade process. See Section 10.7, *Upgrading from a Previous Version of Red Hat Linux* for important information on what you'll need to do if you're upgrading Red Hat Linux.

3. Installing the secure server after installing Red Hat Linux 7.0 — if you previously installed Red Hat Linux 7.0, and at a later date decide that you want to provide the secure server functionality, you can use the RPM Package Manager (RPM), Gnome-RPM or Kpackage to install the secure server packages from a Red Hat Linux CD.

Additionally, if you're upgrading from any version of Apache (including any previous version of the Red Hat Linux Secure Web Server), you'll need to know about certain issues concerning the upgrade process. See Section 10.6, *Upgrading from a Previous Version of Apache* before you begin the installation process, if you're upgrading Apache.

10.4 Choose Which Packages to Install

To install the secure server, you'll need to install three packages at minimum:

`apache`

The `apache` package contains the Apache Web server.

`mod_ssl`

The `mod_ssl` package includes the mod_ssl module, which provides strong cryptography for the Apache Web server via the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols.

`openssl`

The `openssl` package contains the OpenSSL toolkit. The OpenSSL toolkit implements the SSL and TLS protocols and also includes a general purpose cryptography library.

Additionally, other software packages included with Red Hat Linux can add functionality to your secure server (but are not required by the secure server to function):

OpenSSH

The `openssh` package provides the OpenSSH set of network connectivity tools for logging in to and executing commands on a remote machine. OpenSSH tools encrypt all traffic (including passwords), so you can avoid eavesdropping, connection hijacking, and other attacks on the communications between your machine and the remote machine.

The `openssh` package includes the OpenSSH clients: `ssh`, a secure replacement for `rsh`; `slogin`, a secure replacement for `rlogin` (remote login) and `telnet` (communications with another host via the TELNET protocol); and `scp`, a secure replacement for `rcp` (for copying files between machines) and `ftp` (for transferring files between machines).

The `openssh-askpass` package supports the display of a dialog window which prompts for a password during use of the OpenSSH agent with RSA authentication.

The `openssh-askpass-gnome` package contains a GNOME GUI desktop environment dialog window which is displayed when OpenSSH programs prompt for a password. If you're running GNOME and using OpenSSH utilities, you should install this package.

The `openssh-server` package contains the `sshd` secure shell daemon and man page. The secure shell daemon is the server side of the OpenSSH suite, and must be installed on your host if you want to allow SSH clients to connect to your host.

The `openssh-clients` package contains the client programs needed to make encrypted connections to SSH servers.

For more information about OpenSSH, see the OpenSSH website at http://www.openssh.com/.

Stunnel

The `stunnel` package provides the Stunnel SSL wrapper. Stunnel supports the SSL encryption of TCP connections, so it can provide encryption for non-SSL aware daemons and protocols (e.g., POP, IMAP, LDAP) without requiring any changes to the daemon's code.

`apache-devel`

The `apache-devel` package contains the Apache include files, header files and the APXS utility. You will need all of these if you intend to load any extra modules, other than the modules provided with this product. Please see Section 12.2, *Adding Modules to Your Server* for more information on loading modules into your Red Hat Linux Secure Web Server using Apache's DSO functionality.

If you do not intend to load other modules into your Red Hat Linux Secure Web Server, you do not need to install this package.

`apache-manual`

The `apache-manual` package contains the Apache Project's *Apache 1.3 User's Guide* in HTML format. This manual is also available on the Web at http://www.apache.org/docs/.

`openssl-devel`

The `openssl-devel` package contains the static libraries and include file necessary for compiling applications with support for various cryptographic algorithms and protocols. You need to install this package only if you're developing applications which include SSL support — you don't need this package to use SSL.

10.5 Installing the Red Hat Linux Secure Web Server During the Installation of Red Hat Linux

If you're installing Red Hat Linux and the Red Hat Linux Secure Web Server at the same time, follow the instructions provided in Chapter 15, *Installing Red Hat Linux via the GUI* to begin the installation of Red Hat Linux. Follow the instructions until you get to the point where you need to choose an installation class: workstation, server or custom.

1. If you choose a server-class installation, the secure server packages (`apache`, `mod_ssl` and `openssl` will be selected automatically. The `stunnel` and `openssh` packages, which provide security-related functionalities, will also be selected.
2. If you choose a workstation-class installation, the secure server packages and the security-related packages will not be automatically selected for installation, but you can choose to install them during the package selection customization process.
3. If you choose a custom-class installation, since you have complete control over what packages are installed, you'll need to select the secure server packages and any security-related packages you want.

Once you've chosen an installation class, continue following the installation instructions to partition and configure your system. When you reach the section on selecting package groups, or components, select the **Web Server** package group. **Web Server** includes the `apache` and `mod_ssl` packages that you must install to run the secure server. Since `openssl` is a dependency for the `mod_ssl` package, `openssl` will also be chosen for installation.

If you'd like to install any of the additional security-related packages described in Section 10.4, *Choose Which Packages to Install*, you'll need to identify their packages to the installation program. To do this, choose **Select individual packages** on the same **Package Group Selection** screen.

Select the security-related packages that you want to install according to the instructions provided in the *Official Red Hat Linux Installation Guide*. To help you find them, a table of their locations is provided as Table 10–1, *Security Packages*.

After making sure that the packages you need are selected, continue with the installation process.

10.6 Upgrading from a Previous Version of Apache

If you're upgrading Red Hat Linux and Apache (including any version of the Red Hat Linux Secure Web Server), you'll need to be aware of two issues:

- In the version of Apache included in Red Hat Linux 7.0, the `DocumentRoot` is `/var/www/html`.
- You may have customized your Apache configuration file (`httpd.conf`); you probably want to know what will happen to your customizations during the upgrade process (read on).

10.6.1 Where is the DocumentRoot?

Basically, the `DocumentRoot` is the directory on your system which holds most of the Web pages served by your Apache Web server. The `DocumentRoot` is set by a configuration directive in Apache's configuration file, `httpd.conf`. If you're unfamiliar with the `DocumentRoot` configuration directive, see Section 12.1.28, *`DocumentRoot`* for a more detailed explanation.

In previous versions of the Apache Web server shipped with Red Hat Linux, the `DocumentRoot` was `/home/httpd/html`. In the default (non-secure) version of Apache's configuration file, the `DocumentRoot` is `/usr/local/apache/htdocs`. It is also possible that you (or a predecessor) used an entirely different `DocumentRoot`. The important point is — in Red Hat Linux 7.0 the `DocumentRoot` is now, by default, `/var/www/html`.

Does this matter to you? It does, if you used a previous version of Apache to serve Web pages. Any Web pages that were previously served from a different `DocumentRoot` will not be found (or served) by the Apache shipped with Red Hat Linux 7.0 in its default configuration. You'll need to take one of the following steps:

Move all of the files in the old `DocumentRoot` (`/home/httpd/html`, `/usr/local/apache/htdocs`, or wherever) to the new one (`/var/www/html`).

or

Edit the Apache configuration file and change all references to the `DocumentRoot` back to the old directory path.

The solution you choose depends upon your system's configuration. Generally, if you automount `/home` on your system, you won't want to have your `DocumentRoot` in `/home`. On the other hand, if you don't have much space in `/var`, then you probably won't want your `DocumentRoot` in `/var`. You, or your system administrator, will have to decide the best solution based on your system's configuration and your Web server's needs. The Red Hat Linux Secure Web Server's default configuration is intended to address the needs of most webmasters; unfortunately, we can't configure it for everyone's individual situation.

10.6.2 What Happens to My Old Configuration File?

If you had another version of Apache installed and you customized its configuration files, the configuration files will be saved in their directory with an extension of `.rpmsave` during the installation of Apache. If you had another version of Apache installed but you never altered its configuration files, they will be written over during the installation of this product.

After installing Apache, you can cut and paste your customizations from your old Apache configuration file (`httpd.conf.rpmsave`) into the newly installed `httpd.conf` configuration file for your secure server.

10.7 Upgrading from a Previous Version of Red Hat Linux

If you're already running a previous version of Red Hat Linux on your system, you'll have to upgrade to Red Hat Linux 7.0 (instead of performing a full installation). If you need to upgrade, you must choose **Upgrade** instead of choosing an installation class (server, workstation or custom). Follow the instructions contained in the *Official Red Hat Linux Installation Guide* on how to upgrade your system. During the upgrade, you'll need to make sure that the secure server packages are selected by the installation program.

When you perform an upgrade to your Red Hat Linux system, the installation program checks to see what packages are already installed. Those packages will automatically be updated to the versions included in Red Hat Linux 7.0 during the upgrade process. Obviously, however, if you don't have a particular package already installed, the installation program will not install that package — unless you customize your upgrade.

If you're upgrading from the US/Canada version of Red Hat Linux Professional, you'll need to customize your upgrade and choose the secure server packages for installation. You may already have `apache` installed, but `mod_ssl` and `openssl` will not be installed (they were not included in Red Hat Linux before Red Hat Linux 7.0). You'll need to customize the upgrade to choose at least `mod_ssl` and `openssl`. See Section 10.7.1, *Customizing Your Upgrade to Install the Secure Server* for instructions on finding the packages you'll need to choose.

If you're upgrading from the International version of Red Hat Linux Professional, and you had the `apache`, `mod_ssl` and `openssl` packages installed, then the installation program will select and upgrade these programs automatically.

If you're upgrading from the International version of Red Hat Linux Professional, but you did not have the `apache`, `mod_ssl` or `openssl` packages installed, then you'll need to customize your upgrade and choose these packages for installation. See Section 10.7.1, *Customizing Your Upgrade to Install the Secure Server* for instructions on finding the packages you'll need to choose.

If you're upgrading from the Red Hat Linux Secure Web Server version 1.0 or 2.0 and you want to use your old key and certificate, you'll need to move them to the

right places. See Section 11.1, *Using Pre-existing Keys and Certificates* for more information on what to do with keys and certificates used with Red Hat Linux Secure Web Server versions 1.0 and 2.0.

10.7.1 Customizing Your Upgrade to Install the Secure Server

If you need to customize the upgrade process, follow the upgrading instructions contained in the *Official Red Hat Linux Installation Guide*; basically, choose **Upgrade** as your **Install Type** and then select **Customize packages to be upgraded.** Then you'll need to select the packages to upgrade, as described in the Official Red Hat Linux Installation Guide. To help you in your selection, Table 10–1, *Security Packages* provides the location of each secure server-related package and whether it is optional.

Table 10–1 Security Packages

Package Name	Located in Group	Optional?
`apache`	System Environment/Daemons	no
`mod_ssl`	System Environment/Daemons	no
`openssl`	System Environment/Libraries	no
`apache-devel`	Development/Libraries	yes
`apache-manual`	Documentation	yes
`openssh`	Applications/Internet	yes
`openssh-askpass`	Applications/Internet	yes
`openssh-askpass-gnome`	Applications/Internet	yes
`openssh-clients`	System Environment/Daemons	yes
`openssh-server`	System Environment/Daemons	yes
`openssl-devel`	Development/Libraries	yes
`stunnel`	Applications/Internet	yes

10.8 Installing the Secure Server After Installation of Red Hat Linux

If you installed Red Hat Linux 7.0 without installing the secure server packages, and then at a later date decide that you want to install the secure server, you can. The easiest way to do this is to use RPM, Gnome-RPM or Kpackage to install the RPM packages included on the Red Hat Linux CD.

The Red Hat Linux Secure Web Server is provided in RPM (RPM Package Manager) format. RPM is a software packaging system which makes it easy to install, uninstall, upgrade and query software packages. If you always use RPM to install software, RPM will keep track of the packages installed on your system and the files that they include.

10.8.1 Stop Any Running Web Server Processes

Before you begin this process, if you are running any Web server on your system, you must stop the server process before installing the Red Hat Linux Secure Web Server. If you are running an Apache Web server, stop the server process by issuing the appropriate command or commands, as root, from the following list:

```
/etc/rc.d/init.d/httpsd stop
/etc/rc.d/init.d/httpd stop
```

If you're running an Apache-based secure Web server, use the first command to stop the server process. If you're running a regular (non-secure) Apache Web server, use the second command. If you're running both, use both commands.

10.8.2 Using Gnome-RPM or Kpackage

If you're running GNOME or KDE, you can use a GUI program like Gnome-RPM or Kpackage to install the secure server packages. Alternatively, you can use RPM.

More information on how to use Gnome-RPM is included in Chapter 6, *Gnome-RPM* and in the *Official Red Hat Linux Getting Started Guide*. Instructions on how to use Kpackage are included on the *Kpackage Handbook* Web page at http://www.general.uwa.edu.au/u/toivo/kpackage/.

10.8.3 Using RPM

The Red Hat Linux Secure Web Server packages are provided in RPM format, so you can install the packages using RPM.

Mounting the CD-ROM

To begin the installation process, you must first mount the CD-ROM. Place the appropriate Red Hat Linux CD in your CD-ROM drive. As root, type the following command to mount the CD:

```
mount /mnt/cdrom
```

This command will work if you have an entry in your `/etc/fstab` file for the CD-ROM drive. If for some reason you get an error message after this command, try:

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

to mount the CD-ROM drive. Also, on your system, you or the system administrator may allow users (in addition to root) to mount the CD-ROM drive. You'll need to be root to use RPM to install the packages, in any case.

Once you've mounted the CD-ROM drive, the next step is to `cd` to the directory on the CD that contains the RPMs.

```
cd /mnt/cdrom/RedHat/RPMS
```

Once you're there, use RPM commands to install the packages you want. You'll need to install `apache`, `openssl` and `mod_ssl`.

For example, to install the `apache` package, become root (if you aren't already) and type in the following command:

```
# rpm -Uvh apache-1.3.12-3.i386.rpm
apache                  ##############################
```

The `apache` package will be installed. You'll need to repeat the previous command with each package you want to install.

Please Note

Complete instructions on how to use RPM are included in the Chapter 5, *Package Management with RPM*. (A condensed version of the same instructions is included in the *Official Red Hat Linux Getting Started Guide*.) RPM is a powerful and capable package management system. Please check the complete instructions on using RPM, if you have any questions about using it.

Once you have finished installing your packages, you'll need to unmount your CD-ROM. First, use `cd ..` to move one level above the `/mnt/cdrom` directory. Then, type `umount /mnt/cdrom` to unmount the CD-ROM. Type `eject /mnt/cdrom` and the CD-ROM drive will open so that you can remove the CD.

After you've installed the necessary packages, the next step is to create your key and obtain a certificate. Please continue to Chapter 11, *Obtaining a Certificate for your Secure Server* to create your key and certificate.

10.9 Finding Help and Documentation

If you followed the steps outlined in Chapter 10, *Installing the Red Hat Linux Secure Web Server* but you experienced a problem, the first thing you should do is check the Red Hat Errata section of the Red Hat website at http://www.redhat.com/support/errata.

If you purchased an Official Red Hat product which included support, you are entitled to technical support. Be sure to visit the Red Hat Support website at http://www.redhat.com/support to register for support.

10.9.1 Useful Sources of Information

Other sources of information about Apache and mod_ssl are available, including the following:

- The Tips, FAQs and HOWTO documents provided on the Red Hat website at http://www.redhat.com/support/docs/howto/

- The Red Hat Linux Apache Centralized Knowledgebase at http://www.redhat.com/support/docs/faqs/RH-apache-FAQ/book1.html
- Documentation was installed on your machine along with the Red Hat Linux Secure Web Server. After you've successfully installed the Red Hat Linux Secure Web Server, but before you've changed your home page, you can access the HTML documentation on your machine from your domain's URL (http://your_domain)
- The Apache website provides complete documentation for the Apache Web server at http://www.apache.org/docs/
- The mod_ssl website (http://www.modssl.org/ is the definitive source for information about mod_ssl. The website includes a wealth of documentation, including a *User Manual* at http://www.modssl.org/docs/.
- The redhat-secure-server mailing list. You can subscribe to this mailing list at http://www.redhat.com/community/list_subscribe.html

  You can also subscribe to the redhat-secure-server mailing list by e-mailing `redhat-secure-server-request@redhat.com` and include the word "subscribe" (without the quotation marks) in the Subject line.
- After installation of a package, you can find documentation for each package (if any exists) in `/usr/share/doc/<package_name> <version_number>/`.

10.10 How to Uninstall the Red Hat Linux Secure Web Server

If you need to uninstall your Red Hat Linux Secure Web Server, use these commands (as root): `rpm -e mod_ssl`, `rpm -e openssl` and `rpm -e apache`.

11 Obtaining a Certificate for your Secure Server

This chapter will guide you through the process of securing your server by obtaining and installing a certificate.

To make your customers feel safe doing business with you over the Web, your Web server needs to be secure. Secure servers use the **Secure Sockets Layer** (SSL) protocol, which encrypts the data sent back and forth between a browser and the server. When your browser is communicating using SSL, you'll see the https: prefix before the Uniform Resource Locator (URL) in the navigation bar.

Customers feel more comfortable when making purchases from websites if they know that their transactions are secure, but secure servers aren't used only for electronic commerce. A secure server may also be used to transmit sensitive data such as sales figures to sales people on the road or to business partners over the Internet.

A secure server uses a certificate to identify itself to Web browsers. You can generate your own certificate (called a "self-signed" certificate) or you can get a certificate from a Certificate Authority or CA. A certificate from a reputable CA guarantees that a website is associated with a particular company or organization.

If your server will be used for e-commerce, you'll probably want to purchase a certificate from a CA. A CA's certificate provides two advantages: (usually) browsers will recognize it automatically and the CA guarantees the identity of the organization responsible for the website. Self-signed certificates will not be automatically accepted by a user's browser — the user will be asked by the browser if they want to accept the certificate and create the secure connection.

When you use a CA-signed certificate, you guarantee the identity of the organization running the server. For example, if the certificate says the website is Red Hat's, and the user trusts the CA, then there is no reason to doubt that any files or programs downloaded from that site really are from Red Hat.

The first step is to create a public and private key pair. Then you'll need to either create a certificate request (CSR) to send to a CA or you'll need to create your own

self-signed certificate. This chapter provides instructions on how to get signed certificates from VeriSign (http://www.verisign.com or http://www.verisign.com/offer/redhat/ for details on a VeriSign discount to Red Hat customers) and Thawte (http://www.thawte.com), and how to generate your own certificate.

Please Note

You can get signed certificates from any CA that you choose, and not just the CAs that are mentioned in this manual. However, VeriSign is offering a discount on certificates to Red Hat customers. See http://www.verisign.com/offer/redhat for details on VeriSign's discount.

Once you have a self-signed certificate or a signed certificate from the CA of your choice, you'll learn how to install it on your Red Hat Linux Secure Web Server.

11.1 Using Pre-existing Keys and Certificates

If you already have an existing key and certificate (for example, if you're installing the Red Hat Linux Secure Web Server to replace another company's secure Web server product), you will probably be able to use your existing key and certificate with the Red Hat Linux Secure Web Server. In the following two situations, you will not be able to use your existing key and certificate:

- if you're changing your IP address or domain name
- if you have a certificate from VeriSign and you're changing your server software

You can't use your old key and certificate if you are changing your IP address or domain name. Certificates are issued for a particular IP address and domain name pair. You will need to get a new certificate if you're changing your IP address or domain name.

VeriSign is a widely used CA. If you already have a VeriSign certificate for another purpose, you may have been considering using your existing VeriSign certificate with

your new Red Hat Linux Secure Web Server. However, you won't be allowed to, because VeriSign issues certificates for one particular server software and IP address/domain name combination.

If you change either of those parameters (for example, if you previously used another secure Web server product and now you want to use the Red Hat Linux Secure Web Server), the VeriSign certificate you obtained to use with the previous configuration will not work with the new configuration. You will need to obtain a new certificate.

If you have an existing key and certificate that you can use, you will not have to follow the instructions contained in Chapter 11, *Obtaining a Certificate for your Secure Server*. You do need to move and rename the files which contain your key and certificate.

Move your existing key file to:

```
/etc/httpd/conf/ssl.key/server.key
```

Move your existing certificate file to:

```
/etc/httpd/conf/ssl.crt/server.crt
```

After you've moved your key and certificate, skip to Section 11.10, *Testing Your Certificate*.

If you're upgrading from the Red Hat Secure Web Server versions 1.0 and 2.0, your old key (`httpsd.key`) and certificate (`httpsd.crt`) will be located in `/etc/httpd/conf/`. You'll need to move and rename your key and certificate, so that the Red Hat Linux Secure Web Server can use them. Use the following two commands to move and rename your key and certificate files:

```
mv /etc/httpd/conf/httpsd.key /etc/httpd/conf/ssl.key/server.key
mv /etc/httpd/conf/httpsd.crt /etc/httpd/conf/ssl.crt/server.crt
```

Then start your Red Hat Linux Secure Web Server as described in Section 11.11, *Starting and Stopping Apache*. You should not need to get a new certificate, if you are upgrading from a previous version of the Red Hat Linux Secure Web Server.

11.2 A General Overview of Web Server Security

Your Red Hat Linux Secure Web Server provides security using a combination of the Secure Sockets Layer (SSL) protocol and (in most cases) CA-approved digital certificates. SSL handles the encrypted communications and the mutual authentication between browsers and your Red Hat Linux Secure Web Server. The CA-approved digital certificate provides authentication for your Red Hat Linux Secure Web Server (the CA puts its reputation behind its certification of your organization's identity).

Encryption depends upon the use of keys (think of them as secret encoder/decoder rings in data format). In conventional or symmetric cryptography, both ends of the transaction have the same key, which they use to decode each other's transmissions. In public or asymmetric cryptography, two keys co-exist: a public key and a private key. A person or an organization keeps their private key a secret, and publishes their public key. Data encoded with the public key can only be decoded with the private key; data encoded with the private key can only be decoded with the public key.

You'll use public cryptography to create a public and private key pair. In most cases, you'll send your certificate request (including your public key), proof of your company's identity and payment to a CA. The CA verifies the certificate request and your identity, and then sends back a certificate for your Red Hat Linux Secure Web Server.

Alternatively, you can create your own self-signed certificate. Note, however, that self-signed certificates should not be used in production environments. See Section 11.3, *Types of Certificates* for more information on the differences between self-signed and CA-signed certificates.

11.3 Types of Certificates

If you installed your Red Hat Linux Secure Web Server using the Red Hat Linux installation program, a random key and a test certificate are generated and put into the appropriate directories. Before you begin using your secure server, however, you'll need to generate your own key and obtain a certificate which correctly identifies your server.

You need a key and a certificate to operate your Red Hat Linux Secure Web Server — you can either generate a self-signed certificate or purchase a CA-signed certificate from a CA. What are the differences between the two?

A CA-signed certificate provides two important capabilities for your server:

- Browsers will (usually) automatically recognize the certificate and allow a secure connection to be made, without prompting the user.
- When a CA issues a signed certificate, they are guaranteeing the identity of the organization that is providing the Web pages to the browser.

You can generate a self-signed certificate for your Red Hat Linux Secure Web Server, but be aware that a self-signed certificate will not provide the same functionalities as a CA-signed certificate. A self-signed certificate will not be automatically recognized by users' browsers, and a self-signed certificate does not provide any guarantee for the identity of the organization that is providing the website. A CA-signed certificate provides both of these important capabilities for a secure server. If your secure server will be used in a production environment, you'll probably need a CA-signed certificate.

If your secure server is being accessed by the public at large, your Red Hat Linux Secure Web Server needs a certificate signed by a CA, so that people who visit your website can rely that the website is owned by the organization who claims to own it. Before signing a certificate, a CA verifies that the organization requesting the certificate was actually who they claimed to be.

Most Web browsers that support SSL have a list of CAs whose certificates they will automatically accept. If a browser encounters a certificate whose authorizing CA is not in the list, the browser will ask the user to choose whether to accept or decline the connection.

The process of getting a certificate from a CA is fairly easy. A quick overview is as follows:

- Create an encryption private and public key pair.
- Create a certificate request based on the public key. The certificate request contains information about your server and the company hosting it.

- Send the certificate request, along with documents proving your identity, to a CA.
- When the CA is satisfied that you are indeed who you claim to be, they will send you a digital certificate.
- Install this certificate on your Web server, and begin handling secure transactions.

11.4 Deciding on a Certificate Authority

We can't tell you which certificate authority to choose. Your decision may be based on your past experiences, or on the experiences of your friends or colleagues, or purely on monetary factors. We will guide you through the process of getting a digital certificate from VeriSign and from Thawte, but you can choose a different CA.

Many other CAs exist. Click on the **Security** button on your Navigator toolbar or on the padlock icon at the bottom left of the screen, then click on **Signers** to see a list of certificate signers from whom your browser will accept certificates. You can also search the Web for CAs. The process of getting a certificate from a different CA will be similar to the processes described in this manual.

11.4.1 VeriSign Certificate Packages

VeriSign offers discounts on its certificate offerings to Red Hat customers. To take advantage of this offer, refer to http://www.verisign.com/offer/redhat/.

VeriSign and Thawte offer several tiers of server certificate options, as listed below. Check the appropriate website for completely up-to-date information.

Commerce Site Services

Commerce Site (with 40-bit encryption and Payflow Pro online payment management services)

Commerce Site Pro (with 128-bit encryption and Payflow Pro online payment management services)

Secure Site Services

Secure Site (with 40-bit encryption)

Secure Site Pro (with 128-bit encryption)

VeriSign's 128-bit SSL IDs enable the world's strongest encryption technology for Web servers on both domestic and export versions of Microsoft and Netscape browsers.

VeriSign's solutions each include an SSL server certificate (or "Server ID") plus other features, including:

- The NetSure Protection Plan, an extended warranty program that protects VeriSign customers against up to $250,000 of economic loss resulting from the theft, corruption, impersonation, or loss of use of a certificate.
- The VeriSign Secure Site Seal, allowing visitors to check your Server ID's information and status in real time.
- Commerce Site Services and Secure Site Services also include Keynote Site Performance Measurement Services, Netcraft E-Commerce Security Analysis, and Qualys' network scanning service to determine your site's vulnerabilities.
- Training Discounts toward the VeriSign course "Building Secure Web Servers."
- An audited authentication process, which ensures that VeriSign verifies the identity of every site equipped with a Server ID.

For more information about VeriSign's Server ID solutions, see http://www.verisign.com/server/index.html. For more information about Thawte SSL Server Certificates, see http://www.thawte.com/certs/server/contents.html.

11.5 Proving Your Organization's Identity to a CA

When you request a signed certificate from a CA, you'll need to prove that your organization has the right to conduct business using your organization's name. CAs are very specific about their requirements for proving your identity, and you'll need to check with the CA of your choice to see what their requirements are.

In some cases, copies of the following documents will need to be mailed or faxed to the CA, and your certificate will not be issued until the documents have been received and verified by the CA.

11.5.1 Proving Your Organization's Identity to VeriSign

The easiest way to prove to VeriSign that your organization has the right to do business is to provide your Dun & Bradstreet (D-U-N-S) number. If you don't have a D-U-N-S number, you can request one from the Dun & Bradstreet website at http://www.dnb.com/dunsno/whereduns.htm.

If you don't know whether you have a D-U-N-S number, you can find out if you have one from VeriSign at https://digitalid.verisign.com/dnb_query.htm.

If you don't have a D-U-N-S number and you don't want to get one, see http://www.verisign.com/server/enroll/globalpreparing.html#proof for the documentation you'll need to provide to VeriSign.

11.5.2 Proving Your Organization's Identity to Thawte

See http://www.thawte.com/certs/server/docs.html for a list of what Thawte requires to prove your organization's identity. At the time this document was written, you needed to provide the following:

1. Proof of organizational name
2. Proof of the right to a domain name

"Proof of Organizational Name" means that you have to prove your right to use your company's or organization's name. This proof may consist of a copy of your official company's registration documents or a copy of your certificate of incorporation in your state or country. A number of other documents can also be used to prove your organizational name; see http://www.thawte.com/certs/server/right_name.html for many more examples.

"Proof of the Right to a Domain Name" is unnecessary if your domain is registered exactly to the company name included in your certificate request (which you haven't created yet). In most cases, this will be the case. Run a `whois` on your domain to see your domain's registration information. See http://www.thawte.com/certs/server/right_domain.html for more information on how you might prove your right to a domain name, if your certificate request will not exactly match the information returned by a `whois` on your domain.

Once you've gathered the information you'll need to prove your organization's identity to a CA, you can go on to creating a key and certificate request.

11.6 Generating a Key

First, you'll need to remove the key and certificate that were generated during the installation. `cd` to the `/etc/httpd/conf` directory.

Use the following commands to remove the two files:

```
rm ssl.key/server.key
```

and

```
rm ssl.crt/server.crt
```

The first step towards creating a certificate is to create your own random key. Type in the following command, which will generate your key:

```
make genkey
```

Your system will display a message similar to the following:

```
umask 77 ; \
/usr/bin/openssl genrsa -des3 1024 > /etc/httpd/conf/ssl.key/server.key
Generating RSA private key, 1024 bit long modulus
.......++++++
.....................................................................++++++
e is 65537 (0x10001)
Enter PEM pass phrase:
```

You now need to type in a password. For best security, your password should be at least eight characters, should include numbers or punctuation, and should not be a word in a dictionary. Also, remember that your password is case sensitive.

Please Note

You will need to remember and enter this password every time you start your Red Hat Linux Secure Web Server, so don't forget it.

You will be asked to re-type the password, to verify that it's correct. Once you've typed it in correctly, a file called `server.key`, containing your key, will be created.

Note that if you don't want to type in a password every time you start your Red Hat Linux Secure Web Server, you will need to use the following two commands instead of `make genkey` to create the key. Both of these commands should be typed in entirely on one line.

Use the following command:

```
/usr/bin/openssl genrsa 1024 > /etc/httpd/conf/ssl.key/server.key
```

to create your key. Then use this command:

```
chmod go-rwx /etc/httpd/conf/ssl.key/server.key
```

to make sure that the permissions are set correctly on your key.

After you use the above commands to create your key, you will not need to use a password to start your Red Hat Linux Secure Web Server.

Disabling the password feature for your secure Web server is a security risk. We DO NOT recommend that you disable the password feature for your Red Hat Linux Secure Web Server.

The problems associated with not using a password are directly related to the security maintained on the host machine. For example, if an unscrupulous individual compromises the regular UNIX security on the host machine, that person could obtain your private key (the contents of your `server.key` file). The key could be used to "spoof" Web pages that will appear to be coming from your site.

If UNIX security practices are scrupulously being followed for the host computer (i.e., all operating system patches and updates are installed as soon as they're available, no unnecessary or risky services are operating, and so on), the Red Hat Linux Secure Web Server's password may seem unnecessary. However, since your Red Hat

Linux Secure Web Server shouldn't need to be re-booted very often, the extra security provided by entering a password is a worthwhile effort in most cases.

The `server.key` file should be owned by the root user on your system and should not be accessible to any other user. Make a backup copy of this file and keep the backup copy in a safe, secure place. You need the backup copy because if you ever lose the `server.key` file after using it to create your certificate request, your certificate will no longer work and the CA will not be able to help you. Your only option would be to request (and pay for) a new certificate.

If you're going to purchase a certificate from a CA, continue to Section 11.7, *Generating a Certificate Request to Send to a CA*. If you're generating your own self-signed certificate, continue to Section 11.9, *Creating a Self-Signed Certificate*.

11.7 Generating a Certificate Request to Send to a CA

Once you've created a key, the next step is to generate a certificate request which you will need to send to the CA of your choice. Type in the following command:

```
make certreq
```

Your system will display the following output and will ask you for your password (unless you disabled the password option):

```
umask 77 ; \
/usr/bin/openssl req -new -key /etc/httpd/conf/ssl.key/server.key
-out /etc/httpd/conf/ssl.csr/server.csr
Using configuration from /usr/share/ssl/openssl.cnf
Enter PEM pass phrase:
```

Type in the password that you chose when you were generating your key. Your system will display some instructions and then ask for a series of responses from you. Your inputs will be incorporated into the certificate request. The display, with example responses, will look like this:

```
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a
DN.
```

```
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:North Carolina
Locality Name (eg, city) []:Durham
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Test Company
Organizational Unit Name (eg, section) []:Testing
Common Name (eg, your name or your server's hostname) []:test.mydomain.com
Email Address []:admin@mydomain.com
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
```

The default answers appear in brackets `[]` immediately after each request for input. For example, the first information required is the name of the country where the certificate will be used, shown like the following:

```
Country Name (2 letter code) [AU]:
```

The default input, in brackets, is **AU**. To accept the default, just press [Enter], or fill in your countries' two letter code.

You will have to type in the rest of the inputs (`State or Province Name`, `Locality Name`, `Organization Name`, `Organizational Unit Name`, `Common Name`, and `Email address`). All of these should be self-explanatory, but you need to follow these guidelines:

- Do not abbreviate the locality or state. Write them out (for example, St. Louis should be written out as Saint Louis).
- If you're sending this CSR to a CA, be very careful to provide correct information for all of the fields, but especially for the `Organization Name` and the `Common Name`. CAs check the information provided in the CSR to determine whether your organization is responsible for what you provided as the `Common Name`. CAs will reject CSRs which include information they perceive as invalid.

- For `Common Name`, make sure you type in the *real* name of your Red Hat Linux Secure Web Server (a valid DNS name) and not any aliases which the server may have.
- The `Email Address` should be the e-mail address for the webmaster or system administrator.
- Avoid any special characters like @, #, &, !, etc. Some CAs will reject a certificate request which contains a special character. So, if your company name includes an ampersand (&), spell it out as "and" instead of "&."
- Don't use either of the extra attributes (`A challenge password` and `An optional company name`). To continue without entering these fields, just press [Enter] to accept the blank default for both inputs.

When you've finished entering your information, a file named `server.csr` will be created. `server.csr` is your certificate request, ready to send to your CA.

11.8 Buying a Certificate

After you've generated a CSR, you need to send it to a CA.

11.8.1 Purchasing a Certificate From VeriSign

First, you'll need to decide which certificate you'd like to purchase. See Section 11.4.1, *VeriSign Certificate Packages* for a description of VeriSign's certificate products.

Once you've decided on the one you want, go to http://www.verisign.com/server/. Select the certificate that you want to buy.

The purchase process is somewhat standardized. In this document, we'll go step-by-step through the process of purchasing a Secure Site certificate, but you should be able to use the instructions to purchase another type of certificate.

1. For Secure Site certificates, you can choose a **Two Year Option** (if that's what you want) and confirm the location of your secure server. Click on **Continue** when you're finished.

2. The next page is **Preparing for Enrollment**. This page provides a summary of the information that you'll need to provide to VeriSign. Read through this page and make sure that you have the necessary information ready before you continue with the application process. When you're finished, click on the **Continue** button at the bottom of the page.

3. The next page is **CSR: Wizard: Verify Distinguished Name**. If you have not already generated a key and CSR according to the instructions provided in Section 11.6, *Generating a Key* and Section 11.7, *Generating a Certificate Request to Send to a CA*, do so now. Then select **I have already prepared a CSR for this enrollment** and click on **Continue**.

4. The next step, **Submit CSR**, is shown as Figure 11–1, *Submitting a Certificate Request to VeriSign*. Select **Red Hat** from the **Server Software Vendor** list.

Figure 11–1 Submitting a Certificate Request to VeriSign

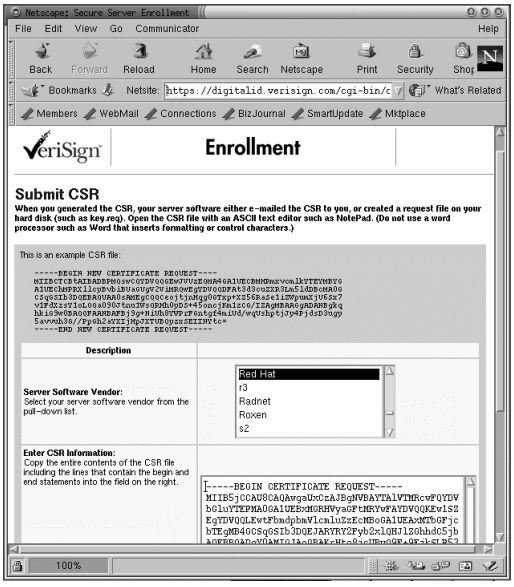

Paste the contents of your CSR into the **Enter CSR Information** text box. To cut and paste the contents of the file (if you're using X), first `cd` to `/etc/httpd/conf/ssl.csr`. Display the contents of `server.csr` with

the `cat server.csr` command. Highlight the contents of the file by clicking and dragging with your left mouse button. Left click on the text box on the Web page. Click the middle mouse button to paste the highlighted text.

When you're copying and pasting the CSR, be careful not to copy any extra empty or white spaces before or after the text (including the `-----BEGIN CERTIFICATE REQUEST-----` and `-----END CERTIFICATE REQUEST-----` lines). CAs have been known to reject CSRs which include those unwelcome spaces.

After you've successfully pasted in the CSR, click on **Continue.**

5. The next step is to **Provide Proof of Right.** This means that you need to prove to VeriSign that your organization is legitimate. VeriSign first attempts to match the organizational name you provided in the Dun & Bradstreet database. If your organization was found, select it. If your organization was not found, select **My company and/or my company's correct address is not displayed in this list.** Click on **Continue.**

   The easiest way to prove your organization's identity to VeriSign is to provide them with your D-U-N-S number, but there are other ways if you don't have a D-U-N-S number or you don't want to use one. Refer to the instructions provided by VeriSign if you need to prove your organization's identity with something other than a D-U-N-S number. You'll need the proof, ready for submission to VeriSign, before you can apply for a certificate. Once you have the required documents, continue with the enrollment process.

6. After selecting the correct organization from the Dun & Bradstreet database list and clicking on **Continue**, the next page is **Confirm Domain Registration.** On this page, VeriSign is checking to see if your domain is registered to your organization. For more information on registering a domain name, see the InterNIC FAQ at http://www.internic.net/faq.html and/or ask your network administrator.

   Your domain name must be registered to your organization. So, the **Organization name listed in domain registry** should be the same as the **Organization name you entered.** If they are not the same, you'll probably need to create a new CSR which includes the correct information.

In most cases, the two fields will be the same, so you can select **These organization names match** and then click on **Continue**.

7. The next page should congratulate you on passing VeriSign's initial validation checks. Click on **Continue**.

8. The next page, **Complete Application**, is shown as Figure 11–2, *Application for VeriSign Certificate*.

Figure 11–2 Application for VeriSign Certificate

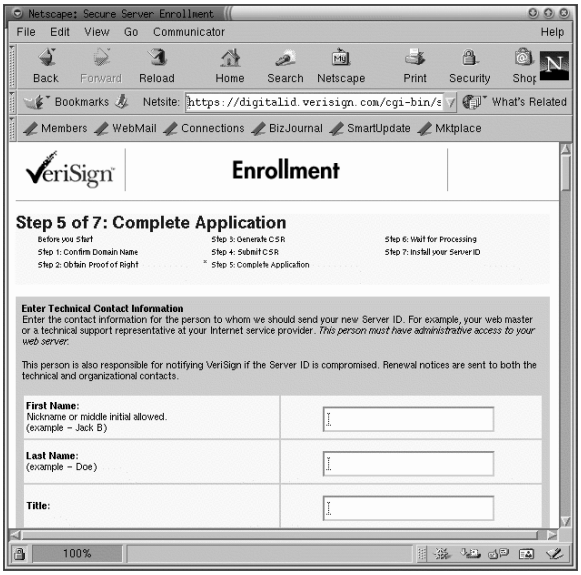

Fill in the **Enter Technical Contact Information** section with information about your Red Hat Linux Secure Web Server's administrator or webmaster.

Fill in the **Enter Organizational Contact Information** section with the appropriate information, according to the instructions provided by VeriSign.

Fill in the **Enter Billing Contact Information** with information for the person who will be contacted for billing purposes.

Type a "challenge phrase" and a "reminder question" into the area provided. You may be asked for your challenge phrase if you ever need support from VeriSign, so be sure to record it and keep it someplace safe.

Indicate how you are going to pay for your certificate.

Read the subscriber agreement at the bottom of the page. After you've read the agreement, click on the **Continue** button at the bottom of the page. Your application will be submitted.

After you've successfully completed your enrollment form and your information and payment has been provided to VeriSign, they will authenticate your organization's identity and issue your certificate. When your application has been approved, they will send your certificate by e-mail to the technical and organizational contacts you provided.

Save the certificate VeriSign sends you in the file `server.crt` in `/etc/httpd/conf/ssl.crt/`. Follow the steps outlined in Section 11.10, *Testing Your Certificate* to install your certificate.

11.8.2 Purchasing a Certificate From Thawte

To purchase a certificate from Thawte, follow these instructions:

1. Point your browser to http://www.thawte.com/certs/server/request.html, where Thawte provides an overview of the necessary steps.

2. The first thing you need to do is gather the documents that they require, as discussed in both Section 11.5.2, *Proving Your Organization's Identity to Thawte* and the aforementioned Web page.

3. The next step is to generate a key and a certificate signing request (CSR). If you followed the instructions contained in Section 11.6, *Generating a Key* and Section 11.7, *Generating a Certificate Request to Send to a CA*, you already have a key (`/etc/httpd/conf/ssl.key/server.key`) and a CSR (`/etc/httpd/conf/ssl.csr/server.csr`). If you did not already create your key and certificate request, do so now using the instructions provided in this document.

4. Go to Thawte's **Buy a Certificate** Web page at https://www.thawte.com/cgi/server/step1.exe. Select **SSL Server Certificate**. Click on the **Next** button at the bottom of the page.

5. The next page is **Server Cert Enrollment**. Paste the contents of your `/etc/httpd/conf/ssl.csr/server.csr` file into the **Certificate Signing Request (CSR)** text box, shown in Figure 11–3, *Thawte Enrollment Form*.

Figure 11–3 Thawte Enrollment Form

To cut and paste the contents of the file (if you're using X), first `cd` to `/etc/httpd/conf/ssl.csr`. Display the contents of `server.csr` with the `cat server.csr` command. Highlight the contents of the file by clicking and dragging with your left mouse button. Left click on the text box on the Web page. Click the middle mouse button to paste the highlighted text.

When you're copying and pasting the CSR, be careful not to copy any extra empty or white spaces before or after the text (including the `-----BEGIN CERTIFICATE REQUEST-----` and `-----END CERTIFICATE REQUEST-----`

lines). CAs have been known to reject CSRs which include those unwelcome spaces.

6. Choose **Red Hat Secure Server** from the **Web Server Software** pull-down menu.
7. Choose how you want to pay for the certificate.
8. Click on **Next** at the bottom of the page.
9. The next page displays an **Analysis of Certificate Signing Request**, shown as Figure 11–4, *Analysis of CSR*.

Figure 11–4 Analysis of CSR

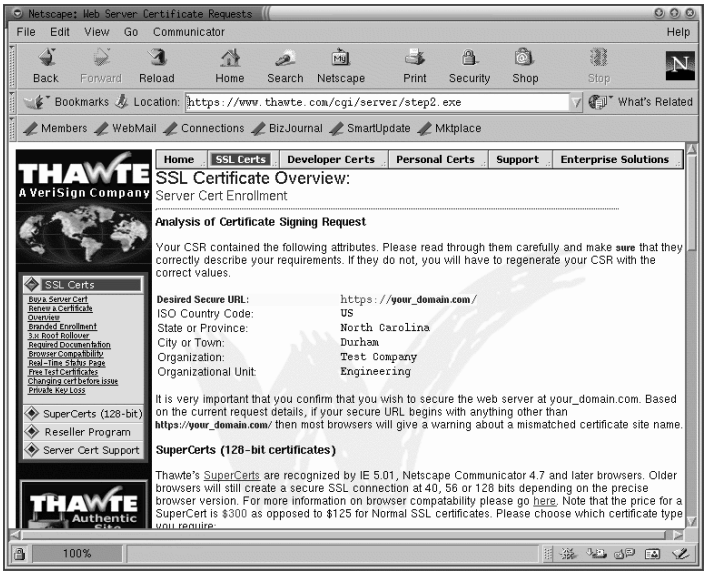

Scroll down the page to **Background Information**, where you need to select a description for your organization from the pull-down menu, or type your own description into the text box provided.

10. If you have a D-U-N-S number, type it into the text box under **DUNS Number**.

11. Review Thawte's **Subscriber Agreement**. Fill in the required information for the person in your organization who will be authorize the **Subscriber Agreement**, as described in Section 11.5.2, *Proving Your Organization's Identity to Thawte*.

12. Under **Technical Contact/Webmaster**, fill in contact information for your Red Hat Linux Secure Web Server's administrator or webmaster.

13. Click on the **Next** button at the bottom of the page.

14. The next page, also entitled **Server Cert Enrollment**, is the last page of their enrollment form and is shown as Figure 11–5, *Thawte CSR Application*. From the first pull-down menu, choose the currency in which you are going to pay Thawte.

Figure 11–5 Thawte CSR Application

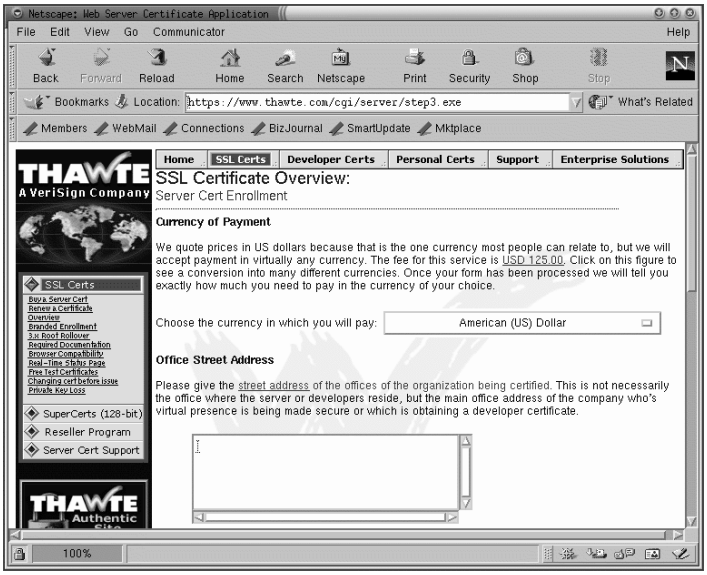

15. Type the street address for your organization into the **Office Street Address** text box.

16. Type in your organization's fax number into the text box under **Office Fax Number**.

17. From the pull-down menu under **Nearest Thawte Office**, choose the Thawte office closest to your organization.

18. Type a password or challenge phrase into the text box under **Privacy Protection Password**. After you've submitted your application, you'll be able to check on its status on the Web.

