

Red Hat Linux 7.1

The Official Red Hat Linux Customization Guide

ISBN: N/A

 Red Hat, Inc.

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

© 2001 Red Hat, Inc.

rhl-cg(EN)-7.1-Print-RHI (2001-03-09T14:34-0500)

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

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

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

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

Linux is a registered trademark of Linus Torvalds.

Motif and UNIX are registered trademarks of The Open Group.

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

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

Windows is a registered trademark of Microsoft Corporation.

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

FireWire is a trademark of Apple Computer Corporation.

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

Printed in Canada, Ireland, and Japan

Contents

Red Hat Linux 7.1

Introduction	ix
Document Conventions	ix
Using the Mouse	xiii
Copying and Pasting Text With X	xiii
More to Come	xiii
Sign Up for Support	xiv
Part I Installation-Related Reference	15
Chapter 1 Configuring a Dual-Boot System	17
1.1 If Your Computer Already Has An Operating System	17
1.2 Setting Up a Dual-Boot Environment	19
1.3 Partitioning with FIPS	21
Chapter 2 Kickstart Installations	27
2.1 What are Kickstart Installations?	27
2.2 How Do You Perform a Kickstart Installation?	27
2.3 Starting a Kickstart Installation	29
2.4 The Kickstart File	30
2.5 Kickstart Configurator	31
2.6 Kickstart Options	36
Chapter 3 Rescue Mode	55
3.1 What is Rescue Mode?	55
Chapter 4 Software RAID Configuration	59
Part II Network-Related References	63

Chapter 5	Controlling Access to Services	65
5.1	Additional Resources	67
Chapter 6	Anonymous FTP	69
Chapter 7	OpenSSH	71
7.1	Why Use OpenSSH?	71
7.2	Configuring an OpenSSH Server	71
7.3	Configuring an OpenSSH Client	72
7.4	Additional Resources	77
Chapter 8	Network File System (NFS)	79
8.1	Why Use NFS?	79
8.2	Mounting NFS Filesystems	79
8.3	Exporting NFS Filesystems	81
8.4	Additional Resources	81
Chapter 9	Samba	83
9.1	Why Use Samba?	83
9.2	Configuring Samba	83
9.3	Connecting to a Samba Share	84
9.4	Using Samba with Windows NT 4.0 and Windows 2000	84
9.5	Additional Resources	85
Part III	System Configuration	87
Chapter 10	Gathering System Information	89
10.1	System Processes	89
10.2	Memory Usage	91
10.3	Filesystems	92
10.4	Sysreport	94
10.5	Additional Resources	95

Chapter 11 Apache Configuration	97
11.1 Basic Settings.....	98
11.2 Default Settings.....	99
11.3 Virtual Hosts Settings	107
11.4 Server Settings	112
11.5 Performance Tuning.....	114
11.6 Saving Your Settings	116
11.7 Additional Resources	117
Chapter 12 BIND Configuration	119
12.1 Adding a Forward Master Zone	120
12.2 Adding a Reverse Master Zone	122
12.3 Adding a Slave Zone	124
Chapter 13 Printer Configuration	127
13.1 Adding a Local Printer	129
13.2 Adding a Remote UNIX Printer	131
13.3 Adding a Samba (SMB) Printer	133
13.4 Adding a Novell NetWare (NCP) Printer	135
13.5 Adding a JetDirect Printer	136
13.6 Printing a Test Page	138
13.7 Creating Printer Aliases.....	138
13.8 Modifying Existing Printers	138
13.9 Additional Resources	139
Chapter 14 Linuxconf	141
14.1 Starting Linuxconf	141
14.2 Linuxconf User Interfaces.....	141
14.3 Gnome-Linuxconf Interface.....	142
14.4 Enabling Web-Based Linuxconf Access	143
14.5 Adding a User Account	144
14.6 Modifying a User Account	149
14.7 Changing a User's Password.....	150
14.8 Changing the Root Password	150