19. Click on **Next** at the bottom of the page.

20. The next page will indicate that your submission is complete. This page provides you with a tracking number for your application, so that you can monitor its status over the Web.

21. After Thawte receives your documentation and payment, your certificate should be issued by e-mail. When you receive your certificate, save it into the `/etc/httpd/conf/ssl.crt/server.crt` file. See Section 11.10, *Testing Your Certificate*, for instructions on installing your certificate.

11.9 Creating a Self-Signed Certificate

You can create your own self-signed certificate. Please note that a self-signed certificate will not provide the security guarantees provided by a CA-signed certificate. See Section 11.3, *Types of Certificates* for more details about certificates.

If you'd like to make your own self-signed certificate, you'll first need to create a random key using the instructions provided in Section 11.6, *Generating a Key*. Once you have a key, use the following command:

```
make testcert
```

You'll see the following output and you'll be prompted for your password (unless you generated a key without a password):

```
umask 77 ; \
/usr/bin/openssl req -new -key /etc/httpd/conf/ssl.key/server.key
-x509 -days 365 -out /etc/httpd/conf/ssl.crt/server.crt
Using configuration from /usr/share/ssl/openssl.cnf
Enter PEM pass phrase:
```

After you enter your password (or without a prompt if you created a key without a password), you'll be asked for more information. The computer's output and a set of inputs looks like the following (you'll need to provide the correct information for your organization and host):

```
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a
DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:North Carolina
Locality Name (eg, city) []:Durham
Organization Name (eg, company) [Internet Widgits Pty Ltd]:My Company, Inc.
Organizational Unit Name (eg, section) []:Documentation
Common Name (eg, your name or your server's hostname) []:myhost.mydomain.com
Email Address []:myemail@mydomain.com
```

After you provide the correct information, a self-signed certificate will be created and placed in `/etc/httpd/conf/ssl.crt/server.crt`. You'll need to re-start your secure server after generating the certificate. See Section 11.11, *Starting and Stopping Apache* for instructions on stopping and starting your secure Web server.

11.10 Testing Your Certificate

When the secure server is installed by the Red Hat Linux installation program, a random key and a generic certificate are installed, for testing purposes. You can connect to your secure server using this certificate. For any purposes other than testing, however, you need to get a certificate from a CA or generate a self-signed certificate. See Section 11.3, *Types of Certificates* if you need more information on the different types of certificates available.

If you've followed the instructions provided in this guide to either purchase a certificate from a CA or generate a self-signed certificate, you should have a file named `/etc/httpd/conf/ssl.key/server.key`, containing your key, and a file named `/etc/httpd/conf/ssl.crt/server.crt`, containing your test certificate. If your key and certificate are somewhere else, move them to these directories. If you changed any of the default locations or filenames for the Red Hat Linux Secure Web Server in your Apache configuration files, you should put these two files in the appropriate directory, based on your modifications.

Now stop and start your server as described in Section 11.11, *Starting and Stopping Apache*. If your key file is encrypted, you will be asked for the password. Type in your password and your server should start.

Point your Web browser to your server's home page. The URL to access your Red Hat Linux Secure Web Server will look like this:

```
https://your_domain
```

Please Note

Note the "s" after "http." The https: prefix is used for secure HTTP transactions. If the connection is made, you should see a dialog box indicating that your browser must be configured to accept the test certificate.

If you're using a CA-signed certificate from a well-known CA, your browser will probably automatically accept the certificate (without prompting you) and create the secure connection. Your browser will not automatically recognize a test or a self-signed certificate, because the certificate is not signed by a CA. If you're not using a certificate from a CA, follow the instructions provided by your browser to accept the certificate. You can just accept the defaults by clicking **Next** until the dialogs are finished.

Once your browser accepts the certificate, your Red Hat Linux Secure Web Server will show you a default home page as shown in Figure 11–6, *The Default Home Page*.

Figure 11–6 The Default Home Page

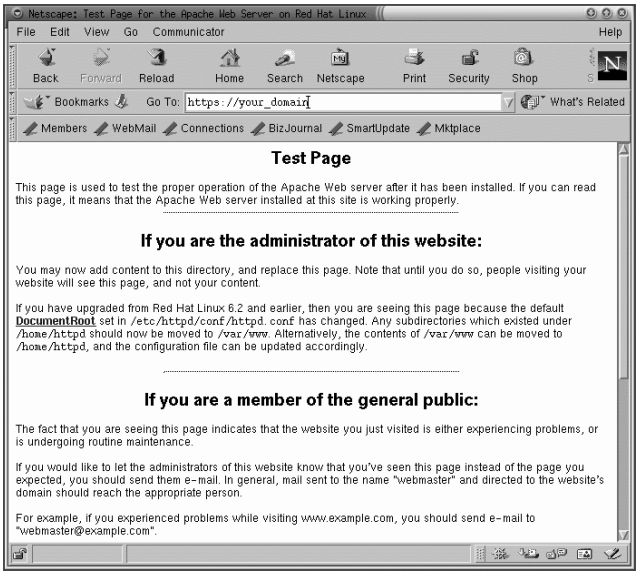

11.11 Starting and Stopping Apache

During the installation process, a Bourne shell (sh) script named `httpd` was installed into `/etc/rc.d/init.d`. To manually stop and start the server, run `httpd` with either `stop` or `start` as an argument.

To start your server, type the command:

```
/etc/rc.d/init.d/httpd start
```

You will be prompted to fill in your password. After you type it in, your server will start.

To stop your server, type the command:

```
/etc/rc.d/init.d/httpd stop
```

The command `restart` is a shorthand way of stopping and then starting your server. `restart` does explicitly stop and then start your server, so you will be prompted for your password. `restart` looks like the following:

```
/etc/rc.d/init.d/httpd restart
```

If you just finished editing something in your `httpd.conf` file, you don't need to explicitly stop and start your server. Instead you may use the `reload` command. The benefit of using `reload` is that you will not need to type in your password. Your password will remain cached across reloads, but it will not be cached between stops and starts. `reload` looks like the following:

```
/etc/rc.d/init.d/httpd reload
```

Your server (the `httpd` process) will start automatically when your machine boots. Be aware that you'll be prompted for the secure server's password after the machine boots, unless you generated a key for your secure server without password protection.

11.12 Accessing Your Secure Server

To access your secure server, use a URL like this:

https://your_domain

Note that URLs which are intended to connect to your Red Hat Linux Secure Web Server should begin with the https: protocol designator instead of the more common http: protocol designator. https: is the protocol designator for secure HTTP protocol communications.

So your non-secure server can be accessed using an URL like this:

http://your_domain

The standard port for secure Web communications is port 443. The standard port for non-secure Web communications is port 80. The Red Hat Linux Secure Web Server default configuration listens on both of the two standard ports. Therefore, you won't need to specify the port number in a URL (the port number is assumed).

However, if you configure your server to listen on a non-standard port (i.e., anything besides 80 or 443), you'll need to specify the port number in every URL which is intended to connect to the server on the non-standard port.

For example, you may have configured your server so that you have a virtual host running non-secured on port 12331. Any URLs intended to connect to that virtual host must specify the port number in the URL. The following URL example will attempt to connect to a non-secure Web server listening on port 12331:

```
http://your_domain:12331
```

Some of the example URLs used in this manual may need to be changed, depending upon whether you're accessing your Red Hat Linux Secure Web Server or your non-secure Web server. Please view all URLs in this manual as general examples, and not as explicit instructions that will work under all circumstances.

12 Configuring Your Secure Server

The default configuration of the Red Hat Linux Secure Web Server should work for most users. You may never need to change any of Apache's configuration directives. If you do want to change any of the default configuration options, you'll need to know what some of the options are, and know where to find them. This chapter covers the configuration options available to you.

After you've installed the Red Hat Linux Secure Web Server, the Apache Web server documentation is available at http://*your_domain*/manual/ or you can use the Apache documentation available on the Web at http://www.apache.org/docs/. The Apache Web server documentation contains a full list and complete descriptions of all of Apache's configuration options. For your convenience, short descriptions of the configuration directives used by your Red Hat Linux Secure Web Server are provided in this manual.

When you are looking through your Web server's configuration file, be aware that your default configuration includes both a non-secure and a secure Web server. The secure Web server runs as a virtual host, which is configured in the `httpd.conf` configuration file. For more information about virtual hosts, see Section 12.3, *Using Virtual Hosts*.

Please Note

We do not include FrontPage extensions, because the Microsoft(TM) license prohibits the inclusion of the extensions in a third party product.

12.1 Configuration Directives in `httpd.conf`

The Apache Web server's configuration file is `/etc/httpd/conf/httpd.conf`. The `httpd.conf` file is well-commented and somewhat self-explanatory. The default configuration of the Red Hat Linux Secure Web Server will work for most people, so you probably

won't need to change the directives in `httpd.conf`. However, you may want to be familiar with the most important configuration options.

The empty `srm.conf` and `access.conf` files are also in the `/etc/httpd/conf` directory. `srm.conf` and `access.conf` were formerly used, along with `httpd.conf`, as configuration files for Apache.

If you need to configure your Red Hat Linux Secure Web Server, you simply edit `httpd.conf`, and then either reload, or stop and start your Red Hat Linux Secure Web Server. How to reload, stop and start your server is covered in Section 11.11, *Starting and Stopping Apache*.

Before you edit `httpd.conf`, you should first copy the original file to something like `httpd.confold` (or to any name you want). Then, if you make a mistake while you're editing the configuration file, you'll have a backup to start over with.

If you do make a mistake, and your Red Hat Linux Secure Web Server doesn't work correctly, the first place to look is at what you just edited in `httpd.conf`. Make sure that you didn't make a typo. The next place to look is your Red Hat Linux Secure Web Server's error log (`/var/log/httpd/error_log`). The error log may not be easy to interpret, depending on your level of experience. If you've just experienced a problem, however, the last entries in the error log should provide some clues about what has happened.

The next sections provide short descriptions of the directives which are included in `httpd.conf`, in the order that you'll find them. These descriptions are not exhaustive. If you need more information, please refer to the Apache documentation provided in HTML format at http://*your_domain*/manual/ or to the Apache group documentation at http://www.apache.org/docs/. For more information about mod_ssl directives, refer to the documentation included in HTML format as http://*your_domain*/manual/mod/mod_ssl/, or see the mod_ssl *User Manual* at http://www.modssl.org/docs/2.6/.

12.1.1 `ServerType`

Your `ServerType` can be either `inetd` or `standalone`. By default, your Red Hat Linux Secure Web Server is set to `ServerType standalone`.

ServerType standalone means that the server is started once and then that server handles all of the connections. ServerType inetd means that for every HTTP connection, a new instance of the server is started. Each server instance handles the connection and exits when the connection is ended. As you can probably imagine, using inetd is very inefficient. Another problem is that inetd may not work correctly, according to the Apache group. And finally, since Red Hat Linux 7.0 uses xinetd, additional configuration will be needed to get xinetd to start the server. For these reasons, you'll want to leave your Red Hat Linux Secure Web Server's ServerType set to standalone.

12.1.2 ServerRoot

The ServerRoot is the top-level directory which will contain the server's files. Both your secure and non-secure servers are set to use a ServerRoot of /etc/httpd.

12.1.3 LockFile

LockFile sets the path to the lockfile used when the Apache server is compiled with either USE_FCNTL_SERIALIZED_ACCEPT or USE_FLOCK_SERIALIZED_ACCEPT. LockFile should normally be left at its default value.

12.1.4 PidFile

PidFile names the file in which the server records its process ID (pid). Your Red Hat Linux Secure Web Server is set to record its pid in /var/run/httpd.pid.

12.1.5 ScoreBoardFile

The ScoreBoardFile stores internal server process information, which is used for communication between the parent server process and its child processes. Your Red Hat Linux Secure Web Server's ScoreBoardFile is set to /var/run/httpd.scoreboard.

12.1.6 ResourceConfig

The ResourceConfig directive instructs the server to read the file after ResourceConfig for more directives. The ResourceConfig directive is

commented out, because your Web server only uses `httpd.conf` for configuration directives.

12.1.7 AccessConfig

The `AccessConfig` directive instructs the server to read the file named after `AccessConfig` for more directives, after it has read the file named by `ResourceConfig`. The `AccessConfig` directive is commented out, because your Web server only uses `httpd.conf` for configuration directives.

12.1.8 Timeout

`Timeout` defines, in seconds, the amount of time that your server will wait for receipts and transmissions during communications. Specifically, `Timeout` defines how long your server will wait to receive a GET request, how long it will wait to receive TCP packets on a POST or PUT request and how long it will wait between ACKs responding to TCP packets. `Timeout` is set to 300 seconds, which is appropriate for most situations.

12.1.9 KeepAlive

`KeepAlive` sets whether your server will allow persistent connections (i.e., more than one request per connection). `KeepAlive` can be used to prevent any one client from consuming too much of the server's resources. By default, `KeepAlive` is set to `on`, which means that your server allows persistent connections. You could set it to `off`, which would disable persistent connections. See Section 12.1.10, *`MaxKeepAliveRequests`* for a related way to limit requests per connection.

12.1.10 MaxKeepAliveRequests

This directive sets the maximum number of requests allowed per persistent connection. The Apache Group recommends a high setting, which will improve your server's performance. `MaxKeepAliveRequests` is set to `100` by default, which should be appropriate for most situations.

12.1.11 `KeepAliveTimeout`

`KeepAliveTimeout` sets the number of seconds your server will wait for a subsequent request, after a request has been served, before it closes the connection. Once a request has been received, the `Timeout` directive applies instead.

12.1.12 `MinSpareServers` and `MaxSpareServers`

The Apache Web server dynamically adapts to the perceived load by maintaining an appropriate number of spare server processes based on the traffic. The server checks the number of servers waiting for a request and kills some if there are more than `MaxSpareServers` or creates some if the number of servers is less than `MinSpareServers`.

Your server's default `MinSpareServers` is `5`; your server's default `MaxSpareServers` is `20`. These default settings should be appropriate in almost all situations. You should not increase the `MinSpareServers` to a very large number, since that will create a heavy processing load on your server even when traffic is light.

12.1.13 `StartServers`

`StartServers` sets how many server processes are created upon startup. Since your Web server dynamically kills and creates server processes based on traffic load, you won't ever need to change this parameter. Your Web server is set to start eight server processes at startup.

12.1.14 `MaxClients`

`MaxClients` sets a limit on the total number of server processes (i.e., simultaneously connected clients) that can run at one time. You want to keep `MaxClients` at a high number (your server's default is set to 150), because no one else will be allowed to connect once that number of simultaneously connected clients is reached. You can't set `MaxClients` to higher than 256 without recompiling Apache. The main reason for having `MaxClients` is so that a runaway Web server doesn't crash your operating system.

12.1.15 MaxRequestsPerChild

`MaxRequestsPerChild` sets the total number of requests each child server process serves before the child dies. The main reason for setting `MaxRequestsPerChild` is to avoid long-lived process induced memory leaks. The default `MaxRequestsPerChild` for your server is `100`.

12.1.16 Listen

The `Listen` command identifies the ports on which your Red Hat Linux Secure Web Server will accept incoming requests. Your Red Hat Linux Secure Web Server is set to listen to port 80 for non-secure Web communications and (in virtual host tags that define the secure server) to port 443 for secure Web communications.

If you set Apache to listen to a port under 1024, the `httpd` process will need to start as root. For port 1024 and above, `httpd` can start as a regular user.

`Listen` can also be used to specify particular IP addresses over which the server will accept connections.

12.1.17 BindAddress

`BindAddress` is a way of specifying which IP addresses your server will listen to. You should use the `Listen` directive instead if you need this functionality. `BindAddress` is not used by your Web server; by default it is commented out in `httpd.conf`.

12.1.18 LoadModule

`LoadModule` is used to load in Dynamic Shared Object (DSO) modules. More information on the Red Hat Linux Secure Web Server's DSO support, including exactly how to use the `LoadModule` directive, can be found in Section 12.2, *Adding Modules to Your Server*. Note that the order of the modules is important, so don't move them around.

12.1.19 IfDefine

The <IfDefine> and </IfDefine> tags surround configuration directives that are applied if the "test" stated in the <IfDefine> tag is true; the directives are ignored if the test is false.

The test in the <IfDefine> tags is a parameter-name (e.g., HAVE_PERL). If the parameter is defined (i.e., provided as an argument to the server's start-up command), then the test is true. In this case, when your Red Hat Linux Secure Web Server is started, the test is true and the directives contained in the IfDefine tags are applied.

By default, <IfDefine HAVE_SSL> tags surround the virtual host tags for your secure server. <IfDefine HAVE_SSL> tags also surround the LoadModule and AddModule directives for the ssl_module.

12.1.20 ClearModuleList

The ClearModuleList directive is located immediately before the long list of AddModule directives. ClearModuleList erases the server's built-in list of active modules. Then the list of AddModule directives re-creates the list, immediately after ClearModuleList.

12.1.21 AddModule

AddModule is the directive used to create a complete list of all available modules. You will use the AddModule directive if you add your own module in as a DSO. For more information on how AddModule is used for DSO support, see Section 12.2, *Adding Modules to Your Server*.

12.1.22 ExtendedStatus

The ExtendedStatus directives controls whether Apache generates basic (off) or detailed server status information (on), when the server-status handler is called. Server-status is called using Location tags; more information on calling server-status is included in Section 12.1.30, *Location*.

12.1.23 Port

Normally, Port defines the port that your server is listening to. Your Red Hat Linux Secure Web Server, however, is listening to more than one port by default, since the Listen directive is also being used. When Listen directives are in effect, your server listens at all of those ports. See the description of the Listen directive for more information about Listen.

The Port command is also used to specify the port number used to construct a canonical name for your server. See Section 12.1.40, *UseCanonicalName* for more information about your server's canonical name.

12.1.24 User

The User directive sets the userid used by the server to answer requests. User's setting determines the server's access. Any files inaccessible to this user will also be inaccessible to your website's visitors. The default for User is apache.

The User should only have privileges so that it can access files which are supposed to be visible to the outside world. The User is also the owner of any CGI processes spawned by the server. The User should not be allowed to execute any code which is not intended to be in response to HTTP requests.

Please Note

Unless you know exactly what you're doing, don't set the User to root. Using root as the User will create wide, gaping security holes for your Red Hat Linux Secure Web Server.

The parent httpd process first runs as root during normal operations, but is then immediately handed off to the apache user. The server must start as root because it needs to bind to a port below 1024 (the default port for secure Web communications is port 443; the default port for non-secure Web communications is port 80). Ports below 1024 are reserved for system use, so they can't be used by anyone but root.

Once the server has attached itself to its port, however, it hands the process off to the User before it accepts any connection requests.

12.1.25 Group

The Group directive is similar to the User. The Group sets the group under which the server will answer requests. The default Group is also apache.

12.1.26 ServerAdmin

ServerAdmin should be the e-mail address of the Red Hat Linux Secure Web Server's administrator. This e-mail address will show up in error messages on server-generated Web pages, so users can report a problem by sending e-mail to the server administrator. ServerAdmin is set by default to root@localhost.

Typically, a good way to set up ServerAdmin is to set it to webmaster@your_domain.com. Then alias webmaster to the person responsible for the Web server in /etc/aliases. Finally, run /usr/bin/newaliases to add the new alias.

12.1.27 ServerName

You can use ServerName to set a hostname for your server which is different from your host's real name. For example, you might want to use www.your_domain.com when your server's real name is actually foo.your_domain.com. Note that the ServerName must be a valid Domain Name Service (DNS) name that you have the right to use (don't just make something up).

If you do specify a ServerName, be sure its IP address and server name pair are included in your /etc/hosts file.

12.1.28 DocumentRoot

The DocumentRoot is the directory which contains most of the HTML files which will be served in response to requests. The default DocumentRoot for both the non-secure and secure Web servers is /var/www/html. For example, the server might receive a request for the following document:

```
http://your_domain/foo.html
```

The server will look for the following file in the default directory:

```
/var/www/html/foo.html
```

If you want to change the DocumentRoot so that it isn't shared by the secure and the non-secure Web servers, see Section 12.3, *Using Virtual Hosts*.

12.1.29 Directory

<Directory /path/to/directory> and </Directory> tags are used to enclose a group of configuration directives that are meant to apply only to that directory and all of its subdirectories. Any directive which is applicable to a directory may be used within <Directory> tags. <File> tags can be used in the same way, to apply to a specific file.

By default, very restrictive parameters are applied to the root directory, using the Options (see Section 12.1.31, *Options*) and AllowOverride (see Section 12.1.32, *AllowOverride*) directives. Under this configuration, any directory on your system which needs more permissive settings has to be explicitly given those settings.

Using Location tags, the DocumentRoot (referred to as "/") is defined to have less rigid parameters, so that HTTP requests can be served from it.

The cgi-bin directory is set up to allow the execution of CGI scripts, with the ExecCGI option. If you need to execute a CGI script in another directory, you'll need to set ExecCGI for that directory. For example, if your cgi-bin is /var/www/cgi-bin, but you want to execute CGI scripts from within /home/my_cgi_directory, add an ExecCGI directive to a set of Directory directives like the following to your httpd.conf file:

```
<Directory /home/my_cgi_directory>
    Options +ExecCGI
</Directory>
```

To allow CGI script execution in /home/my_cgi_directory, you'll need to take a few extra steps besides setting ExecCGI. You'll also need to have the AddHandler directive uncommented to identify files with the .cgi extension as CGI

scripts. See Section 12.1.66, *AddHandler* for instructions on setting AddHandler. Permissions for CGI scripts, and the entire path to the scripts, must be set to 0755. Finally, the owner of the script and the owner of the directory must be the same user.

12.1.30 Location

<Location> and </Location> tags allow you to specify access control based on the URL.

The first use of Location tags is to configure Options and provide other configuration guidelines for the DocumentRoot. These configuration directives, located within the <Location "/"> and </Location> tags, are necessary to provide access to the documents located in the DocumentRoot.

The next use of Location tags is located within IfModule mod_perl.c tags. These configuration directives are in effect if the mod_perl.so DSO is loaded. See Section 12.2, *Adding Modules to Your Server* for more information about adding modules to Apache.

The Location tags name the /var/www/perl directory (an Alias for /perl) as the directory from which Perl scripts will be served. If a document is requested with an URL containing /perl in the path, your Web server will look in /var/www/perl/ for the appropriate Perl script.

Several other <Location> options are commented out in your httpd.conf file. If you want to enable the functionality they provide, you'll need to uncomment the appropriate section of directives.

Immediately after the Perl directives discussed previously, your httpd.conf file includes a section of directives for enabling HTTP PUT (e.g., Netscape Gold's publish feature, which can post Web pages to a Web server). If you want to allow HTTP PUT, you'll need to uncomment this entire section:

```
#LoadModule put_module          modules/mod_put.so
#AddModule mod_put.c
#
#Alias /upload /tmp
#<Location /upload>
#    EnablePut On
```

```
#    AuthType Basic
#    AuthName Temporary
#    AuthUserFile /etc/httpd/conf/passwd
#    EnableDelete Off
#    umask 007
#    <Limit PUT>
#       require valid-user
#    </Limit>
#</Location>
```

If you want to allow people connecting from your domain to see server status reports, you should uncomment the next section of directives:

```
#<Location /server-status>
#    SetHandler server-status
#    Order deny,allow
#    Deny from all
#    Allow from .your_domain.com
#</Location>
```

You must replace `.your_domain.com` with your second level domain name.

If you want to provide server configuration reports (including installed modules and configuration directives) to requests from inside your domain, you'll need to uncomment the following lines:

```
#<Location /server-info>
#    SetHandler server-info
#    Order deny,allow
#    Deny from all
#    Allow from .your_domain.com
#</Location>
```

Again, you must fill in `.your_domain.com`.

The next section of directives use `Location` tags to allow access to the documentation in `/usr/share/doc` (for example, with a URL like http://*your_domain*/doc/whatever.html). These directives only allow this access to requests made from the localhost.

Another use of the Location tags is a commented-out section which is intended to track attacks on your Web server which exploit an old bug from pre-Apache 1.1 days. If you want to track these requests, uncomment the following lines:

```
#<Location /cgi-bin/phf*>
#    Deny from all
#    ErrorDocument 403 http://phf.apache.org/phf_abuse_log.cgi
#</Location>
```

If these lines are uncommented, your Web server will redirect any requests which end in /cgi-bin/phf* to a logging CGI script run by the Apache Group.

12.1.31 Options

The Options directive controls which server features are available in a particular directory. For example, under the restrictive parameters specified for the root directory, Options is set to only FollowSymLinks. No features are enabled, except that the server is allowed to follow symbolic links in the root directory.

By default, in your DocumentRoot directory, Options is set to include Indexes, Includes and FollowSymLinks. Indexes permits the server to generate a directory listing for a directory if no DirectoryIndex (i.e., index.html, etc.) is specified. Includes means that server-side includes are permitted. FollowSymLinks allows the server to follow symbolic links in that directory.

You'll also need to include Options statements for directories within virtual hosts directives, if you want your virtual hosts to recognize those Options.

For example, server side includes are already enabled inside the /var/www/html directory, because of the Options Includes line within the Location "/" directives section. However, if you want a virtual host to recognize that server side includes are allowed within /var/www/html, you'll need to include a section like the following within your virtual host's tags:

```
<Directory /var/www/html>
Options Includes
</Directory>
```

12.1.32 AllowOverride

The `AllowOverride` directive sets whether or not any `Options` can be overridden by the declarations in an `.htaccess` file. By default, both the root directory and the `DocumentRoot` are set to allow no `.htaccess` overrides.

12.1.33 Order

The `Order` directive simply controls the order in which `allow` and `deny` directives are evaluated. Your server is configured to evaluate the `Allow` directives before the `deny` directives for your `DocumentRoot` directory.

12.1.34 Allow

`Allow` specifies which requester can access a given directory. The requester can be `all`, a domain name, an IP address, a partial IP address, a network/netmask pair, etc. Your `DocumentRoot` directory is configured to `Allow` requests from `all` (i.e., anyone).

12.1.35 deny

`Deny` works just like `allow`, but you're specifying who is denied access. Your `DocumentRoot` isn't configured to `deny` requests from anyone.

12.1.36 UserDir

`UserDir` is the name of the subdirectory within each user's home directory where they should place personal HTML files which are to be served by the Web server. By default, the subdirectory is `public_html`. For example, the server might receive the following request:

```
http://your_domain/~username/foo.html
```

The server would look for the file:

```
/home/username/public_html/foo.html
```

In the above example, `/home/username` is the user's home directory (note that the default path to users' home directories may be different on your system).

Make sure that the permissions on the users' home directories are set correctly. Users' home directories must be set to 0755. The read (r) and execute (x) bits must be set on the users' `public_html` directories (0755 will work). Files that will be served in users' `public_html` directories must be set to at least 0644.

12.1.37 `DirectoryIndex`

The `DirectoryIndex` is the default page served by the server when a user requests an index of a directory by specifying a forward slash (/) at the end of the directory name.

For example, when a user requests the page http://*your_domain*/*this_directory*/, they are going to get either the `DirectoryIndex` page if it exists, or a server-generated directory listing. The default for `DirectoryIndex` is `index.html index.htm index.shtml index.cgi`. The server will try to find any one of these four files, and will return the first one it finds. If it doesn't find any of these files and if `Options Indexes` is set for that directory, the server will generate and return a listing, in HTML format, of the subdirectories and files in the directory.

12.1.38 `AccessFileName`

`AccessFileName` names the file which the server should use for access control information in each directory. By default, your Web server is set to use `.htaccess`, if it exists, for access control information in each directory.

Immediately after the `AccessFileName` directive, a set of `Files` tags apply access control to any file beginning with a `.ht`. These directives deny Web access to any `.htaccess` files (or other files which begin with `.ht`) for security reasons.

12.1.39 `CacheNegotiatedDocs`

By default, your Red Hat Linux Secure Web Server asks proxy servers not to cache any documents which were negotiated on the basis of content (i.e., they may change over time or because of the input from the requester). If you uncomment `CacheNegotiatedDocs`, you are disabling that function and proxy servers will be allowed to cache the documents from then on.

12.1.40 UseCanonicalName

UseCanonicalName is set by default to on. UseCanonicalName allows the server to construct an URL that references itself, using ServerName and Port. When the server refers to itself in response to requests from clients, it uses this URL. If you set UseCanonicalName to off, the server will instead use the value that came in the request from the client to refer to itself.

12.1.41 TypesConfig

TypesConfig names the file which sets the default list of MIME type mappings (filename extensions to content types). The default TypesConfig file is /etc/mime.types. Instead of editing /etc/mime.types, the recommended way to add MIME type mappings is to use the AddType directive.

12.1.42 DefaultType

DefaultType sets a default content type for the Web server to use for documents whose MIME types can't be determined. Your Web server defaults to assume a plain text content type for any file with an indeterminate content type.

12.1.43 IfModule

<IfModule> and </IfModule> tags surround directives that are conditional. The directives contained within the IfModule tags are processed under one of two conditions. The directives are processed if the module contained within the starting <IfModule> tag is compiled in to the Apache server. Or, if an "!" (an exclamation point) is included before the module name, the directives are processed only if the module in the starting <IfModule> tag is *not* compiled in.

The mod_mime_magic.c file is included in these IfModule tags. The mod_mime_magic module can be compared to the UNIX file command, which looks at a few bytes of a file's contents, then uses "magic numbers" and other hints in order to figure out the MIME type of the file.

If the mod_mime_magic module is compiled in to Apache, these IfModule tags tell the mod_mime_magic module where the hints definition file is: share/magic in this case.

The mod_mime_magic module is not compiled in by default. If you would like to use it, see Section 12.2, *Adding Modules to Your Server*, for instructions on how to add modules to your server.

12.1.44 HostnameLookups

`HostnameLookups` can be set to `on` or `off`. If you allow `HostnameLookups` (by setting it to `on`), your server will automatically resolve the IP address for each connection which requests a document from your Web server. Resolving the IP address means that your server will make one or more connections to the DNS in order to find out the hostname that corresponds to a particular IP address.

Generally, you should leave `HostnameLookups` set to `off`, because the DNS requests add a load to your server and may slow it down. If your server is busy, the effects of `HostnameLookups` may be quite noticeable.

`HostnameLookups` are also an issue for the Internet as a whole. All of the individual connections made to look up each hostname add up. Therefore, for your own Web server's benefit, as well as for the good of the Internet as a whole, you should leave `HostnameLookups` set to `off`.

12.1.45 ErrorLog

`ErrorLog` names the file where server errors are logged. As this directive indicates, the error log file for your Web server is `/var/log/httpd/error_log`.

The error log is a good place to look if your Web server ever generates any errors or fails and you aren't sure what happened.

12.1.46 LogLevel

`LogLevel` sets how verbose the error messages in the error logs will be. `LogLevel` can be set (from least verbose to most verbose) to `emerg`, `alert`, `crit`, `error`, `warn`, `notice`, `info` or `debug`. Your Red Hat Linux Secure Web Server's `LogLevel` is set to `warn` (a happy medium).

12.1.47 LogFormat

The `LogFormat` directives in your `httpd.conf` file set up a format for the messages in your access log; hopefully, this format will make your access log more readable.

12.1.48 CustomLog

`CustomLog` identifies the log file and the log file format. In your Red Hat Linux Secure Web Server's default configuration, `CustomLog` defines the log file in which accesses to your Web server are recorded: `/var/log/httpd/access_log`. You'll need to know the location of this file if you want to generate any access-based server performance statistics for your Web server.

`CustomLog` also sets the log file format to common. The common logfile format looks like this:

```
remotehost rfc931 authuser [date] "request" status bytes
```

remotehost

The remote hostname. If the hostname is not available from DNS, or if `HostnameLookups` is set to `Off`, then *`remotehost`* will be the IP address of the remote host.

rfc931

Not used. You'll see a `-` in the log file in its place.

authuser

If authentication was required, this is the username with which the user identified him or herself. Usually, this isn't used, so you'll see a `-` in its place.

[date]

The date and time of the request.

"request"

The request string exactly as it came from the browser or client.

status

The HTTP status code which was returned to the browser or client.

bytes

The size of the document.

The `CustomLog` command can be used to set up specific log files to record referers (the URL for the Web page which linked to a page on your Web server) and/or agents (the browsers used to retrieve Web pages from your Web server). The relevant `CustomLog` lines are commented out, as shown, but you should uncomment them if you want those two log files:

```
#CustomLog /var/log/httpd/referer_log referer
#CustomLog /var/log/httpd/agent_log agent
```

Alternatively, you can also set the `CommonLog` directive to use a `combined` log by uncommenting the following line:

```
#CustomLog /var/log/httpd/access_log combined
```

A `combined` log will add the referer and agent fields to the end of the common log fields. If you want to use a `combined` log, you'll need to comment out the `CustomLog` directive setting your access log to the common logfile format.

12.1.49 `ServerSignature`

The `ServerSignature` directive adds a line containing the Apache server version and the `ServerName` of the serving host to any server-generated documents (for example, error messages sent back to clients). `ServerSignature` is set to `on` by default. You can change it to `off`, so no signature line will be added, or you can change it to `EMail`. `EMail` will add a `mailto:ServerAdmin` HTML tag to the signature line.

12.1.50 `Alias`

The `Alias` setting allows directories to be outside the `DocumentRoot` directory and yet still accessible to the Web server. Any URL ending in the alias will automatically resolve to the alias' path. By default, one alias is already set up. An `icons`

directory can be accessed by the Web server, but the directory is not in the `DocumentRoot`. The `icons` directory, an alias, is actually `/var/www/icons/`, not `/var/www/html/icons/`.

12.1.51 ScriptAlias

The `ScriptAlias` setting defines where CGI scripts (or other types of scripts) can be found. Generally, you don't want to leave CGI scripts within the `DocumentRoot`. If CGI scripts are in `DocumentRoot`, they could potentially be viewed as text documents. Even if you don't care if people can see (and then use) your CGI scripts, revealing how they work creates opportunities for unscrupulous people to exploit any security holes in the script, and may create a security risk for your server. By default, the `cgi-bin` directory is a `ScriptAlias` of `/cgi-bin/`, and is actually located in `/var/www/cgi-bin/`.

Your `/var/www/cgi-bin` directory has `Options ExecCGI` set, meaning that execution of CGI scripts is permitted within that directory.

See Section 12.1.66, *`AddHandler`* and Section 12.1.29, *`Directory`* for instructions on how to execute CGI scripts in directories other than the `cgi-bin`.

12.1.52 Redirect

When a Web page is moved, `Redirect` can be used to map the old URL to a new URL. The format is as follows:

```
Redirect /path/foo.html http://new_domain/path/foo.html
```

So, if an HTTP request is received for a page which used to be found at http://*your_domain*/*path*/foo.html, the server will send back the new URL (http://*new_domain*/*path*/foo.html) to the client, which should attempt to fetch the document from the new URL.

12.1.53 IndexOptions

`IndexOptions` controls the appearance of server generated directing listings, by adding icons and file descriptions, etc. If `Options Indexes` is set (see Section 12.1.31, *`Options`*), your Web server may generate a directory listening when your Web server receives an HTTP request like the following:

http://*your_domain*/*this_directory*/

First, your Web server looks in that directory for a file from the list after the `DirectoryIndex` directive (e.g., `index.html`). If your Web server doesn't find one of those files, it creates an HTML directory listing of the subdirectories and files in the directory. You can modify the appearance of this directory listing using certain directives in `httpd.conf`, including `IndexOptions`.

Your default configuration sets `FancyIndexing` on. If `FancyIndexing` is turned on, clicking on the column headers in the directory listing will sort the order of the display by that header. Another click on the same header will switch from ascending to descending order and back. `FancyIndexing` also shows different icons for different files, depending upon file extensions. If you use the `AddDescription` directive and turn `FancyIndexing` on, then a short description of a file will be included in the server generated directory listing.

`IndexOptions` has a number of other parameters which can be set to control the appearance of server generated directories. Parameters include `IconHeight` and `IconWidth`, to make the server include HTML `HEIGHT` and `WIDTH` tags for the icons in server generated Web pages; `IconsAreLinks`, for making the icons act as part of the HTML link anchor along with the filename, and others.

12.1.54 `AddIconByEncoding`

This directive names icons which will be displayed by files with MIME encoding, in server generated directory listings. For example, by default, your Web server shows the `compressed.gif` icon next to MIME encoded x-compress and x-gzip files in server generated directory listings.

12.1.55 `AddIconByType`

This directive names icons which will be displayed next to files with MIME types in server generated directory listings. For example, your server is set to show the icon `text.gif` next to files with a mime-type of "text," in server generated directory listings.

12.1.56 AddIcon

`AddIcon` tells the server which icon to show in server generated directory listings for certain file types or for files with certain extensions. For example, your Web server is set to show the icon `binary.gif` for files with `.bin` or `.exe` extensions.

12.1.57 DefaultIcon

`DefaultIcon` names the icon to show in server generated directory listings for files which have no other icon specified. `unknown.gif` is the `DefaultIcon` for those files by default.

12.1.58 AddDescription

You can use `AddDescription` to show text that you specify for certain files, in server generated directory listings (you'll also need to enable `FancyIndexing` as an `IndexOptions`). You can name specific files, wildcard expressions or file extensions to specify the files which this directive should apply to. For example, you could use the following line:

```
AddDescription "A file that ends in .ni" .ni
```

In server generated directory listings, all files with extensions of `.ni` would have the description `A file that ends in .ni` after the filename. Note that you'll also need `FancyIndexing` turned on.

12.1.59 ReadmeName

`ReadmeName` names the file which (if it exists in the directory) will be appended to the end of server generated directory listings. The Web server will first try to include the file as an HTML document and then try to include it as plain text. By default, `ReadmeName` is set to `README`.

12.1.60 HeaderName

`HeaderName` names the file which (if it exists in the directory) will be prepended to the start of server generated directory listings. Like `ReadmeName`, the server will try to include it as an HTML document if possible, or in plain text if not.

12.1.61 `IndexIgnore`

`IndexIgnore` lists file extensions, partial filenames, wildcard expressions or full filenames. The Web server will not include any files which match any of those parameters in server generated directory listings.

12.1.62 `AddEncoding`

`AddEncoding` names filename extensions which should specify a particular encoding type. `AddEncoding` can also be used to instruct some browsers (not all) to uncompress certain files as they are downloaded.

12.1.63 `AddLanguage`

`AddLanguage` associates filename extensions with specific content languages. This directive is mostly useful for content negotiation, when the server returns one of several documents based on the client's language preference as set in their browser.

12.1.64 `LanguagePriority`

`LanguagePriority` allows you to set precedence for different languages in which to serve files, which will be in effect if the client expressed no preference for language in their browser.

12.1.65 `AddType`

Use the `AddType` directive to define MIME type and file extension pairs. For example, if you are using PHP4, your Web server is using the `AddType` directive to make your Web server recognize files with PHP extensions (`.php4`, `.php3 .phtml .php`) as PHP MIME types.

The following `AddType` line tells your server to recognize the `.shtml` file extension (for server side includes):

```
AddType text/html .shtml
```

You'll need to include the above line within the virtual host tags for any virtual hosts which should allow server side includes.

12.1.66 AddHandler

`AddHandler` maps file extensions to specific handlers. For example, the `cgi-script` handler can be used matched with the extension `.cgi` to automatically treat a file ending with `.cgi` as a CGI script. This will work, even for files outside of the `ScriptAlias` directory, as long as you follow the instructions provided here.

You have a CGI `AddHandler` line in your `httpd.conf` file:

```
AddHandler cgi-script .cgi
```

You'll have to uncomment the line. Then Apache will execute CGI scripts for files ending in `.cgi`, even if they are outside of the `ScriptAlias`, which is set by default to locate your `/cgi-bin/` directory in `/var/www/cgi-bin/`.

You'll also need to set `ExecCGI` as an `Options` for any directory containing a CGI script. See Section 12.1.29, *`Directory`* for more information about setting `ExecCGI` for a directory. Additionally, you'll need to make sure the permissions are set correctly for the CGI scripts and the directories containing CGI scripts. CGI scripts and the entire directory path to the scripts must be set to 0755. Finally, the owner of the directory and the owner of the script file must be the same user.

You'll need to add the same `AddHandler` line to your `VirtualHost` setup, if you're using virtual hosts and you want them to also recognize CGI scripts outside the `ScriptAlias`.

In addition to CGI scripts, your Web server also uses `AddHandler` to process server-parsed HTML and imagemap files.

12.1.67 Action

`Action` allows you to specify a MIME content type and CGI script pair, so that whenever a file of that media type is requested, a particular CGI script will be executed.

12.1.68 MetaDir

`MetaDir` specifies the name of a directory where your Web server should look for files containing meta information (extra HTTP headers) to include when serving documents.

12.1.69 MetaSuffix

`MetaSuffix` specifies the filename suffix for the file that contains meta information (extra HTTP headers), which should be located in the `MetaDir` directory.

12.1.70 ErrorDocument

By default, in the event of a problem or error, your Web server outputs a simple (and usually cryptic) error message back to the requesting client. Instead of using the default, you can use `ErrorDocument` to configure your Web server so that it outputs a customized message or redirects the client to a local or external URL. `ErrorDocument` simply associates a HTTP response code with a message or a URL which will be sent back to the client.

12.1.71 BrowserMatch

The `BrowserMatch` directive allows your server to define environment variables and/or take appropriate actions based on the User-Agent HTTP header field, which identifies the client's browser. By default, your Web server uses `BrowserMatch` to deny connections to specific browsers with known problems and also to disable keepalives and HTTP header flushes for browsers that are known to have problems with those actions.

12.1.72 ProxyRequests

If you uncomment the `IfModule` tags surrounding the `ProxyRequests` et al, your Apache server will also function as a proxy server. You'll also need to load the `mod_proxy` module. For instructions on how to load in modules, see Section 12.2, *Adding Modules to Your Server*.

12.1.73 ProxyVia

The `ProxyVia` command controls whether or not an HTTP Via: header line is sent along with requests or replies which go through the Apache proxy server. The Via: header will show the hostname if `ProxyVia` is set to `On`, the hostname and Apache version for `Full`, any Via: lines will be passed along unchanged for `Off`, and Via: lines will be removed for `Block`.

12.1.74 Cache Directives

A number of cache directives are commented out in the proxy `IfModule` tags mentioned above. If you are using the proxy server functionality and you want to also enable the proxy cache, you should uncomment the cache directives as described. The default settings for your cache directives should be appropriate for most configurations.

`CacheRoot` sets the name of the directory which will contain cached files. The default `CacheRoot` is `/var/cache/httpd`.

`CacheSize` sets how much space the cache can use, in KB. The default `CacheSize` is `5` KB.

`CacheGcInterval` sets a number of hours. After that number of hours, files in the cache will be deleted if the cache is using more space than allowed by `CacheSize`. The default for `CacheGcInterval` is four hours.

Cached HTML documents will be retained (without a reload from the originating Web server) in the cache for a maximum number of hours set by `CacheMaxExpire`. The default is `24` hours.

The `CacheLastModifiedFactor` affects the creation of an expiry (expiration) date for a document which did not come from its originating server with its own expiry set. The default `CacheLastModifiedFactor` is set to `0.1`, meaning that the expiry date for such documents equals one-tenth of the amount of time since the document was last modified.

`CacheDefaultExpire` is the expiry time in hours for a document that was received using a protocol that doesn't support expiry times. The default is set to one hour.

Any document that is retrieved from a host and/or domain that matches one set in `NoCache` will not be cached. If you know of hosts or domains from which you don't want to cache documents, uncomment `NoCache` and set their domains or hostnames here.

12.1.75 `NameVirtualHost`

You'll need to use the `NameVirtualHost` directive for the IP address (and port number if necessary) of any name-based virtual hosts you're setting up. The name-based virtual hosts configuration is used when you want to set up different virtual hosts for different domains, but you don't have (or don't want to use) different IP addresses for all of the different domain names for which your Web server serves documents.

Please Note

You can't use name-based virtual hosts with your secure server. Any name-based virtual hosts you set up will only work with non-secure HTTP connections and not with SSL connections.

You can't use name-based virtual hosts with your secure server because the SSL handshake (when the browser accepts the secure Web server's authenticating certificate) occurs before the HTTP request which identifies the correct name-based virtual host. In other words, authentication occurs before there is any identification of different name-based virtual hosts. If you want to use virtual hosts with your secure server, you'll need to use IP address-based virtual hosts.

If you're using name-based virtual hosts, uncomment the `NameVirtualHost` configuration directive and add the correct IP address for your server after `NameVirtualHost`. Then add more information about the different domains using the `Virtual Host` tags which surround the `ServerName` for each virtual host, plus any other configuration directives which are only applicable to that virtual host.

12.1.76 VirtualHost

<VirtualHost> and </VirtualHost> tags surround any configuration directives which are intended to apply to a virtual host. Most configuration directives can be used within virtual host tags, and then they only apply to that particular virtual host.

A set of commented out VirtualHost tags surround some example configuration directives and placeholders for the information you'd need to fill in to set up a virtual host. Please see Section 12.3, *Using Virtual Hosts*, for more information about virtual hosts.

12.1.77 SetEnvIf

The Apache configuration directive SetEnvIf is used to disable HTTP keepalive and to allow SSL to close the connection without a close notify alert from the client browser. This setting is necessary for certain browsers that don't reliably shut down the SSL connection.

12.1.78 SSL Configuration Directives

The SSL directives in your server's httpd.conf file are included to enable secure Web communications using SSL and TLS.

For more information on SSL directives, please point your browser to http://*your_domain*/manual/mod/mod_ssl/. More information on SSL directives is also available at http://www.modssl.org/docs/2.6/ssl_reference.html/, a chapter in a Web document about mod_ssl by Ralf Engelschall. The same document, the mod_ssl *User Manual*, begins at http://www.modssl.org/docs/2.6/ and is a great reference source for mod_ssl (of course) and for Web cryptography in general. This manual provides general information about securing your Web server in Chapter 11, *Obtaining a Certificate for your Secure Server*.

Please Note

Don't modify your SSL directives unless you're absolutely sure about what you're doing. For the vast majority of Red Hat Linux Secure Web Servers, the SSL directives are configured appropriately as installed.