14.9	Disabling a User Account	150
14.10	Enabling a User Account	151
14.11	Deleting a User Account	151
14.12	Groups	153
14.13	Filesystems	157
14.14	Network Configuration with Linuxconf	161
14.15	Finding Your Way Through Linuxconf	167
14.16	Additional Resources	168
Chapter 15 Control Panel		169
15.1	Network Configurator	170
15.2	Time and Date	175
Chapter 16 Building a Custom Kernel		177
16.1	The 2.4 Kernel	177
16.2	Building a Modularized Kernel	177
16.3	Making an initrd Image	181
16.4	Building a Monolithic Kernel	182
16.5	Loading Kernel Modules	182
Part IV Package Management		185
Chapter 17 Package Management with RPM		187
17.1	RPM Design Goals	187
17.2	Using RPM	188
17.3	Checking a Package's Signature	194
17.4	Impressing Your Friends with RPM	196
17.5	Additional Resources	198
Chapter 18 Gnome-RPM		201
18.1	Starting Gnome-RPM	202
18.2	The Package Display	204
18.3	Installing New Packages	205

18.4	Configuration.....	207
18.5	Package Manipulation.....	213
Chapter 19 Red Hat Network.....		219
Part V Appendixes.....		221
Appendix A Getting Started with Gnu Privacy Guard.....		223
A.1	An Introduction to GnuPG	223
A.2	Generating a Keypair.....	224
A.3	Generating a Revocation Certificate.....	226
A.4	Exporting your Public Key	227
A.5	Importing a Public Key	230
A.6	What Are Digital Signatures?	231
A.7	Additional Resources	231

Introduction

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

The *Official Red Hat Linux Customization Guide* contains information on how to customize your Red Hat Linux system to fit your needs. If you are looking for step-by-step, task-oriented guides for configuring and customizing your system, this is the guide for you. This manual discusses many beginner and intermediate topics such as the following:

- Setting up a Network Interface Card (NIC)
- Configuring a dual-boot system
- Configuring Samba shares
- Managing your software with RPM

This manual is divided into the following main categories:

- Installation-Related Reference
- Network-Related Reference
- System Configuration
- Package Management
- Advanced Topics

This guide assumes you have a basic understanding of your Red Hat Linux system. If you need reference material which covers more basic issues, please refer to the *Official Red Hat Linux Getting Started Guide*. If you need more advanced documentation, please refer to the *Official Red Hat Linux Reference Guide*.

HTML and PDF versions of all the Official Red Hat Linux manuals are available online at <http://www.redhat.com/support/manuals/>.

Document Conventions

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

command

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

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

filename

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

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

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

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

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

application

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

Use Netscape Navigator to browse the Web.

[key]

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

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

[key]-[combination]

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

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

text found on a GUI interface

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

On the GNOME **Control Center** screen, you can customize your GNOME window manager.

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

top level of a menu on a GUI screen or window

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

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

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

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

button on a GUI screen or window

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

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

computer output

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

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

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

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

prompt

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

```
$
#
[stephen@maturin stephen]$
leopard login:
```

user input



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

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

Another example, with the word **root** displayed as something the user needs to type in:

If you need to log in as root when you first log into your system, and you are using the graphical login screen, at the Login prompt, type **root**. At the Password prompt, type in the root password.

glossary entry

A word that appears in the glossary will be shown in the body of the document in this style. For example:

The lpd **daemon** handles printing requests.

In this case, the style of the word **daemon** should indicate to you that a definition of the term is available in the glossary.

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

Note

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



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

WARNING

If you choose not to partition manually, a server-class installation will remove all existing partitions on all installed hard drives. Don't choose this installation class unless you're sure you have no data you need to save.

Using the Mouse

Red Hat Linux is designed to use a three-button mouse. If you have a two-button mouse, you should have selected three-button emulation during the installation process. If you're using three-button emulation, pressing both mouse buttons at the same time equates to pressing the missing third (middle) button.