12.2 Adding Modules to Your Server

Since Apache 1.3 supports DSOs, you can easily load Apache modules or compile in your own modules to your Red Hat Linux Secure Web Server. DSO support means that modules may be loaded at runtime. Since the modules are only loaded as necessary, they won't use any memory unless they're loaded and less memory will be needed overall.

The Apache Group provides complete DSO Documentation at http://www.apache.org/docs/dso.html. After installation of your server, you can also check http://*your_domain*/manual/mod/ for documentation on Apache modules in HTML format (if you installed the `apache-manual` package). A "quick and dirty" description of how to load modules is provided next, but if you need more details, check the URLs provided.

For your Red Hat Linux Secure Web Server to use a dynamically shared module, that module must have a `LoadModule` line and an `AddModule` line in `httpd.conf`. By default, many modules have these two lines already included in `httpd.conf`, but a few of the less commonly used modules are commented out. The commented out modules were included during compilation, but they are not loaded by default.

If you need to use one of those non-loaded modules, look in the `httpd.conf` file to see all the available modules. Each of the available modules has a corresponding `LoadModule` line. To show you an example, the `LoadModule` section begins with these seven lines:

```
#LoadModule mmap_static_module modules/mod_mmap_static.so
LoadModule vhost_alias_module modules/mod_vhost_alias.so
LoadModule env_module         modules/mod_env.so
LoadModule config_log_module  modules/mod_log_config.so
```

```
LoadModule agent_log_module   modules/mod_log_agent.so
LoadModule referer_log_module modules/mod_log_referer.so
#LoadModule mime_magic_module  modules/mod_mime_magic.so
```

Most of the lines are not commented out, indicating that each associated module was compiled in and is loaded in by default. The first line is commented out, which means that the corresponding module (`mmap_static_module`) was compiled in but not loaded.

To make your Red Hat Linux Secure Web Server load an unloaded module, first uncomment the corresponding `LoadModule` line. For example, if you wanted to make your Red Hat Linux Secure Web Server load in the `mime_magic_module`, change that `LoadModule` line from the original:

```
#LoadModule mime_magic_module modules/mod_mime_magic.so
```

Uncomment the previous line so that it reads:

```
LoadModule mime_magic_module modules/mod_mime_magic.so
```

Next, you need to uncomment the corresponding line from the `AddModule` section in `httpd.conf`. To continue with our previous example, uncomment the `mod_mime_magic` line. The original (default) line looks like the following:

```
#AddModule mod_mime_magic.c
```

The uncommented line should read:

```
AddModule mod_mime_magic.c
```

Once you've uncommented the `LoadModule` and `AddModule` lines for the module that you want to load in, stop and start your Web server, as covered in Section 11.11, *Starting and Stopping Apache*. After starting, the module should be loaded in to your Red Hat Linux Secure Web Server.

If you have your own module, you can add it to the `httpd.conf` file so that it is compiled in and loaded as a DSO. If you want to do this, you need to install the `apache-devel` package, as covered in Chapter 10, *Installing the Red Hat Linux Secure Web Server*. You need the `apache-devel` package because it installs the

include files, the header files and the APache eXtenSion (APXS) support tool. APXS uses the include files and the header files to compile your module so that it will work with Apache.

If you've written your own module or are borrowing someone else's, you should be able to use APXS to compile your module sources outside the Apache source tree, without needing to tweak any compiler and/or linker flags. If you need more information on APXS, please see the Apache documentation at http://www.apache.org/docs/dso.html.

Once you've compiled your module using APXS, put your module into `/usr/lib/apache`. Then your module needs both a `LoadModule` line and an `AddModule` line in the `httpd.conf` file, just as described previously for Apache's own modules. After the `LoadModule` list in `httpd.conf`, add a line for the shared object file for your module like the following:

```
LoadModule foo_module  modules/mod_foo.so
```

Note that you'll need to change the name of the module and the name of your shared object file as appropriate.

At the end of the `AddModule` list in `httpd.conf`, add a line for the source code file for your module like the following:

```
AddModule mod_foo.c
```

Note that you'll need to change the name of the source code file as appropriate.

Once you've completed the previous steps, stop and start your Web server as outlined in Section 11.11, *Starting and Stopping Apache*. If you've done everything correctly, and your module is correctly coded, your Web server should find your module and load it in as it starts.

12.2.1 The mod_ssl Security Module

The mod_ssl security portion of the Red Hat Linux Secure Web Server is provided as a Dynamic Shared Object (DSO). This means that the Apache Web server can be re-compiled by users if the EAPI extension patch from the mod_ssl security module is

applied to Apache. Follow the instructions for building mod_ssl into Apache included with the mod_ssl documentation, but add the following flag:

```
--with-eapi-only
```

The complete command line should look like the following:

```
./configure [userflags] --with-eapi-only
```

Then build and install Apache.

Please Note

Red Hat cannot support re-compiled versions of the Apache Web server. Installation of the shipped version is supported, but if you re-compile Apache, you're on your own. Please don't re-compile Apache unless you know exactly what you're doing.

12.3 Using Virtual Hosts

You can use Apache's virtual hosts capability to run different servers for different IP addresses, different host names or different ports on the same machine. If you're interested in using virtual hosts, complete information is provided in the Apache documentation on your machine or on the Web at http://www.apache.org/docs/vhosts/.

Please Note

You can't use name-based virtual hosts with your Red Hat Linux Secure Web Server, because the SSL handshake (when the browser accepts the secure Web server's certificate) occurs before the HTTP request which identifies the appropriate name-based virtual host. If you want to use name-based virtual hosts, they will only work with your non-secure Web server.

Virtual hosts are configured within the `httpd.conf` file, as described in Section 12.1, *Configuration Directives in `httpd.conf`*. Please review that section before you start to change the virtual hosts configuration on your machine.

12.3.1 The Red Hat Linux Secure Web Server Virtual Host

The default configuration of your Red Hat Linux Secure Web Server runs a non-secure and a secure server. Both servers use the same IP address and host name, but they listen on different ports, and the secure server is a virtual host. This configuration enables you to serve both secure and non-secure documents in the most efficient manner possible. As you may know, secure HTTP transmissions take more time than non-secure, because a lot more information is being passed back and forth during secure transactions. So using your secure server for non-secure Web traffic is not a good idea.

The configuration directives for your secure server are contained within virtual host tags in the `httpd.conf` file. If you need to change something about the configuration of your secure server, you'll need to change the configuration directives inside virtual host tags in the `httpd.conf` file. If you want to enable certain features (for example, server side includes) for your secure server, they will need to be enabled within the virtual host tags that define your secure server.

The non-secure Web server is configured as the "non-virtual" host in the `httpd.conf` file. In other words, the non-secure Web server's configuration options are outside of the virtual host tags in `httpd.conf`. If you want to

change something about your non-secure Web server, you'll need to change the configuration directives in `httpd.conf` outside of the virtual host tags.

By default, both the secure and the non-secure Web servers share the same `DocumentRoot`, a configuration directive specified in `httpd.conf`. In other words, the secure and the non-secure Web server look in the same place for the HTML files that they provide in response to requests. By default, the `DocumentRoot` is set to `/var/www/html`.

To change the `DocumentRoot` so that it is no longer shared by both the secure server and the non-secure server, change one of the `DocumentRoot` directives in `httpd.conf`. The `DocumentRoot` outside the virtual host tags defines the `DocumentRoot` for your non-secure Web server. The `DocumentRoot` within the virtual host tags that define your secure server is (obviously) for your secure server.

If for some reason you want to disable the non-secure Web server on your machine, you can. Your secure server listens on port 443, the default port for secure Web communications, while your non-secure Web server listens on port 80, the default port for non-secure Web communications. To stop the non-secure Web server from accepting connections, in `httpd.conf`, find the line which reads:

```
Port 80
```

Change the above line so that it reads:

```
Port 443
```

Then comment out the `Listen 80` line, so that instead of:

```
Listen 80
```

Change the above line so that it reads:

```
#Listen 80
```

After these two steps, your Red Hat Linux Secure Web Server will be accepting connections on port 443, the default port for secure Web communications. However, your server will not accept connections on port 80, the default port for non-secure communications, so the non-secure Web server will be effectively disabled.

12.3.2 Setting Up Virtual Hosts

Most people will probably use their Red Hat Linux Secure Web Server as it is configured. Therefore, they'll be using the built-in virtual hosts capability, but they won't have to do any manipulation of the virtual hosts directives in `httpd.conf`. However, if you would like to use the virtual hosts capability for some other reason, you can.

To create a virtual host, you'll need to alter the virtual host lines, provided as an example, in `httpd.conf`, or create your own virtual host section. (Remember that name-based virtual hosts won't work with your secure server — you'll need to use IP address-based virtual hosts if you need SSL-enabled virtual hosts. Your non-secure server, however, will support both IP address and name-based virtual hosts.)

The virtual host example lines read as follows:

```
#<VirtualHost ip.address.of.host.some_domain.com>
#    ServerAdmin webmaster@host.some_domain.com
#    DocumentRoot /www/docs/host.some_domain.com
#    ServerName host.some_domain.com
#    ErrorLog logs/host.some_domain.com-error_log
#    CustomLog logs/host.some_domain.com-access_log common
#</VirtualHost>
```

Uncomment all of the lines (remove the # from the beginning of each line). Then add the correct information for your machine and/or your virtual host to each line.

In the first line, change `ip.address.of.host.some_domain.com` to your server's IP address. Change the `ServerName` to a *valid* DNS name to use for the virtual host. (In other words, don't just make something up. Ask your system administrator if you don't know how to get a valid domain name.)

You'll also need to uncomment one of the `NameVirtualHost` lines in `httpd.conf`:

```
#NameVirtualHost 12.34.56.78:80
#NameVirtualHost 12.34.56.78
```

Uncomment one of the lines and change the IP address to the IP address (and port if necessary) for that virtual host.

Many other configuration directives can be placed between the virtual host tags, depending upon why you're setting up a virtual host.

If you set up a virtual host and want it to listen on a non-default port (80 is the default port for non-secure Web communications; 443 is the default port for secure Web communications), you'll need to set up a virtual host for that port and add a `Listen` directive to `httpd.conf`, corresponding to that port.

To have a virtual host work specifically for that port, add the port number to the first line of the virtual host configuration. The first line should look something like the following:

```
<VirtualHost ip_address_of_your_server:12331>
```

This line would create a virtual host that listens on port 12331. Substitute the port number you want to use for `12331` in the previous example.

Underneath the `Listen` lines in `httpd.conf`, add a line like the following, which will instruct your Web server to listen on port 12331:

```
Listen 12331
```

You must restart your server to start a new virtual host.

Much more complete information about creating and configuring both name-based and IP address-based virtual hosts is provided on the Web at http://www.apache.org/docs/vhosts/index.html. Please check the Apache Group's virtual host documentation for more details on using virtual hosts.

Part III Installation-Related Reference

13 Preparing for a Text Mode Installation

13.1 Things You Should Know

Before attempting to install Red Hat Linux, you should collect some system information in order to 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 hardware supported by Red Hat Linux can be found at http://www.redhat.com/hardware. It's a good idea to check your hardware against this list before proceeding.

Please Note

You can perform a text mode installation of Red Hat Linux 7.0 by following instructions in this chapter and in Chapter 14, *Installing Red Hat Linux via Text Mode*. However, if you're installing from a CD-ROM, you might prefer to use the graphical installation mode, which offers both ease of use and a flexible, custom-class installation mode. For more information on graphical installations, turn to the *Official Red Hat Linux Installation Guide*.

Tip

At the end of the *Before You Begin* chapter in the *Official Red Hat Linux Installation Guide* is a table for you to fill out with your specific system requirements, which will help you keep up with any information you will need during your installation.

13.1.1 Basic Hardware Configuration

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

- hard drive(s) -- Specifically, the number, size, and type. If you have more than one, it's helpful to know which one is first, second, and so on. It is also good to know if your drives are IDE or SCSI. 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 ATAPI) are the most common type 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's a good idea to collect this information anyway, just to be sure.

Learning About Your Hardware with Windows

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

Figure 13–1 Windows System Properties

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

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

Figure 13–2 Device Manager Under Windows 95

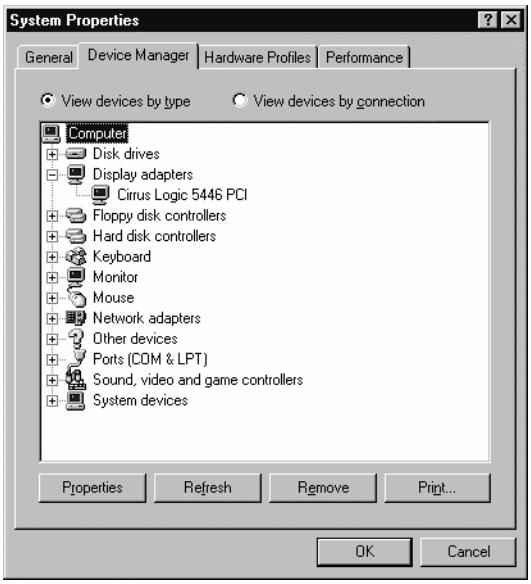

- **Disk drives** -- You will find the type (IDE or SCSI) of hard drive here. (IDE drives will normally include the word "IDE," while SCSI drives won't.)
- **Hard disk controllers** -- You can get more information about your hard drive controller here.
- **CDROM** -- Here is where you'll find out about any CD-ROM drives connected to your computer.

Please Note

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

- **Mouse** -- The type of mouse present on your computer can be found here.
- **Display adapters** -- If you're interested in running the X Window System, you should write down the information you find here.
- **Sound, video and game controllers** -- If your computer has sound capabilities, you'll find more information about that here.
- **Network adapters** -- Here you'll find additional information on your computer's network card (if you have one).
- **SCSI controllers** -- If your computer uses SCSI peripherals, you'll find additional info on the SCSI controller here.

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

Please Note

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

13.1.2 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.)

13.1.3 Network-related Information

If you're 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 -- Yet another set of four dot-separated numbers. For instance, `10.0.2.254`.
- one or more name server IP addresses -- One or more sets of dot-separated numbers. For example, `10.0.2.1` might be the address of a name server.
- domain name -- The name your organization uses. For instance, Red Hat has a domain name of `redhat.com`.
- hostname -- The name assigned to your individual system. A computer might be named `pooh`, for instance.

Please Note

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

14 Installing Red Hat Linux via Text Mode

This release of Red Hat Linux features a graphical, mouse-based installation program, documented in the *Official Red Hat Linux Installation Guide*. But you can also install Red Hat Linux with a text mode, keyboard-based installation program. This chapter explains how to use it. Here are some recommendations:

- If you're new to Linux installations, read the *Official Red Hat Linux Installation Guide* first. Although it focuses on the graphical installation process, most of the concepts apply to the text mode installation as well. After that, you'll find that Chapter 13, *Preparing for a Text Mode Installation*, *Preparing for a Text Mode Installation*, will give you more in-depth information regarding those aspects of installing Red Hat Linux that do not apply to the graphical installation process.

  Additionally, you'll find Appendix B, *An Introduction to Disk Partitions*, *An Introduction to Disk Partitions*, helpful, as it discusses disk partition resizing (crucial if you plan to install Linux on a disk where another operating system is currently installed).

- If you'll need PCMCIA support to perform the installation (for example, you're installing on a laptop equipped with a PCMCIA card), you must make a PCMCIA boot disk. The *Official Red Hat Linux Installation Guide* explains how to do this.

- If you plan to install over a network (via NFS, FTP, or HTTP), you must make a network boot disk. The *Official Red Hat Linux Installation Guide* explains how to do this.

- If you've never used the text mode installation program, or need a refresher on its user interface, read the next section.

- To begin installation without further delay, turn to Section 14.2, *Starting the Installation Program*.

14.1 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 14–1, *Installation Program Widgets as seen in* **Configure TCP/IP** and Figure 14–2, *Installation Program Widgets as seen in Disk Druid* illustrate the screens you'll see.

Figure 14–1 Installation Program Widgets as seen in Configure TCP/IP

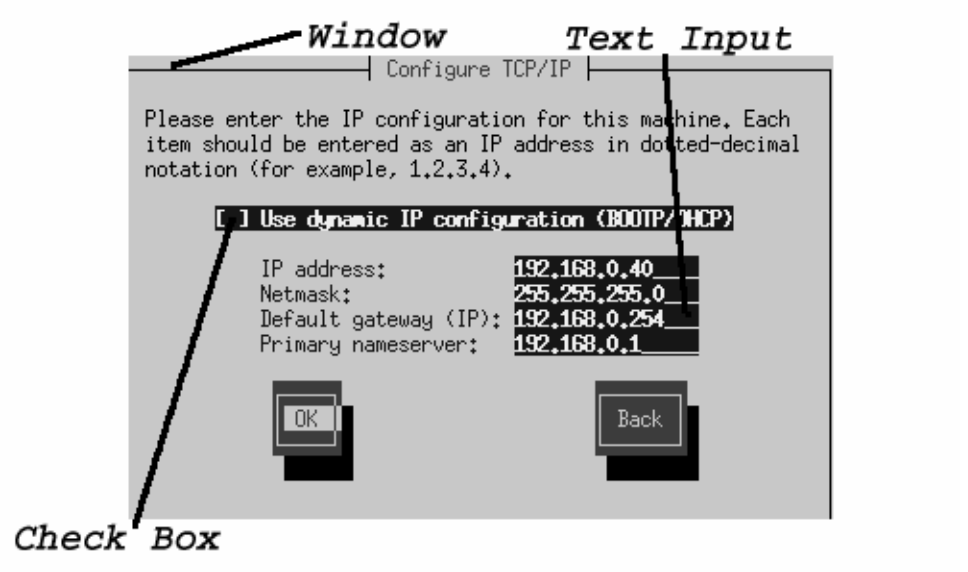

Figure 14–2 Installation Program Widgets as seen in Disk Druid

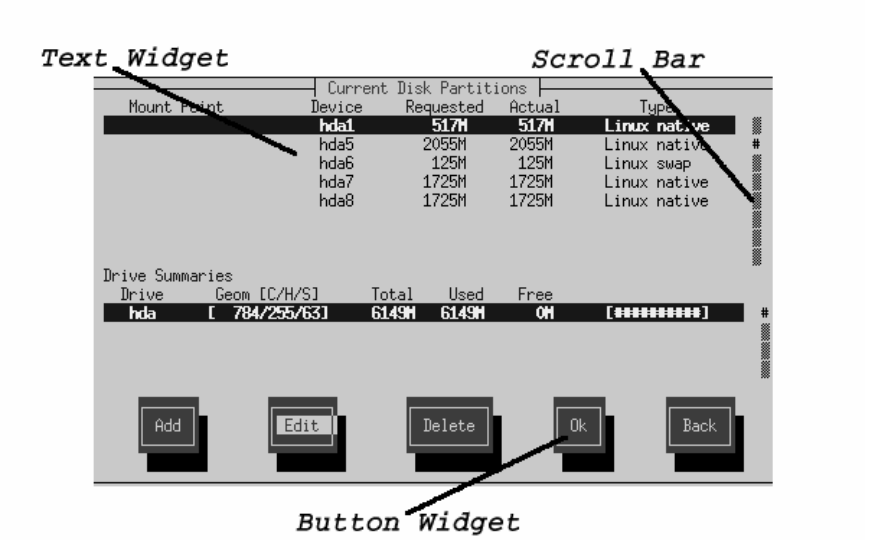

Here's a list of the most important widgets shown in Figure 14–1, *Installation Program Widgets as seen in* **Configure TCP/IP** and Figure 14–2, *Installation Program Widgets as seen in* Disk Druid :

- 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.
- Check Box -- Check boxes 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 check box, 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 check boxes. 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 14–1, *Installation Program Widgets as seen in* **Configure TCP/IP**, the cursor is positioned on the **OK** button. Figure 14–2, *Installation Program Widgets as seen in Disk Druid* shows the cursor on the **Edit** button.

14.1.1 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. Most screens display along the bottom a summary of available cursor positioning keys.

To "press" a button, position the cursor over the button (using [Tab], for instance) 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 **check box**, move the cursor to the check box and press [Space] to select an item. To deselect, press [Space] a second time.

Pressing [F12] accepts the current values and proceeds to the next dialog; it is 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).

14.2 Starting the Installation Program

There are several ways to start the installation. You can:

- insert the diskette included in your boxed set (or a PCMCIA boot or network boot disk that you've created) into the primary diskette drive and reboot your computer
- insert the Red Hat Linux CD into the drive and reboot, if your computer can boot from the CD-ROM drive
- boot MS-DOS, and start a program in the `dosutils` directory of the Red Hat Linux CD-ROM named `autoboot.bat` (this will work from DOS only; it will not work from a DOS window started from Windows)

While the installation program loads, messages will scroll on your screen. When the installation program has loaded, this prompt appears:

```
boot:
```

14.2.1 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.

14.2.2 Text Mode Boot Options

If you press [Enter] at the `boot` prompt, or if you take no action within the first minute after the `boot` prompt appears, the graphical installation program, as explained in the *Official Red Hat Linux Installation Guide*, starts. Pressing one of the help screen function keys as described above disables this autostart feature.

To start the text mode installation program explained here, before pressing [Enter], type:

```
boot: text
```

If the installation program does not properly detect your hardware, you may need to restart the installation in "expert" mode. To start an expert mode installation, type:

```
boot: text expert
```

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

Please Note

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

Note that the command to start a serial installation has changed. If you must perform the installation in serial mode, type:

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

Where <*device*> should be the device you are using (such as ttyS0 or ttyS1).

Other options that may be entered with the boot command include passing options to the kernel. For example, to instruct the kernel to use all the RAM in a 128 MB system, enter:

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

To explicitly request a dialog where you can configure additional devices (such as ISA devices) include the 'isa' directive:

```
boot: linux isa
```

14.3 Choosing a Language

Using the [Up] and [Down] arrow keys, select the appropriate language for both the installation program and the system default, and press [Enter] (Figure 14–3, *Selecting a Language*).

Figure 14–3 Selecting a Language

A scroll bar appears to the right of the list. This indicates that there are more entries than can be displayed in the available space at once. You'll see other scroll bars like this throughout the installation program.

14.4 Selecting a Keyboard Type

Next, choose a keyboard type (Figure 14–4, *Selecting a Keyboard Type*).

Figure 14–4 Selecting a Keyboard Type

After selecting the appropriate keyboard type, press [Enter]; the keyboard type you select will be loaded automatically both for the remainder of the installation process and each time you boot your Red Hat Linux system.

Tip

If you wish to change your keyboard type after you have installed your Red Hat Linux system, become root and type either `/usr/sbin/kbdconfig` or `setup` at the root prompt.

14.5 Selecting an Installation Method

If you booted directly from the Red Hat Linux CD-ROM, you'll see the **Welcome** dialog: turn to Section 14.8, *Welcome*.

Otherwise, an **Installation Method** dialog appears. The choices presented in the dialog vary depending on the type of diskette you booted from (the one in your boxed set, or a network or PCMCIA boot disk that you created).

Figure 14–5, *Installation Method Dialog* shows all available choices.

Figure 14–5 Installation Method Dialog

Please Note

If you are performing a network installation and are copying the files from the Red Hat Linux CD-ROM or an FTP site, 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.

Red Hat Linux can be installed via any of the following:

Local CDROM

If you booted from the diskette in your boxed set and are installing with the Red Hat Linux CD-ROM in a local drive. If you choose this method, you'll next see the **Welcome** dialog: turn to Section 14.8, *Welcome*.

NFS Image

If you booted from a network or PCMCIA boot disk and are installing from an NFS Image server which is exporting the Red Hat Linux CD-ROM (or a mirror image of one.) If you choose this method, you'll next see the **NFS Setup** dialogs: turn to Section 14.7, *Installing over a Network.*

FTP

If you booted from a network or PCMCIA boot disk and are installing directly from an FTP server. If you select this method, you'll next see the **FTP Setup** dialogs: turn to Section 14.7, *Installing over a Network.*

HTTP

If you booted from a network or PCMCIA boot disk and are installing directly from an HTTP Web server. If you select this method, you'll next see the **HTTP Setup** dialogs: turn to Section 14.7, *Installing over a Network.*

Hard Drive

If you booted from the diskette in your boxed set and are installing from the Red Hat Linux files that you have previously copied to a local hard drive. If you select this method, you'll next see the **Select Partition** dialog: turn to Section 14.6, *Identify Disk Partition to Install From*.

14.6 Identify Disk Partition to Install From

The **Select Partition** screen (Figure 14–6, *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 you are installing from.

Figure 14–6 Selecting Partition Dialog for Hard Drive Installation

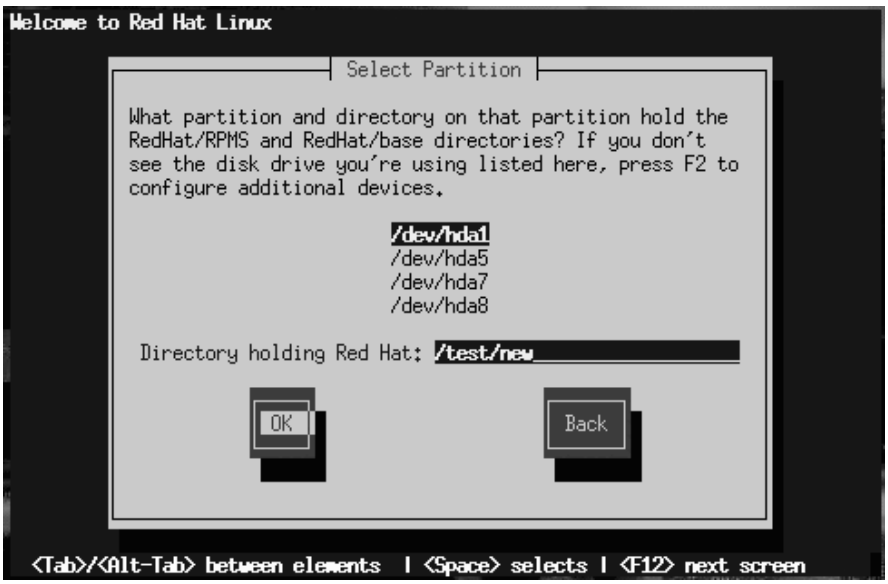

Enter the device name of the partition containing the `RedHat` directory tree. There is also a field labelled **Directory holding Red Hat**. If the `RedHat` directory is not in the root directory of that partition, enter the path to the `RedHat` directory (for example, if the `RedHat` directory is at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

After you've identified the disk partition, you'll next see the **Welcome** dialog: Turn now to Section 14.8, *Welcome*.

14.7 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 14.14, *Configuring a Network Connection* and then return here.

14.7.1 NFS Setup

The NFS setup dialog (Figure 14–7, *NFS Setup Dialog*) applies only if you are installing from an NFS server (that is, if you booted from a network or PCMCIA boot disk and selected **NFS Image** in the **Installation Method** dialog).

Figure 14–7 NFS Setup Dialog

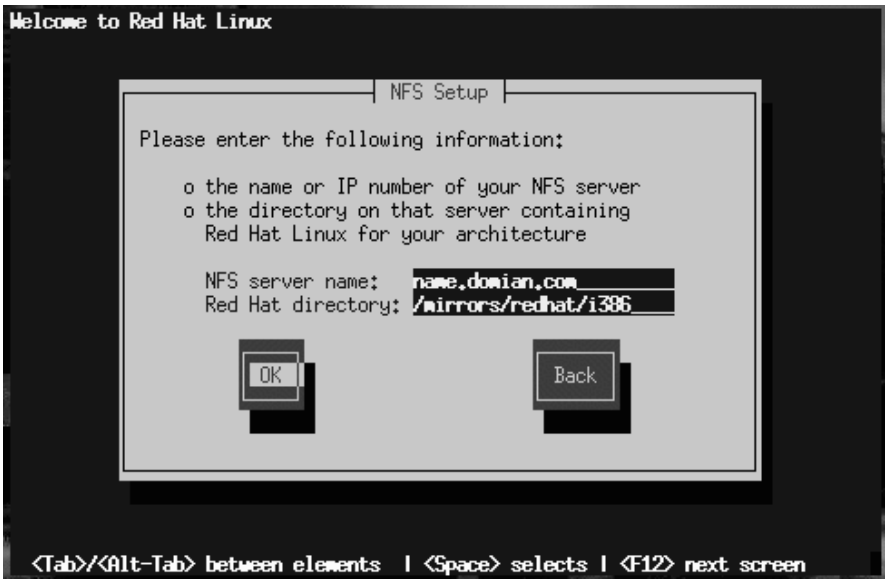

Enter the fully-qualified domain name or IP address of your NFS server, and the name of the exported directory that contains the Red Hat Linux installation files or CD. For example, if you're installing from a host named `eastcoast` in the domain `redhat.com`, enter `eastcoast.redhat.com` in the **NFS 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/i386/RedHat`, enter `/mirrors/redhat/i386`.

Please 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'll see the **Welcome** dialog: Turn now to Section 14.8, *Welcome*.

14.7.2 FTP Setup

The **FTP Setup** screen (Figure 14–8, *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 you are installing from.

Figure 14–8 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 there which contains the `RedHat` installation files for your architecture. For example, if the FTP site contains the directory `/pub/mirrors/redhat/i386/RedHat`, enter **`/pub/mirrors/redhat/i386`**.

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

Next you'll see the **Welcome** dialog: Turn now to Section 14.8, *Welcome*.

14.7.3 HTTP Setup

The **HTTP Setup** screen (Figure 14–9, *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 you are installing from.

Figure 14–9 HTTP Setup Dialog

Enter the name or IP address of the HTTP site you are installing from, 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 `/pub/mirrors/redhat/i386/RedHat`, enter **`/pub/mirrors/redhat/i386`**.

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

Next you'll see the **Welcome** dialog: Turn now to Section 14.8, *Welcome*.

14.8 Welcome

Once you've selected your language and keyboard, and indicated the installation method, Figure 14–10, *Welcome dialog* appears. Press **OK** to continue.

Figure 14–10 Welcome dialog

14.9 Upgrading or Installing

The **Installation Type** dialog (Figure 14–11, *Installation Type Dialog*) presents you with five choices, described below.

Figure 14–11 Installation Type Dialog

- **Install Workstation** -- Install on a system that will be used primarily as a workstation. Load the GNOME (and/or KDE) GUI and configure the system to start GNOME (or KDE) as the desktop default. The installation program deletes all data in any existing Linux partitions, decides how to partition the disk for the new version, and chooses which software packages to load.

WARNING

Do not choose this method if you're sharing a disk with Windows NT; if you do, you will be unable to boot Windows NT. LILO will write over NT's boot loader and you will be unable to boot NT. You must perform a custom-class installation and configure LILO so that it is not installed on the Master Boot Record (MBR).

To create a dual-boot environment on a system that currently has NT, you must install LILO on the first sector of the root partition, not the MBR. Please be sure to create a boot disk. In a case such as this, you will either need to use the boot disk, or configure the NT system loader to boot LILO from the first sector of the root partition. Be sure to check out http://www.linux-doc.org/HOWTO/mini/Linux+NT-Loader.html for more information on setting up LILO and NT.

Below are the minimum recommended disk space requirements for a workstation-class installation.

– Workstation choosing GNOME -- 900M

– Workstation choosing KDE -- 900M

– Workstation choosing both GNOME and KDE -- 1.1G

If you plan to choose all group packages (for example, GNOME is one package group), as well as select additional individual packages, you may want to allow yourself 1.7G or more of disk space. This is also allow for some room where additional data may be written.

- **Install Server System** -- Install on a system that will be used primarily as a server. The X Window System is not configured and no GUI starts when the system boots. The Installation program deletes *all* data in *all* existing partitions of any kind, decides how to partition the disk for the new version, and chooses which software packages to load.

  Below are the recommended disk space requirements for a server-class installation.

  – Server (minimum) -- 450M

  – Server (choosing everything) -- 1G

  If you plan to choose all group packages, as well as select additional individual packages, you may want to allow yourself 1.7G or more of disk space. This is also allow for some room where additional data may be written.

- **Install Custom System** -- Perform a custom installation. You make all decisions regarding disk partitioning and initialization, which software packages to install, and how to configure the X Window System and the user interface.

  Below are the recommended disk space requirements for a custom-class installation.

  – Custom (minimum) -- 250M

  – Custom (choosing everything) -- 1.7G

- **Upgrade Existing System** -- Upgrade an earlier version of Red Hat Linux (3.0.3 or later) without deleting any existing data. The installation program updates the modular 2.2.x kernel and all currently installed software packages.

14.9.1 Upgrading

If you choose to upgrade and the installation program detects more than one installed Linux version on the system, you'll be asked which version to upgrade. After you indicate this, or if there's only one installed Linux version on the system, the installation program probes your existing system to determine which software packages require updating and presents the **Customize Packages to Upgrade** dialog (Figure 14–12, *Customize Packages to Upgrade Dialog*).

Figure 14–12 Customize Packages to Upgrade Dialog

If you answer **No**, the installation program starts upgrading existing packages.

Answer **Yes** if you want to add to or remove items from the list of individual packages to be upgraded. The package selection dialog is seen in Section 14.20.1, *Selecting Individual Packages*. The upgrade starts when you finish making your changes.

Please Note

Some upgraded packages may require that other packages are also installed for proper operation. The upgrade procedure takes care of these **dependencies**, but in doing so it may need to install additional packages which are not on your existing system.

The upgrade process preserves existing configuration files by renaming them using a `.rpmsave` extension (e.g., `sendmail.cf.rpmsave`) and leaves a log telling

what actions it took in `/tmp/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.

The next dialog you'll see is Figure 14–40, *Package Installation Status Dialog*. This dialog remains on the screen until the upgrade is complete.

14.10 Automatic Partitioning

If you choose a workstation- or server-class installation, Figure 14–13, *Automatic Partitioning Dialog* appears.

Figure 14–13 Automatic Partitioning Dialog

If you select **Continue** and press **OK**, the installation program partitions your disk and decides which software packages to install. Next, you'll see the **Hostname** dialog explained in Section 14.13, *Naming Your Computer*.

If you select **Manually partition** or choose to perform a custom-class installation, Disk Druid will begin. The disk partitioning dialogs described in the next section will

appear, showing you any current partitions on your system. It is up to you to indicate the partitions and mount points to be used for installation of this version of Red Hat Linux.

Installing Red Hat Linux over another installation of Linux (including Red Hat Linux) does *not* preserve any information (files or data) from the prior installation. Make sure you save any important files! If you are worried about saving the current data on your existing Red Hat Linux system (without making a backup on your own), you should consider performing an upgrade instead.

14.11 Partitioning Your Disk for Red Hat Linux

If you have not yet planned how you will set up your partitions, turn to Appendix B, *An Introduction to Disk Partitions*. There you'll find an introduction to basic disk partitioning concepts. As a bare minimum, you'll need an appropriately-sized root partition, and a swap partition of at least 16 MB.

Figure 14–14, *Disk Setup Dialog* shows the two disk partitioning applications that are available for you to use.

If you will be using fdisk to partition your drive, please see Section 14.11.12, *Using fdisk* for those instructions. If you select Disk Druid, continue reading below.

Figure 14–14 Disk Setup Dialog

The following sections describe the layout of Figure 14–15, *Disk Druid Main Screen* and how to use its buttons to set up partitions. If you're already familiar with Disk Druid, you can partition your disk and skip to Section 14.11.11, *Choose Partitions to Format Dialog*.

You use the disk partitioning dialogs to tell the installation program where to install Red Hat Linux (Figure 14–15, *Disk Druid Main Screen*).

Figure 14–15 Disk Druid Main Screen

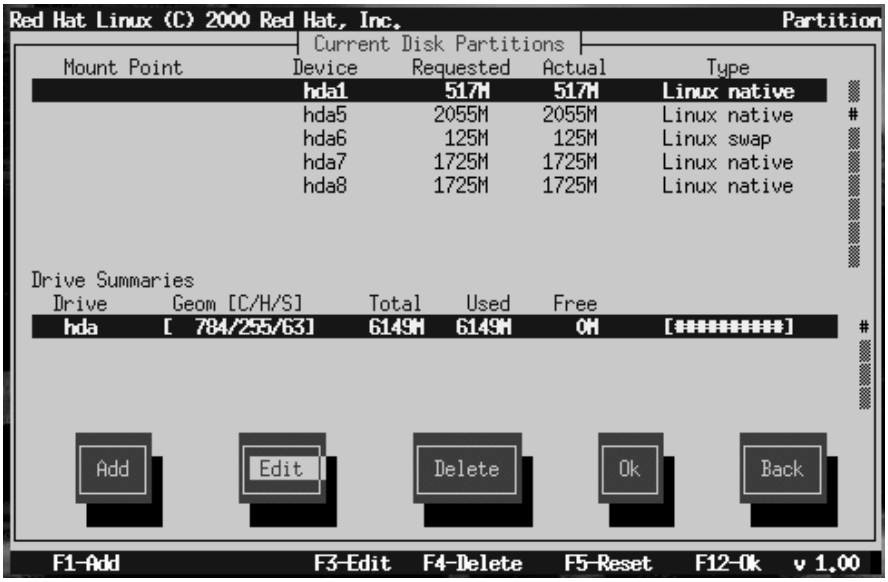

14.11.1 The Current Disk Partitions Section

Each line in the **Current Disk Partitions** section represents a disk partition. In this example (Figure 14–15, *Disk Druid Main Screen*), there are Linux existing partitions.

Note the scroll bar to the right, which indicates that there may be more partitions than can be displayed at one time. If you use the [Up] and [Down] arrow keys, you can see if there are any additional partitions. Each line in this section has five different fields:

- **Mount Point** -- Indicates where the partition will be mounted when Red Hat Linux is installed (such as `/`, `/boot`, or `swap`).
- **Device** -- Displays specific hard drive and partition information.
- **Requested** -- Shows the partition's initial size.
- **Actual** -- Shows the partition's current size.
- **Type** -- Shows the partition's type.

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

14.11.2 The Drive Summaries Section

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

- **Drive** -- Shows the hard disk's device name.
- **Geom [C/H/S]** -- 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** -- Shows the total available space on the hard disk.
- **Used** -- Shows how much of the hard disk's space is currently allocated to partitions.
- **Free** -- Shows how much of the hard disk's space is still unallocated.
- **Bar Graph** -- Presents a visual representation of the space currently used on the hard disk. The more pound signs there are between the square braces, the less free space there is. In Figure 14–15, *Disk Druid Main Screen*, the bar graph shows no free space.

Please Note

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. This is described more completely in Section 14.11.5, *Adding a Partition.*

14.11.3 Disk Druid's Buttons

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

- **Add** -- Request a new partition. Selecting this button causes a dialog box to appear containing fields that must be filled in.
- **Edit** -- Modify the mount point of the partition currently highlighted in the **Current Disk Partitions** section. Selecting this button will cause a dialog box to appear allowing you to change the name of the mount point.
- **Delete** -- Delete the partition currently highlighted in the **Current Disk Partitions** section. Selecting this button will cause a dialog box to appear asking you to confirm the deletion.
- **OK** -- Confirm that changes made to your system's partitions to be written to disk. You will be asked to confirm the changes before Disk Druid rewrites your hard disk partition table(s). In addition, any mount points you've defined are passed to the installation program, and will eventually be used by your Red Hat Linux system to define the filesystem layout.
- **Back** -- Abort without saving any changes you've made. When this button is selected, the installation program will take you back to the previous screen, so you can start over.

14.11.4 Handy Function Keys

Use the [F5] (Reset) function key to discard all changes you may have made while in Disk Druid, and return the list of partitions to those read from the partition table(s) on your hard disk(s). When selected, you'll be asked to confirm whether you want to discard the changes. Note that any mount points you've specified will be lost, and will need to be reentered.

Please Note

You will need to dedicate at least one partition to Red Hat Linux, and optionally more. This is discussed more completely in Section B.1.8, *How Many Partitions?*.

14.11.5 Adding a Partition

To add a new partition, select the **Add** button and press [Space] or [Enter]. The **Edit New Partition** dialog (Figure 14–16, *Edit New Partition Dialog*) appears.

Figure 14–16 Edit New Partition Dialog

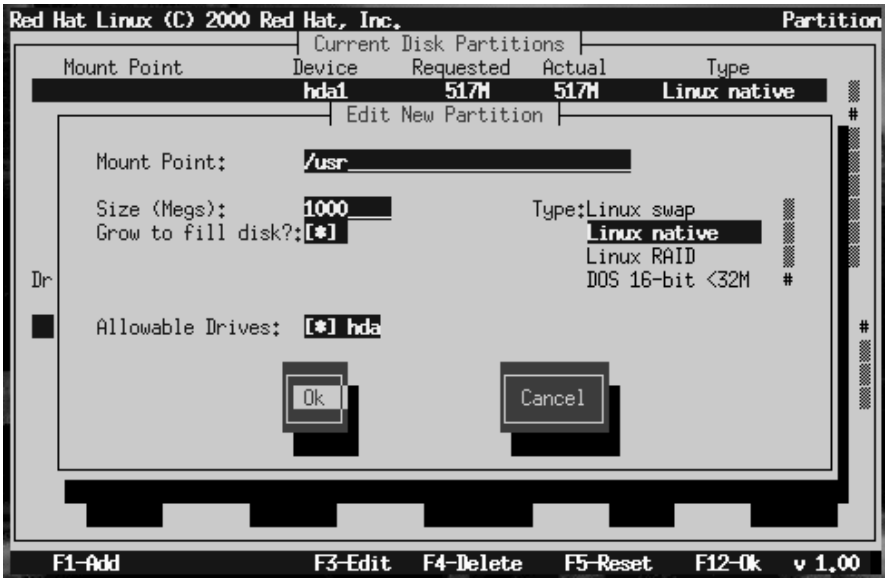

The screen contains the following fields:

- **Mount Point** -- Highlight this field and enter the partition's mount point. For example, if this partition should be the root partition, enter `/`; enter `/usr` for the `/usr` partition, and so on.

- **Size (Megs)** -- In this field, enter the size (in megabytes) of the partition. Note that this field starts with a "1" in it, meaning that unless you change it, you'll end up with a 1 MB partition. Delete it using the [Backspace] key, and enter the desired partition size.
- **Grow to fill disk?** -- This check box indicates whether the size you entered in the previous field is to be considered the partition's exact size, or its minimum size. Press [Space] to select this option. When selected, the partition will grow to fill all available space on the hard disk. In this case, the partition's size will expand and contract as other partitions are modified. If you make more than one partition growable, the partitions will compete for the available free space on the disk.
- **Type** -- This field contains a list of different partition types. Select the appropriate partition type by using the [Up] and [Down] arrow keys.
- **Allowable Drives** -- This field contains a list of the hard disks installed on your system, with a check box for each. If a hard disk's box is checked, then this partition may be created on that hard disk. By using different check box settings, you can direct Disk Druid to place partitions as you see fit, or let Disk Druid decide where partitions should go.
- **OK** -- Select this button and press [Space] when you are satisfied with the partition's settings, and wish to create it.
- **Cancel** -- Select this button and press [Space] when you don't want to create the partition.

14.11.6 Recommended Partitioning Scheme

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

- A swap partition (at least 16MB) -- 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 your computer has 16MB of RAM or less, you *must* create a swap partition. Even if you have more memory, a swap partition is still recommended. The minimum size of your swap partition should be equal to your computer's RAM, or 16MB (whichever is larger). In Disk Druid, the partition field for swap should look similar to:

```
<Swap>  hda2  125M   125M  Linux swap
```

- A /boot partition (16MB, maximum) -- The partition mounted on /boot contains the operating system kernel (which allows your system to boot Red Hat Linux), along with files used during the bootstrap process. Due to the limitations of most PC BIOSes, creating a small partition to hold these files is a good idea. This partition should be no larger than 16MB. In Disk Druid, the partition field for /boot should look similar to:

  ```
  /boot  hda1  16M   19M  Linux native
  ```

- A root partition (900MB-1.7GB) -- This is where "/" (the root directory) resides. In this setup, all files (except those stored in /boot) reside on the root partition. A 850MB root partition will permit the equivalent of a workstation-class installation (with *very* little free space), while a 1.7GB root partition will let you install every package. In Disk Druid, the partition field for / should look similar to:

  ```
  /  hda2  900M   3669M   Linux native
  ```

14.11.7 Problems When Adding a Partition

Please Note

If you are having problems adding a partition, turn to Appendix B, *An Introduction to Disk Partitions*, *An Introduction to Disk Partitions*, to find a solution.

If you attempt to add a partition and Disk Druid can't carry out your request, you'll see a dialog box listing any partitions that are currently unallocated, along with the reason they could not be allocated, as in Figure 14–17, *Unallocated Partitions Dialog*. Select the **OK** button, and press [Space] to continue. Note that the unallocated partition(s) are also displayed on Disk Druid's main screen (though you may have to scroll the **Current Disk Partitions** section to see them).

Figure 14–17 Unallocated Partitions Dialog

14.11.8 Editing a Partition

To change a partition's mount point, highlight the partition in the **Current Disk Partitions** section, select the **Edit** button, and press [Space]. The **Edit Partition** dialog is similar to the one shown in Figure 14–16, *Edit New Partition Dialog*. The difference here is that all fields except the mount point are read-only. To modify any other value, delete the partition and add it again with the new values.

Figure 14–18 Edit Partition Dialog

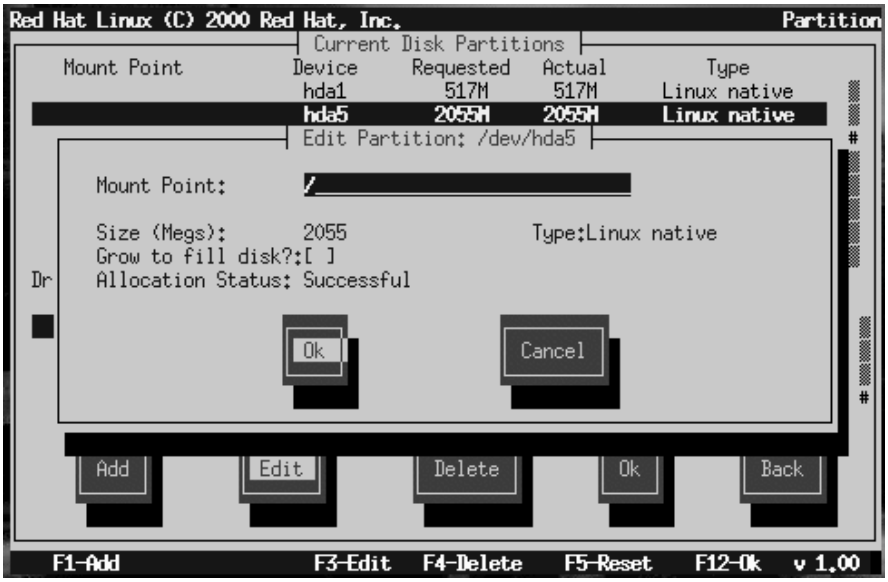

14.11.9 Deleting a Partition

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

14.11.10 When You're Finished

Once you've configured your partitions and entered your mount points, you screen should look something like Figure 14–19, *Current Disk Partitions Dialog*.