In this document, if you are instructed to click with the mouse on something, that means click the left mouse button. If you need to use the middle or right mouse button, that will be explicitly stated. (This will be reversed if you've configured your mouse to be used by a left handed person.)

The phrase "drag and drop" may be familiar to you. If you're instructed to drag and drop an item on your GUI desktop, click on something and hold the mouse button down. While continuing to hold down the mouse button, drag the item by moving the mouse to a new location. When you've reached the desired location, release the mouse button to drop the item.

Copying and Pasting Text With X

Copying and pasting text is easy using your mouse and the X Window System. To copy text, simply click and drag your mouse over the text to highlight it. To paste the text somewhere, click the middle mouse button in the spot where the text should be placed.

More to Come

The *Official Red Hat Linux Customization Guide* is part of Red Hat's growing commitment to provide useful and timely support to Red Hat Linux users. As new tools and applications are released, this guide will be expanded to include them.

Send in Your Feedback

If you spot a typo in the *Official Red Hat Linux Customization Guide*, or if you've thought of a way to make this manual better, we'd love to hear from you! Please submit a report in Bugzilla (<http://www.redhat.com/bugzilla>) against the component `rhl-cg`.

Be sure to mention the manual's identifier:

```
rh1-cg(EN)-7.1-Print-RHI (2001-03-09T14:34-0500)
```

If you mention this manual's identifier, we'll know exactly which version of the guide you have.

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

Sign Up for Support

If you have an official edition of Red Hat Linux 7.1, 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.
- Red Hat Network — Easily update your packages and receive security notices that are customized for your system. Go to <http://www.redhat.com/network> for more details.
- Priority FTP access — No more late-night visits to congested mirror sites. Owners of Red Hat Linux 7.1 receive free access to priority.redhat.com, Red Hat's preferred customer FTP service, offering high bandwidth connections day and night.
- *Under the Brim: The Official Red Hat E-Newsletter* — Every month, get the latest news and product information directly from Red Hat.

To sign up, go to <http://www.redhat.com/apps/activate/>. You'll find your Product ID on a black, 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 *Getting Technical Support Appendix* in the *Official Red Hat Linux x86 Installation Guide*.

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

The Red Hat Documentation Team

Part I Installation-Related Reference

1 Configuring a Dual-Boot System

This chapter explains how to install Red Hat Linux on a computer that currently runs another operating system, such as Microsoft Windows, and how to create a dual-boot environment.

1.1 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 you are uncomfortable with disk partitioning? You can install Red Hat Linux on your system without creating any Linux partitions¹ by performing a **partitionless** installation. During a partitionless installation, the installation program will install Red Hat Linux on an existing, formatted Windows partition. You will 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 what you want to do, refer to the *Official Red Hat Linux x86 Installation Guide* for those instructions.

- Do you want to install Red Hat Linux and then have the option of booting either Red Hat Linux or your other operating system? A workstation- or custom-class installation can be performed so that Red Hat Linux is installed on your system, without affecting the other operating system. In fact, a workstation-class installation will accomplish this by default. In a custom-class installation, you can install LILO (the Linux LOader) to boot Linux and the other operating system.

Install the other operating system first and then install Red Hat Linux. The Red Hat Linux installation program will usually detect the other operating system and automatically configure LILO to boot either Red Hat Linux or the other operating system. The *Official Red Hat Linux x86 Installation Guide* provides instructions on installing and configuring LILO. After the installation, whenever you start the computer, you can indicate whether you want to start Red Hat Linux or the other operating system.

Remember to back up all important information before configuring your system to boot more than one operating system. Be sure to create a boot disk for both operating systems in case the boot loader fails to recognize both of them.

¹ A partition is a physical division on a hard drive.