Figure 14–19 Current Disk Partitions Dialog

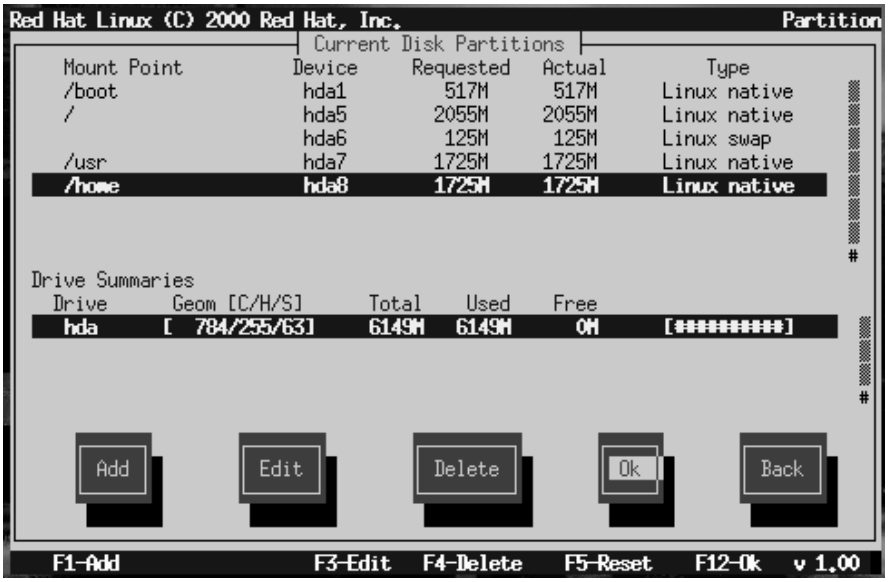

14.11.11 Choose Partitions to Format Dialog

Next, select which partitions you want to format (Figure 14–20, *Choose Partitions to Format*). You must format all newly created partitions and other partitions that contain old data (assuming they don't contain data you wish to keep).

Figure 14–20 Choose Partitions to Format

CAUTION

If partitions such as `/home` or `/usr/local` already exist and contain data you wish to keep, do not select these for formatting.

When you have selected the partitions to format, press [Space]. If you wish to check for bad blocks while formatting each filesystem (recommended for those with older disk drives), select **Check for bad blocks during format**. Select **OK**, and press [Space].

Please Note

Selecting **check for bad blocks** may dramatically increase your total installation time. Since most newer hard drives are quite large in size, checking for bad blocks may take a while depending on the size of your hard drive.

14.11.12 Using fdisk

If you chose a custom-class installation, you also chose which disk partitioning application to use. This section only applies if you opted to use fdisk.

Once you've selected fdisk, you'll be presented with the **Partition Disks** dialog box (Figure 14–21, *Disk Setup Dialog*). In this box is a list of every disk on your computer.

Using [Tab], and the [Up] and [Down] arrow keys, highlight the disk you'd like to partition, select **Edit**, and press [Space].

You will then enter fdisk and can partition the disk you selected. Repeat this process for each disk you want to partition. When you're finished, select **Done**.

Figure 14–21 Disk Setup Dialog

An Overview of fdisk

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

- The command for help is m.
- To list the current partition table, use the p command (see Figure 14–22, *Sample Output from fdisk*).
- To add a new partition, use n.
- Linux fdisk creates partitions of type **Linux native** by default. When you create a swap partition, don't forget to change it to type **Linux swap** using the t command. The value for the **Linux swap** type is 82. For other partition types, use the l command to see a list of partition types and values.
- Linux allows up to four (4) partitions on one disk. If you wish to create more than that, one (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's a good idea to write down which partitions (e.g., `/dev/hda2`) are meant for which filesystems (e.g., `/usr`) as you create each one.

Please Note

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

Figure 14–22 Sample Output from fdisk

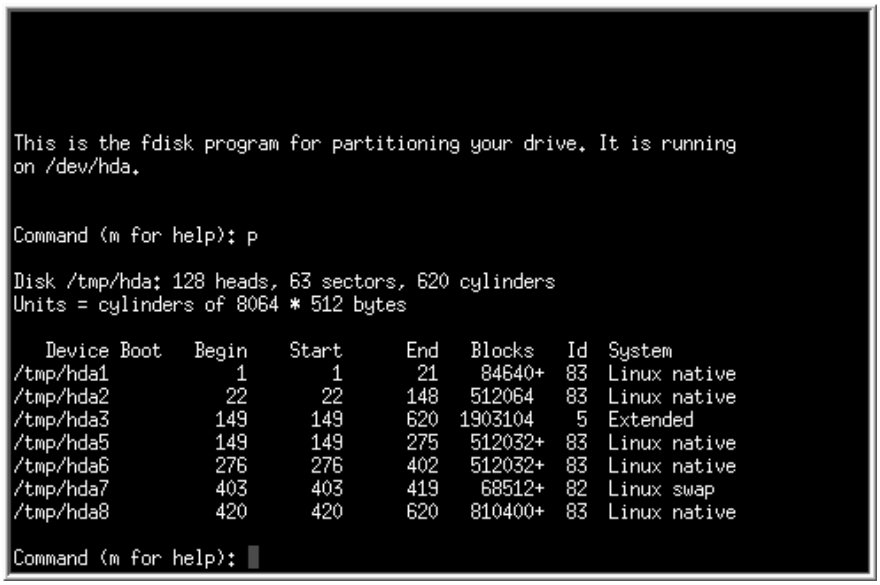

Changing the Partition Table

When you are finished partitioning your disks, press **Done**; you may see a message indicating that the installation program needs to reboot. This is a normal occurrence

after changing a disk's partition data; it usually happens if you created, changed or deleted any extended partitions. After you press **OK**, your machine will reboot and the installation will begin again. Repeat the same installation steps you performed earlier until you reach the **Partition Disks** dialog; then simply choose **Done**.

14.12 Installing LILO

The LInux LOader (LILO) lets you specify at boot time whether to start Linux or another operating system. (If you are performing a workstation- or server-class installation, LILO is configured automatically in the Master Boot Record [MBR].) If you are performing a custom-class installation, the **LILO Installation** dialogs let you indicate how or whether to install LILO.

The **Choosing LILO in LILO Configuration Dialog** dialog (Figure 14–23, *Choosing LILO in LILO Configuration Dialog*) lets you add default options to the LILO boot command or choose to not install LILO at all. Any options you enter will be passed to the Linux kernel at boot time.

Figure 14–23 Choosing LILO in LILO Configuration Dialog

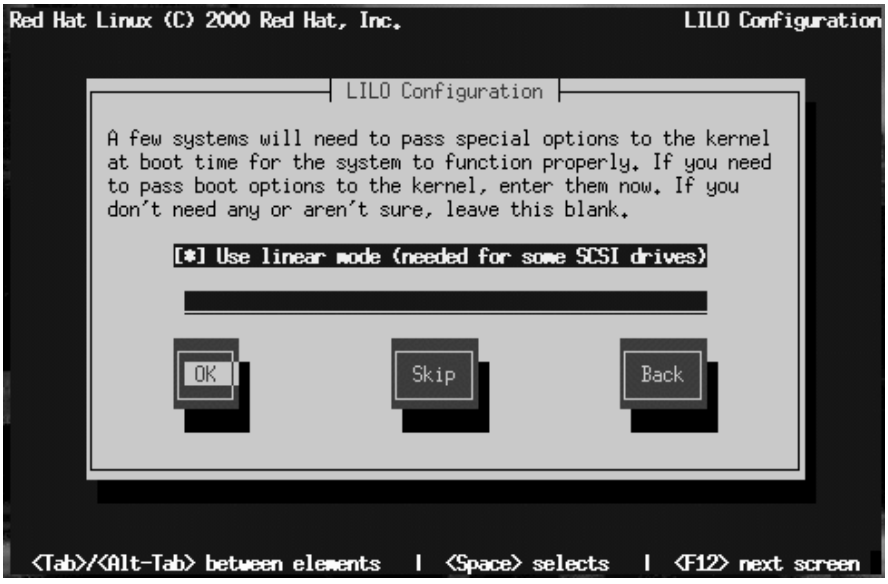

In Section 13.1.1, *Basic Hardware Configuration*, you were asked to review your computer's BIOS settings. In reviewing the BIOS settings, if you determined your system does not access a hard drive in linear mode, deselect this option. **Use linear mode** is selected by default. Select **OK** and press [Space] to continue.

If you do not wish to install LILO, press **Skip**.

WARNING

If you choose Skip, you will not be able to boot your Red Hat Linux system directly, and will need to use another boot method (such as a boot disk). Use this option only if you are sure you have another way of booting your Red Hat Linux system!

14.12.1 Alternatives to LILO

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

Boot Disk

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

LOADLIN

LOADLIN can load Linux from MS-DOS; unfortunately, it requires a copy of the Linux kernel (and an initial RAM disk, if you have a SCSI adapter) to be available on an MS-DOS partition. The only way to accomplish this is to boot your Red Hat Linux system using some other method (e.g., from LILO on a diskette) and then copy the kernel to an MS-DOS partition. LOADLIN is available from ftp://metalab.unc.edu/pub/Linux/system/boot/dualboot/ and associated mirror sites.

SYSLINUX

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

Commercial Bootloaders

Some commercial bootloaders are able to boot Linux. However, these products still require LILO to be installed in your Linux boot partition).

14.12.2 LILO Configuration

Choose where you want to install LILO (Figure 14–24, *Installing LILO in LILO Configuration Dialog*).

Figure 14–24 Installing LILO in LILO Configuration Dialog

You may install LILO in one of two places:

WARNING

To create a dual-boot environment on a system that currently has NT, you must install LILO on the first sector of the root partition, not the MBR. Please be sure to create a boot disk. In a case such as this, you will either need to use the boot disk, or configure the NT system loader to boot LILO from the first sector of the root partition. Be sure to check out http://www.linuxdoc.org/HOWTO/mini/Linux+NT-Loader.html for more information on setting up LILO and NT.

The Master Boot Record (MBR)

The recommended place to install LILO, unless the MBR already starts another OS loader, such System Commander or OS/2's Boot Manager. The MBR is a special area on your hard drive that is automatically loaded by your computer's BIOS, and is the earliest point at which LILO can take control of the boot process. If you install LILO in the MBR, when your machine boots, LILO will present a `boot:` prompt. You can then boot Red Hat Linux or any other operating system you configure LILO to boot.

The first sector of your boot partition

Recommended if you are already using another boot loader on your system (such as OS/2's Boot Manager). In this case, your other boot loader will take control first. You can then configure that boot loader to start LILO (which will then boot Red Hat Linux).

Select the location where you wish to install LILO and press **OK**.

Finally the installation program lets you set the default operating system and specify boot labels, as in Figure 14–25, *Selecting Bootable Partitions in LILO Configuration Dialog*.

Figure 14–25 Selecting Bootable Partitions in LILO Configuration Dialog

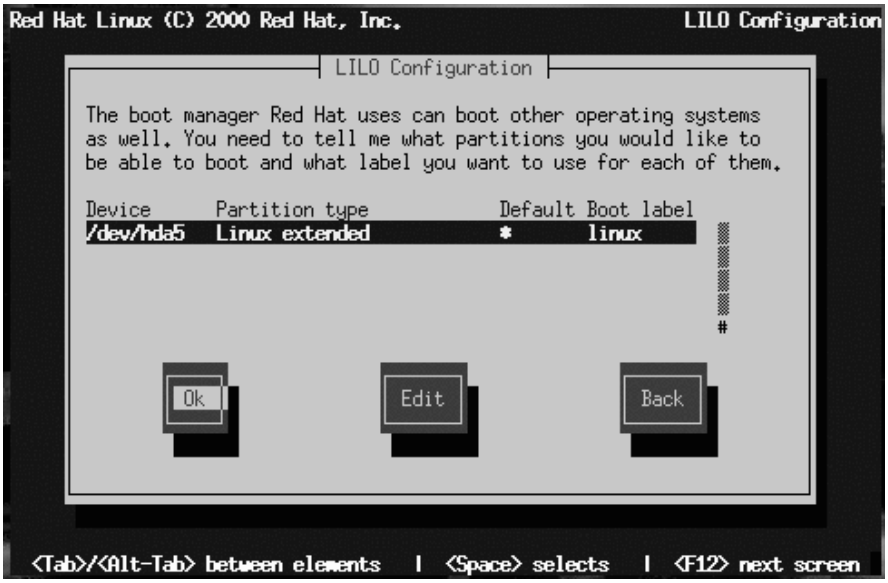

Every partition that is bootable is listed, including partitions used by other operating systems. The **Boot label** column will be filled in with the word `linux` on the partition holding your Red Hat Linux system's root filesystem. Other partitions may also have boot labels (such as a `dos` boot label for Windows 95/98 partitions).

To add a boot label for a partition (or change an existing boot label), use the arrow keys to highlight the desired partition. Then use the [Tab] key to select the **Edit** button, and press [Space]. You'll then see a small dialog box permitting you to enter/modify the partition's boot label. Press **OK** when done.

Tip

The contents of the **Boot label** column show what you will need to enter at LILO's `boot:` prompt in order to boot the desired operating system. Should you forget the boot labels defined on your system, press [Tab] at LILO's `Boot:` prompt to display a list of defined boot labels.

There is also a column labeled **Default**. Only one partition will contain an asterisk under that column. The partition marked as the default will be the partition LILO will boot if there is no user input during the boot process. Initially the root partition for your Red Hat Linux installation will be selected as the default. If you'd like to change this, use the arrow keys to highlight the partition you'd like to make the default, and press [F2]. The asterisk will move to the selected partition. When you've finished, select **OK**, and press [Space].

14.12.3 SMP Motherboards and LILO

If the installer detects a symmetric multi-processor motherboard on your system, it will automatically create two `lilo.conf` entries: `linux`, the default, boots the system in SMP mode and `linux-up` boots the system in uni-processor mode.

14.13 Naming Your Computer

The **Hostname Configuration** dialog prompts you to enter a name for your computer (called a **hostname**).

Figure 14–26 Hostname Configuration Dialog

If you have a network card in your computer, enter a **fully-qualified domain name** in this format:

hostname.domain.name

In this example, *hostname* is the name you've chosen for your computer and *domain.name* is the TCP/IP domain. (A domain name may contain more nodes: for example, *eastcoast.mainserver.redhat.com* and *westcoast.mainserver.redhat.com*.)

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

Tip

To change your hostname once you have rebooted your system, first become root. In a terminal at the root prompt, type `hostname` *newname*, where *newname* is what you want the hostname to be. If you just want to have the hostname echoed onto the display, type `hostname` and it will display the system's hostname.

14.14 Configuring a Network Connection

If you are installing over the network, you have performed your network configuration at the beginning of the installation process and do not need to complete this information again. If you are installing via local media and have a network card in place, please continue with this section.

Figure 14–27, *Network Configuration Dialog* appears only if your computer has a network card. (If there is more than one network card, this dialog configures the primary card.)

Figure 14–27 Network Configuration Dialog

You have two choices in this dialog:

- Select **Use bootp/dhcp**. In this case, an existing server on your LAN dynamically supplies network-related information needed to add this system to the network at connect-time. (Note that if you do this, the remaining fields in this dialog will be disabled, as DHCP and BOOTP essentially "fill in the blanks" for you.)
- Enter static network information in the fields provided. In this case, the network information you supply is assigned permanently to this computer.

14.15 Configuring Your Mouse

Next, the installation program probes for a mouse (Figure 14–28, *Mouse Selection Dialog*). Use the [Up] and [Down] arrow keys to confirm or change the selection as required.

Figure 14–28 Mouse Selection Dialog

If no mouse is detected, you'll have to select one manually.

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. On the other hand, if the connector is round, you have a PS/2 mouse. If you are installing Red Hat Linux on a laptop computer, in most cases the pointing device will be PS/2 compatible.

If you cannot find an exact match, select one of the **Generic** entries, based on your mouse's number of buttons, and its interface.

The **Emulate 3 Buttons** check box allows you to emulate a three-button mouse if your mouse only has two buttons. If you select this check box, you can simulate the third, "middle" button by pressing both mouse buttons simultaneously. If you have a two-button mouse, checking this box is a good idea, since the X Window System is easiest to use with a three-button mouse.

If you've selected a mouse with a serial interface, highlight the appropriate serial port, select **OK**, and press [Space].

Tip

To change your mouse configuration after you have booted your Red Hat Linux system, become root and use the `/usr/sbin/mouseconfig` command.

To configure your mouse as a left-handed mouse after you have booted your Red Hat Linux system, open a terminal and type `gpm -B 321`.

14.16 Configuring the Time Zone

Next, enter your system's time zone (see Figure 14–29, *Time Zone Selection Dialog*).

Figure 14–29 Time Zone Selection Dialog

If you wish to set the hardware (CMOS) clock to GMT (Greenwich Mean Time, also known as UTC, or Universal Coordinated Time), select **Hardware clock set to GMT**. Setting it to GMT means your system will properly handle daylight-saving time, if your time zone uses it.

WARNING

If your computer uses another operating system, setting the clock to GMT may cause the other operating system to display the incorrect time. Also keep in mind that if more than one operating system is allowed to automatically change the time to compensate for daylight saving time, it is likely that the time will be improperly set.

Select your time zone from the list and press [Enter].

Tip

To change your time zone configuration after you have booted your Red Hat Linux system, use the `/usr/sbin/timeconfig` command.

14.17 Setting a Root Password

The **Root Password** dialog prompts you to set a **root password** for your system. You'll use the root password to log into your Red Hat Linux system to perform system administration functions.

Figure 14–30 Root Password Dialog

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 poor passwords. Good passwords mix numerals with upper and lower case letters and do not contain dictionary words: **Aard387vark** or **420BMttNT**, for example. Remember that the password is case-sensitive. Write down this password and keep it in a secure place.

14.18 Creating a User Account

Next, you can create a user account for yourself which is for your everyday use (as in Figure 14–31, *Add User Dialog*). If you do not create a user account, after installation you'll have to log in as the **root** user (also known as the **superuser**). Root has complete access to the entire system. Logging in as the root user is best done *only* to perform system maintenance or administration. For instructions on how to create or modify user accounts after installation, refer to Chapter 3, *System Configuration* or the *Getting Started* chapter in the *Official Red Hat Linux Getting Started Guide*.

Figure 14–31 Add User Dialog

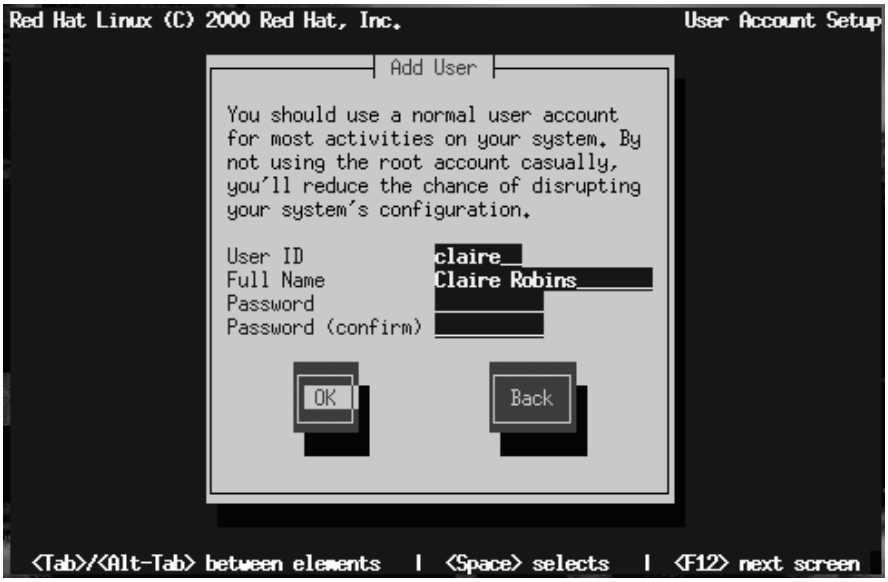

If you choose to create a user account, the account directory will be created under the path `/home` (for example, `/home/claire`). Passwords are case-sensitive and must contain at least six characters.

After you add a user account for yourself, the **User Account Setup** dialog (Figure 14–32, *User Account Setup Dialog*) appears, giving you the opportunity to create additional accounts. Select **Add** to do so or **OK** to continue.

Figure 14–32 User Account Setup Dialog

If you are performing a workstation- or server-class installation, your next task is to confirm your video hardware: turn to Section 14.21, *Configuring Your Video Adapter*. Otherwise, continue reading.

14.19 Authentication Configuration

If you are performing a custom installation, your next step is to configure the type of password authentication your Red Hat Linux system will use (see Figure 14–33, *Authentication Configuration Dialog*). You will also have the opportunity to configure NIS support; If you are unsure as to whether or not you should do this, ask your network administrator.

Figure 14–33 Authentication Configuration Dialog

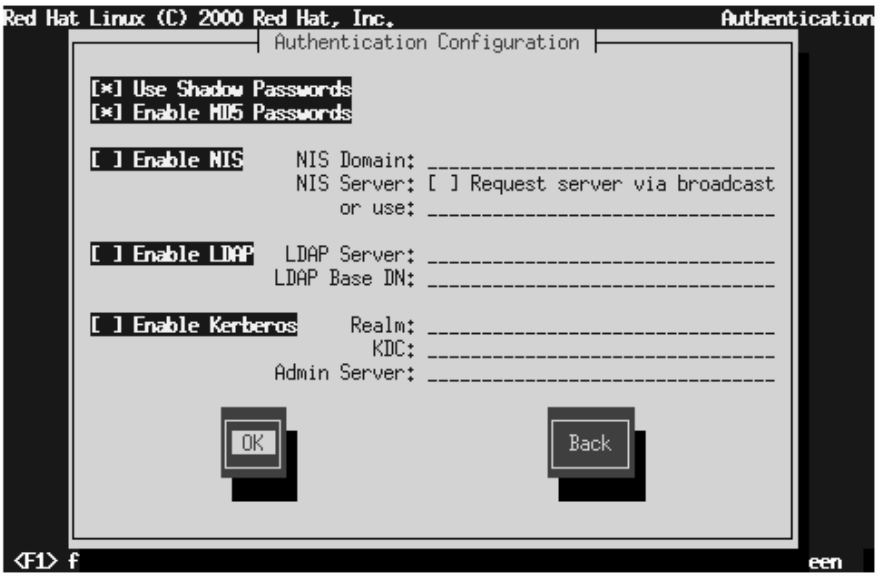

Here's a brief explanation of the authentication password prompts:

- **Use Shadow Passwords** -- provides a very secure method of retaining passwords for you. The password filed in the `/etc/passwd` file is replaced by `/etc/shadow` which is readable only by root.
- **Enable MD5 Passwords** -- allows a long password to be used (up to 256 characters) instead of the standard eight letters or less.
- **Enable NIS** -- allows you to run a group of computers in the same Network Information Service domain with a common password and group file. There are two options here to choose from:
  - **NIS Domain** -- this option allows you to specify which domain or group of computers your system will belong to.

  - **NIS Server** -- this option 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** -- 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 Chapter 7, *Lightweight Directory Access Protocol (LDAP)*. There are two options to choose from here:

  - **LDAP Server** -- this option allows you to access a server running the LDAP protocol.

  - **LDAP Base DN** -- this option allows you to look up user information by its *Distinguished Name* (DN).

- **Enable Kerberos** -- Kerberos is a secure system for providing network authentication services. For more information about Kerberos, see Chapter 8, *Using Kerberos 5 on Red Hat Linux*. 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 very 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.

Please Note

To configure the NIS option, you must be connected to an NIS network. If you are unsure whether you are connected to an NIS network, please ask your network administrator.

Unless you are setting up NIS, you will notice that both shadow passwords and MD5 passwords are selected. We recommend you use both to make your machine as secure as possible.

14.20 Select Packages to Install

Figure 14–34, *Package Group Selection Dialog* appears only if you're performing a custom-class installation. Use this dialog to select the types of software packages you wish to install.

Figure 14–34 Package Group Selection Dialog

Selecting **Everything** (which can be found at the end of the component list) installs all packages included with Red Hat Linux. Selecting every package will require close to 1.7GB of free disk space.

14.20.1 Selecting Individual Packages

You can also select or deselect individual packages in an application category. To do this, check the **Select individual packages** check box. When you select **OK**, a screen

like Figure 14–35, *Selecting Individual Packages in Package Group Selection Dialog* appears.

Figure 14–35 Selecting Individual Packages in Package Group Selection Dialog

To see the list of packages in the **Applications/Editors** category, for example, use the arrow keys to select this group and press [Enter] or [Space]. When you do, the **+** sign (closed) changes to a **-** sign (open) and a list of packages in this category appears. An ***** appears beside currently selected packages. Press the [Enter] or [Space] keys to select or deselect packages.

To see a description of a particular package, make sure that package is highlighted and press [F1].

Note that some packages that are required for every Red Hat Linux system (such as the kernel and certain libraries) do not appear in the package selection dialogs.

14.20.2 Unresolved Package Dependencies

Many software packages, in order to work correctly, require that other software packages are also installed on your system. For example, many of the graphical 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.

Figure 14–36 Package Dependencies Dialog

After you have finished selecting packages to install, the installation program checks the list of selected packages for dependencies. If any package requires another package which you have not selected to install, the program presents a list of these **unresolved dependencies** and gives you the opportunity to resolve them (see Figure 14–36, *Package Dependencies Dialog*). If you simply press **OK**, the program will resolve them automatically by adding all required packages to the list of selected packages.

14.21 Configuring Your Video Adapter

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

Figure 14–37 Video Card Confirmation

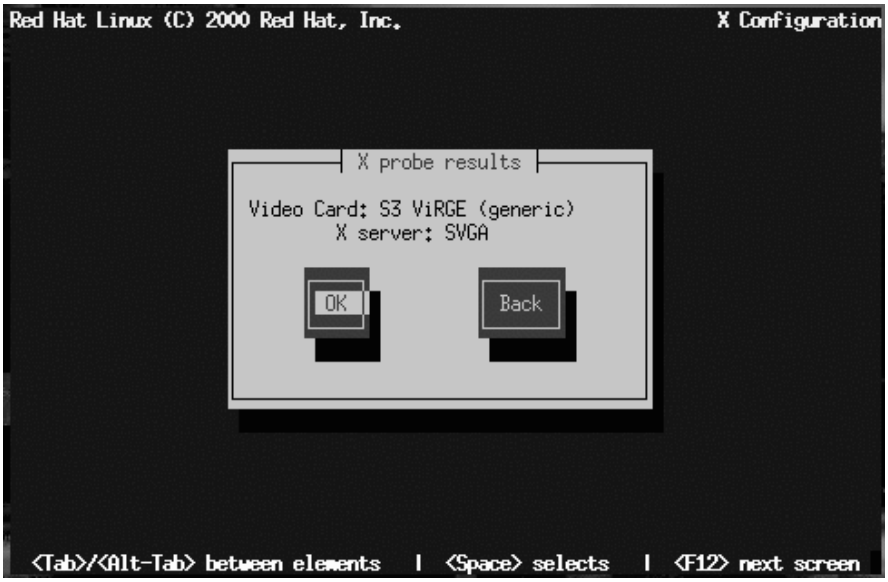

If the installation program cannot detect your video adapter, you'll see Figure 14–38, *Video Card Selection Dialog*. In this case, select your video card from the list using the arrow keys and [Space].

Figure 14–38 Video Card Selection Dialog

14.22 Package Installation

Figure 14–39, *Installation to Begin Dialog* appears when the installation program is ready to format partitions and load software packages. The installation program names the log file (`/tmp/install.log`) for this Red Hat Linux installation. Select **OK** and press [Space] to continue.

Figure 14–39 Installation to Begin Dialog

While software packages are being installed, a screen like Figure 14–40, *Package Installation Status Dialog* appears.

Figure 14–40 Package Installation Status Dialog

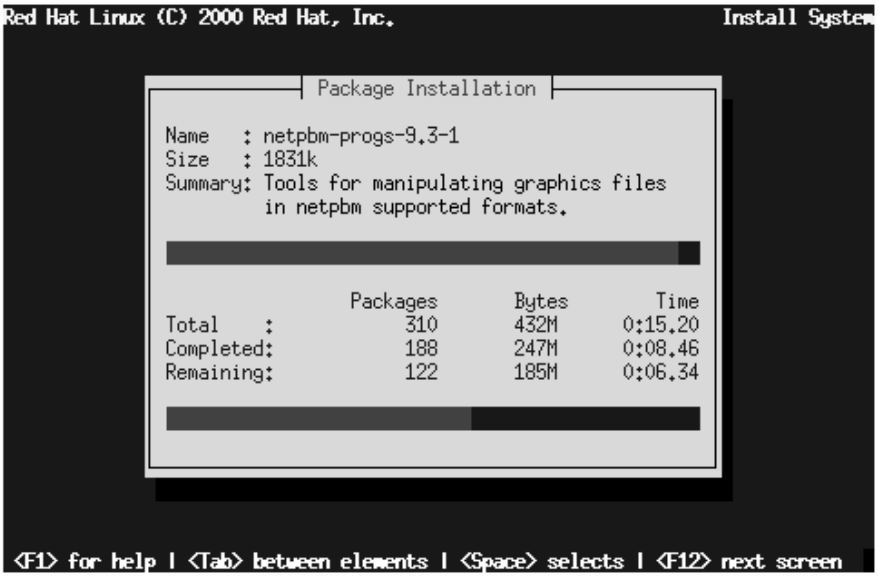

As each software package is installed, the top part of the **Package Installation** screen displays its name, size, and a summary description of its function.

In the bottom part of the **Package Installation** screen, the **Total** row shows the total number of packages to be installed, the total size of all packages, and the amount of time required to install all packages. As each package begins and completes installation, the completed and remaining rows are updated.

If you are performing a server-class installation, you are almost finished; turn to Section 14.25, *Finishing Up*.

14.23 Creating a Boot Disk

If you're performing a custom-class installation, the **Bootdisk** dialog (see Figure 14–41, *Creating a Boot Disk*) gives you the opportunity to create a customized boot disk for your Red Hat Linux system.

Figure 14–41 Creating a Boot Disk

A boot disk can be handy in these situations:

- Use It Instead of LILO -- You can use a boot disk instead of LILO. This is handy if you're trying Red Hat Linux for the first time, and you'd feel more comfortable if the boot process for your other operating system is left unchanged. With a boot disk, going back to your other operating system is as easy as removing the boot disk and rebooting.
- Use It When Another Operating System Overwrites LILO -- Other operating systems may not be as flexible as Red Hat Linux when it comes to supported boot methods. Quite often, installing or updating another operating system can cause the master boot record (originally containing LILO) to be overwritten, making it impossible to boot your Red Hat Linux installation. The boot disk can then be used to boot Red Hat Linux so you can reinstall LILO.

Select **Yes** and press [Space] to create a boot disk. Next, you'll be prompted to insert a blank, formatted diskette.

Figure 14–42 Boot Disk Dialog

For instructions on how to create a boot disk after the installation, read the `mkbootdisk` man page. Be sure you create a new boot disk if you make any changes to your kernel.

14.24 Configuring the X Window System

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

First, `Xconfigurator` presents a list of monitors (see Figure 14–43, *Monitor Setup Dialog*). If your monitor is listed, select it and press [Enter]. Otherwise, select **Custom.**

Figure 14–43 Monitor Setup Dialog

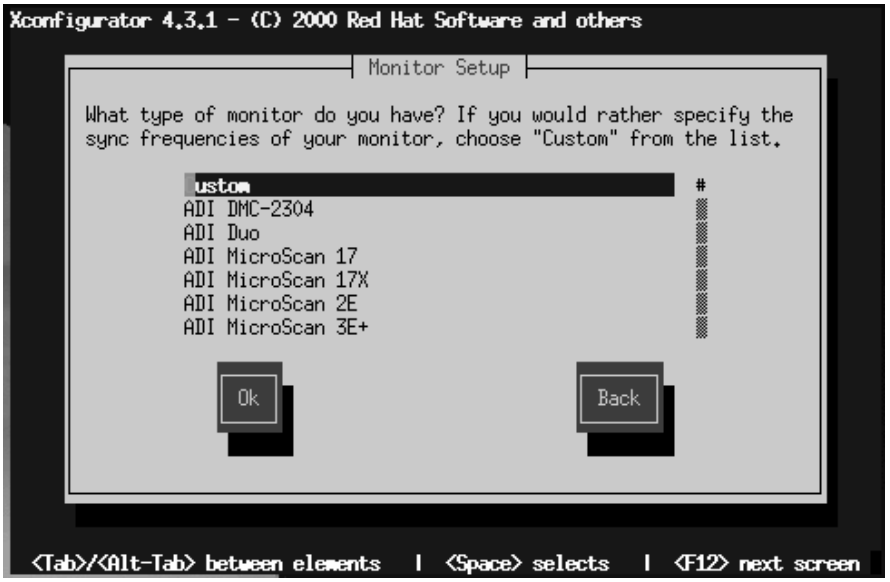

If you select a monitor from the list, you will next enter your video memory, see Figure 14–48, *Video Memory*.

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 14–44 Custom Monitor Setup

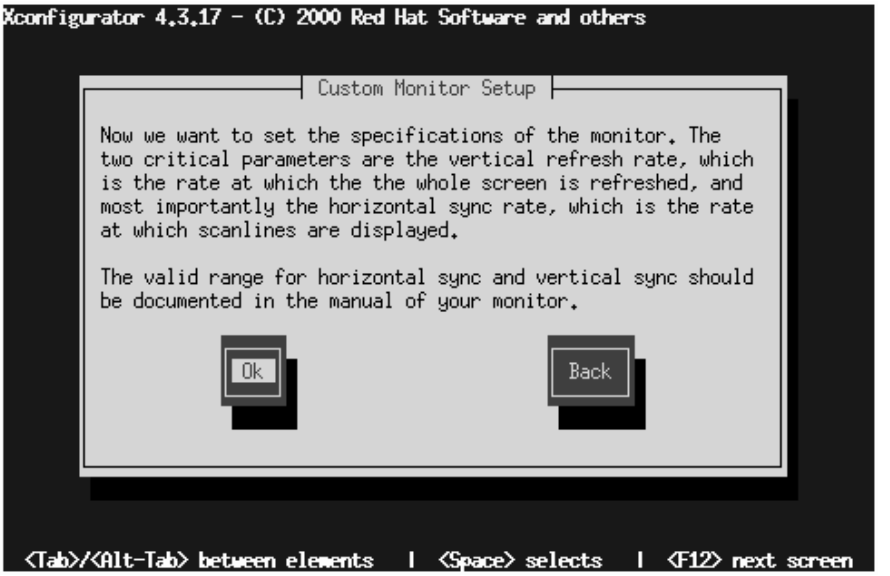

CAUTION

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 14–45, *Custom Monitor Setup - Horizontal Sync*). Be sure *not* to choose a range outside of your monitor's capacity. For custom settings, refer to Figure 14–47, *Custom Monitor Information*.

Figure 14–45 Custom Monitor Setup - Horizontal Sync

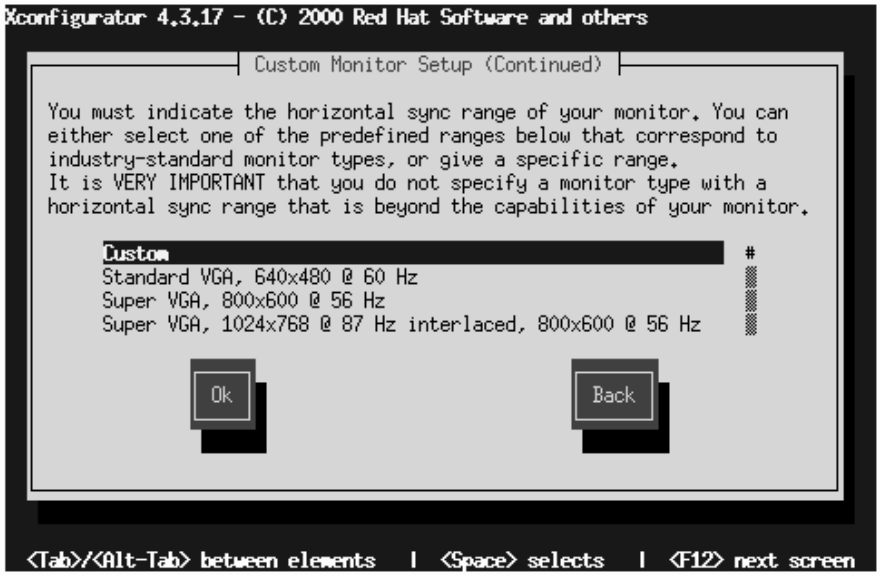

Next, choose a vertical sync range.

Figure 14–46 Custom Monitor Setup - Vertical Sync

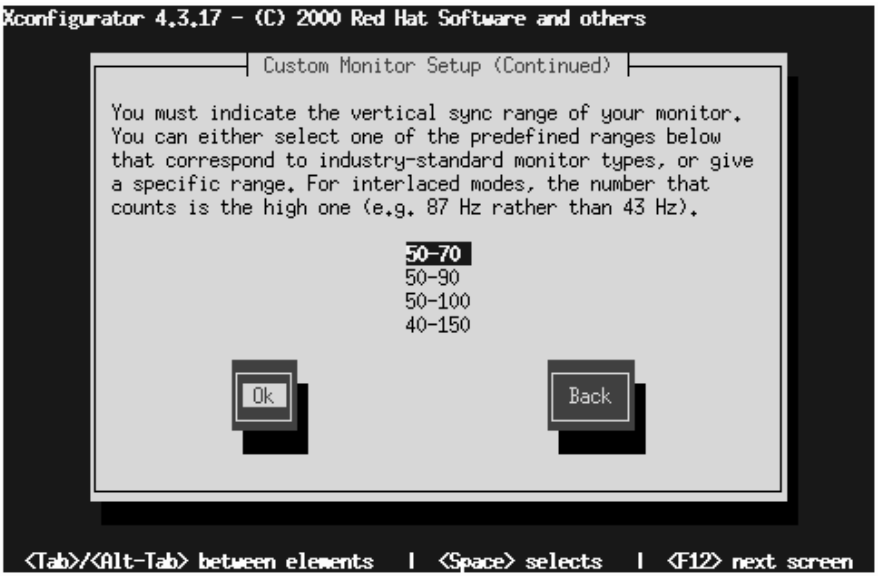

If you chose to enter custom monitor settings, please enter both the horizontal and vertical sync rages for your monitor (see Figure 14–47, *Custom Monitor Information*).

Figure 14–47 Custom Monitor Information

After you've selected your monitor (and its settings), you must choose the amount of memory present on your video card (see Figure 14–48, *Video Memory*).

Figure 14–48 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 14–49 Clockchip Configuration Dialog

The **Probe for Clocks** screen (see Figure 14–50, *Probe for Clocks*) prompts you to allow the installation program to detect what video modes your video card and monitor are capable of using.

You should select **Probe** for best results. If you allow the installation program to detect these modes for you, your choices for video modes (see Figure 14–51, *Select Video Modes*) may be reduced based on your video card and monitor capabilities.

However, if a previous attempt to probe video modes ended with a bad result (such as having to reboot your system), you should choose **Skip** and then choose your preferred video modes (see Figure 14–52, *Select Video Modes*).

Figure 14–50 Probe for Clocks

To choose your preferred display mode (see Figure 14–51, *Select Video Modes* and Figure 14–52, *Select Video Modes*), select one or more modes by pressing [Space].

Figure 14–51 Select Video Modes

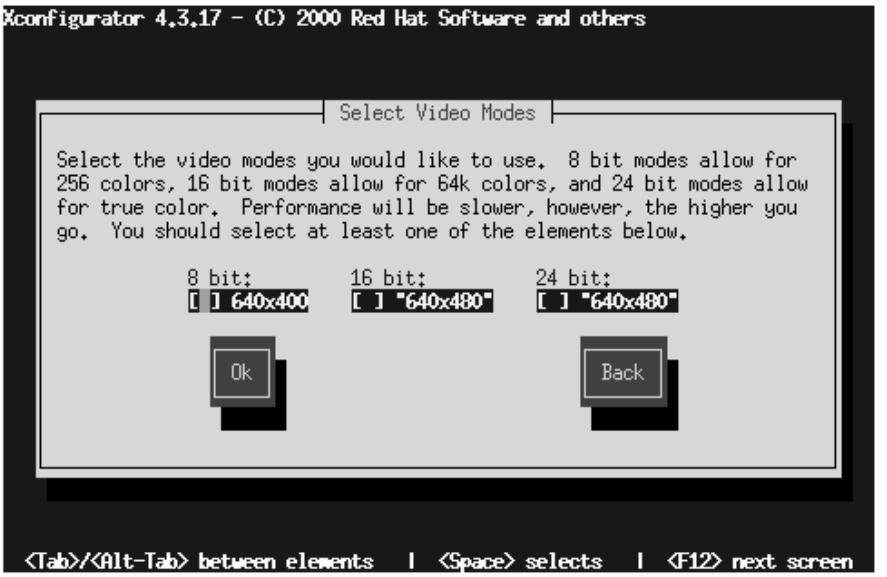

Tip

Minimum recommended settings for video modes are 16 bit for most applications (video modes set higher than 16 bit tend to run more slowly) in 1024x768.

Figure 14–52 Select Video Modes

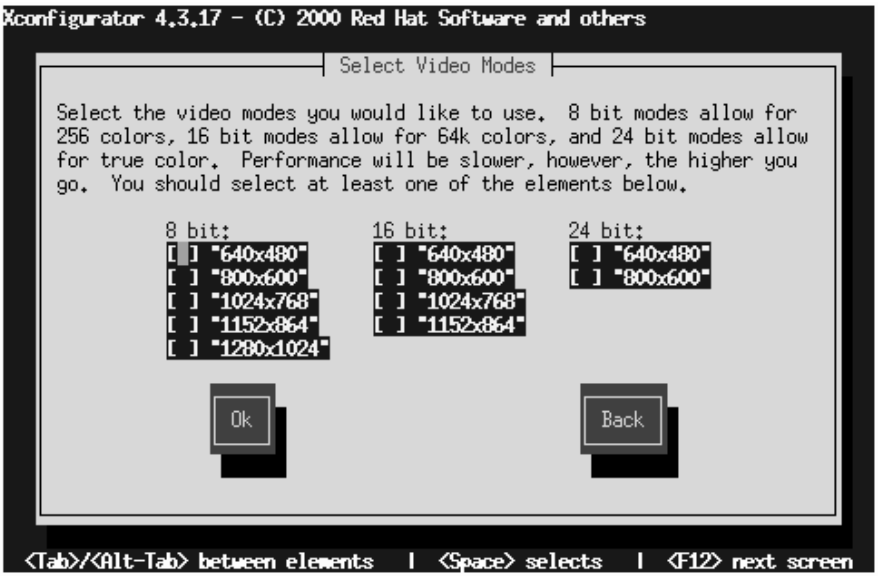

Do not select a mode that exceeds the capacity of your monitor.

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

Figure 14–53 Test Your X Configuration

If you select **OK**, you'll have the opportunity to go back and enter different configuration values should there be an error. If there is no error, 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'll see the completion screen.

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

14.25 Finishing Up

When finished, the installation program displays the **Complete** dialog (as in Figure 14–54, *Compete Dialog*), telling you to reboot your computer. Remove any diskette

from the diskette drive (unless you skipped the LILO installation, in which case you must use the boot disk created during the installation), or the Red Hat Linux CD if your system booted from the CD-ROM.

Figure 14–54 Compete Dialog

Select **OK** to reboot your newly installed system. After your computer's normal power-up sequence has completed, you should see LILO's GUI prompt, at which you can do any of the following:

- Press [Enter] -- Causes LILO's default boot entry (as seen in Figure 14–25, *Selecting Bootable Partitions in LILO Configuration Dialog*) to be booted.
- Select a Boot Label, followed by [Enter] -- Causes LILO to boot the operating system corresponding to the boot label. (Press [?] at the LILO text `boot:` prompt for a list of valid boot labels.)
- Wait -- After LILO's timeout period, (which, by default, is five seconds) LILO will automatically boot the default boot entry.

Do whatever is appropriate to boot Red Hat Linux. You should see one or more screens of messages scroll by. Eventually, you should see either a graphical login screen, or a `login:` prompt.

Congratulations! Your Red Hat Linux installation is complete!

15 Installing Red Hat Linux via the GUI

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

15.1 The Installation Program User Interface

If you've used a **graphical user interface (GUI)** before, you'll 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.

Please Note

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

```
boot: text
```

For text mode installation instructions, please refer to Chapter 14, *Installing Red Hat Linux via Text Mode*.

15.1.1 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 giving you a way to enter commands from a shell prompt. It presents this information on five **virtual consoles**, among which you can switch using a single keystroke.

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 15–1, *Console, Keystrokes, and Contents* for a listing of the virtual consoles, keystrokes to switch to them, and their contents.

Table 15–1 Console, Keystrokes, and Contents

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

Generally, there's no reason to leave the default console (virtual console #7) unless you are attempting to diagnose installation problems. But if you get curious, feel free to look around.

15.2 Starting the Installation Program

Now it's time to begin installing Red Hat Linux. To start the installation, you must first boot the installation program. Please make sure you have all the resources you'll need for the installation. If you've already read through *Official Red Hat Linux Installation Guide*, and followed the instructions, you should be ready to begin.

Please Note

Occasionally, some hardware components require a **driver disk** during the installation. A driver disk adds support for hardware that is not otherwise supported by the installation program. The driver disk could be produced by Red Hat, it could be a disk you make yourself, or it could be a disk that a hardware vendor includes with a piece of hardware.

If a screen appears prompting you to insert a driver disk, and you have a vendor supplied driver disk, please do so. Another source for finding driver disks is at http://www.redhat.com/support/errata/. For more information about driver disks, refer to Appendix C, *Driver Disks*.

15.2.1 Booting the Installation Program

Please Note

If you need to create a boot disk, please refer to section *Step 6 - How Do You Want to Start the Installation?* in the *Official Red Hat Linux Installation Guide*.

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

Tip

To change your BIOS settings, you will need to take note of the instructions given when your computer first begins to boot. Often you will see a line of text telling you to press the [Del] key to enter the BIOS settings. Once you have done whatever process is needed to enter your computer's BIOS, you can then change the boot order to allow your computer to boot from the CD-ROM drive or diskette drive first when bootable software is detected. For more information, please refer to the documentation that came with your system.

There are four possible boot methods:

- *Bootable CD-ROM* -- your machine supports a bootable CD-ROM drive and you want to perform a local CD-ROM installation.
- *Local boot disk* -- your machine will not support a bootable CD-ROM and you want to install from a local CD-ROM or a hard drive.
- *Network boot disk* -- use to install from NFS, FTP and HTTP installation methods.
- *PCMCIA boot disk* -- use in cases where you need PCMCIA support, but your machine does not support booting from the CD-ROM drive *or* if you need PCMCIA support in order to make use of the CD-ROM drive on your system. This boot disk offers you all installation methods (CD-ROM, hard drive, NFS, FTP, and HTTP).

After a short delay, a screen containing the `boot:` prompt should appear. The screen contains information on a variety of boot options. Each boot option also has one or more help screens associated with it. To access a help screen, press the appropriate function key as listed in the line at the bottom of the screen.

You should keep two things in mind:

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

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

Expert mode can be entered using the following boot command:

```
boot: linux expert
```

If you do not wish to perform a CD-ROM GUI installation, you can choose to perform a text mode installation by using the following boot command:

```
boot: text
```

For text mode installation instructions, please refer to Chapter 14, *Installing Red Hat Linux via Text Mode*.

The command to start a **serial installation** has changed. If you need to perform the installation in serial mode, type:

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

Where <*device*> should be the device you are using (such as ttyS0 or ttyS1).

To explicitly request a dialog where you can configure additional devices (such as ISA devices) include the 'isa' directive:

```
boot: linux isa
```

Please Note

If the **Mouse Not Detected** screen (see Figure 15–1, *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 15–4, *Mouse Configuration*).

Figure 15–1 Mouse Not Detected

For text mode workstation-class installation instructions, please refer to the *Official Red Hat Linux Installation Guide*.

For text mode custom installation instructions, please refer to Chapter 14, *Installing Red Hat Linux via Text Mode*.

Please Note

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

Options can also be passed to the kernel.

For example, to instruct the kernel to use all the RAM in a 128MB system, enter:

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

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

If you do need to specify boot options to identify your hardware, please make note of them -- they will be needed during the LILO configuration portion of the installation (please see Section 15.16, *Installing LILO* for more information).

Booting without diskettes

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

If you have MS-DOS installed on your system, you can boot directly from the CD-ROM drive without using a boot disk.

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

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

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

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

15.3 Selecting an Installation Method

Next, you will be asked what type of installation method you wish to use. You can install Red Hat Linux via the following basic methods:

CD-ROM

If you have a CD-ROM drive and the Red Hat Linux CD-ROM. Requires a boot disk, a bootable CD-ROM or a PCMCIA boot disk.

Hard Drive

If you copied the Red Hat Linux files to a local hard drive. Refer to Chapter 14, *Installing Red Hat Linux via Text Mode* for hard drive installation instructions. Requires a boot disk or a PCMCIA 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. Requires a network or PCMCIA boot disk. Refer to Chapter 14, *Installing Red Hat Linux via Text Mode* for network installation instructions. Please note: NFS installations may also be performed in GUI mode.

FTP

If you are installing directly from an FTP server. Requires a network or PCMCIA boot disk. Refer to Chapter 14, *Installing Red Hat Linux via Text Mode* for FTP installation instructions.

HTTP

If you are installing directly from an HTTP Web server. Requires a network or PCMCIA boot disk. Refer to Chapter 14, *Installing Red Hat Linux via Text Mode* for HTTP installation instructions.

15.4 Beginning the Installation

If you are planning to install via CD-ROM using the graphical interface, please read on.

Please Note

If you'd rather perform a text mode installation, reboot your system and at the `boot:` prompt, type **`text`**. Refer to Chapter 14, *Installing Red Hat Linux via Text Mode* for further instructions.

15.4.1 Installing from CD-ROM

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

The installation program will then probe your system and attempt to identify your CD-ROM drive. It will start by looking for an IDE (also known as ATAPI) CD-ROM drive. If found, you will continue to the next stage of the installation process (see Section 15.5, *Language Selection*).

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

SCSI

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

Other

If your CD-ROM drive is neither an IDE nor a SCSI, it's an "other." Sound cards with proprietary CD-ROM interfaces are good examples of this CD-ROM

type. The installation program presents a list of drivers for supported CD-ROM drives -- choose a driver and, if necessary, specify any driver options.

Tip

A partial list of optional parameters for CD-ROM drives can be found in Appendix A, *General Parameters and Modules*.

What If the IDE CD-ROM Was Not Found?

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

- `a` - First IDE controller, master
- `b` - First IDE controller, slave
- `c` - Second IDE controller, master
- `d` - Second IDE controller, slave

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

Once identified, you will be asked to insert the Red Hat Linux CD into your CD-ROM drive. Select **OK** when you have done so. After a short delay, the next dialog box will appear.

After booting, the installation program begins by displaying the language screen.

Please Note

If you wish to abort the installation process at this time, simply reboot your machine then eject the boot diskette or CD-ROM. You can safely cancel the installation at any point before the **About to Install** screen, see Section 15.23, *Preparing to Install*.

15.5 Language Selection

Using your mouse, select the language you would prefer to use for the installation and as the system default (see Figure 15–2, *Language Selection*).

Figure 15–2 Language Selection

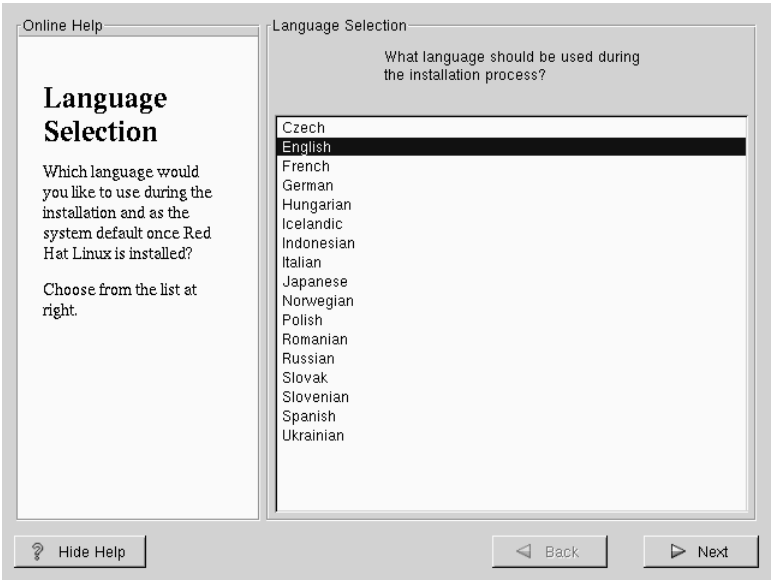

15.6 Keyboard Configuration

Choose the model that best fits your system (see Figure 15–3, *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 post-installation, become **root** and use the `/usr/sbin/kbdconfig` command, or 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 15–3 Keyboard Configuration

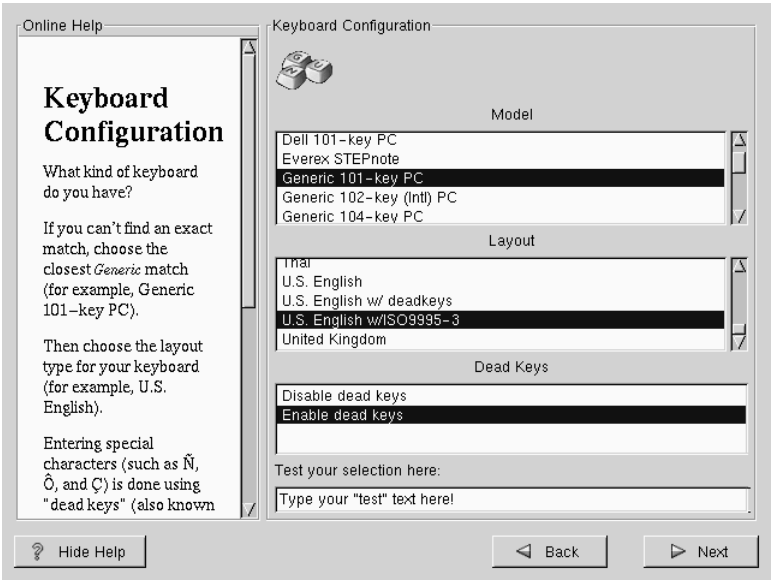

15.7 Mouse Configuration

Choose the correct mouse type for your system. If an exact match cannot be found, choose a mouse type that you are sure is compatible with your system (see Figure 15–4, *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 are installing Red Hat Linux on a laptop computer, in most cases the pointing device will be PS/2 compatible.

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 15–4 Mouse Configuration

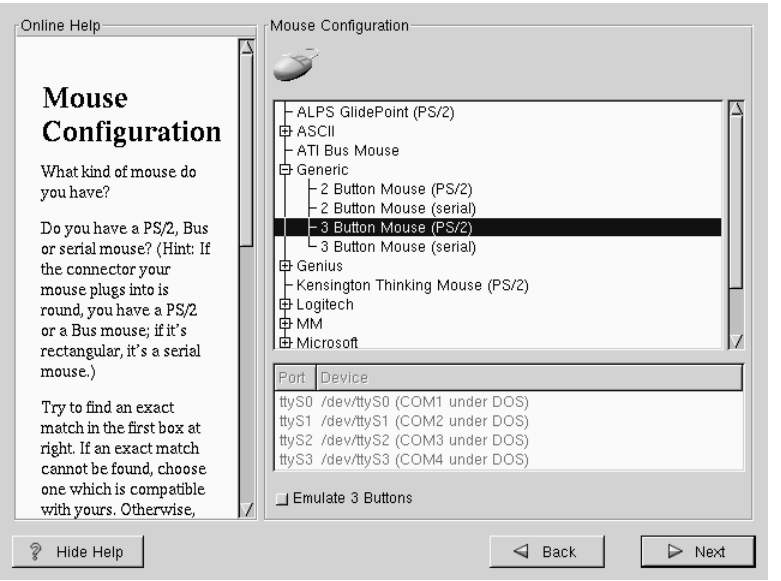

If you have a PS/2 or a Bus mouse, you do not need to pick a port and device. If you have a serial mouse, you should choose the correct port and device that your serial mouse is on.

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

Tip

To change your mouse configuration post-installation, become root. You can then use the `/usr/sbin/mouseconfig` command from the shell prompt.

To configure your mouse as a left-handed mouse, you can reset the order of the mouse buttons. This can be done after you have booted your Red Hat Linux system, by typing `gpm -B 321` at the shell prompt.

15.8 Welcome to Red Hat Linux

The **Welcome** screen (see Figure 15–5, *Welcome to Red Hat Linux*) does not prompt you for any installation 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.

Figure 15–5 Welcome to Red Hat Linux

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

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

15.9 Install Options

Please Note

`Feature:` Red Hat Linux 7.0 has a installation method known as a "partitionless" installation. If your system has a FAT (DOS/Windows) partition with sufficient free space, you can install Red Hat Linux without repartitioning your hard drive. This method is perfect for people who are new to Linux, and would like to try Red Hat Linux with a minimum of disruption to their computer. To learn more about this method, refer to *Installing Without Partitioning* in the *Official Red Hat Linux Installation Guide*.

Choose whether you would like to perform a full installation or an upgrade (see Figure 15–6, *Choosing Install or Upgrade*).

In the top right-hand corner of the **Install Type** screen there is a box you may select if you wish to partition using fdisk. Note that fdisk is not as intuitive to use as Disk Druid and is not selected by default. If you have not used fdisk before, you should read about both fdisk and Disk Druid to determine which will best suit your needs.

Figure 15–6 Choosing Install or Upgrade

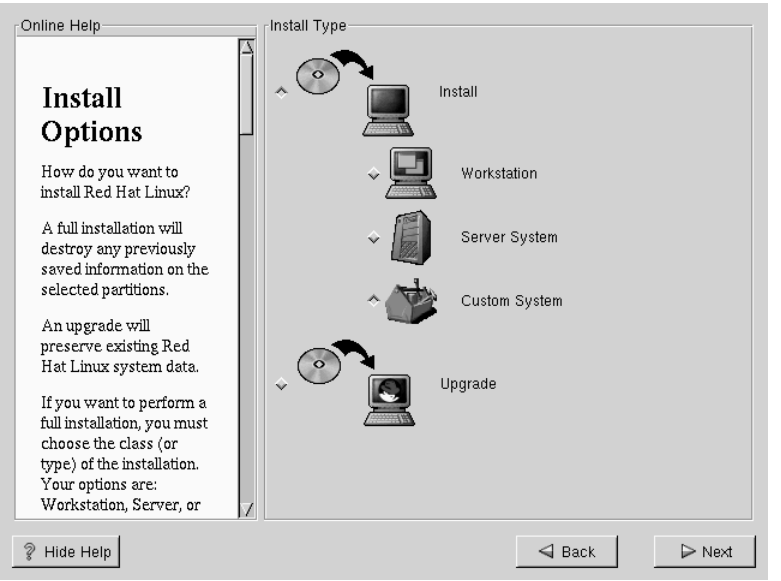

To perform an upgrade, please refer to *Upgrading Your Current System* in the *Official Red Hat Linux Installation Guide*.

15.10 Continuing the Installation

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

WARNING

Installing Red Hat Linux over another installation of Linux (including Red Hat Linux) does *not* preserve any information (files or data) from a prior installation. Make sure you save any important files! If you are worried about saving the current data on your existing system (without making a backup on your own), you should consider performing an upgrade instead.

In choosing a full installation, you must also choose the class of the installation. Your options include: **Workstation**, **Server** or **Custom**.

Most suitable for new users, the workstation-class installation will install your choice of a GNOME or KDE desktop environment (or both), and the X Window System.

WARNING

Do not choose this method if you're sharing a disk with Windows NT; if you do, you will be unable to boot Windows NT. LILO will write over NT's boot loader and you will be unable to boot NT. You must perform a custom-class installation and configure LILO so that it is not installed on the Master Boot Record (MBR).

To create a dual-boot environment on a system that currently has NT, you must install LILO on the first sector of the root partition, not the MBR. Please be sure to create a boot disk. In a case such as this, you will either need to use the boot disk, or configure the NT system loader to boot LILO from the first sector of the root partition. Be sure to check out http://www.linuxdoc.org/HOWTO/mini/Linux+NT-Loader.html for more information on setting up LILO and NT.

WARNING

A workstation-class installation will erase *all information* in *all Linux-related partitions* from *every one* of your computer's hard drive(s).

Below are the minimum recommended disk space requirements for a workstation-class installation.

- Workstation choosing GNOME -- 900M
- Workstation choosing KDE -- 900M
- Workstation choosing both GNOME and KDE -- 1.1G

If you plan to choose all group packages (for example, GNOME is one package group), as well as select additional individual packages, you may want to allow yourself 1.7G or more of disk space. This is also allow for some room where additional data may be written.

Please Note

Unlike previous workstation-class installations, performing a Red Hat Linux 7.0 workstation-class installation will not install the network daemon xinetd (inet services). Not installing xinetd results in a more secure installation; however,in-bound[1], network-related services such as finger, telnet, talk, and FTP will not work. If you require these types of services, please go back and choose a server- or a custom-class installation.

A server-class installation is most appropriate for you if you'd like your system to function as a Linux-based server, and you don't want to heavily customize your system configuration.

Below are the recommended disk space requirements for a server-class installation.

- Server (minimum) -- 450M
- Server (choosing everything) -- 1G

If you plan to choose all group packages, as well as select additional individual packages, you may want to allow yourself 1.7G or more of disk space. This is also allow for some room where additional data may be written.

[1] In-bound services mean that you can telnet, for instance, out, but that other systems cannot telnet into your system.

WARNING

A server-class installation will erase *all partitions* (both Linux and non-Linux) from *every one* of your computer's hard drive(s).

The *custom-class installation* allows you the most flexibility during your installation. The workstation-class and server-class installations automatically go through the installation process for you and omit certain steps. During a custom-class installation, it is up to *you* how disk space should be partitioned. You have complete control over the packages that will be installed on your system. You can also determine whether you'll use LILO (the LInux LOader) to boot your system. Unless you have prior Linux experience, you should not select the custom-class installation method.

Below are the recommended disk space requirements for a custom-class installation.

- Custom (minimum) -- 250M
- Custom (choosing everything) -- 1.7G

If you would like to know what steps are omitted by not performing a custom-class installation please refer to *Step 7 - Which Installation Type is Best For You?* in the *Official Red Hat Linux Installation Guide*.

15.11 Automatic 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.

The **Automatic Partitioning** screen is only seen when performing a workstation- or server-class installation. If you are performing a custom-class installation, or choose to manually partition, please refer to Section 15.13, *Partitioning Your System*.

In this screen, you can choose to continue with this installation, to partition manually, or use the **Back** button to choose a different installation method (see Figure 15–7, *Automatic Partitioning*).

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

Figure 15–7 Automatic Partitioning

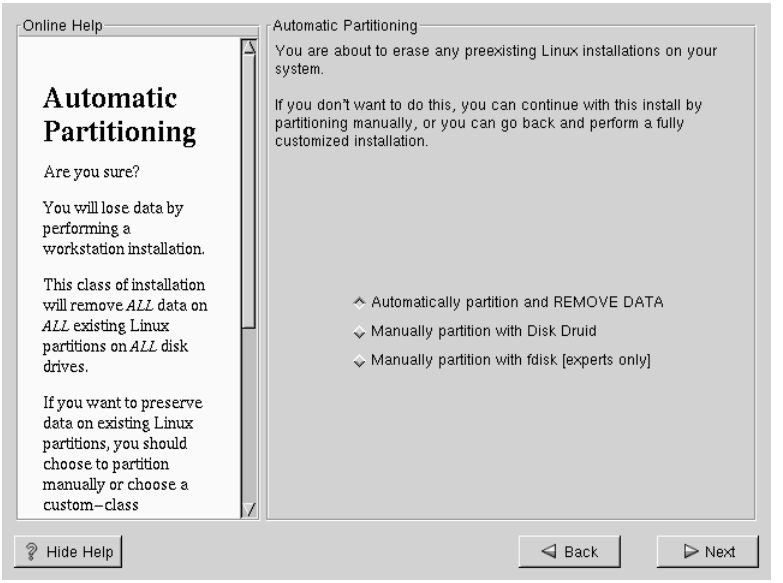

CAUTION

A workstation-class installation will remove all data on all currently existing Linux partitions.

If you do not want Red Hat Linux to be installed on your master boot record (MBR) or if you want to use a boot manager other than LILO, do not choose this installation method.

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

If you have another OS on your system that you wish to keep installed, if you do not want Red Hat Linux to be installed on your master boot record (MBR), or if you want to use a boot manager other than LILO, do not choose this installation method.

If you are unsure how you want your system to be partitioned, please read Appendix B, *An Introduction to Disk Partitions*.

15.12 Manual Partitioning

`Manual Partitioning` allows you to perform an installation while partitioning your drive(s) yourself. If you do not feel comfortable with partitioning your system, it is recommended that you *do not* choose to partition manually and instead choose **Automatic Partitioning** by clicking the **Back** and choosing to perform a workstation- or server-class installation.

Figure 15–8 Manual Partitioning

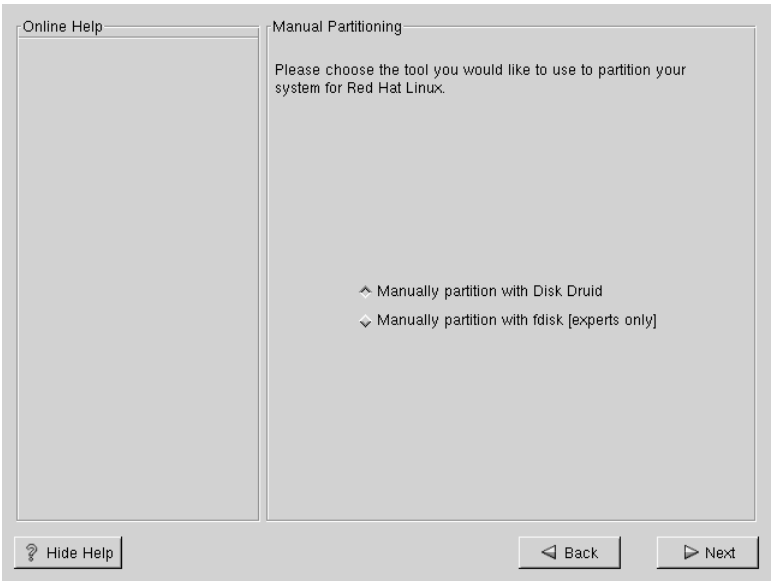

15.13 Partitioning Your System

If you are performing a workstation- or server-class installation and you chose *not* to partition manually, please skip to Section 15.17, *Network Configuration*.

At this point, it's necessary to let the installation program know where it should install Red Hat Linux. This is done by defining mount points for one or more disk partitions in which Red Hat Linux will be installed. You may also need to create and/or delete partitions at this time (refer to Figure 15–9, *Partitioning with Disk Druid*).

Please Note

If you have not yet planned how you will set up your partitions, refer to Appendix B, *An Introduction to Disk Partitions*. As a bare minimum, you'll need an appropriately-sized root partition, and a swap partition of at least 16 MB.

Figure 15–9 Partitioning with Disk Druid

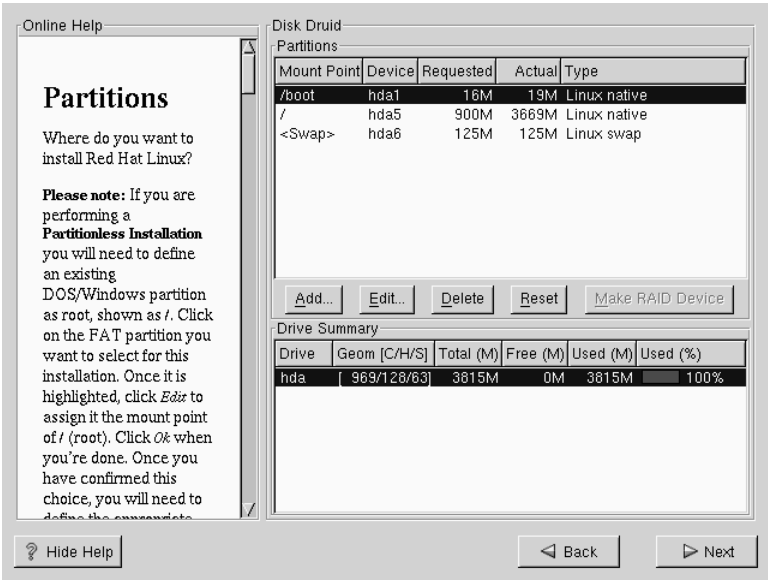

The partitioning tool used in Red Hat Linux 7.0 is Disk Druid. With the exception of certain esoteric situations, Disk Druid can handle the partitioning requirements for a typical Red Hat Linux installation.

15.13.1 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 said to be mounted at this location. This field indicates where the partition will be mounted. If a partition exists, but is "not set" you need to define its mount point. Double-click on the partition or use the **Edit** key.

Device:

This field displays the partition's device name.

Requested:

This field shows the partition's original size. To re-define the size, 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 (such as Linux Native or DOS).

15.13.2 Recommended Partitioning Scheme

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

- A swap partition (at least 16MB) -- 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 your computer has 16MB of RAM or less, you *must* create a swap partition. Even if you have more memory, a swap partition is still recommended. The minimum size of your swap partition should be equal to your computer's RAM, or 16MB (whichever is larger). In Disk Druid, the partition field for `swap` should look similar to:

```
<Swap>  hda2  125M   125M  Linux swap
```

- A `/boot` partition (16MB, maximum) -- The partition mounted on `/boot` contains the operating system kernel (which allows your system to boot Red Hat Linux), along with files used during the bootstrap process. Due to the limitations

of most PC BIOSes, creating a small partition to hold these files is a good idea. This partition should be no larger than 16MB. In Disk Druid, the partition field for `/boot` should look similar to:

```
/boot  hda1  16M   19M  Linux native
```

- A `root` partition (900MB-1.7GB) -- This is where "`/`" (the root directory) resides. In this setup, all files (except those stored in `/boot`) reside on the root partition. A 850MB root partition will permit the equivalent of a workstation-class installation (with *very* little free space), while a 1.7GB root partition will let you install every package. In Disk Druid, the partition field for `/` should look similar to:

```
/  hda2  900M   3669M   Linux native
```

15.13.3 Problems When Adding a Partition

If you attempt to add a partition and Disk Druid can't carry out your request, you'll 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.

15.13.4 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.

Free:

This field shows how much of the hard disk's space is still unallocated.

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. That is done using the **Allowable Drives** field in Section 15.13.6, *Adding Partitions*.

15.13.5 Disk Druid's Buttons

These buttons control Disk Druid's actions. They are used to add and delete partitions, and to change partition attributes. There are also buttons that are used to accept the changes you've made, or to exit Disk Druid. Let's 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'll 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 Appendix E, *RAID (Redundant Array of Independent Disks)*.

15.13.6 Adding Partitions

To add a new partition, select the **Add** button. A dialog box will appear (see Figure 15–10, *Adding a Partition*).

Please Note

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

Figure 15–10 Adding a Partition

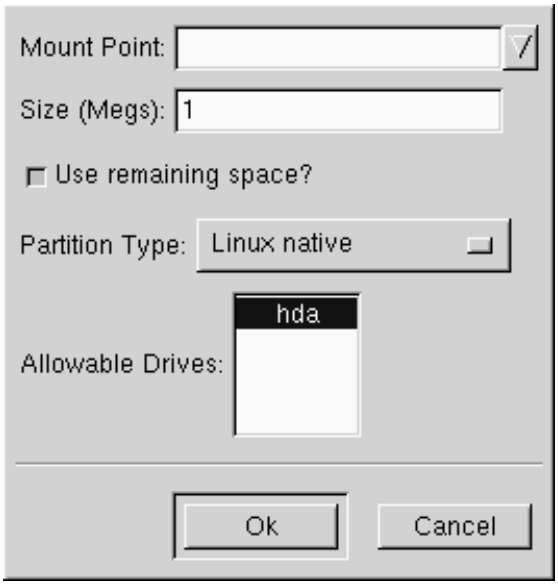

- **Mount Point**: Highlight and enter the partition's mount point. For example, if this partition should be the root partition, enter `/`; enter `/boot` for the `/boot` partition, and so on. You can also use the pull-down menu to choose the correct mount point for your partition.
- **Size (Megs)**: Enter the size (in megabytes) of the partition. Note this field starts with a "1" in it; unless changed you'll end up with a 1 MB partition.
- **Use remaining space**: This check box indicates if 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 expand and contract 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 check box settings, you can direct Disk Druid to place partitions as you see fit, or let Disk Druid decide where partitions should go.
- **Ok:** Select **Ok** once you're satisfied with the settings, and wish to create the partition.
- **Cancel:** Select **Cancel** if you don't want to create the partition.

15.13.7 Editing Partitions

To edit a partition, select the **Edit** button or double-click on the existing partition (see Figure 15–11, *Editing a Partition*).

Figure 15–11 Editing a Partition

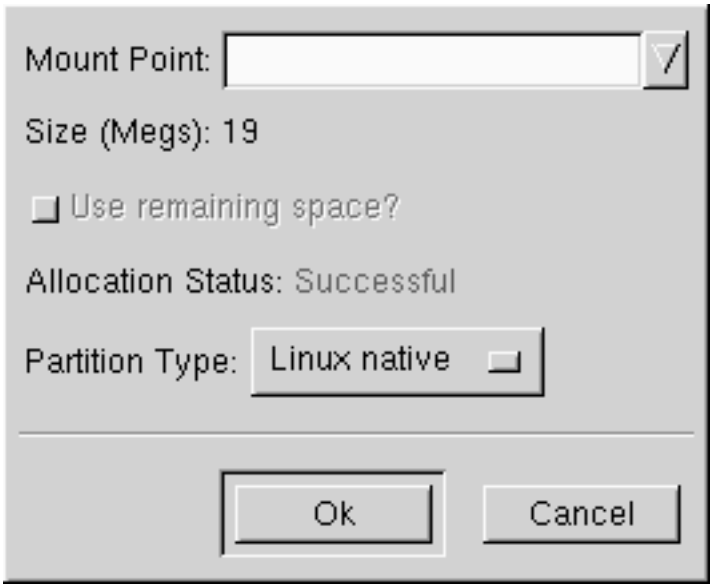

Please Note

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

15.13.8 Deleting a Partition

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

Skip to Section 15.15, *Choose Partitions to Format* for further installation instructions.

15.14 Partitioning with fdisk

Unless you have previously used fdisk and understand how it works, we do not recommend that you use it. Disk Druid is an easier and friendlier partitioning tool for those new to partitioning their system. To exit fdisk click **Back** to return to the previous screen, deselect fdisk, and then click **Next**.

This section applies only if you chose to use fdisk to partition your system. If are not using fdisk, please skip to Section 15.11, *Automatic Partitioning* for automatic partitioning or Section 15.13, *Partitioning Your System* for partitioning with Disk Druid.

If you have chosen to use fdisk, the next screen (see Figure 15–12, *fdisk*) will prompt you to select a drive to partition using fdisk.

Figure 15–12 fdisk

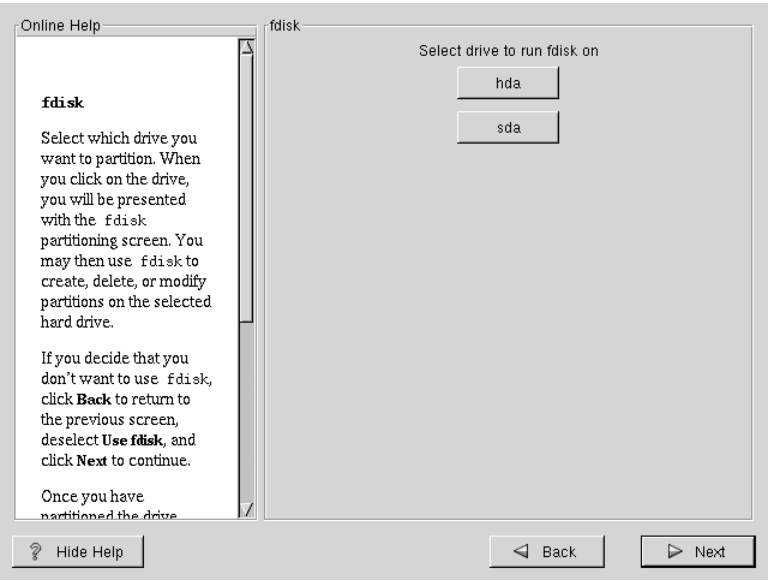

Once you have chosen which drive to partition, you will be presented with the fdisk command screen (see Figure 15–13, *Partitioning with fdisk*). If you are unsure as to what command you should use, type [m] at the prompt for help. Please refer to the Chapter 14, *Installing Red Hat Linux via Text Mode* for an overview of fdisk. When you've 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.

Figure 15–13 Partitioning with fdisk

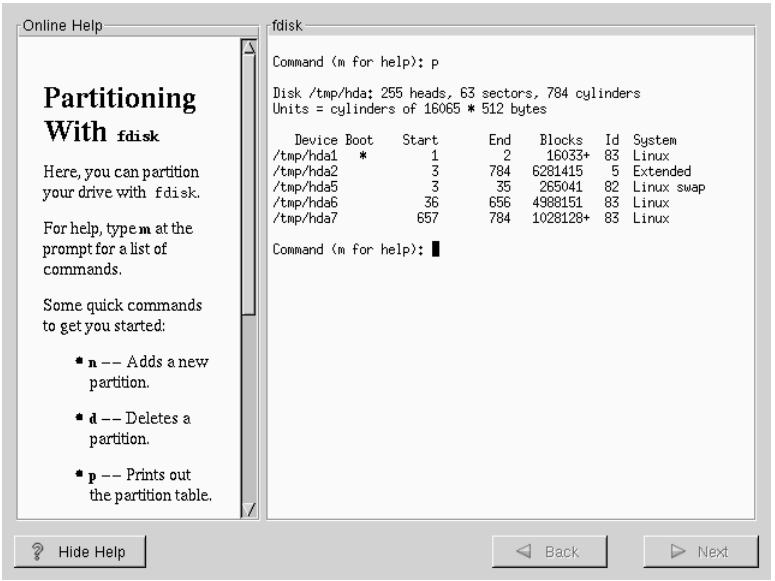

After you have partitioned your drive(s), click **Next**. You will then use Disk Druid to assign **mount points** to your partitions.

You will not be able to add new partitions using Disk Druid, but you will be able to edit mount points for those you have already created.

Figure 15–14 Editing with Disk Druid

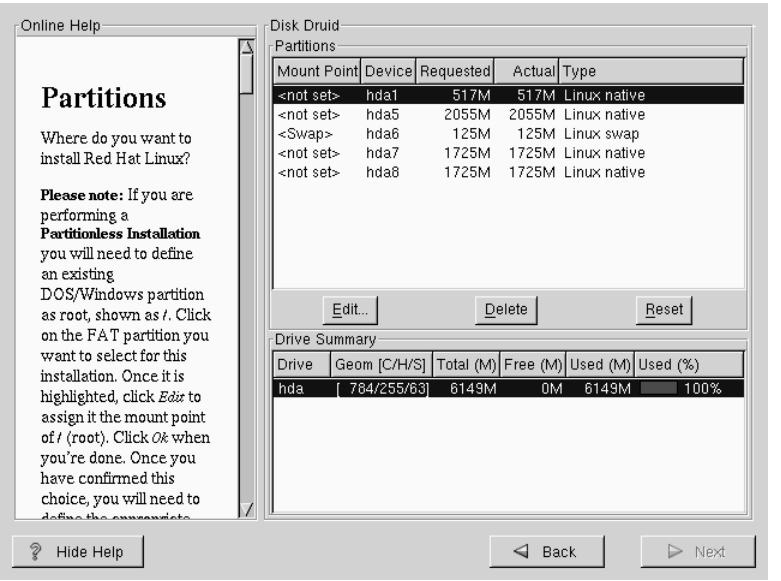

Skip to Section 15.15, *Choose Partitions to Format* for further installation instructions.

15.15 Choose Partitions to Format

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 15–15, *Choosing Partitions to Format*).

Figure 15–15 Choosing Partitions to Format

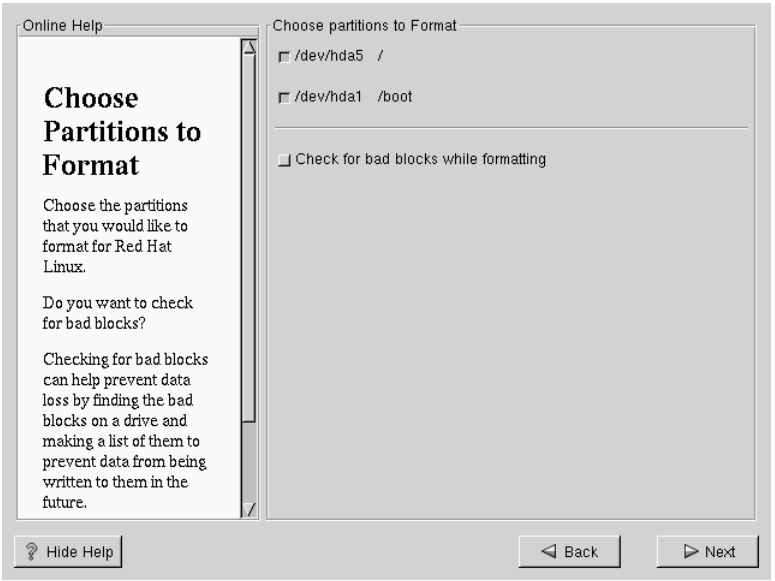

If you wish to check for bad blocks while formatting each filesystem, please make sure to select the **check for bad blocks** option.

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.

Please Note

Selecting **check for bad blocks** may dramatically increase your total installation time. Since most newer hard drives are quite large in size, checking for bad blocks may take a while depending on the size of your hard drive.

15.16 Installing LILO

If you're performing a workstation- or server-class installation, please skip ahead to Section 15.18, *Time Zone Configuration*.

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

The master boot record (MBR)

The recommended place to install LILO, unless the MBR already starts another operating system loader, such System Commander or OS/2's Boot Manager. The master boot record is a special area on your hard drive that is automatically loaded by your computer's BIOS, and is the earliest point at which LILO can take control of the boot process. If you install LILO in the MBR, when your machine boots, LILO will present a `boot:` prompt. You can then boot Red Hat Linux or any other operating system you configure LILO to boot.

The first sector of your root partition

Recommended if you are already using another boot loader on your system (such as OS/2's Boot Manager). In this case, your other boot loader will take control first. You can then configure that boot loader to start LILO (which will then boot Red Hat Linux).

If you choose to install LILO, please select where you would like LILO to be installed on your system (see Figure 15–16, *LILO Configuration*). If your system will use only Red Hat Linux you should choose the master boot record (MBR). For systems with Win95/98, you also should install LILO to the MBR so that LILO can boot both operating systems.

If you have Windows NT (and you want to install LILO) you should choose to install LILO on the first sector of the root partition, not the MBR. Please be sure to create a boot disk. In a case such as this, you will either need to use the boot disk, or configure the NT system loader to boot LILO from the first sector of the root partition. Be sure to check out http://www.linuxdoc.org/HOWTO/mini/Linux+NT-Loader.html for more information on setting up LILO and NT.

CAUTION

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

The `Use linear mode` button is selected by default. In most cases, linear mode should be enabled; if your computer cannot use linear mode to access your hard drives, deselect this option.

Figure 15–16 LILO Configuration

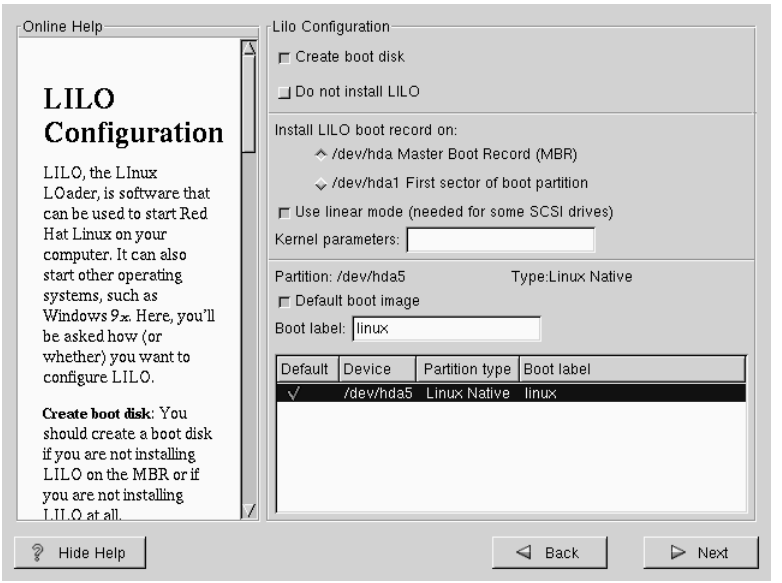

If you wish to add default options to the LILO boot command, enter them into the kernel parameters field. Any options you enter will be passed to the Linux kernel every time it boots.

Bootable Partition -- Every bootable partition is listed, including partitions used by other operating systems. The "Boot label" column will be filled in with the word `linux` on the partition holding your Red Hat Linux system's root filesystem. Other partitions may also have boot labels. If you would like to add boot labels for other partitions (or change an existing boot label), click once on the partition to select it. Once selected, you can change the boot label.

Please Note

The "Boot label" column lists what you must enter at LILO's `boot:` prompt in order to boot the desired operating system. However, if you forget the boot labels defined on your system, you can always press [Tab] at LILO's `boot:` prompt to display a list of defined boot labels.

15.16.1 Configuring LILO

- **Create boot disk** -- The **Create boot disk** option is checked by default. If you do not want to create a boot disk, you should deselect this option. However, we strongly urge you to create a boot disk. A boot disk can be handy for a number of reasons:
  - For use instead of LILO -- You can use a boot disk instead of LILO. This is handy if you're trying Red Hat Linux for the first time, and you'd feel more comfortable if the boot process for your other operating system is left unchanged. With a boot disk, going back to your other operating system is as easy as removing the boot disk and rebooting. If you would rather use a boot disk instead of LILO, make sure you enable the **Do not install LILO** option.
  - For use if another operating system overwrites LILO -- Other operating systems may not be as flexible as Red Hat Linux when it comes to supported boot methods. Quite often, installing or updating another operating system can cause the master boot record (originally containing LILO) to be overwritten, making it impossible to boot your Red Hat Linux installation. The boot disk can then be used to boot Red Hat Linux so you can reinstall LILO.

- **Do not install LILO** -- if you have Windows NT installed on your system, you may not want to install LILO. If you choose not to install LILO for this reason, make sure that you have chosen to create a boot disk; otherwise you will not be able to boot Linux. You can also choose to skip LILO if you do not want to write LILO to your hard drive.

Tip

To use the boot disk with rescue mode, you have several options:

- Using the CD-ROM to boot, type `linux rescue` at the `boot:` prompt.
- Using the network boot disk, type `linux rescue` at the `boot:` prompt. You will then be prompted to pull the rescue image from the network.
- Using the boot disk included with the Red Hat Linux boxed set, type `linux rescue` at the `boot:` prompt. You then pick an installation method and choose a valid installation tree to load from.

For more information regarding rescue mode, refer to Chapter 2, *System Administration*.

15.16.2 Alternatives to LILO

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

Boot Disk

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

LOADLIN

You can load Linux from MS-DOS. Unfortunately, it requires a copy of the Linux kernel (and an initial RAM disk, if you have a SCSI adapter) to be available on an MS-DOS partition. The only way to accomplish this is to boot your Red Hat Linux system using some other method (e.g., from LILO on a diskette) and then copy the kernel to an MS-DOS partition. `LOADLIN` is available from ftp://metalab.unc.edu/pub/Linux/system/boot/dualboot/ and associated mirror sites.

SYSLINUX

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

Some commercial bootloaders

For example, System Commander and Partition Magic, which are able to boot Linux (but still require LILO to be installed in your Linux root partition).

15.16.3 SMP Motherboards and LILO

This section is specific to SMP motherboards only. If the installer detects an SMP motherboard on your system, it will automatically create two **lilo.conf** entries, rather than the usual single entry.

One entry will be called **linux** and the other will be called **linux-up**. The *linux* will boot by default. However, if you have trouble with the SMP kernel, you can elect to boot the *linux-up* entry instead. You will retain all the functionality as before, but you will only be operating with a single processor.

15.17 Network Configuration

If you have a network card and have not already configured your networking information, you now have the opportunity to configure networking (as shown in Figure 15–17, *Network Configuration*).

Choose your device type and whether you would like to configure using DHCP. If you have multiple Ethernet devices, each device will keep the information you have

provided. You may switch between devices, for example eth0 and eth1, and the information you give will be specific to each device. If you select **Activate on boot**, your network interface 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, please contact your network administrator.

Figure 15–17 Network Configuration

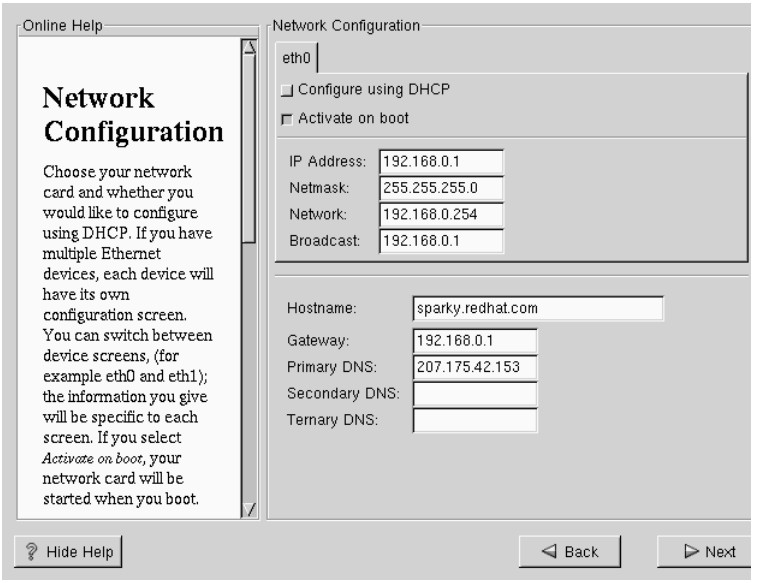

Tip

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

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

15.18 Time Zone Configuration

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

Figure 15–18 Configuring Time Zone

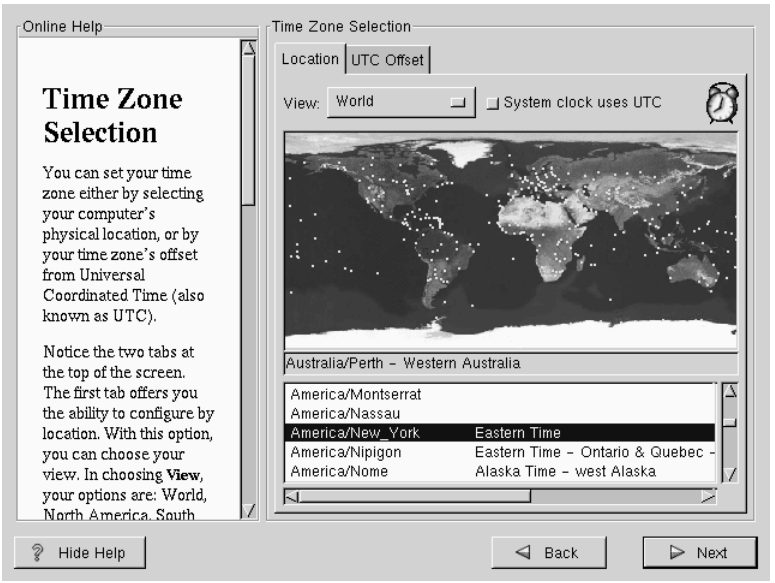

Notice the two tabs at the top of the screen (see Figure 15–18, *Configuring Time Zone*). The first tab offers you the ability to configure by location. With this option, you can choose your view. In choosing **view**, your options are: **World**, **North America**, **South America**, **Pacific Rim**, **Europe**, **Africa**, and **Asia**.

From the interactive map, you can also click on a specific city, as indicated by the yellow dots; a red **x** will appear indicating your selection. You can also scroll through a list and choose your desired time zone.

The second tab offers you the ability to use the UTC offset. UTC presents you with a list of offsets to choose from, as well as an option to set daylight saving time.

For both tabs, there is the option of selecting **System Clock uses UTC**. Please select this if you know that your system is set to UTC.

Tip

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

15.19 Account Configuration

The **Account Configuration** screen allows you to set your root password. Additionally, you can set up user accounts for you to log into once the installation is complete (see Figure 15–19, *Account Creation*).

Figure 15–19 Account Creation

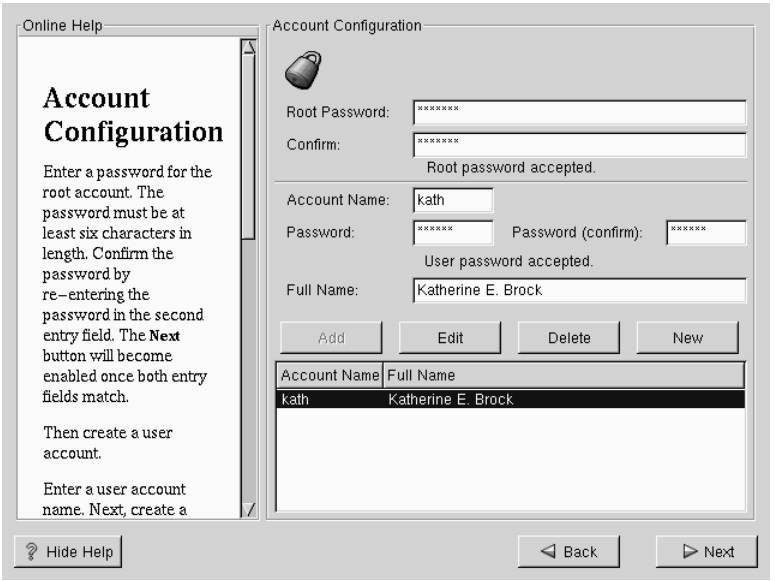

15.19.1 Setting the Root Password

The installation program will prompt you to set a **root password** for your system.

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 poor passwords. Good passwords mix numerals with upper and lower case letters and do not contain dictionary words: **Aard387vark** or **420BMttNT**, for example. Remember that the password is case-sensitive. Write down this password and keep it in a secure place.

Please Note

The **root** user (also known as the **superuser**) has complete access to the entire system; for this reason, logging in as the root user is best done *only* to perform system maintenance or administration.

15.19.2 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 other accounts.

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

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 or no longer want.

15.20 Authentication Configuration

If you are performing a workstation-class installation, please skip ahead to Section 15.22, *GUI X Configuration Tool*.

If you are performing a server-class installation, please skip ahead to Section 15.23, *Preparing to Install*.

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

Unless you are setting up **NIS** authentication, you will notice that both **MD5** and **shadow** passwords are selected (see Figure 15–20, *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 unsure whether you are connected to an NIS network, please ask your system administrator.

Figure 15–20 Authentication Configuration

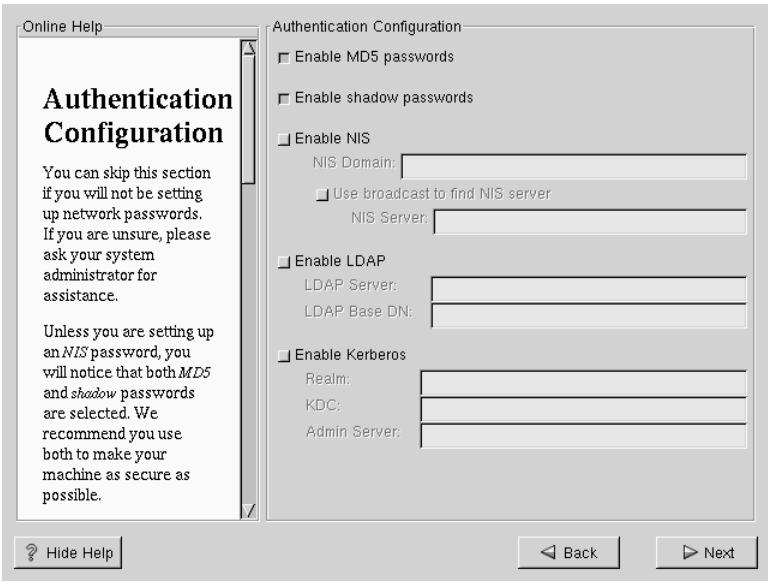

- **MD5 Password** -- allows a long password to be used (up to 256 characters), instead of the standard eight letters or less.
- **Shadow Password** -- provides a secure method of retaining passwords. The passwords are stored in `/etc/shadow`, which is readable only by root.
- **Enable NIS** -- allows you to run a group of computers in the same Network Information Service domain with a common password and group file. There are two options to choose from here:
  - **NIS Domain** -- this option allows you to specify which domain or group of computers your system belongs to.

  - **NIS Server** -- this option 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** -- 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 Chapter 7, *Lightweight Directory Access Protocol (LDAP)*. There are two options to choose from here:

  - **LDAP Server** -- this option allows you to access a server running the LDAP protocol.

  - **LDAP Base DN** -- this option allows you to look up user information by its *Distinguished Name* (DN).

- **Enable Kerberos** -- Kerberos is a secure system for providing network authentication services. For more information about Kerberos, see Chapter 8, *Using Kerberos 5 on Red Hat Linux*. 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 very 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.

15.21 Package Group Selection

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

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 check box beside it (see Figure 15–21, *Package Group Selection*).

Figure 15–21 Package Group Selection

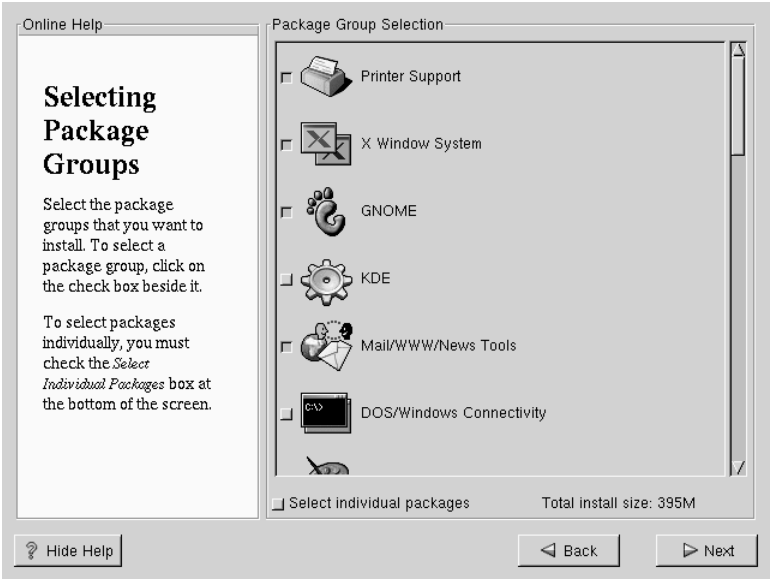

Select each component you wish to install. Selecting `Everything` (which can be found at the end of the component list) installs all packages included with Red Hat Linux. Selecting every package will require close to 1.7GB of free disk space.

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

15.21.1 Selecting Individual Packages

After selecting the components you wish to install, you can select or deselect individual packages. The installation program presents a list of the packages in that group, which you can select or deselect using your mouse (see Figure 15–22, *Selecting Individual Packages*).

Figure 15–22 Selecting Individual Packages

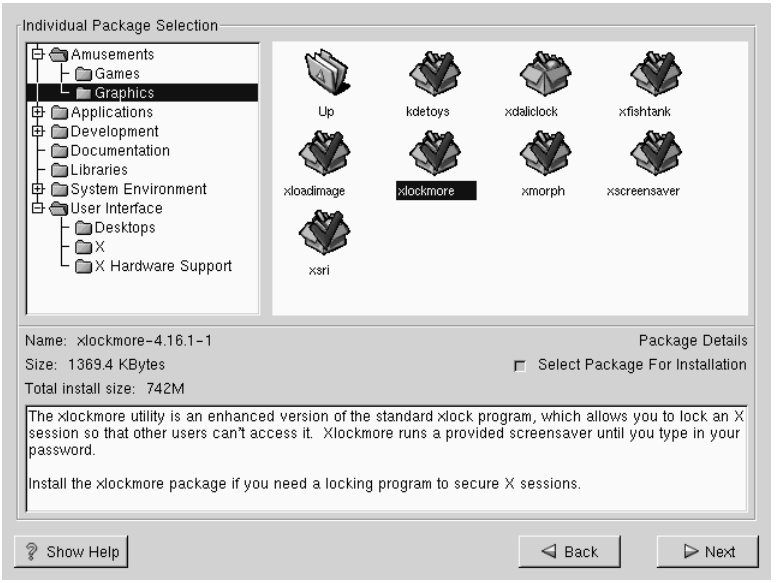

On the left side of the screen you will see a directory listing of various package groups. When you expand this list (double-click to select it) and double-click on a single directory, the list of packages available for installation will appear on the right.

To select an individual package, double-click on it, or click on it once to highlight it and click on the **Select Package For Installation** button below. A red check mark will appear on any of the packages you have selected for installation.

To read information about a particular package before choosing it for installation, left-click on it once to highlight it, and the information will appear at the bottom of the screen along with the name and size of the package.

Please 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.

15.21.2 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 15–23, *Unresolved Dependencies*).

The **Unresolved Dependencies** screen will only appear if you are missing certain packages that are needed by your selected packages. Under the list of missing packages, there is an **Install packages to satisfy dependencies** check box at the bottom of the screen which 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 15–23 Unresolved Dependencies

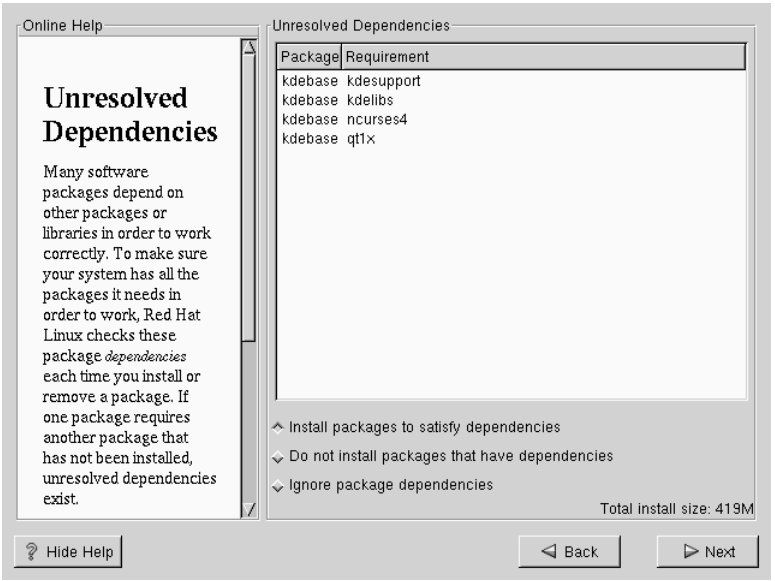

15.22 GUI X Configuration Tool

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 15.24, *Installing Packages*.

15.22.1 Configuring Your Monitor

Xconfigurator, the X Window System configuration tool, first 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 15–24 Monitor Selection

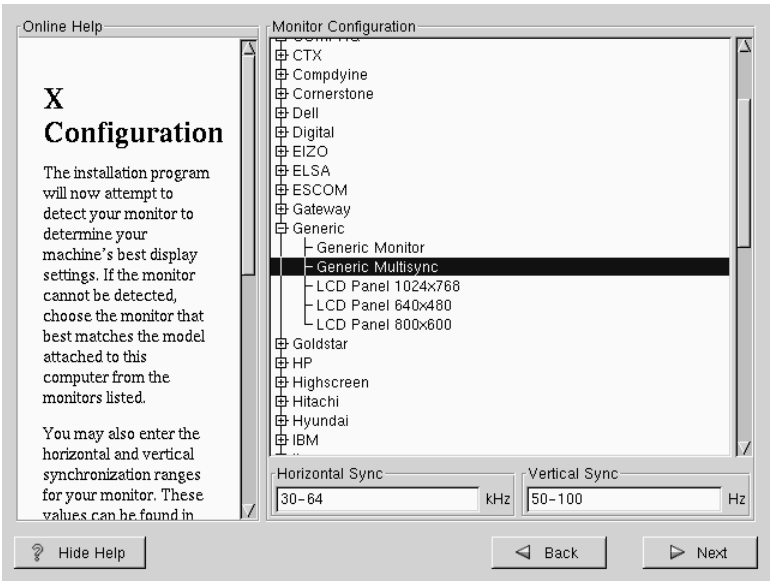

If your monitor does not appear on the list, select the most appropriate **Generic** model available. If you do select a **Generic** monitor, Xconfigurator will suggest horizontal and vertical sync ranges. These values are generally available in the documentation which accompanies your monitor, or from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

CAUTION

Do not select a monitor *similar* to your monitor unless you are certain that the monitor you are selecting does not exceed the capabilities of your monitor. Doing so may overclock your monitor and damage or destroy it.

Also presented are the horizontal and vertical ranges that Xconfigurator suggests.

Click **Next** when you have finished configuration of your monitor.

15.22.2 Video Hardware Configuration

Next, Xconfigurator will probe for any video hardware you have (see Figure 15–25, *Videocard Setup*). Failing that, Xconfigurator will present a list of video cards and monitors for you to select from.

If your video card does not appear on the list, **XFree86** may not support it. However, if you have technical knowledge about your card, you may choose **Unlisted Card** and attempt to configure it by matching your card's video chipset with one of the available X servers.

Figure 15–25 Videocard Setup

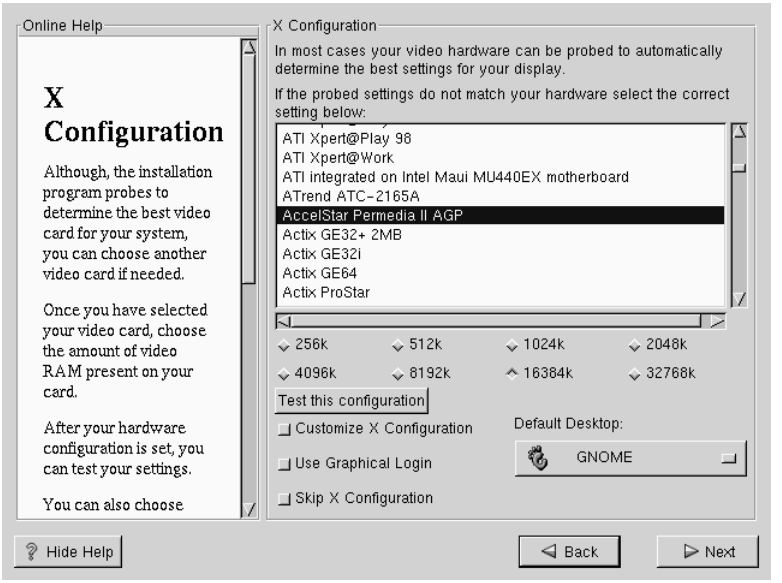

Next, Xconfigurator prompts you for the amount of video memory installed on your video card. If you are not sure, please consult the documentation accompanying your video card. You will not damage your video card by choosing more memory than is available, but the XFree86 server may not start correctly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

Once your hardware has been determined, you can test the configuration settings. We recommend that you do test your configuration to make sure that the resolution and color is what you want to work with.

If you would like to customize the X configuration, please make sure the **Customize X Configuration** button is selected. If you choose to customize, you will be presented with another screen that lets you select what your resolution should be (see Figure 15–26, *X Customization*). Again, you will have the option of testing the configuration.

Be sure to select either GNOME or KDE as your desktop default, if you installed one or both of them.

Figure 15–26 X Customization

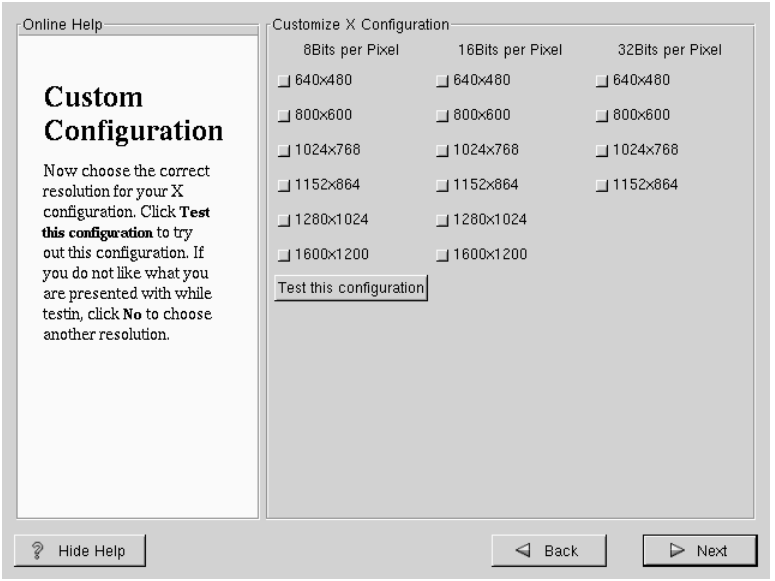

You may also choose to **Skip X Configuration** if you would rather configure X after the install or not at all.

15.23 Preparing to Install

You will now see a screen preparing you for the installation of Red Hat Linux (see Figure 15–27, *Ready to Install*).

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.

Figure 15–27 Ready to Install

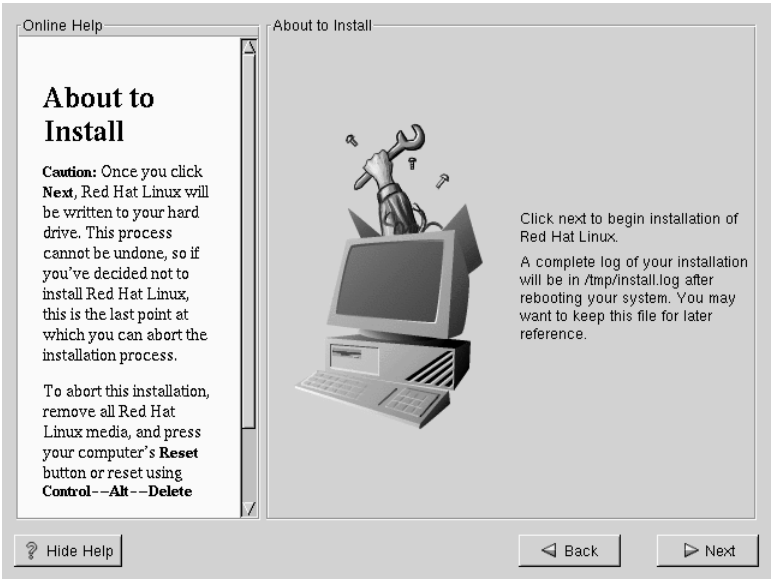

15.24 Installing Packages

At this point there's nothing left for you to do until all the packages have been installed (see Figure 15–28, *Installing Packages*). How quickly this happens depends on the number of packages you've selected, and your computer's speed.

Figure 15–28 Installing Packages

15.25 Boot Disk Creation

If you chose to to create a boot disk, you should now insert a blank, formatted diskette into your floppy drive (see Figure 15–29, *Creating Your Boot Disk*).

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

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

Figure 15–29 Creating Your Boot Disk

15.26 Installation Complete

Congratulations! Your Red Hat Linux 7.0 installation is now complete!

The installation program will prompt you to prepare your system for reboot (see Figure 15–30, *Installation Complete*). Don't forget to remove any diskette in the floppy drive or CD in the CD-ROM drive. If you did not install LILO, you'll need to use your boot disk now.

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

- Press [Enter] -- Causes LILO's default boot entry to be booted.
- Select a boot label, followed by [Enter] -- Causes LILO to boot the operating system corresponding to the boot label. (Press [?] at the LILO text `boot:` for a list of valid boot labels.)

- Do nothing -- After LILO's timeout period, (which, by default, is five seconds) LILO will automatically boot the default boot entry.

Figure 15–30 Installation Complete

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

Part IV Appendixes

A General Parameters and Modules

This appendix is provided to illustrate *some* of the possible parameters that may be needed by certain drivers. It should be noted that, in most cases, these additional parameters are unnecessary. Also included is a list of network hardware and the associated modules required by that hardware.

Please keep in mind that if a device you are attempting to use requires one of these parameters, and support for that device is *not* compiled into the kernel, the traditional method of adding the parameter to the LILO boot command will not work. Drivers loaded as modules require that these parameters are specified when the module is loaded. The Red Hat Linux installation program gives you the option to specify module parameters when a driver is loaded.

A.1 A Note About Kernel Drivers

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

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

After the installation is complete you may want to rebuild a kernel that includes support for your specific hardware configuration. See Section 2.8, *Building a Custom Kernel* for information on building a customized kernel. Note that, in most cases, a custom-built kernel is not necessary.

A.2 CD-ROM Module Parameters

Please Note

Not all of the cards that are listed are supported. Please check the hardware compatibility list on Red Hat's World Wide Web site at http://www.redhat.com/support/hardware to make sure your card is supported.

One of the more commonly used parameters, the `hdX=cdrom` parameter, *can* be entered at the boot prompt, as it deals with support for IDE/ATAPI CD-ROMs, which is part of the kernel.

In the tables below, most modules without any parameters listed are either able to auto-probe to find the hardware, or require you to manually change settings in the module source code, and recompile.

Table A–1 Hardware Parameters

Hardware	Module	Parameters
ATAPI/IDE CD-ROM Drives		hdX=cdrom
Aztech CD268-01A, Orchid CD-3110, Okano/Wearnes CDD110, Conrad TXC, CyCDROM CR520, CyCDROM CR540 (non-IDE)	`aztcd.o`	aztcd=io_port
Sony CDU 31A or 33A CD-ROM	`cdu31a.o`	cdu31a=io_port,IRQ[,PAS] cdu31a_port=base_addr cdu31a_irq=irq

Hardware	Module	Parameters
Philips/LMS CDROM drive 206 with cm260 host adapter card	cm206.o	cm206=io_port,IRQ
Goldstar R420 CD-ROM	gscd.o	gscd=io_port
ISP16, MAD16, or Mozart sound card CD-ROM interface (OPTi 82C928 and OPTi 82C929) with Sanyo/Panasonic, Sony, or Mitsumi drives	isp16.o	isp16=io_port, IRQ, dma,drive_type isp16_cdrom_base=io_port isp16_cdrom_irq=IRQ isp16_cdrom_dma=dma isp16_cdrom_type=drive_type
Mitsumi CD-ROM, Standard	mcd.o	mcd=io_port,IRQ
Mitsumi CD-ROM, Experimental	mcdx.o	mcdx=io_port_1, IRQ_1, io_port_n, IRQ_n
Optics storage 8000 AT CD-ROM "Dolphin" drive; Lasermate CR328A	optcd.o	optcd=io_port
SB Pro 16 compatible	sbpcd.o	sbpcd=io_port,sb_pro_Setting
Sanyo CDR-H94A	sjcd.o	sjcd=io_port sjcd_base=io_port
Sony CDU-535 & 531 (some Procomm drives)	sonycd535.o	sonycd535=io_port

Here are some examples of these modules in use:

Table A–2 Hardware Parameters Configuration Examples

Configuration	Example
ATAPI CD-ROM, jumpered as master on 2nd IDE channel	hdc=cdrom

Configuration	Example
non-IDE Mitsumi CD-ROM on port 340, IRQ 11	mcd=0x340,11
Three non-IDE Mitsumi CD-ROM drives using the experimental driver, io ports 300, 304, and 320 with IRQs 5, 10 and 11	mcdx=0x300,5,0x304,10,0x320,11
Sony CDU 31 or 33 at port 340, no IRQ	cdu31=0x340,0 cdu31_port=0x340 cdu31a_irq=0
Aztech CD-ROM at port 220	aztcd=0x220
Panasonic-type CD-ROM on a SoundBlaster interface at port 230	sbpcd=0x230,1
Phillips/LMS cm206 and cm260 at IO 340 and IRQ 11	cm206=0x340,11
Goldstar R420 at IO 300	gscd=0x300
Mitsumi drive on a MAD16 soundcard at IO Addr 330 and IRQ 1, probing DMA	isp16=0x330,11,0,Mitsumi
Sony CDU 531 at IO address 320	sonycd535=0x320

Please Note

Most newer Sound Blaster cards come with IDE interfaces. For these cards, you do not need to use `sbpcd` parameters, only use `hdx` parameters.

A.3 SCSI parameters

Table A–3 SCSI Parameters

Hardware	Module	Parameters
NCR53c810/820/720, NCR53c700/710/700-66	`53c7,8xx.o`	
AM53/79C974 PC-SCSI Driver Qlogic PCI-Basic	`AM53C974.o`	AM53C974=host-scsi-id, target-scsi-id,max-rate, max-offset
Most Buslogic (now Mylex) cards with "BT" part number	`BusLogic.o`	BusLogic_Options=option,option,… (See `README.BusLogic` in …`drivers/scsi/`)
	`NCR53c406a.o`	ncr53c406a=io_port[, IRQ[, FASTPIO]] ncr53c406a io=io_port irq=IRQ fastpio=FASTPIO
Advansys SCSI Cards	`advansys.o`	
Adaptec AHA 152x	`aha152x.o`	aha152x=io_base, IRQ, scsi_id, reconnect, parity
Adaptec AHA 1542	`aha1542.o`	aha1542=io_base,buson,busoff,dmaspeed
Adaptec AHA 1740	`aha1740.o`	

Hardware	Module	Parameters
Adaptec AHA-274x, AHA-284x, AHA-29xx, AHA-394x, AHA-398x, AHA-274x, AHA-274xT, AHA-2842, AHA-2910B, AHA-2920C, AHA-2930/U/U2, AHA-2940/W/U/UW/AU/, U2W/U2/U2B/, U2BOEM, AHA-2944D/WD/UD/UWD, AHA-2950U2/W/B, AHA-3940/U/W/UW/, AUW/U2W/U2B, AHA-3950U2D, AHA-3985/U/W/UW, AIC-777x, AIC-785x, AIC-786x, AIC-787x, AIC-788x , AIC-789x, AIC-3860	`aic7xxx.o`	aic7xxx=string
Data Technology Corp DTC3180/3280	`dtc.o`	

Hardware	Module	Parameters
DTP SCSI host adapters (EATA/DMA) PM2011B/9X ISA, PM2021A/9X ISA, PM2012A, PM2012B, PM2022A/9X EISA, PM2122A/9X, PM2322A/9X, SmartRAID PM3021, PM3222, PM3224	`eata.o`	eata=port0,port1,port2,… options eata io_port=port0,port1,port2,… option=value
DTP SCSI Adapters PM2011, PM2021, PM2041, PM3021, PM2012B, PM2022, PM2122, PM2322, PM2042, PM3122, PM3222, PM3332, PM2024, PM2124, PM2044, PM2144, PM3224, PM3334	`eata_dma.o`	
DTP EATA-PIO boards	`eata_pio.o`	

Hardware	Module	Parameters
Future Domain TMC-16x0- based cards TMC-1800, TMC-18C50, TMC-18C30, TMC-36C70, Future Domain TMC-1650, TMC-1660, TMC-1670, TMC-1680, TMC-1610M/MER/MEX, TMC-3260 (PCI), Quantum ISA-200S, ISA-250MG, Adaptec AHA-2920A (PCI) (NOT AHA-2920C)	`fdomain.o`	fdomain=io_base, IRQ[, ADAPTER_ID]
NCR5380 and NCR53c400 cards	`g_NCR5380.o`	ncr5380=io_port,IRQ,dma ncr53c400=io_port,IRQ ncr5380 io=io_port irq=IRQ dma=dma ncr53c400 io=io_port irq=IRQ
GDT ISA/EISA/PCI Disk Array Controller	`gdth.o`	gdth=IRQ0,IRQ1,IRQ2,… options:values
IOMEGA MatchMaker parallel port SCSI adapter	`imm.o`	
Always IN2000 ISA SCSI card	`in2000.o`	in2000=setup_string:value in2000 setup_string=value
Initio INI-9X00U/UW SCSI host adapters	`initio.o`	
AMI MegaRAID 418, 428, 438, 466, 762	`megaraid.o`	

Hardware	Module	Parameters
NCR SCSI controllers with 810/810A/815/ 825/825A/860/875/876/895 chipsets	`ncr53c8xx.o`	ncr53c8xx=option1:value1,option2:value2,… ncr53c8xx="option1:value1 option2:value2…"
Pro Audio Spectrum/Studio 16	`pas16.o`	pas16=port,irq
IOMEGA PPA3 parallel port SCSI host adapter	`ppa.o`	
Perceptive Solutions PSI-240I EIDE	`psi240i.o`	
QLogic Fast SCSI FASXXX ISA/VLB/PCMCIA	`qlogicfas.o`	
QLogic ISP2100 SCSI-FCP	`qlogicfc.o`	
QLogic ISP1020 Intelligent SCSI cards IQ-PCI, IQ-PCI-10, IQ-PCI-D	`qlogicisp.o`	
Seagate ST01/ST02	`seagate.o`	controller_type=1 base_address=base_addr irq=irq
Future Domain TMC-885, TMC-950	`seagate.o`	controller_type=2 base_address=base_addr irq=irq

Hardware	Module	Parameters
Cards with the sym53c416 chipset	`sym53c416.o`	sym53c416=PORT-BASE[,IRQ] sym53c416 io=PORTBASE irq=IRQ
Trantor T128/T128F/T228 SCSI Host Adapter	`t128.o`	
Tekram DC390 and other AMD53C974A based PCI SCSI adapters	`tmscsim.o`	tmscsim=ID,SPEED
UltraStor 14F/34F SCSI host adapters (14F, 24F, 34F)	`u14-34f.o`	u14-34f=io_port1,io_port2,… io_port10 u14-34f io_port=io_port1,io_port2, … io_port10
UltraStor 14F, 24F, and 34F	`ultrastor.o`	
WD7000-FASST2,WD7000-ASC, WD7000-AX/MX/EX	`wd7000.o`	wd7000=IRQ,dma,io_port wd7000 io=io_port irq=IRQ dma=dma

Here are some examples of these modules in use:

Table A–4 SCSI Parameters Configuration Examples

Configuration	Example
Adaptec AHA1522 at port 330, IRQ 11, SCSI ID 7	aha152x=0x330,11,7
Adaptec AHA1542 at port 330	bases=0x330
Future Domain TMC-800 at CA000, IRQ 10	controller_type=2 base_address=0xca000 irq=10

When a parameter has commas, make sure you do *not* put a space after a comma.

A.4 Ethernet parameters

Table A–5 Ethernet Module Parameters

Hardware	Module	Parameters
3Com 3c501	`3c501.o`	3c501=io_port,IRQ
3Com 3c503 and 3c503/16	`3c503.o`	3c503=io_port,IRQ 3c503 io=io_port_1,io_port_n irq=IRQ_1,IRQ_n
3Com EtherLink Plus (3c505)	`3c505.o`	3c505=io_port,IRQ,DMA 3c505 io=io_port_1,io_port_n irq=IRQ_1,IRQ_2 dma=dma_1,dma_n
3Com EtherLink 16	`3c507.o`	3c507=io_port,IRQ 3c507 io=io_port irq=IRQ
3Com EtherLink III	`3c509.o`	3c509=IRQ
3Com ISA EtherLink XL "Corkscrew"	`3c515.o`	
3Com EtherLink PCI III/XL Vortex (3c590, 3c592, 3c595, 3c597) Boomerang (3c900, 3c905, 3c595)	`3c59x.o`	
Apricot 680x0 VME, 82596 chipset	`82596.o`	82596=IRQ 82596 irq=IRQ
Ansel Communications AC3200 EISA	`ac3200.o`	ac3200=io_port,IRQ ac3200 io=io_port_1,io_port_n irq=IRQ_1,IRQ_n

Hardware	Module	Parameters
Alteon AceNIC Gigabit Ethernet driver	`acenic.o`	acenic=trace,link acenic trace=trace link=val
Allied Telesis AT1700	`at1700.o`	at1700=io_port,IRQ at1700 io=io_port irq=IRQ
Tangent ATB-II, Novel NL-10000, Daystar Digital LT-200, Dayna DL2000, DaynaTalk PC (HL), COPS LT-95, Farallon PhoneNET PC II, III	`cops.o`	cops=io_port,IRQ cops io=io_port irq=IRQ
Modular driver for the COSA or SRP synchronous serial card	`cosa.o`	cosa=io_port,IRQ,dma
Crystal LAN CS8900/CS8920	`cs89x0.o`	cs89x0=io_port,IRQ,MEDIA_TYPE cs89x0 io=io_port irq=IRQ media=TYPE
EtherWORKS DE425 TP/COAX EISA, DE434 TP PCI, DE435/450 TP/COAX/AUI PCI DE500 10/100 PCI Kingston, LinkSys, SMC8432, SMC9332, Znyx31[45], and Znyx346 10/100 cards with DC21040 (no SROM), DC21041[A], DC21140[A], DC21142, DC21143 chipsets	`de4x5.o`	de4x5=io_port de4x5 io=io_port de4x5 args='ethX[fdx] autosense=MEDIA_STRING'

Hardware	Module	Parameters
D-Link DE-600 Ethernet Pocket Adapter	de600.o	
D-Link DE-620 Ethernet Pocket Adapter	de620.o	de620 io=io_port irq=IRQ bnc=1 utp=1
DIGITAL DEPCA & EtherWORKS DEPCA, DE100, DE101, DE200 Turbo, DE201Turbo DE202 Turbo TP/BNC, DE210, DE422 EISA	depca.o	depca=io_port,IRQ depca io=io_port irq=IRQ
Digi Intl. RightSwitch SE-X EISA and PCI	dgrs.o	
Cabletron E2100 series ethercards	e2100.o	e2100=io_port,IRQ e2100 io=io_port irq=IRQ
Intel i82595 ISA EtherExpressPro10/10+ driver	eepro.o	eepro=io_port,IRQ,mem eepro io=io_port irq=IRQ mem=mem
Intel i82557/i82558 PCI EtherExpressPro driver	eepro100.o	
Intel EtherExpress 16 (i82586)	eexpress.o	eexpress=io_port,IRQ eexpress io=io_port irq=IRQ
SMC EtherPower II 9432 PCI (83c170/175 EPIC series)	epic100.o	
Racal-Interlan ES3210 EISA Network Adapter	es3210.o	es3210=io_port,IRQ,mem es3210 io=io_port irq=IRQ mem=mem

Hardware	Module	Parameters
ICL EtherTeam 16i/32	`eth16i.o`	eth16i=io_port,mediatype eth16i ioaddr=io_port mediatype=type
EtherWORKS 3 (DE203, DE204 and DE205)	`ewrk3.o`	ewrk=io_port,IRQ ewrk io=io_port irq=IRQ
Fujitsu FMV-181/182/183/184	`fmv18x.o`	fmv18x=io_port,IRQ fmv18x io=io_port irq=IRQ
Modular driver for the Comtrol Hostess SV11	`hostess_sv11.o`	hostess_sv11=io_port, IRQ, DMABIT hostess_sv11 io=io_port irq=IRQ dma=DMABIT
HP PCLAN/plus	`hp-plus.o`	hp-plus=io_port,IRQ hp-plus io=io_port irq=IRQ
HP LAN Ethernet	`hp.o`	hp=io_port,IRQ hp io=io_port irq=IRQ
100VG-AnyLan Network Adapters HP J2585B, J2585A, J2970, J2973, J2573 Compex ReadyLink ENET100-VG4, FreedomLine 100/VG	`hp100.o`	hp100=io_port,name hp100 hp100_port=io_port hp100_name=name
IBM Token Ring 16/4	`ibmtr.o`	ibmtr=io_port,IRQ,mem ibmtr io=io_port irq=IRQ mem=mem

Hardware	Module	Parameters
AMD LANCE/PCnet Allied Telesis AT1500, HP J2405A, NE2100, NE2500	`lance.o`	lance=io_port,IRQ,dma lance io=io_port_1,io_port_n irq=IRQ_1,IRQ_2 dma=dma_1,dma_n
Mylex LNE390 EISA cards (LNE390A, LNE390B)	`lne390.o`	lne390=io_port,IRQ,mem lne390 io=io_port irq=IRQ mem=mem
	`ltpc.o`	ltpc=io_port,IRQ ltpc io=io_port irq=IRQ
NE1000 / NE2000 (non-pci)	`ne.o`	ne=io_port,IRQ ne io=io_port irq=IRQ
PCI NE2000 cards RealTEk RTL-8029, Winbond 89C940, Compex RL2000, KTI ET32P2, NetVin, NV5000SC, Via 82C926, SureCom NE34	`ne2k-pci.o`	
Novell NE3210 EISA Network Adapter	`ne3210.o`	ne3210=io_port,IRQ,mem ne3210 io=io_port irq=IRQ mem=mem
MiCom-Interlan NI5010 ethercard	`ni5010.o`	ni5010=io_port,IRQ ni5010 io=io_port irq=IRQ
NI5210 card (i82586 Ethernet chip)	`ni52.o`	ni52=io_port,IRQ ni52 io=io_port irq=IRQ
NI6510, ni6510 EtherBlaster	`ni65.o`	ni65=io_port,IRQ,dma ni65 io=io_port irq=IRQ dma=dma

Hardware	Module	Parameters
AMD PCnet32 and AMD PCnetPCI	`pcnet32.o`	
RedCreek Communications PCI	`rcpci.o`	
RealTek cards using RTL8129 or RTL8139 Fast Ethernet chipsets	`rtl8139.o`	
Sangoma S502/S508 multi-protocol FR	`sdla.o`	
Sangoma S502A, ES502A, S502E, S503, S507, S508, S509	`sdladrv.o`	
SysKonnect Token Ring ISA/PCI Adapter, TR4/16(+) ISA or PCI, TR4/16 PCI, and older SK NET TR4/16 ISA cards	`sktr.o`	sktr=io_port,IRQ,mem sktr io=io_port irq=IRQ mem=mem
SMC Ultra and SMC EtherEZ ISA ethercard (8K, 83c790)	`smc-ultra.o`	smc-ultra=io_port,IRQ smc-ultra io=io_port irq=IRQ
SMC Ultra32 EISA Ethernet card (32K)	`smc-ultra32.o`	
SMC 9000 series of Ethernet cards	`smc9194.o`	smc9194=io_port,IRQ smc9194 io=io_port irq=IRQ ifport=[0,1,2]

Hardware	Module	Parameters
Compaq Netelligent 10/100 TX PCI UTP Compaq Netelligent 10 T PCI UTP Compaq Integrated NetFlex 3/P Compaq Netelligent Dual 10/100 TX PCI UTP Compaq Netelligent Integrated 10/100 TX UTP Compaq Netelligent 10/100 TX Embedded UTP Compaq Netelligent 10 T/2 PCI UTP/Coax Compaq Netelligent 10/100 TX UTP Compaq NetFlex 3/P Olicom OC-2325, OC-2183, OC-2326	`tlan.o`	tlan=io_port,IRQ,aui,debug tlan io=io_port irq=IRQ Other Module Options: speed=10Mbs,100Mbs debug=0x0[1,2,4,8] aui=1 duplex=[1,2]
Digital 21x4x Tulip PCI Ethernet cards SMC EtherPower 10 PCI(8432T/8432BT) SMC EtherPower 10/100 PCI(9332DST) DEC EtherWorks 100/10 PCI(DE500-XA) DEC EtherWorks 10 PCI(DE450) DEC QSILVER's, Znyx 312 etherarray Allied Telesis LA100PCI-T Danpex EN-9400, Cogent EM110	`tulip.o`	

Hardware	Module	Parameters
VIA Rhine PCI Fast Ethernet cards with either the VIA VT86c100A Rhine-II PCI or 3043 Rhine-I D-Link DFE-930-TX PCI 10/100	via-rhine.o	
AT&T GIS (nee NCR) WaveLan ISA Card	wavelan.o	wave-lan=[IRQ,0],io_port,NWID
WD8003 and WD8013 "compatible" ethercards	wd.o	wd=io_port, IRQ, mem,mem_end wd io=io_port irq=IRQ mem=mem mem_end=end
Packet Engines Yellowfin	yellowfin.o	
G-NIC PCI Gigabit Ethernet adapter		
Z8530 based HDLC cards for AX.25	z85230.o	

Here are some examples of these modules in use:

Table A–6 Ethernet Parameter Configuration Examples

Configuration	Example
NE2000 ISA card at IO address 300 and IRQ 11	ne=0x300,11 ether=0x300,11,eth0
Wavelan card at IO 390, autoprobe for IRQ, and use the NWID to 0x4321	wavelan=0,0x390,0x4321 ether=0,0x390,0x4321,eth0

A.4.1 Using Multiple Ethernet Cards

You can use multiple Ethernet cards in one machine. If each card uses a different driver (e.g., a 3c509 and a DE425), you simply need to add `alias` (and possibly `options`) lines for each card to `/etc/conf.modules`. See Section 3.2.2, *Loading Kernel Modules* for more information.

If any two Ethernet cards use the same driver (e.g., two 3c509's or a 3c595 and a 3c905), you will need to either give the two card addresses on the driver's options line (in the case of ISA cards), or (for PCI cards) simply add one `alias` line for each card.

http://metalab.unc.edu/LDP/HOWTO/Ethernet-HOWTO.html For more information about using more than one Ethernet card, see the *Linux Ethernet-HOWTO* at http://metalab.unc.edu/LDP/HOWTO/Ethernet-HOWTO.html. If you installed the `howto` package when you installed Red Hat Linux, you can find it in the file `/usr/share/doc/HOWTO/Ethernet-HOWTO`.

B An Introduction to Disk Partitions

Disk partitions are a standard part of the personal computer landscape, and have been for quite some time. However, with so many people purchasing computers featuring preinstalled operating systems, relatively few people understand how partitions work. This chapter attempts to explain how disk partitions work so you'll find your Red Hat Linux installation is as simple as possible.

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

B.1 Hard Disk Basic Concepts

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

When discussing issues such as disk partitioning, it's important to know a bit about the underlying hardware; unfortunately, it's easy to become bogged down in details. Therefore, let's use a simplified diagram of a disk drive to help us explain what goes on "under the hood." Figure B–1, *An Unused Disk Drive* shows a brand-new, unused disk drive.

Figure B–1 An Unused Disk Drive

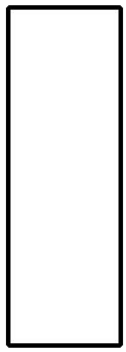

Not much to look at, is it? But if we're talking about disk drives on a basic level, it will do. Let's say that we'd like to store some data on this drive. As things stand now, it won't work. There's something we need to do first…

B.1.1 It's Not What You Write, it's How You Write It

The old-timers in the audience probably got this one on the first try. We need to **format** the drive. Formatting (usually known as "making a **filesystem**" in Linux parlance) writes information to the drive, creating order out of the empty space in an unformatted drive.

Figure B–2 Disk Drive with a Filesystem

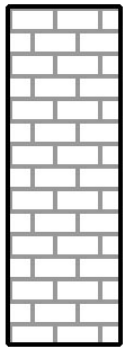

As Figure B–2, *Disk Drive with a Filesystem* implies, the order imposed by a filesystem involves some tradeoffs:

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

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

It's also worth noting that there is no single, universal filesystem; as Figure B–3, *Disk Drive with a Different Filesystem* shows, a disk drive may have one of many different filesystems written on it. As you might guess, different filesystems tend

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

to be incompatible; that is, an operating system that supports one filesystem (or a handful of related filesystem types) may not support another. This last statement is not a hard-and-fast rule, however. For example, Red Hat Linux supports a wide variety of filesystems (including many commonly used by other operating systems), making data interchange easy.

Figure B–3 Disk Drive with a Different Filesystem

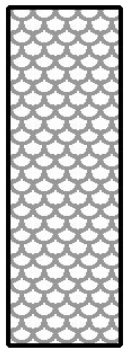

Of course, writing a filesystem to disk is only the beginning. The goal of this process is to actually *store* and *retrieve* data. Let's take a look at our drive after some files have been written to it.

Figure B–4 Disk Drive with Data Written to It

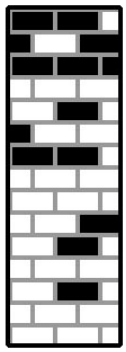

As Figure B–4, *Disk Drive with Data Written to It* shows, 14 of the previously-empty blocks are now holding data. We cannot determine how many files reside on this drive; it may be as few as one or as many as 14, as all files use at least one block. Another important point to note is that the used blocks do not have to form a contiguous region; used and unused blocks may be interspersed. This is known as **fragmentation**. Fragmentation can play a part when attempting to resize an existing partition.

As with most computer-related technologies, disk drives continued to change over time. In particular, they changed in one specific way -- they got bigger. Not bigger in size, but bigger in capacity. And it was this additional capacity that drove a change in the way disk drives were used.

B.1.2 Partitions – Turning One Drive Into Many

As disk drive capacities soared, some people started wondering if having all that space in one big chunk wasn't such a great idea. This line of thinking was driven by several

issues, some philosophical, some technical. On the philosophical side, above a certain size, it seemed that the additional space provided by a larger drive created more clutter. On the technical side, some filesystems were never designed to support larger drives. Or the filesystems *could* support larger drives, but the overhead imposed by the filesystem became excessive.

The solution to this problem was to divide disks into **partitions**. Each partition can be accessed as if it was a separate disk. This is done through the addition of a **partition table**.

Please Note

While the diagrams in this chapter show the partition table as being separate from the actual disk drive, this is not entirely accurate. In reality, the partition table is stored at the very start of the disk, before any filesystem or user data. But for clarity, we'll keep it separate in our diagrams.

Figure B–5 Disk Drive with Partition Table

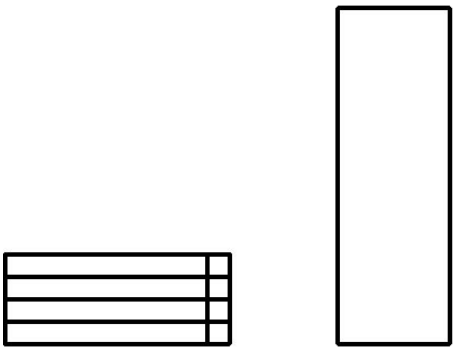

As Figure B–5, *Disk Drive with Partition Table* shows, the partition table is divided into four sections. Each section can hold the information necessary to define a single partition, meaning that the partition table can define no more than four partitions.

Each partition table entry contains several important characteristics of the partition:

- The points on the disk where the partition starts and ends;
- Whether the partition is "active";
- The partition’s type.

Let’s take a closer look at each of these characteristics. The starting and ending points actually define the partition’s size and location on the disk. The "active" flag is used by some operating systems’ boot loaders. In other words, the operating system in the partition that is marked "active" will be booted.

The partition’s type can be a bit confusing. The type is a number that identifies the partition’s anticipated usage. If that statement sounds a bit vague, that’s because the

meaning of the partition type is a bit vague. Some operating systems use the partition type to denote a specific filesystem type, to flag the partition as being associated with a particular operating system, to indicate that the partition contains a bootable operating system, or some combination of the three.

Table B–1, *Partition Types* contains a listing of some popular (and obscure) partition types, along with their numeric values.

Table B–1 Partition Types

Partition Type	Value	Partition Type	Value
Empty	00	Novell Netware 386	65
DOS 12-bit FAT	01	PIC/IX	75
XENIX root	02	Old MINIX	80
XENIX usr	03	Linux/MINUX	81
DOS 16-bit <=32M	04	Linux swap	82
Extended	05	Linux native	83
DOS 16-bit >=32	06	Linux extended	85
OS/2 HPFS	07	Amoeba	93
AIX	08	Amoeba BBT	94
AIX bootable	09	BSD/386	a5
OS/2 Boot Manager	0a	OpenBSD	a6
Win95 FAT32	0b	NEXTSTEP	a7
Win95 FAT32 (LBA)	0c	BSDI fs	b7
Win95 FAT16 (LBA)	0e	BSDI swap	b8
Win95 Extended (LBA)	0f	Syrinx	c7
Venix 80286	40	CP/M	db
Novell?	51	DOS access	e1

Partition Type	Value	Partition Type	Value
Microport	52	DOS R/O	e3
GNU HURD	63	DOS secondary	f2
Novell Netware 286	64	BBT	ff

Now you might be wondering how all this additional complexity is normally used. See Figure B–6, *Disk Drive With Single Partition* for an example.

Figure B–6 Disk Drive With Single Partition

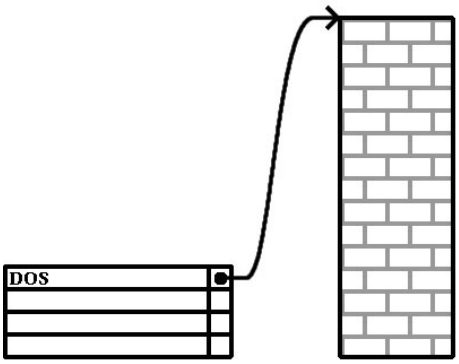

That’s right -- in many cases there is but a single partition spanning the entire disk, essentially duplicating the pre-partitioned days of yore. The partition table has only one entry used, and it points to the start of the partition.

We’ve labeled this partition as being of type "DOS," although as you can see from Table B–1, *Partition Types*, that’s a bit simplistic, but adequate for the purposes of this discussion. This is a typical partition layout for most newly purchased computers with some version of Windows pre-installed.

B.1.3 Partitions within Partitions – An Overview of Extended Partitions

Of course, in time it became obvious that four partitions would not be enough. As disk drives continued to grow, it became more and more likely that a person could configure four reasonably-sized partitions and still have disk space left over. There needed to be some way of creating more partitions.

Enter the extended partition. As you may have noticed in Table B–1, *Partition Types*, there is an "Extended" partition type; it is this partition type that is at the heart of extended partitions. Here's how it works.

When a partition is created and its type is set to "Extended," an extended partition table is created. In essence, the extended partition is like a disk drive in its own right -- it has a partition table that points to one or more partitions (now called **logical partitions**, as opposed to the four **primary partitions**) contained entirely within the extended partition itself. Figure B–7, *Disk Drive With Extended Partition* shows a disk drive with one primary partition, and one extended partition containing two logical partitions (along with some unpartitioned free space).

Figure B–7 Disk Drive With Extended Partition

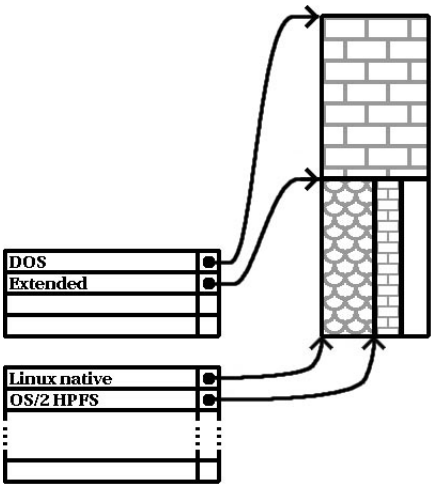

As this figure implies, there is a difference between primary and logical partitions -- there can only be four primary partitions, but there is no fixed limit to the number of logical partitions that can exist. (However, in reality it is probably not a good idea to try to define and use more than 12 logical partitions on a single disk drive.)

Now that we've discussed partitions in general, let's see how to use this knowledge to get Red Hat Linux installed.

B.1.4 Making Room For Red Hat Linux

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

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

Let’s look at each scenario in order.

Please Note

Please keep in mind that the following illustrations are simplified in the interest of clarity, and do not reflect the exact partition layout that you will encounter when actually installing Red Hat Linux.

Using Unpartitioned Free Space

In this situation, the partitions already defined do not span the entire hard disk, leaving unallocated space that is not part of any defined partition. Figure B–8, *Disk Drive with Unpartitioned Free Space* shows what this might look like.

Figure B–8 Disk Drive with Unpartitioned Free Space

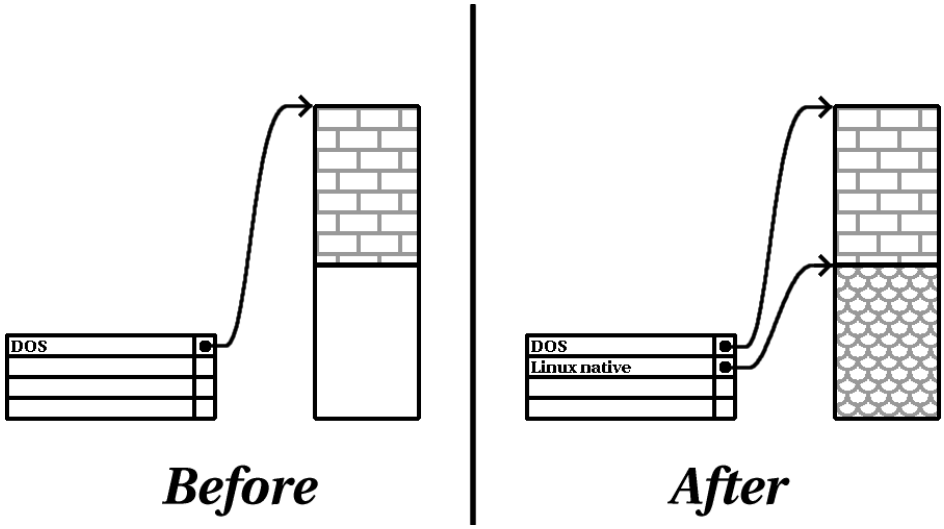

If you think about it, an unused hard disk also falls into this category; the only difference is that *all* the space is not part of any defined partition.

In any case, you can simply create the necessary partitions from the unused space. Unfortunately, this scenario, although very simple, is not very likely (unless you've just purchased a new disk just for Red Hat Linux).

Let's move on to a slightly more common situation.

Using Space From An Unused Partition

In this case, maybe you have one or more partitions that you just don't use any longer. Perhaps you've dabbled with another operating system in the past, and the partition(s) you've dedicated to it never seem to be used anymore. Figure B–9, *Disk Drive With an Unused Partition* illustrates such a situation.

Figure B–9 Disk Drive With an Unused Partition

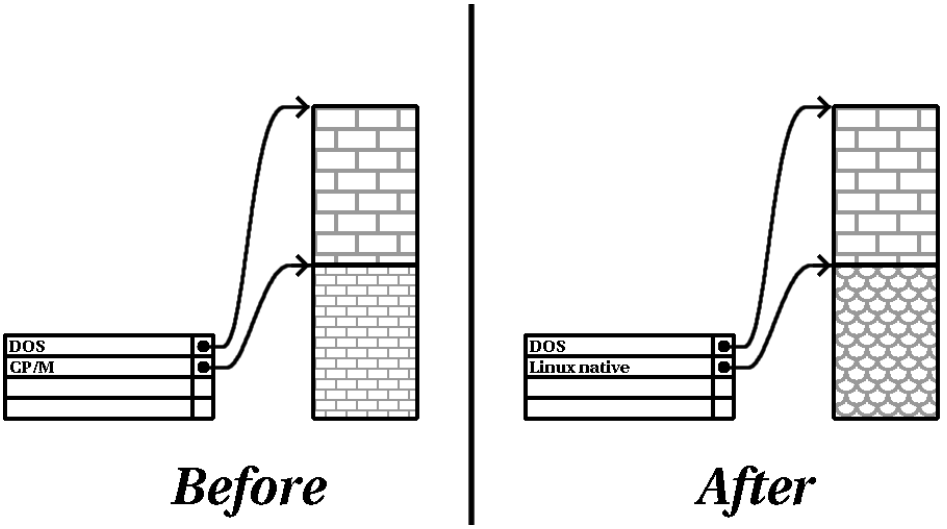

If you find yourself in this situation, you can use the space allocated to the unused partition. You'll first need to delete the partition, and then create the appropriate Linux partition(s) in its place. You can either delete the partition using DOS `fdisk`, or you'll be given the opportunity to do so during a custom-class installation.

Using Free Space From An Active Partition

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

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

Destructive Repartitioning

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

After creating a smaller partition for your existing software, you can reinstall any software, restore your data, and continue with your Red Hat Linux installation. Figure B–10, *Disk Drive Being Destructively Repartitioned* shows this being done.

Figure B–10 Disk Drive Being Destructively Repartitioned

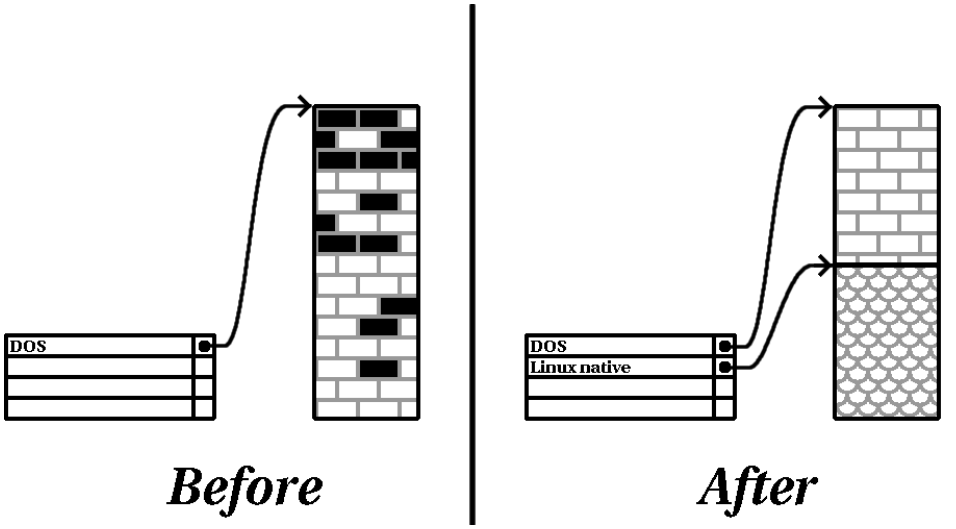

CAUTION

As Figure B–10, *Disk Drive Being Destructively Repartitioned* shows, any data present in the original partition will be lost without proper backup!

Non-Destructive Repartitioning

Here, you run a program that does the seemingly impossible: it makes a big partition smaller without losing any of the files stored in that partition. Many people have found this method to be reliable and trouble-free. What software should you use to perform this feat? There are several disk management software products on the market; you'll have to do some research to find the one that is best for your situation.

While the process of non-destructive repartitioning is rather straightforward, there are a number of steps involved:

- Compress existing data
- Resize partition
- Create new partition(s)

Let’s take a look at each step in a bit more detail.

Compress existing data

As Figure B–11, *Disk Drive Being Compressed* shows, the first step is to compress the data in your existing partition. The reason for doing this is to rearrange the data such that it maximizes the available free space at the "end" of the partition.

Figure B–11 Disk Drive Being Compressed

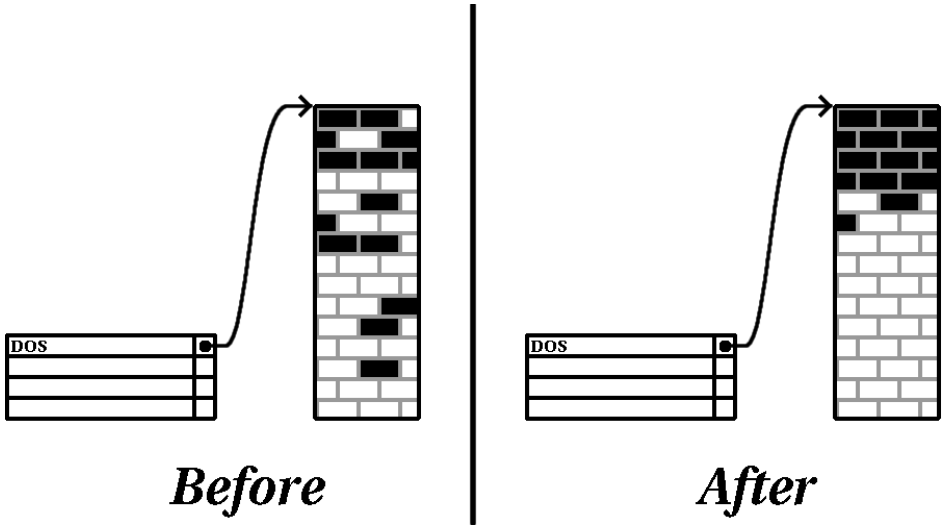

This step is crucial; without it, it is possible that the location of your data could prevent the partition from being resized to the extent desired. Note also that, for one reason or another, some data cannot be moved. If this is the case (and it restricts the size of your new partition(s)), you may be forced to destructively repartition your disk.

Resize partition

Figure B–12, *Disk Drive with Partition Resized* shows the actual resizing process. While the actual end-product of the resizing operation varies depending on the software used, in most cases the newly freed space is used to create an unformatted partition of the same type as the original partition.

Figure B–12 Disk Drive with Partition Resized

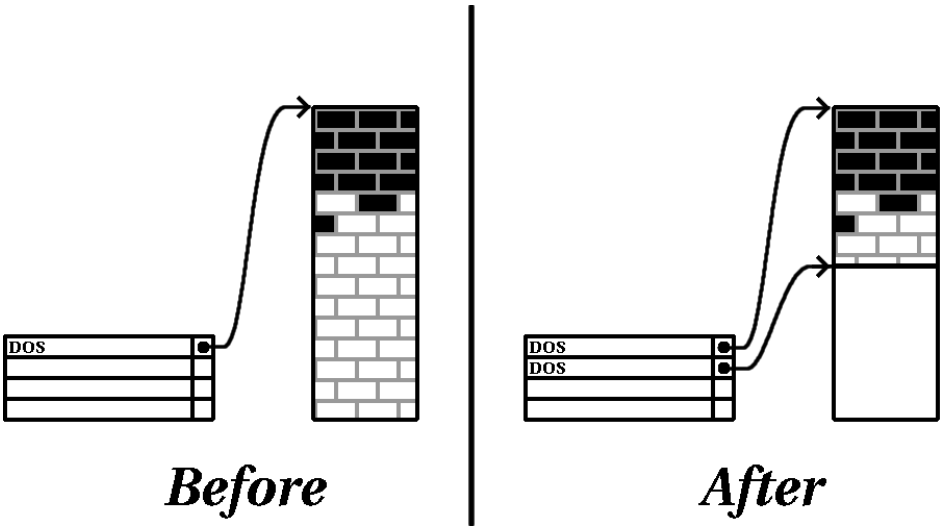

It's important to understand what the resizing software you use does with the newly freed space, so that you can take the appropriate steps. In the case we've illustrated, it would be best to simply delete the new DOS partition, and create the appropriate Linux partition(s).

Create new partition(s)

As the previous step implied, it may or may not be necessary to create new partitions. However, unless your resizing software is Linux-aware, it is likely you'll need to delete the partition that was created during the resizing process. Figure B–13, *Disk Drive with Final Partition Configuration* shows this being done.

Figure B–13 Disk Drive with Final Partition Configuration

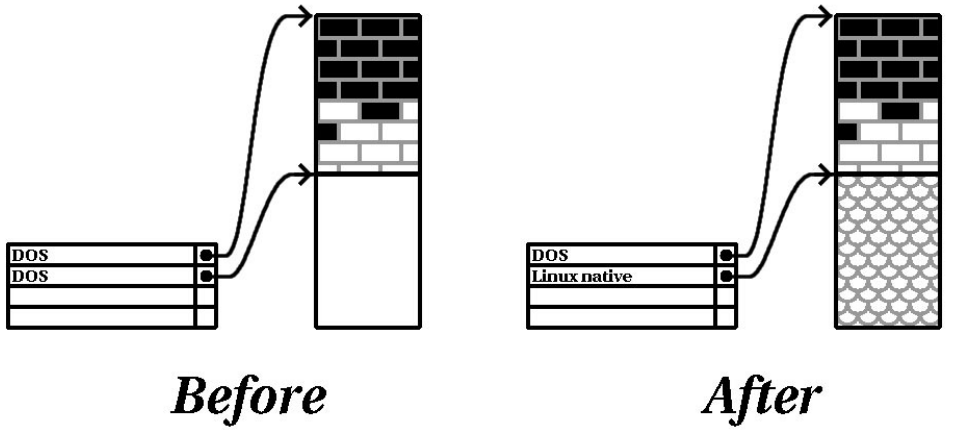

Intel

The following information is specific to Intel-based computers only.

As a convenience to our customers, we provide the `fips` utility. This is a freely available program that can resize FAT (File Allocation Table) partitions. It's included on the Red Hat Linux/Intel CD-ROM in the `dosutils` directory.

Please Note

Many people have successfully used `fips` to repartition their hard drives. However, because of the nature of the operations carried out by `fips`, and the wide variety of hardware and software configurations under which it must run, Red Hat cannot guarantee that `fips` will work properly on your system. Therefore, no installation support whatsoever is available for `fips`; use it at your own risk.

That said, if you decide to repartition your hard drive with `fips`, it is *vital* that you do two things:

- Perform a Backup -- Make two copies of all the important data on your computer. These copies should be to removable media (such as tape or diskettes), and you should make sure they are readable before proceeding.
- Read the Documentation -- Completely read the `fips` documentation, located in the `/dosutils/fipsdocs` subdirectory on Red Hat Linux/Intel CD 1.

Should you decide to use `fips`, be aware that after `fips` runs you will be left with *two* partitions: the one you resized, and the one `fips` created out of the newly freed space. If your goal is to use that space to install Red Hat Linux, you should delete the newly created partition, either by using `fdisk` under your current operating system, or while setting up partitions during a custom-class installation.

B.1.5 Partition Naming Scheme

Linux refers to disk partitions using a combination of letters and numbers which may be confusing, particularly if you're used to the "C drive" way of referring to hard disks and their partitions. In the DOS/Windows world, here is how partitions are named:

- Each partition's type is checked to determine if it can be read by DOS/Windows.
- If the partition's type is compatible, it is assigned a "drive letter." The drive letters start with "C".

- The drive letter can then be used to refer to that partition as well as the filesystem contained on that partition.

Red Hat Linux uses a naming scheme that is more flexible and conveys more information than the approach used by other operating systems. The naming scheme is file-based, with filenames in the form:

```
/dev/xxyN
```

Here's how to decipher the partition naming scheme:

`/dev/`

This string is the name of the directory in which all device files reside. Since partitions reside on hard disks, and hard disks are devices, the files representing all possible partitions reside in `/dev/`.

xx

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

y

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

N

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

Please Note

There is no part of this naming convention that is based on partition type; unlike DOS/Windows, *all* partitions can be identified under Red Hat Linux. Of course, this doesn't mean that Red Hat Linux can access data on every type of partition, but in many cases it is possible to access data on a partition dedicated to another operating system.

Keep this information in mind; it will make things easier to understand when you're setting up the partitions Red Hat Linux requires.

B.1.6 Disk Partitions and Other Operating Systems

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

Intel

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

B.1.7 Disk Partitions and Mount Points

One area that many people new to Linux find confusing is the matter of how partitions are used and accessed by the Linux operating system. In DOS/Windows, it is relatively simple: If you have more than one partition, each partition gets a "drive letter." You then use the drive letter to refer to files and directories on a given partition.

This is entirely different from how Red Hat Linux deals with partitions and, for that matter, with disk storage in general. The main difference is that each partition is used to form part of the storage necessary to support a single set of files and directories. This is done by associating a partition with a directory through a process known as **mounting**. Mounting a partition makes its storage available starting at the specified directory (known as a **mount point**).

For example, if partition `/dev/hda5` were mounted on `/usr`, that would mean that all files and directories under `/usr` would physically reside on `/dev/hda5`. So the file `/usr/share/doc/FAQ/txt/Linux-FAQ` would be stored on `/dev/hda5`, while the file `/etc/X11/gdm/Sessions/Gnome` would not.

Continuing our example, it is also possible that one or more directories below `/usr` would be mount points for other partitions. For instance, a partition (say, `/dev/hda7`) could be mounted on `/usr/local`, meaning that, for example, `/usr/local/man/whatis` would then reside on `/dev/hda7` rather than `/dev/hda5`.

B.1.8 How Many Partitions?

At this point in the process of preparing to install Red Hat Linux, you will need to give some consideration to the number and size of the partitions to be used by your new operating system. The question of "how many partitions" continues to spark debate within the Linux community and, without any end to the debate in sight, it’s safe to say that there are probably as many partition layouts as there are people debating the issue.

Keeping this in mind, we recommend that, unless you have a reason for doing otherwise, you should create the following partitions:

- A swap partition — Swap partitions are used to support virtual memory. In other words, data is written to swap when there is not RAM to hold the data your system is processing. If your computer has 16 MB of RAM or less, you *must* create a swap partition. Even if you have more memory, a swap partition is still recommended. The minimum size of your swap partition should be equal to your computer's RAM, or 16 MB (whichever is larger).
- A `/boot` partition — The partition mounted on `/boot` contains the operating system kernel (which allows your system to boot Red Hat Linux), along with a few other files used during the bootstrap process. Due to the limitations of most PC BIOSes, creating a small partition to hold these files is a good idea. This partition should be no larger than 16MB.

Please Note

Make sure you read Section B.1.9, *One Last Wrinkle: Using LILO* — the information there applies to the `/boot` partition!

- A root partition (/) — The root partition is where `/` (the root directory) resides. In this partitioning layout, all files (except those stored in `/boot`) reside on the root partition. Because of this, it's in your best interest to maximize the size of your root partition. A 900MB root partition will permit the equivalent of a workstation-class installation (with *very* little free space), while a 1.7GB root partition will let you install every package.

B.1.9 One Last Wrinkle: Using LILO

LILO (the LInux LOader) is the most commonly used method to boot Red Hat Linux on Intel-based systems. An operating system loader, LILO operates "outside" of any operating system, using only the Basic I/O System (or BIOS) built into the computer hardware itself. This section describes LILO's interactions with PC BIOSes, and is specific to Intel-compatible computers.

BIOS-Related Limitations Impacting LILO

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

All the data LILO needs to access at boot time (including the Linux kernel) are located in the `/boot` directory. If you follow the partition layout recommended above, or you are performing a workstation- or server-class install, the `/boot` directory will be in a small, separate partition. Otherwise, it will reside in the root partition. In either case, the partition in which `/boot` resides must conform to the following guidelines if you are going to use LILO to boot your Red Hat Linux system:

On First Two IDE Drives

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

On First IDE or First SCSI Drive

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

On First Two SCSI Drives

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

Partition *Completely* Below Cylinder 1023

No matter which of the above configurations apply, the partition that holds `/boot` must be located entirely below cylinder 1023. If the partition holding `/boot` straddles cylinder 1023, you may face a situation where LILO will work initially (because all the necessary information is below cylinder 1023),

but will fail if a new kernel is to be loaded, and that kernel resides above cylinder 1023.

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

Please Note

Disk Druid as well as the workstation- and server-class installs take these BIOS-related limitations into account.

C Driver Disks

C.1 Why Do I Need a Driver Disk?

While the Red Hat Linux installation program is loading, you may see a screen that asks you for a driver disk. The driver disk screen is most often seen in three scenarios:

- if you are running the installation program in `expert mode`
- if you run the installation program by entering `linux dd` at the `boot:` prompt
- if you run the installation program on a computer which does not have any PCI devices

C.1.1 So What Is a Driver Disk Anyway?

A driver disk adds support for hardware that is not otherwise supported by the installation program. The driver disk could be produced by Red Hat, it could be a disk you make yourself, or it could be a disk that a hardware vendor includes with a piece of hardware.

There is really no need to use a driver disk unless you need a particular device in order to install Red Hat Linux. You will most likely use a driver disk for SCSI adapters and NICs, as those are really the only devices which are used during the installation that might require driver disk support. If an unsupported device is not needed to install Red Hat Linux on your system, continue with a regular installation and then add support for the new piece of hardware once the installation is complete.

C.1.2 How Do I Obtain a Driver Disk?

Your best option for finding driver disk information is on Red Hat's website at http://www.redhat.com/support/errata/ under the section called **Bug Fixes**.

If you find a driver disk that is appropriate for your device support needs, create a boot disk using that *`filename`*`.img` file. For instructions on how to make a boot disk, see the *Official Red Hat Linux Installation Guide* section called *Making Installation Diskettes*.

Once you have created your driver disk, boot your system using the diskette as a boot disk and enter either `linux expert` or `linux dd` at the `boot:` prompt.

Please Note

Some driver disk images may also be found in the images directory of your Red Hat Linux CD.

D How to Create a Dual-Boot System

This document explains your options for installing Red Hat Linux on a computer that currently runs another operating system, and how to create a dual-boot environment.

D.1 If Your Computer Already Has An Operating System ...

If the computer you want to install Red Hat Linux on is currently running Windows (or some other operating system), you have an important decision to make. Your choices are:

- Do you want to install Red Hat Linux but feel timid about disk partitioning? You can install Red Hat Linux on your system without creating any Linux partitions by performing a **partitionless** installation. The installation program can install Red Hat Linux on an existing, formatted Windows partition and you'll only need to create a boot disk during the installation to access Red Hat Linux on your system.

  This method is perfect for those who do not want to install Red Hat Linux as the primary OS or as a dual-boot OS on your system. It is a great way of trying out Red Hat Linux without creating Linux partitions on your system.

  If this is your preferred choice, refer to *Official Red Hat Linux Installation Guide*, *Appendix B. Installing Without Partitioning*, for those instructions.

- Do you want to install Red Hat Linux and have the option of booting either Red Hat Linux or your other operating system? Performing a workstation- or custom-class installation will allow Red Hat Linux to install on your system without affecting the other operating system. A workstation-class installation will do this by default. In a custom-class installation, you must install LILO (the LInux LOader) on the MBR (Master Boot Record) -- unless Windows NT is your other OS or if you are already using another boot loader on your system. The *Official Red Hat Linux Installation Guide* has instructions on installing and configuring LILO. After the installation, whenever you start the computer, you will indicate whether you want to start Red Hat Linux or the other operating system.

WARNING

Do not perform a workstation-class installation if you're sharing a disk with Windows NT; if you do, you will be unable to boot Windows NT. LILO will overwrite NT's boot loader and you will be unable to boot NT. You must perform a custom-class installation and configure LILO so that it is not installed on the MBR.

To create a dual-boot environment on a system that currently has NT, you must install LILO on the first sector of the root partition, not the MBR. Please be sure to create a boot disk. In a case such as this, you will either need to use the boot disk, or configure the NT system loader to boot LILO from the first sector of the root partition. Be sure to check out http://www.linux-doc.org/HOWTO/mini/Linux+NT-Loader.html for more information on setting up LILO and NT.

If this is your preferred choice, read Section D.2, *Setting Up a Dual-Boot Environment*

- Do you want Red Hat Linux to be the only operating system on your computer? The Red Hat Linux installation program will remove the other operating system and also any data you have created using that operating system.

Please Note

In order to install Red Hat Linux and keep another OS on your system, there must be sufficient space for Red Hat Linux to be installed on. Otherwise, Red Hat Linux will replace the current OS and files on your system. If you have not partitioned your hard drive to make room for Red Hat Linux or made sure that there is sufficient unpartitioned space available for your installation, Red Hat Linux will install over the existing information by default. It will also happen if you select a server-class installation (note that a server-class installation does not install the X Window System so no GUI environment will be present [1]). So a dual-boot environment is incompatible with this choice.

If this is your preferred choice, first back-up any information on your computer that you want to save or perform a full backup if you think you may want to restore your system to its original configuration), then proceed with installation as explained in the *Official Red Hat Linux Installation Guide*.

D.2 Setting Up a Dual-Boot Environment

Sharing a computer between two operating systems requires dual booting. You can use either operating system on the computer but not both at once. Each operating system boots from and uses its own hard drives or disk partitions (a partition is a physical division on a hard drive).

For clarity, we will assume that the other operating system is Windows. But the general procedures are similar for other operating systems.

[1] A server-class installation is most appropriate for you if you'd like your system to function as a Linux-based server, and you don't want to heavily customize your system configuration.

Please Note

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

Before starting the installation program, you must first make room for Red Hat Linux. Your choices are:

- Add a new hard drive
- Use an existing hard drive or partition
- Create a new partition

D.2.1 Add a New Hard Drive

The simplest way to make room for Red Hat Linux is to add a new hard drive to the computer and then install Red Hat Linux on that drive. For example, if you add a second IDE hard drive to the computer, the Red Hat Linux installation program will recognize it as hdb and the existing drive (the one used by Windows) as hda. (With SCSI hard drives, the newly installed hard drive would be recognized as sdb and the other hard drive as sda.)

If you choose to install a new hard drive for Linux, you don't need to read any further. After starting the Red Hat Linux installation program, just make sure you tell it to install Linux on the newly installed hard drive (hdb, sdb) rather than the one Windows uses.

D.2.2 Use an Existing Hard Drive or Partition

The next simplest way to make room for Linux is to use a hard drive or disk partition that is currently being used by Windows. For example, suppose that Windows

Explorer shows two hard drives, `C:` and `D:`. This could indicate either that the computer has two hard drives, or a single hard drive with two partitions. In either case (assuming it is large enough), you can install Red Hat Linux on the hard drive or disk partition that Windows recognizes as `D:`.

This choice is available to you only if the computer has two or more hard drives or disk partitions.

Please Note

Windows uses letters to refer to removable drives (for example, a ZIP drive) and network storage (virtual drives) as well as for local hard drive space: you cannot install Linux on a removable or network drive.

If a local Windows partition is available that you want to install Linux in, you don't need to read any further. Just do the following:

1. Copy all data you want to save from the selected hard drive or partition (`D:` in this example) to another location.
2. Start the Red Hat Linux installation program and tell it to install Linux in the designated drive or partition -- in this example, in the hard drive or partition that Windows designates as `D:`. Note that Linux distinguishes between hard drives and disk partitions. Thus:
   - If `C:` and `D:` on this computer refer to two separate hard drives, the installation program will recognize them as `hda` and `hdb` (IDE) or `sda` and `sdb` (SCSI). Tell the installation program to use `hdb` or `sdb`.
   - If `C:` and `D:` refer to partitions on a single drive, the installation program will recognize them as `hda1` and `hda2` (or `sda1` and `sda2`). During the partitioning phase of Linux installation, you'll delete the second partition (`hda2` or `sda2`), then partition the unallocated free space for Linux. (You don't have to delete the second partition prior to beginning Linux partitioning. But if you don't, Windows will complain whenever you boot that it cannot read `Drive`

D; and should someone accidentally format D, your Linux system would be destroyed.)

D.2.3 Create a New Partition

The third way to make room for Linux is to create a new partition for Red Hat Linux on the hard drive being used by the other operating system. If Windows Explorer shows only one hard drive (C:), and you don't want to add a new hard drive, you must partition the drive. After partitioning, Windows Explorer will reveal a smaller C: drive; and, when you run the Red Hat Linux installation program, it will partition the remainder of the drive for Linux.

You can use a destructive partitioning program, such as fdisk, to divide the hard drive, but doing so will require you to re-install Windows. (This is probably not your best option.)

A number of non-destructive third-party partitioning programs are available for the Windows operating system. If you choose to use one of these, consult their documentation.

For instructions on how to partition with FIPS, a program that is on the Red Hat Linux CD-ROM, turn to Section D.3, *Partitioning with FIPS*.

D.3 Partitioning with FIPS

As a convenience to our customers, we provide the FIPS utility. This is a freely available program that can resize FAT (File Allocation Table) partitions. It's included on the Red Hat Linux CD-ROM in the dosutils directory.

Please Note

Many people have successfully used FIPS to repartition their hard drives. However, because of the nature of the operations carried out by FIPS, and the wide variety of hardware and software configurations under which it must run, Red Hat cannot guarantee that FIPS will work properly on your system. Therefore, no installation support whatsoever is available for FIPS; use it at your own risk.

That said, if you decide to repartition your hard drive with FIPS, it is vital that you do two things:

- `Perform a Backup` -- Make two copies of all the important data on your computer. These copies should be to removable media (such as tape or diskettes), and you should make sure they are readable before proceeding.
- `Read the Documentation` -- Completely read the FIPS documentation, located in the FIPS directory on Red Hat Linux CD-ROM.

Should use decide to use FIPS, be aware that after FIPS runs you will be left with two partitions: the one you resized, and the one FIPS created out of the newly freed space. If your goal is to use that space to install Red Hat Linux, you should delete the newly created partition, either by using fdisk under your current operating system, or while setting up partitions during a custom-class installation.

The following instructions are a simplified version of the FIPS documentation file, `fips.doc`, located in the FIPS directory (`/dosutils/fips20/*`). These instructions should apply in most instances. If you encounter any problems, see the documentation file.

1. From Windows:
   - Do a full backup.
   - Run scandisk to verify that the hard drive contains no bad clusters.

- Decide how to distribute the available space on the hard drive between the operating systems. Use Windows Explorer to see the free space on the drive. Make a note of the space (in megabytes) that each operating system will have.
- If you don't have one, create a DOS boot disk.

  To create a DOS boot disk, first boot your machine to DOS.

  Next, insert a blank, formatted diskette into the floppy drive.

  Type the following at the command prompt and press [Enter]:

  ```
  FORMAT A: /S
  ```

  If you're using Windows 95, first insert a blank formatted diskette into the floppy drive. Next, go to **Start**/**Run**, and type:

  ```
  FORMAT A: /S
  ```

  The diskette will be formatted, and `COMMAND.COM`, along with the associated hidden files (`IO.SYS`, `MSDOS.SYS`, and `BDLSAPCE.BIN`), will be copied to the diskette.
- Copy the following files on the Red Hat Linux CD-ROM to the DOS boot disk.

  ```
  /mnt/cdrom/dosutils/fips20/fips.exe
  /mnt/cdrom/dosutils/fips20/restorrb.exe
  /mnt/cdrom/dosutils/fips20/errors.txt
  /mnt/cdrom/dosutils/fips20/fips.doc
  /mnt/cdrom/dosutils/fips20/fips.faq
  ```

- Defragment the hard drive.

2. Insert the DOS boot disk into the floppy drive and reboot the system.
3. Start FIPS (type `fips` at the prompt).

When FIPS begins, you'll find a welcome screen similar to the following:

Figure D–1 FIPS Welcome Screen

```
FIPS version 2.0, Copyright (C) 1993/4 Arno Schaefer
FAT32 Support, Copyright (C) 1997 Gordon Chaffee

DO NOT use FIPS in a multitasking environment like Windows, OS/2, Desqview,
Novell Task manager or the Linux DOS emulator; boot from a DOS boot disk first.

If you use OS/2 or a disk compressor, read the relevant sections in FIPS.DOC.

FIPS comes with ABSOLUTELY NO WARRANTY, see file COPYING for details.

This is free software, and you are welcome to redistribute it
under certain conditions; again, see file COPYING for details.

Press any key.
```

When you press a key, a root partition screen similar to the following appears. (Note that, if the computer has more than one hard drive, you'll be asked to select which one you want to partition.)

Figure D–2 FIPS Root Partition Screen

```
Partition table:

     |        |     Start      |      |      End       | Start  |Number of|
Part.|bootable|Head Cyl. Sector|System|Head Cyl. Sector| Sector |Sectors  |  MB
-----+--------+----------------+------+----------------+--------+---------+----
1    |    yes |   0  148      1|   83h|  15  295     63|  149184|   149184|  72
2    |     no |   1    0      1|   06h|  15  139     63|      63|   141057|  68
3    |     no |   0  140      1|   06h|  15  147     63|  141120|     8064|   3
4    |     no |   0    0      0|   00h|   0    0      0|       0|        0|   0

Checking root sector ... OK

Press any key.
```

When you press a key, details about the hard drive, such as the following, will appear.

Figure D–3 FIPS Boot Sector Screen

```
Boot sector:
Bytes per sector: 512
Sectors per cluster: 8
Reserved sectors: 1
Number of FATs: 2
Number of rootdirectory entries: 512
Number of sectors (short): 0
Media descriptor byte: f8h
Sectors per FAT: 145
Sectors per track: 63
Drive heads: 16
Hidden sectors: 63
Number of sectors (long): 141057
Physical drive number: 80h
Signature: 29h

Checking boot sector ... OK
Checking FAT ... OK
Searching for free space ... OK

Do you want to make a backup copy of your root and boot sector before
proceeding? (y/n)
```

You should select **y**, for yes, to make a backup copy of your root and boot sector before proceeding with FIPS.

Next, you'll be presented with the following message:

```
Do you have a bootable floppy disk in drive A: as described in the
documentation? (y/n)
```

Verify that a DOS boot disk is in the floppy drive, and type **y**, for yes. A screen similar to the following will appear, allowing you to resize the partition.

Figure D–4 Partition Resizing Screen

```
Writing file a:\rootboot:000

Enter start cylinder for new partition (33-526)

Use the cursor keys to choose the cylinder, <enter> to continue

Old partition          Cylinder             New partition
 258.9 MB                 33                   3835.8 MB
```

The initial values allocate *all* free space on the disk to the new partition. This is not what you want, because this setting would leave no free space on your Windows partition. Press the [right arrow] to increase the size of the Windows partition and decrease the size of the new (Linux) partition; press the [left arrow] to decrease the size of the Windows partition and increase the size of the Linux partition. When the sizes are what you want, press [Enter]. A verification screen similar to the following appears:

Figure D–5 FIPS Verification Screen

```
First Cluster: 17442
Last Cluster: 65511

Testing if empty ... OK

New partition table:

     |        |     Start      |      |      End       | Start  |Number of|
Part.|bootable|Head Cyl. Sector|System|Head Cyl. Sector| Sector |Sectors  |  MB
-----+--------+----------------+------+----------------+--------+---------+----
1    |    yes |   0  148      1|   83h|  15  295     63|  149184|   149184|  1090
2    |     no |   0  139      1|   06h| 254  521     63| 2233035|  6152995|  3084
3    |     no |   0  140      1|   06h|  15  147     63|  141120|     8064|   3
4    |     no |   0    0      0|   00h|   0    0      0|       0|        0|   0

Checking root sector ... OK

Do you want to continue or reedit the partition table (c/r)?
```

If you answer **r** (to re-edit the partition tables), *Figure 4* reappears, allowing you to change the partition sizes. If you answer **c**, a confirmation screen (*Figure 6*) appears:

Figure D–6 FIPS Confirmation Screen

```
New boot sector:

Boot sector:
Bytes per sector: 512
```

```
Sectors per cluster: 8
Reserved sectors: 1
Number of FATs: 2
Number of rootdirectory entries: 512
Number of sectors (short): 0
Media descriptor byte: f8h
Sectors per FAT: 145
Sectors per track: 63
Drive heads: 16
Hidden sectors: 63
Number of sectors (long): 141057
Physical drive number: 80h
Signature: 29h

Checking boot sector ... OK

Ready to write new partition scheme to disk
Do you want to proceed (y/n)?
```

Answering **y** completes the resizing operation. A harmless error message may occur, stating, in effect that FIPS cannot reboot the system.

After a successful operation, the disk will have two partitions. The first partition (`hda1` or `sda1`) will be used by Windows. We recommend that you start Windows (remember to remove the boot disk from drive `A:`) and run scandisk on drive `C:`.

If you encounter any problems, (for example, Windows will not boot), you can reverse the FIPS resizing operation with the restorrb.exe command, which you copied to your DOS boot disk. In case of any errors, read the FIPS documentation files (`fips.doc` and `fips.faq`), which indicate a number of factors that could cause the resizing operation to fail. If all else fails, you can restore Windows with the backup you made.

The second partition (`hda2` or `sda2`) contains the space that the Red Hat Linux installation program will use. When the Disk Druid screen appears during installation, delete this partition (the installation manual explains how), then proceed with Linux partitioning.

E RAID (Redundant Array of Independent Disks)

E.1 What is RAID?

The basic idea behind RAID is to combine multiple small, inexpensive disk drives into an array which yields performance exceeding that of one large and expensive drive. This array of drives will appear to the computer as a single logical storage unit or drive.

RAID is a method in which information is spread across several disks, using techniques such as **disk striping** (RAID Level 0) and **disk mirroring** (RAID level 1) to achieve redundancy, lower latency and/or higher bandwidth for reading and/or writing to disks, and maximize recoverability from hard-disk crashes.

The underlying concept in RAID is that data may be distributed across each drive in the array in a consistent manner. To do this, the data much first be broken into consistently-sized "chunks" (often 32K or 64K in size, although different sizes can be used). Each chunk is then written to each drive in turn. When the data is to be read, the process is reversed, giving the illusion that multiple drives are actually one large drive.

E.1.1 Who Should Use RAID?

Those of you who need to keep large quantities of data on hand (such as an average administrator) would benefit by using RAID technology. Primary reasons to use RAID include:

- enhanced speed
- increased storage capacity (and more economical)
- greater efficiency in recovering from a disk failure

E.1.2 RAID: Hardware vs. Software

There are two possible approaches to RAID: Hardware RAID and Software RAID.

Hardware RAID

The hardware-based system manages the RAID subsystem independently from the host and presents to the host only a single disk per RAID array.

An example of a hardware RAID device would be one that connects to a SCSI controller and presents the RAID arrays as a single SCSI drive. An external RAID system moves all RAID handling "intelligence" into a controller located in the external disk subsystem. The whole subsystem is connected to the host via a normal SCSI controller and appears to the host as a single disk.

RAID controllers also come in the form of cards that *act* like a SCSI controller to the operating system, but handle all of the actual drive communications themselves. In these cases, you plug the drives into the RAID controller just like you would a SCSI controller, but then you add them to the RAID controller's configuration, and the operating system never knows the difference.

Software RAID

Software RAID implements the various RAID levels in the kernel disk (block device) code. It also offers the cheapest possible solution: Expensive disk controller cards or hot-swap chassis [1]are not required, and software RAID works with cheaper IDE disks as well as SCSI disks. With today's fast CPUs, software RAID performance can excel against hardware RAID.

The MD driver in the Linux kernel is an example of a RAID solution that is completely hardware independent. The performance of a software-based array is dependent on the server CPU performance and load.

E.1.3 Some Features of RAID

For those interested in learning more about what software RAID has to offer, here is a brief list of few of those features:

- Threaded rebuild process
- Fully kernel-based configuration

[1] A hot-swap chassis allow you to remove a hard drive without having to power-down your system.

- Portability of arrays between Linux machines without reconstruction
- Backgrounded array reconstruction using idle system resources
- Hot-swappable drive support
- Automatic CPU detection to take advantage of certain CPU optimizations

Levels and linear support

RAID also offers levels 0, 1, 4, 5, and linear support. These RAID types act as follows:

- *Level 0* -- RAID level 0, often called "striping," is a performance- oriented striped data mapping technique. That means the data being written to the array is broken down into strips and written across the member disks of the array. This allows high I/O performance at low inherent cost but provides no redundancy. Storage capacity of the array is equal to the total capacity of the member disks.
- *Level 1* -- RAID level 1, or "mirroring," has been used longer than any other form of RAID. Level 1 provides redundancy by writing identical data to each member disk of the array, leaving a "mirrored" copy on each disk. Mirroring remains popular due to its simplicity and high level of data availability. Level 1 operates with two or more disks that may use parallel access for high data-transfer rates when reading, but more commonly operate independently to provide high I/O transaction rates. Level 1 provides very good data reliability and improves performance for read-intensive applications but at a relatively high cost[2]. Array capacity is equal to the capacity of one member disk if you use identical disk drives.
- *Level 4* -- Level 4 uses parity[3]concentrated on a single disk drive to protect data. It's better suited to transaction I/O rather than large file transfers. Because the

[2] RAID level 1 is at a high cost because you write the same information to all of the disks in the array, which wastes drive space. For example, you have RAID level 1 set up so that your "/" (root) partition spans across two 4G drives. You have 8G total but are only able to access 4G of that 8G. The other 4G acts like a mirror of the first 4G.

[3] Parity information is calculated based on the contents of the rest of the member disks in the array. This information can then be used to reconstruct data when a disk in the array fails. The reconstructed data can then be used to satisfy I/O requests to the failed disk, and to repopulate the failed disk after it has been repaired or replaced.

dedicated parity disk represents an inherent bottleneck, level 4 is seldom used without accompanying technologies such as write-back caching. Although RAID level 4 is an option in some RAID partitioning schemes, it is not an option allowed in Red Hat Linux RAID installations[4]. Array capacity is equal to the capacity of member disks, minus capacity of one member disk if you use identical disk drives.

- *Level 5* -- The most common type of RAID. By distributing parity across some or all of an array's member disk drives, RAID level 5 eliminates the write bottleneck inherent in level 4. The only bottleneck is the parity calculation process. With modern CPUs and software RAID, that isn't a very big bottleneck. As with level 4, the result is asymmetrical performance, with reads substantially outperforming writes. Level 5 is often used with write-back caching to reduce the asymmetry. Array capacity is equal to the capacity of member disks, minus capacity of one member disk if you use identical disk drives.
- *Linear RAID* -- Linear RAID is a simple grouping of drives to create a larger virtual drive. In linear RAID, the chunks are allocated sequentially from one member drive, going to the next drive only when the first is completely filled. This grouping provides no performance benefit, as it is unlikely that any I/O operations will be split between member drives. Linear RAID also offers no redundancy, and in fact decreases reliability -- if any one member drive fails, the entire array cannot be used. The capacity is total of all member disks.

E.1.4 Creating RAID Partitions

RAID is available in both the GUI and kickstart installation modes. You can use fdisk or Disk Druid to create your RAID configuration, but these instructions will focus mainly on using Disk Druid to complete this task.

Before you can create a RAID device, you must first create RAID partitions, using the following step-by-step instructions.

[4] RAID level 4 takes up the same amount of space as RAID level 5, but level 5 has many advantages which is why level 4 is not supported.

Tip: If You Use fdisk

If you are using fdisk to create a RAID partition, bear in mind that instead of creating a partition as type 83, which is Linux native, you must create the partition as type fd (Linux RAID) and that partitions within a given RAID array should span identical cylinders on drives for best performance.

- Create a partition. In Disk Druid, choose **Add** to create a new partition (see Figure E–1, *Creating a New RAID Partition*).

Figure E–1 Creating a New RAID Partition

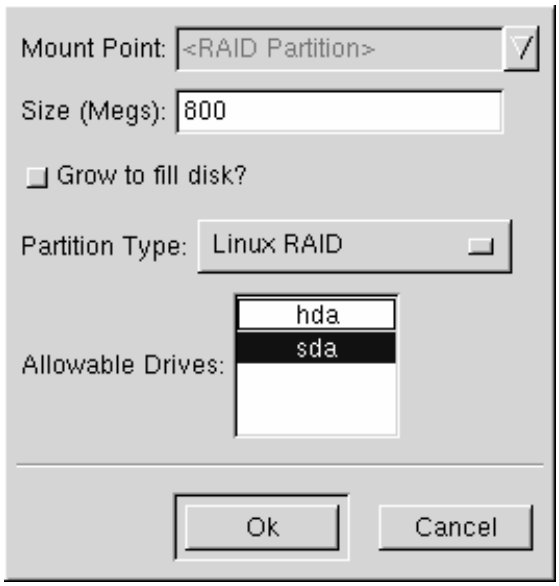

- You will not be able to enter a mount point (you will be able to do that once you've created your RAID device).

- Enter the size that you want the partition to be.
- Select **Grow to fill disk** if you want the partition to grow to fill all available space on the hard disk. In this case, the partition's size will expand and contract as other partitions are modified. If you make more than one partition grow-able, the partitions will compete for the available free space on the disk.
- Enter the partition type as RAID.
- Finally, for **Allowable Drives**, select the drive that RAID will be created on. If you have multiple drives, all drives will be selected here and you must deselect those drives which will *not* have RAID array on it.

Continue these steps to create as many partitions as needed for your RAID setup.

Figure E–2 RAID Partitions

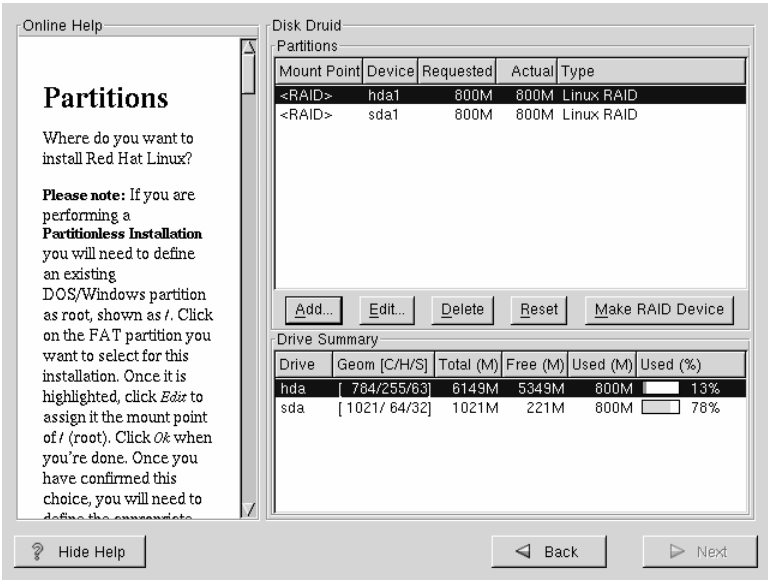

Once you have all of your partitions created as RAID partitions, select the **Make RAID Device** button on the Disk Druid main partitioning screen (see Figure E–2, *RAID Partitions*).

Next, Figure E–3, *Making a RAID Device* will appear which will allow you to make a RAID device.

Figure E–3 Making a RAID Device

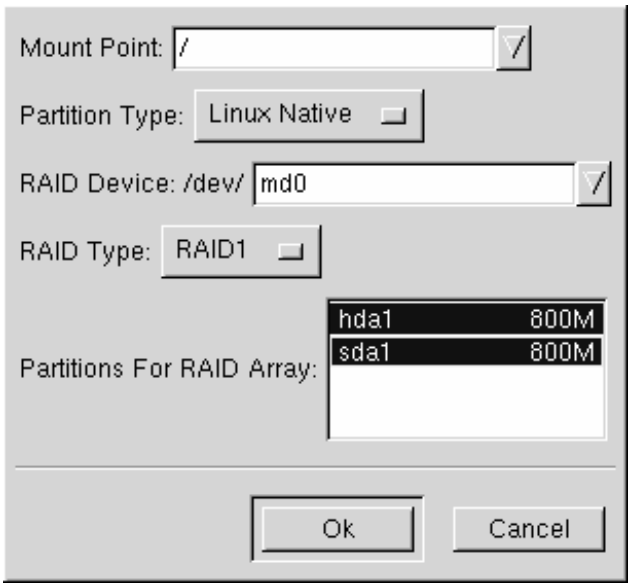

- First, enter a mount point.
- Next, make sure the partition type is set as **Linux Native** (which will be the default).
- Choose your RAID device. You should choose `md0` for your first device, `md1` for your second device, and so on, unless you have a specific reason to make it something else. Raid devices range from md0 to md7, and each may only be used once.
- Choose your RAID type. You can choose from **RAID 0**, **RAID 1**, and **RAID 5**.

Please Note

If you are making a RAID partition of `/boot`, you must choose RAID level 1 and it must use one of the first two drives (IDE first, SCSI second). If you are not creating a RAID partition of `/boot`, and are making a RAID partition of `/`, it must be RAID level 1 and it must use one of the first two drives (IDE first, SCSI second).

- Finally, select which partitions will go into this RAID array (as in Figure E–4, *Creating a RAID Array*) and then click **Next**.

Figure E–4 Creating a RAID Array

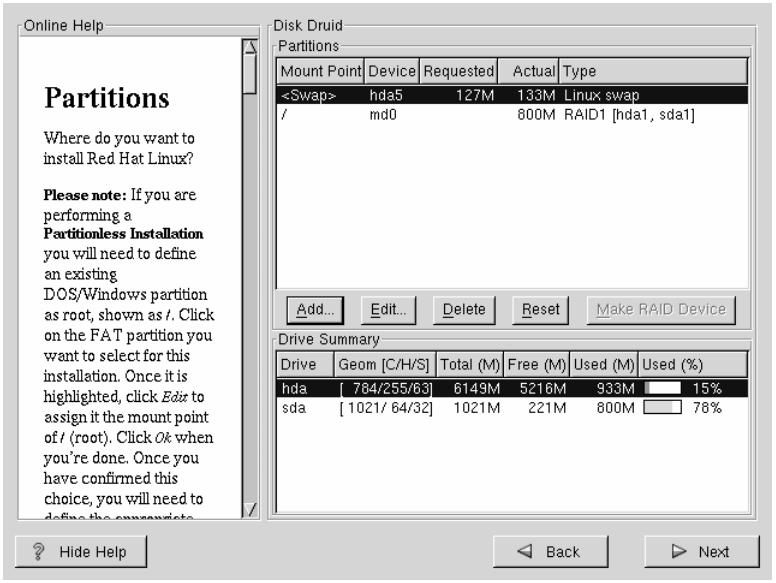

- From here, you can continue with your installation process. Refer back to the *Official Red Hat Linux Installation Guide* for further instructions.

F Kickstart Installations

F.1 What are Kickstart Installations

Due to the need for automated installation, Red Hat has created the kickstart installation method. With this method, a system administrator can create a single file containing the answers to all the questions that would normally be asked during a typical Red Hat Linux installation.

Kickstart files can be kept on single server system, and read by individual computers during the installation. The kickstart installation method is powerful enough that often a single kickstart file can be used to install Red Hat Linux on multiple machines, making it ideal for network and system administrators.

Kickstart lets you automate most of a Red Hat Linux installation, including:

- Language selection
- Network configuration
- Keyboard selection
- Boot loader installation (LILO)
- Disk partitioning
- Mouse Selection
- X Window System configuration

F.2 How Do You Perform a Kickstart Installation?

Kickstart installations can be performed using a local CD-ROM, a local hard drive, NFS, FTP or HTTP installation methods.

To use kickstart mode, you must first create a kickstart file (`ks.cfg`), and make it available to the Red Hat Linux installation program.

F.2.1 Where to Put A Kickstart File

A kickstart file must be placed in one of two locations:

- on a boot disk
- on a network

Normally a kickstart file is copied to the boot disk, or made available on the network. The network-based approach is most commonly used, as most kickstart installations tend to be performed on networked computers.

Let's take a more in-depth look at where the kickstart file may be placed.

To perform a diskette-based kickstart installation, the kickstart file must be named `ks.cfg`, and reside in the boot disk's top-level directory. Note that the Red Hat Linux boot disks are in MS-DOS format, making it easy to copy the kickstart file under Linux using the `mcopy` command:

```
mcopy ks.cfg a:
```

(or, you can also use Windows). You can also mount the MS-DOS boot disk and `cp` the file over. Although there's no technological requirement for it, most diskette-based kickstart installations install Red Hat Linux from a local CD-ROM.

Network installations using kickstart are quite common and are recommended over floppy, because system administrators can easily automate the installation on many networked computers quickly and painlessly. In general, the approach most commonly used is for the administrator to have both a BOOTP/DHCP server and an NFS server on the local network. The BOOTP/DHCP server is used to give the client system its networking information, while the actual files used during the installation are served by the NFS server. Often, these two servers run on the same physical machine, but there is no requirement for this.

To do a network-based kickstart installation, you must have a BOOTP/DHCP server on your network, and it must include configuration information for the machine on which you are attempting to install Red Hat Linux. The BOOTP/DHCP server will be used to give the client its networking information as well as the location of the kickstart file.

If a kickstart file is specified by the BOOTP/DHCP server, the client system will attempt an NFS mount of the file's path, and will copy the specified file to the client, using it as the kickstart file. The exact settings required vary depending on the BOOTP/DHCP server you use.

Here's an example of a line from the `dhcpd.conf` file for the DHCP server shipped with Red Hat Linux:

```
filename "/usr/new-machine/kickstart/";
next-server blarg.redhat.com;
```

Note that you should use `filename` for the kickstart file's name (or the directory in which the kickstart file resides) and `next-server` to set the NFS server name.

If the filename returned by the BOOTP/DHCP server ends with a slash ("/"), then it is interpreted as a path only. In this case, the client system mounts that path using NFS, and searches for a specially-named file. The filename the client searches for is:

```
<ip-addr>-kickstart
```

The *<ip-addr>* section of the filename should be replaced with the client's IP address in dotted decimal notation. For example, the filename for a computer with an IP address of 10.10.0.1 would be `10.10.0.1-kickstart`.

Note that if you don't specify a server name, then the client system will attempt to use the server that answered the BOOTP/DHCP request as its NFS server. If you don't specify a path or filename, the client system will try to mount `/kickstart` from the BOOTP/DHCP server, and will try to find the kickstart file using the same *<ip-addr>*`-kickstart` filename as described above.

F.3 Starting a Kickstart Installation

To begin a kickstart installation, you must boot the system from a Red Hat Linux boot disk, and enter a special boot command at the boot prompt. If the kickstart file resides on the boot disk, the proper boot command would be:

```
boot: linux ks=floppy
```

If, on the other hand, the kickstart file resides on a server, the appropriate boot command would be:

```
boot: linux ks
```

Anaconda looks for a kickstart file if the `ks` command line argument is passed to the kernel. It can take a number of forms:

`ks=floppy`

The installation program looks for the file `ks.cfg` on a VFAT filesystem on the floppy in drive /dev/fd0.

`ks=hd:`*`<device>`*`/`*`<file>`*

The installation program will mount the filesystem on <*device*> (which must be VFAT or ext2), and look for the kickstart configuration file as <*file*> in that filesystem (for example, `ks=hd:sda3/mydir/ks.cfg`).

`ks=file:/`*`<file>`*

The installation program will try to read the file <*file*> from the filesystem; no mounts will be done. This is normally used if the kickstart file is already on the `initrd` image.

`ks=nfs:`*`<server:>`*`/`*`<path>`*

The installation program will look for the kickstart file on the NFS server <*server*>, as file <*path*>. The installation program will use DHCP to configure the Ethernet card.

`ks=cdrom:/`*`<path>`*

The installation program will look for the kickstart file on CD-ROM, as file <*path*>.

`ks`

If `ks` is used alone, the installation program will configure the Ethernet card in the system using DHCP. The system will use the "bootServer" from the DHCP response as an NFS server to read the kickstart file from (by default, this is

the same as the DHCP server). The name of the kickstart file is one of the following:

- If DHCP is specified and the "bootfile" begins with a `/`, that file is looked for on the NFS server.
- If DHCP is specified and the "bootfile" begins with something other then a `/`, that file is looked for in the `/kickstart` directory on the NFS server.
- If DCHP did not specify a "bootfile," then the installation program tries to read the file `/kickstart/1.2.3.4-kickstart`, where *1.2.3.4* is the numeric IP address of the machine being installed.

F.4 The Kickstart File

Now that you have some background information on kickstart installations, let's take a look at the kickstart file itself. The kickstart file is a simple text file, containing a list of items, each identified by a keyword. You can create it by editing a copy of the `sample.ks` file found in the `/doc` directory of the Red Hat Linux CD-ROM, or you can create it from scratch. You should be able to edit it with any text editor or word processor that can save files as ASCII text.

First, some ground rules to keep in mind while creating your kickstart file:

- Items must be specified *in order*. That order is:

```
<command section>
<any combination of %pre, %post, %packages>
<installclass>
```

- Items that aren't required can be omitted.
- Omitting any required item will result in the installation program prompting the user for an answer to the related item, just as during a typical installation. Once the answer is given the installation will continue unattended (unless it comes across another missing item).
- Lines starting with a pound sign ("#") are treated as comments, and are ignored.

- For kickstart *upgrades*, the following items are required:
  - language
  - installation method
  - device specification (if device is needed to perform installation)
  - keyboard setup
  - the `upgrade` keyword
  - LILO configuration

  If any other items are specified for an upgrade, those items will be ignored (note that this includes package selection).
- Kickstart files are split into three sections: commands, package list, and scripts. The file must be of the form:
  - *<kickstart commands>*
  - %packages
  - *<package list>*
  - %post
  - *<post script>*

The order matters; it can't be random. The post section goes to the end of the file and ends the file, no marker is necessary to end the file other than the post section itself.

F.5 Kickstart Commands

The following commands can be placed in a kickstart file.

F.5.1 `auth` – Authentication Options

`auth` (required)

Sets up the authentication options for the system. It's similar to the `authconfig` command that can be run after the install. By default, passwords are normally encrypted and are not shadowed.

`--enablemd5`

Use md5 encryption for user passwords.

`--enablenis`

Turns on NIS support. By default, `--enablenis` uses whatever domain it finds on the network. A domain should almost always be set by hand (via `--nisdomain`).

`--nisdomain`

NIS domain name to use for NIS services.

`--nisserver`

Server to use for NIS services (broadcasts by default).

`--useshadow`

Use shadow passwords.

`--enableldap`

Turns on LDAP support in `/etc/nsswitch.conf`, allowing your system to retrieve information about users (UIDs, home directories, shells, etc.) from an LDAP directory. Use of this option requires that the `nss_ldap` package be installed. You must also specify a server and a base DN.

`--enableldapauth`

Use LDAP as an authentication method. This enables the `pam_ldap` module for authentication and password-changing using an LDAP directory. Use of this option requires that the `nss_ldap` package be installed. You must also specify a server and a base DN.

`--ldapserver=`

The name of the LDAP server used if you use specified either --enableldap or --enableldapauth. This option is set in the `/etc/ldap.conf` file.

`--ldapbasedn=`

The DN (distinguished name) in your LDAP directory tree under which user information is stored. This option is set in the `/etc/ldap.conf` file.

`--enablekrb5`

Use Kerberos 5 for authenticating users. Kerberos itself has no notion of home directories, UIDs, or shells, so if you enable Kerberos you'll still need to enable LDAP, NIS, or Hesiod if you want to avoid having to use the `/usr/sbin/useradd` command to make their accounts known to this workstation. Use of this option requires the `pam_krb5` package to be installed.

`--krb5realm`

The Kerberos 5 realm your workstation belongs to.

`--krb5kdc`

The KDC (or KDCs) that serve requests for the realm. If you have multiple KDCs in your realm, separate their names with commas ([,]).

`--krb5adminserver`

The KDC in your realm that is also running kadmind. This server, which can only be run on the master KDC if you have more than one, handles password- changing and other administrative requests.

`--enablehesiod`

Enable Hesiod support for looking up user home directories, UIDs, and shells. More information on setting up and using Hesiod on your network is in `/usr/share/doc/glibc-2.x.x/README.hesiod`, which is included in the `glibc` package. Hesiod is an extension of DNS

that uses DNS records to store information about users, groups, and various other items.

`--hesiodlhs`

The Hesiod LHS ("left-hand side") option, set in `/etc/hesiod.conf`. This option is used by the Hesiod library to determine the name to search DNS for when looking up information, similar to LDAP's use of a base DN.

`--hesiodrhs`

The Hesiod RHS ("right-hand side") option, set in `/etc/hesiod.conf`. This option is used by the Hesiod library to determine the name to search DNS for when looking up information, similar to LDAP's use of a base DN.

Tip

To look up user information for "jim", the Hesiod library looks up *jim.passwd<LHS><RHS>*, which should resolve to a TXT record that looks like what his passwd entry would look like (`jim:*:501:501:Jungle Jim:/home/jim:/bin/bash`). For groups, the situation is identical, except *jim.group<LHS><RHS>* would be used.

Looking up users and groups by number is handled by making "501.uid" a CNAME for "jim.passwd", and "501.gid" a CNAME for "jim.group". Note that the LHS and RHS do not have periods [.] put in front of them when the library determines the name to search for, so the LHS and RHS most often begin with periods.

F.5.2 `clearpart` – Removing partitions based on partition type

`clearpart` (optional)

Removes partitions from the system, prior to creation of new partitions. By default, no partitions are removed.

`--linux`

Erases Linux (type 0x82, 0x83, and 0xfd [RAID]) partitions

`--all`

Erases all partitions from the system.

F.5.3 device `--opts`

`device` (optional)

On most PCI systems, the installation program will autoprobe for Ethernet and SCSI cards properly. On older systems, and some PCI systems, kickstart needs a hint to find the proper devices, however. The device command tells Anaconda to install extra modules. It is of the form:

```
device <type> <moduleName> --opts <options>
```

<*type*> should be one of "scsi" or "eth", and <*moduleName*> is the name of the kernel module which should be installed.

`--opts`

Options to pass to the kernel module. Note that multiple options may be passed if put in quotes. For example:

```
--opts "aic152x=0x340 io=11"
```

F.5.4 driver disk

`driverdisk` (optional)

During kickstart, driver disks can be used by copying the contents of a driver disk to the root directory of a partition on the system's hard drive and using the `driverdisk` command to tell the installation program where to look for it.

```
driverdisk <partition> [--type <fstype>]
```

<*partition*> is the partition containing the driver disk.

`--type`

Filesystem type (for example, VFAT or ext2).

F.5.5 `install`

`install` (optional)

Tells the system to install a fresh system rather than upgrade an existing system. This is the default mode.

F.5.6 Installation methods

You must use one of the these four commands to specify what type of kickstart is being done:

NFS

Install from the NFS server specified.

- --server <*server*>

  Server from which to install (hostname or IP).

- --dir <*dir*>

  Directory containing the Red Hat installation tree.

For example:

```
nfs --server <server> --dir <dir>
```

CD-ROM

Install from the first CD-ROM drive on the system.

For example:

```
cdrom
```

hard drive

Install from a Red Hat installation tree on a local drive, which must be either VFAT or ext2.

- --partition <*partition*>

  Partition to install from (such as, sdb2).

- --dir <*dir*>

  Directory containing the Red Hat installation tree.

For example:

```
harddrive --partition <partition> --dir <dir>
```

URL

Install from a Red Hat installation tree on a remote server via FTP or HTTP.

For example:

```
url --ulr http://<server>/<dir>
```

F.5.7 keyboard

keyboard (required)

Sets system keyboard type. Here's the list of available keyboards on i386 and Alpha machines:

```
azerty, be-latin1, be2-latin1, fr-latin0, fr-latin1, fr-pc, fr,
wangbe, ANSI-dvorak, dvorak-l, dvorak-r, dvorak, pc-dvorak-latin1,
tr_f-latin5, trf, bg, cf, cz-lat2-prog, cz-lat2, defkeymap,
defkeymap_V1.0, dk-latin1, dk. emacs, emacs2, es, fi-latin1, fi,
gr-pc, gr, hebrew, hu101, is-latin1, it-ibm, it, it2, jp106,
la-latin1, lt, lt.l4, nl, no-latin1, no, pc110, pl, pt-latin1,
pt-old, ro, ru-cp1251, ru-ms, ru-yawerty, ru, ru1, ru2, ru_win,
se-latin1, sk-prog-qwerty, sk-prog, sk-qwerty, tr_q-latin5, tralt,
trf, trq, ua, uk, us, croat, cz-us-qwertz, de-latin1-nodeadkeys,
de-latin1, de, fr_CH-latin1, fr_CH, hu, sg-latin1-lk450,
sg-latin1, sg, sk-prog-qwertz, sk-qwertz, slovene
```

Here's the list for SPARC machines:

```
sun-pl-altgraph, sun-pl, sundvorak, sunkeymap, sunt4-es,
sunt4-no-latin1, sunt5-cz-us, sunt5-de-latin1, sunt5-es,
sunt5-fi-latin1, sunt5-fr-latin1, sunt5-ru, sunt5-uk, sunt5-us-cz
```

F.5.8 language

lang (required)

Sets the default language for the installed system. The language you specify will be used during the installation as well as to configure any language-specific aspect of the installed system. For example, to set the language to English, the kickstart file should contain the following line:

```
lang en_US
```

Valid languages codes are:

```
cs_CZ, en_US, fr_FR, de_DE, hu_HU, is_IS, id_ID, it_IT,
ja_JP.ujis, no_NO, pl_PL, ro_RO, sk_SK, sl_SI, es_ES,
ru_RU.KOI8-R, uk_UA
```

F.5.9 lilo

lilo (required)

Specifies how the boot loader should be installed on the system. By default, LILO installs on the MBR of the first disk, and installs a dual-boot system if a DOS partition is found (the DOS/Windows system will boot if the user types **dos** at the `LILO:` prompt).

--append <*params*>

Specifies kernel parameters.

--linear

Use the `linear` LILO option; this is only for backwards compatibility (and linear is now used by default).

--location

Specifies where the LILO boot record is written. Valid values are **mbr** (default), **partition** (installs the boot loader on the first sector of the partition containing the kernel), or none, which prevents any bootloader from being installed.

F.5.10 lilocheck

lilocheck (optional)

If this is present, the installation program checks for LILO on the MBR of the first hard drive, and reboots the system if it is found -- No installation is done in this case. This can prevent the kickstart from reinstalling an already installed system.

F.5.11 mouse

mouse (required)

Configures the mouse for the system, both in GUI and text modes. Options are:

--device <*dev*>

Device the mouse is on (such as --device ttyS0)

--emulthree

If present, the X Window System uses simultaneous left+right mouse buttons to emulate the middle button (should be used on two button mice).

After options, the mouse type may be specified as one of the following:

```
alpsps/2, ascii, asciips/2, atibm, generic, generic3,
genericps/2, generic3ps/2, geniusnm, geniusnmps/2,
geniusnsps/2, thinking, thinkingps/2, logitech,
logitechcc, logibm, logimman, logimmanps/2, logimman+,
logimman+ps/2, microsoft, msnew, msintelli, msintellips/2,
msbm, mousesystems, mmseries, mmhittab, sun, none
```

If the mouse command is given without any arguments, or it is omitted, the installation program will attempt to autodetect the mouse (which works for most modern mice).

F.5.12 network

network (optional)

Configures network information for the system. If it is not given, and the kickstart installation does not require networking (in other words, it's not installed over NFS), networking is not configured for the system. If the installation does require networking, Anaconda assumes that the install should be done over eth0 via a dynamic IP address (BOOTP/DHCP), and configures the final, installed system to dynamically determine its IP address. The network command configures the networking information for the installation for network kickstarts as well as for the final, installed system.

--bootproto

One of **dhcp**, **bootp**, or **static** (defaults to DHCP, and **dhcp** and **bootp** are treated the same). Must be **static** for static IP information to be used.

--device <*device*>

Used to select a specific ethernet device for installation. Note, using --device <*device*> will not be effective unless the kickstart file is a local file (such as ks=floppy), since the installation program will configure the network to find the kickstart file. Example:

```
network --bootproto dhcp --device eth0
```

--ip

IP address for machine to be installed.

--gateway

Default gateway as an IP address.

--nameserver

Primary name server, as an IP address.

--netmask

Netmask for the installed system

`--hostname`

Hostname for the installed system

There are three different methods of network configuration:

- DHCP
- BOOTP
- static

The DHCP method uses a DHCP server system to obtain its networking configuration. As you might guess, the BOOTP method is similar, requiring a BOOTP server to supply the networking configuration.

The static method requires that you enter all the required networking information in the kickstart file. As the name implies, this information is static, and will be used during the installation, and after the installation as well.

To direct a system to use DHCP to obtain its networking configuration, use the following line:

```
network --bootproto dhcp
```

To direct a machine to use BOOTP to obtain its networking configuration, use the following line in the kickstart file:

```
network --bootproto bootp
```

The line for static networking is more complex, as you must include all network configuration information on one line. You'll need to specify:

- IP address
- netmask
- gateway IP address
- name server IP address

Here's an example static line:

```
network --bootproto static
--ip 10.0.2.15
--netmask 255.255.255.0
--gateway 10.0.2.254
--nameserver 10.0.2.1
```

Please Note

The entire `network` configuration *must* appear on one line! We've wrapped it here to make it easier to read.

There are two restrictions you must keep in mind should you use the static method:

- All static networking configuration information must be specified on *one* line; you cannot wrap lines using a backslash, for example.
- You can only specify one name server here. However, you can use the kickstart file's `%post` section (described in Section F.5.24, *`%post` -- Post-Installation Configuration Section*) to add more name servers, if needed.

F.5.13 `partition`

`part` (required for installs, ignored for upgrades)

Create a partition on the system. Partition requests are of the form:

```
part <mntpoint> --size
<size> [--grow] [--onpart
<partc>] [--ondisk
<disk>] [--onprimary
<N>] [--asprimary <N>]
```

The <*mntpoint*> is where the partition will be mounted, and must be of one of the following forms:

/<*mntpoint*>

(i.e. `/`, `/usr`, `/home`)

swap

The partition will be used as swap space.

raid.<*id*>

The partition will be used for software RAID (see the `raid` command later).

--size <*size*>

Sets the minimum size for the partition

--grow

Tells the partition to grow to fill available space (if any), or up to maximum size setting.

--maxsize <*size*>

Sets the maximum partition size when the partition is set to grow.

--noformat

Tells the installation program not to format the partition, for use with the `--onpart` command.

--onpart <*part*> or --usepart <*part*>

Tells the installation program to put the partition on the *already existing* device <*part*>. For example, `partition /home --onpart hda1` will put `/home` on `/dev/hda1`, which must already exist.

--ondisk <*disk*>

Forces the partition to be created on a particular disk. For example, `--ondisk sdb` will put the partition on the second disk on the system.

--onprimary <*N*>

Forces the partition to be created on primary partition <*N*> or fail. <*N*> can be 1 through 4.

--asprimary <*N*>

Forces auto allocation as a primary partition <*N*> or fail. <*N*> can be 1 through 4.

--bytes-per-inode=<*N*>

<*N*> represents the number of bytes per inode on the filesystem when it is created. It must be given in decimal format. This option is useful for applications where you want to increase the number of inodes on the filesystem.

--type=<*X*>

Sets partition type to <*X*>, where <*X*> is a numerical value.

All partitions created will be formatted as part of the installation process unless `--noformat` and `--onpart` are used.

Please Note

If `--clearpart` is used in the `ks.cfg` file, then `--onpart` cannot be used on a logical partition.

Please Note

If partitioning fails for any reason, diagnostic messages will appear on VC 3.

F.5.14 `raid`

`raid` (optional)

Assembles a software RAID device. This command is of the form:

```
raid <mntpoint> --level <level> --device
<mddevice><partitions*>
```

The <*mntpoint*> is the location to mount the RAID filesystem. If it is `/`, the RAID level must be 1 unless a boot partition (`/boot`) is present in which the `/boot` partition has to be level 1 and the root (`/`) partition can be any of the available types. The <*partitions**> (which denotes that multiple partitions can be listed) lists the RAID identifiers to add to the RAID array.

--level <*level*>

RAID level to use (0, 1, or 5).

--device <*mddevice*>

Name of the RAID device to use (such as md0 or m1). RAID devices range from md0 to md7, and each may only be used once.

Here's an example of how to create a RAID level 1 partition for `/`, and a RAID level 5 for `/usr`, assuming there are three SCSI disks on the system. It also creates three swap partitions, one on each drive.

```
part raid.01 --size 60 --ondisk sda
part raid.02 --size 60 --ondisk sdb
part raid.03 --size 60 --ondisk sdc


part swap --size 128 --ondisk sda part swap --size 128 --ondisk
sdb part swap --size 128 --ondisk sdc


part raid.11 --size 1 --grow --ondisk sda part raid.12 --size 1
--grow --ondisk sdb part raid.13 --size 1 --grow --ondisk sdc


raid / --level 1 --device md0 raid.01 raid.02 raid.03 raid /usr
--level 5 --device md1 raid.11 raid.12 raid.13
```

F.5.15 `reboot`

`reboot` (optional)

Reboot after the installation is complete (no arguments). Normally, kickstart displays a message and waits for the user to press a key before rebooting.

F.5.16 `rootpw`

`rootpw` (required)

usage: rootpw [--iscrypted] <*password*>

Set the system's root password to the <*password*> argument.

--iscrypted

If this is present, the password argument is assumed to already be encrypted.

F.5.17 `skipx`

`skipx` (optional)

If present, X is not configured on the installed system.

F.5.18 `timezone`

`timezone` (required)

timezone [--utc] <*timezone*>

Sets the system time zone to <*timezone*> which may be any of the time zones listed in "timeconfig".

--utc

If present, the system assumes the hardware clock is set to UTC (Greenwich Mean) time.

F.5.19 `upgrade`

`upgrade` (optional)

Tells the system to upgrade an existing system rather than install a fresh system.

F.5.20 xconfig

xconfig (optional)

Configures the X Window System. If this option is not given, the user will need to configure X manually during the installation, if X was installed; this option should not be used if X is not installed on the final system.

--noprobe

Don't probe the monitor.

--card <*card*>

Use card <*card*>; this card name should be from the list of cards in Xconfigurator. If this argument is not provided, Anaconda will probe the PCI bus for the card.

--monitor <*mon*>

Use monitor <*mon*>; this monitor name should be from the list of monitors in Xconfigurator. This is ignored if **--hsync** or **--vsync** is provided; if no monitor information is provided, the monitor is probed via plug-and-play.

--hsync <*sync*>

Specifies the horizontal sync frequency of the monitor.

--vsync <*sync*>

Specifies the vertical sync frequency of the monitor.

--defaultdesktop=(*GNOME* or *KDE*)

Sets the default desktop to either GNOME or KDE (and assumes that GNOME and/or KDE has been installed through %packages).

--startxonboot

Use a graphical login (runlevel 5) for the installed system.

F.5.21 `zerombr` – Partition table initialization

`zerombr` (optional)

If "zerombr" is specified, and "yes" is its sole argument, any invalid partition tables found on disks are initialized. This will destroy all of the contents of disks with invalid partition tables. This command should be used as:

```
zerombr yes
```

No other format is effective.

F.5.22 `%packages` – Package Selection

Use the `%packages` command to begin a kickstart file section that lists the packages you'd like to install (this is for installations only, as package selection during upgrades is not supported).

Packages can be specified by component or by individual package name. The installation program defines several components that group together related packages. See the `RedHat/base/comps` file on any Red Hat Linux CD-ROM for a list of components. The components are defined by the lines that begin with a number followed by a space, and then the component name. Each package in that component is then listed, line-by-line. Individual packages lack the leading number found in front of component lines.

Additionally, there are three other types of lines in the `comps` file you may run across:

Architecture specific (alpha:, i386:, and sparc64:)

If a package name begins with an architecture type, you only need to type in the package name, not the architecture name. For example:

For `i386: netscape-common` you only need to use the `netscape-common` part for that specific package to be installed.

Lines beginning with ?

Lines that begin with a `?`, are specific to the installation program. You do not have to do anything with these type of lines.

Lines beginning with `--hide`

If a package name begins with `--hide`, you only need to type in the package name, minus the `--hide`. For example:

For `--hide KDE Workstation` you only need to use the `KDE Workstation` part for that specific package to be installed.

In most cases, it's only necessary to list the desired components and not individual packages. Note that the `Base` component is always selected by default, so it's not necessary to specify it in the `%packages` section.

Here's an example `%packages` selection:

```
%packages
@ Networked Workstation
@ C Development
@ Web Server
@ X Window System
bsd-games
```

As you can see, components are specified, one to a line, starting with an @ symbol, a space, and then the full component name as given in the `comps` file. Specify individual packages with no additional characters (the `bsd-games` line in the example above is an individual package).

Please Note

You can also direct the kickstart installation to use the workstation- and server-class installations (or choose an everything installation to install all packages). To do this, simply add *one* of the following lines to the `%packages` section:

```
@ Gnome Workstation
@ KDE Workstation
@ Server
       @ Everything
```

F.5.23 %pre – Pre-Installation Configuration Section

You have the option of adding commands to run on the system immediately after the ks.cfg has been parsed. This section must be at the end of the kickstart file (after the commands) and must start with the %pre command. Note, you can access the network in the %pre section; however, **name service** has not been configured at this point, so only IP addresses will work. Here's an example %pre section:

```
%pre

# add comment to /etc/motd
echo "Kickstart-installed Red Hat Linux ‘/bin/date‘" > /etc/motd

# add another nameserver
echo "nameserver 10.10.0.2" >> /etc/resolv.conf
```

This section creates a message-of-the-day file containing the date the kickstart installation took place, and gets around the network command's "one name server only" limitation by adding another name server to /etc/resolv.conf.

Please Note

Note that the pre-install script is not run in the change root enviroment.

F.5.24 %post – Post-Installation Configuration Section

You have the option of adding commands to run on the system once the installation is complete. This section must be at the end of the kickstart file and must start with the %post command. Note, you can access the network in the %post section; however, **name service** has not been configured at this point, so only IP addresses will work. Here's an example %post section:

```
%post

# add comment to /etc/motd
echo "Kickstart-installed Red Hat Linux ‘/bin/date‘" > /etc/motd
```

```
# add another nameserver
echo "nameserver 10.10.0.2" >> /etc/resolv.conf
```

This section creates a message-of-the-day file containing the date the kickstart installation took place, and gets around the `network` command's "one name server only" limitation by adding another name server to `/etc/resolv.conf`.

Please Note

Note that the post-install script is run in a chroot'ed environment; therefore performing tasks such as copying scripts or RPMs from the installation media will not work.

--nochroot

Allows you to specify commands that you would like to run outside of the chroot'ed environment.

--interpreter ***/usr/bin/perl***

Allows you to specify a different scripting language, such as perl.

G Installing and Configuring Tripwire

Tripwire v2.3 software ensures the integrity of critical system files and directories by identifying all changes made to specified system files and directories. Configure Tripwire software to monitor your system in the way that is best for you.

Tripwire software works by comparing files and directories against a baseline. It generates the baseline by taking a **snapshot** of specified files and directories in a known secure state. Tripwire software then compares the current system against the baseline and reports any modifications, additions, or deletions. Use Tripwire software for system security, intrusion detection, damage assessment, and recovery forensics.

While it is recommended that Tripwire be selected and installed during the Red Hat Linux 7.0 installation process, it is possible to install it after your Red Hat Linux system has been installed. The following steps outline this process:

1. Locate the `RedHat/RPMS` directory on the Red Hat Linux 7.0 CD-ROM.
2. Locate the Tripwire binary RPM.
3. Type `rpm -i <name>` (where *<name>* is the name of the Tripwire RPM found in step 2)
4. After installing the Tripwire binary RPM, follow the post-installation instructions outlined below.

We recommend you read the release notes and README file.

G.1 Post-Installation Instructions

The Tripwire binary RPM installs the basic program files needed to run the software. However, this installation does not complete custom configurations that Tripwire 2.3 needs to perform correctly. After you unpack the RPM, you must:

1. Run the configuration script `/etc/tripwire/twinstall.sh` to sign these files. This script walks you through the processes of setting passphrases and signing the Tripwire policy and configuration files.

Please Note

Once encoded and signed, the configuration file should not be renamed or moved.

2. Initialize the Tripwire database file. (`/usr/sbin/tripwire--init`)
3. Run the first integrity check. (`/usr/sbin/tripwire--check`)
4. Edit the configuration file (`twcfg.txt`) with a text editor, if desired.
5. Edit the policy file (`twpol.txt`) with a text editor, if desired.

Please Note

If you plan to modify the policy file, we recommend you do so before running the configuration script. If you modify the policy file after running the configuration script, you must re-run the configuration file before initializing the database file.

G.2 Modifying the Policy File

You can specify how Tripwire software checks your system in the Tripwire policy file (`twpol.txt`). A default policy file is included in the Tripwire software installation. We recommend you tailor this policy file to fit your particular system. Tailoring the policy file greatly increases Tripwire software's ability to ensure the integrity of your system.

Locate the default policy file at `/etc/tripwire/twpol.txt`. An example policy file (located at `/usr/share/doc/tripwire-*/policyguide.txt`) is

included to help you learn the policy language. Read the sample policy file and the comments in the sample policy file to learn the policy language.

After you modify the policy file, follow the post-installation Instructions (run the configuration script). This script signs the modified policy file and renames it to `tw.pol`. This is the active policy file that runs as part of the Tripwire software.

G.3 Selecting Passphrases

Tripwire files are signed or encrypted using site or local keys. These keys are protected by passphrases. When selecting passphrases, the following recommendations apply: Use at least eight alphanumeric and symbolic characters for each passphrase. The maximum length of a passphrase is 1023 characters. Quotes should not be used as passphrase characters.

Assign a unique passphrase for the site key. The site key passphrase protects the site key, which is used to sign Tripwire software configuration and policy files. Assign a unique passphrase for the local key. The local key signs Tripwire database files. The local key may sign the Tripwire report files also.

Store the passphrases in a secure location. There is no way to remove encryption from a signed file if you forget your passphrase. If you forget the passphrases, the files are unusable. In that case you must reinitialize the baseline database.

G.4 Initializing the Database

In Database Initialization mode, Tripwire software builds a database of filesystem objects based on the rules in the policy file. This database serves as the baseline for integrity checks. The syntax for Database Initialization mode is:

```
tripwire --init
```

G.5 Running an Integrity Check

The Integrity Check mode compares the current file system objects with their properties recorded in the Tripwire database [1]. Violations are printed to standard output.

[1] The Tripwire RPM adds a file to the `/etc/cron.daily` directory that will automatically run an integrity check once every day.

The report file is saved and can later be accessed by `twprint`. An email option enables you to send email. The syntax for Integrity Check mode is:

```
tripwire --check
```

G.6 Printing Reports - `twprint` Print Report Mode

The `twprint --print-report` mode prints the contents of a Tripwire report. If you do not specify a report with the `--twrfile` or `-r` command-line argument, the default report file specified by the configuration file **REPORTFILE** variable is used.

Example: On a machine named `LIGHTHOUSE`, the command would be:

```
./twprint -m r --twrfile LIGHTHOUSE-19990622-021212.twr
```

G.7 Updating the Database after an Integrity Check

Database Update mode enables you to update the Tripwire database after an integrity check if you determine that the violations discovered are valid. This update process saves time by enabling you to update the database without having to re-initialize it. It also enables selective updating, which cannot be done through re-initialization. The syntax for Database Update mode is:

```
tripwire --update
```

G.8 Updating the Policy File

Change the way that Tripwire software scans the system by changing the rules in the policy file. You can then update the database without a complete re-initialization. This saves a significant amount of time and preserves security by keeping the policy file synchronized with the database it uses. The syntax for Policy Update mode is:

```
tripwire --update-policy
```

G.9 Testing email functions

Test mode tests the software's email notification system, using the settings currently specified in the configuration file. The syntax for Email Test Reporting mode is:

```
tripwire --test
```

G.10 Tripwire Components

The policy file begins as a text file containing comments, rules, directives, and variables. These dictate the way Tripwire software checks your system. Each rule in the policy file specifies a system object to be monitored. Rules also describe which changes to the object to report, and which to ignore.

System objects are the files and directories you wish to monitor. Each object is identified by an object name. A property refers to a single characteristic of an object that Tripwire software can monitor. Directives control conditional processing of sets of rules in a policy file. During installation, the text policy file is encrypted and renamed, and becomes the active policy file.

The database file is an important component of Tripwire software. When first installed, Tripwire software uses the policy file rules to create the database file. The database file is a baseline snapshot of the system in a known secure state. Tripwire software compares this baseline against the current system to determine what changes have occurred. This is an integrity check.

When you perform an integrity check, Tripwire software produces report files. Report files summarize any changes that violated the policy file rules during the integrity check. You can view the report file in a variety of formats, at varying levels of detail.

The Tripwire configuration file stores system-specific information, such as the location of Tripwire data files. Tripwire software generates some of the configuration file information during installation. The system administrator can change parameters in the configuration file at any time. The configuration file variables **POLFILE**, **DBFILE**, **REPORTFILE**, **SITEKEYFILE**, and **LOCALKEYFILE** specify where the policy file, database file, report files, and site and local key files reside. These variables must be defined,

or the configuration file is invalid. If any of these variables are undefined, an error occurs on execution of Tripwire software and the program exits.

G.11 Tripwire Help

All Tripwire commands support the `--help` option. Example: To get help with Create Configuration File mode, type:

```
./twadmin --help --create-cfgfile
```

The following options illustrate the types of help available in the Tripwire software:

`-?`

Display usage and version information

`--help`

Display all command modes

`--help all`

Display help for all command modes

`--help <mode>`

Display help for current command mode

`--version`

Display version information

G.12 How to Use Tripwire Software

The following flowchart illustrates how Tripwire should be used.

Figure G–1 How to Use Tripwire Software

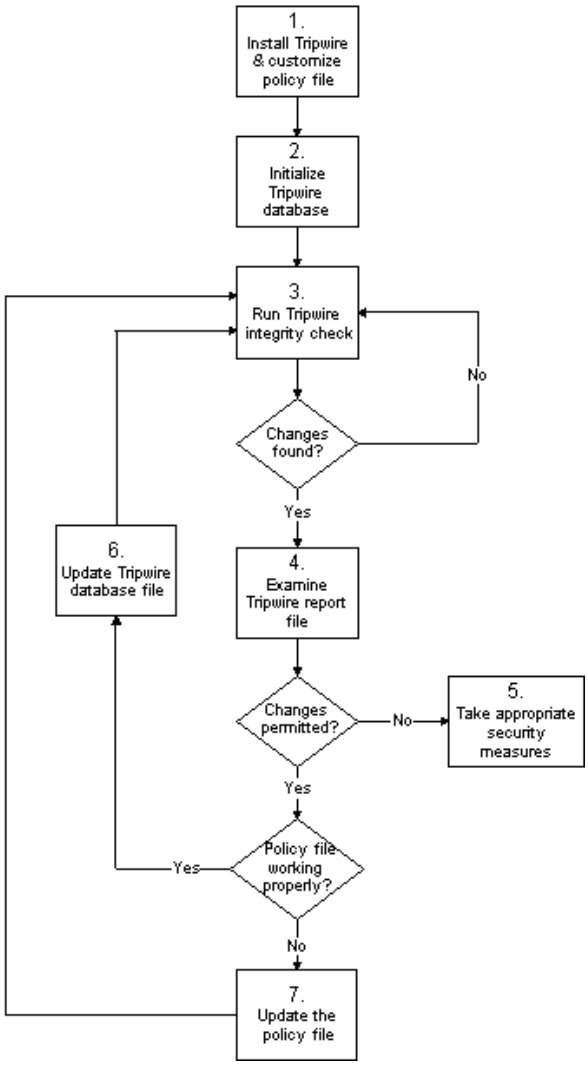

Index

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

W

X