WARNING

The BIOS in some systems cannot access more than the first 1024 cylinders on a hard drive. If this is the case, the `/boot` Linux partition must be located on the first 1024 cylinders of your hard drive for LILO to boot.

WARNING

If you want to dual-boot Red Hat Linux and Windows NT, you should install Windows NT first because it installs its own boot loader on the Master Boot Record (MBR). After installing Windows NT, if you install LILO during the Red Hat Linux installation program, the NT boot loader will be overwritten, but it should add a LILO entry labeled `dos` to boot Windows NT. Remember that a workstation-class installation automatically installs LILO to the MBR. Installing LILO on the MBR to boot Windows NT has been known to fail in some cases. If this is the case, you should perform a custom-class installation and install LILO on the first sector of the root partition instead of the MBR.

If you install LILO on the first sector of the root partition, be sure to create a boot disk. You will need to either use the boot disk to boot Red Hat Linux or configure the NT boot loader to boot LILO from the first sector of the root partition. For more information on configuring the NT boot loader, refer to <http://www.linux-doc.org/HOWTO/mini/Linux+NT-Loader.html>.

If this what you want to do, read Section 1.2, *Setting Up a Dual-Boot Environment*.

- Do you want Red Hat Linux to be the only operating system on your computer? Choose a server-class installation, choose a workstation-class installation and manually delete the DOS (Windows) partitions, or choose a custom-class installation and delete the existing DOS (Windows) partitions.
-

Note

In order to install Red Hat Linux and keep another OS on your system, you must have sufficient space on which Red Hat Linux will be installed. 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. This will also happen if you select a server-class installation. Unless you have sufficient room on your hard drive for Red Hat Linux, you cannot install it.

If this is what you want to do, 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 the installation as explained in the *Official Red Hat Linux x86 Installation Guide*.

1.2 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.

For clarity, we will assume that the other operating system is Windows. But the general procedures are similar for other operating systems.

Note

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

Before starting the installation program, you must first make room for Red Hat Linux. Your choices are as follows:

- Add a new hard drive.
 - Use an existing hard drive or partition.
-

- Create a new partition.

1.2.1 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`. (For 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, all you need to do is start the Red Hat Linux installation program. After starting the Red Hat Linux installation program, just make sure you tell it to install Linux on the newly installed hard drive (such as `hdb` or `sdb`) rather than the hard drive used by Windows.

1.2.2 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 the hard drive 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.

Note

Windows uses letters to refer to removable drives (for example, a ZIP drive) and network storage (virtual drives) as well as for local hard drive space; you cannot install Linux on a removable or network drive.

If a local Windows partition is available in which you want to install Linux, complete the following steps:

1. Copy all data you want to save from the selected hard drive or partition (`D:` in this example) to another location.
 2. Start the Red Hat Linux installation program and tell it to install Linux in the designated drive or partition — in this example, in the hard drive or partition that Windows designates as `D:`. Note that Linux distinguishes between hard drives and disk partitions. Thus:
 - If `C:` and `D:` on this computer refer to two separate hard drives, the installation program will recognize them as `hda` and `hdb` (IDE) or `sda` and `sdb` (SCSI). Tell the installation program to install on `hdb` or `sdb`.
-

- If C: and D: refer to partitions on a single drive, the installation program will recognize them as hda1 and hda2 (or sda1 and sda2). During the partitioning phase of the Red Hat Linux installation, 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. If you don't, however, Windows will complain whenever you boot that it cannot read Drive D; and should someone accidentally format D, your Linux system would be destroyed.)

1.2.3 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 see 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 1.3, *Partitioning with FIPS*.

1.3 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.

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 the Red Hat Linux CD-ROM.

Should you decide to use FIPS, be aware that after FIPS runs you will be left with two partitions: the one you resized, and the one FIPS created out of the newly freed space. If your goal is to use that space to install Red Hat Linux, you should delete the newly created partition, either by using `fdisk` under your current operating system, or while setting up partitions during a custom-class installation.

The following instructions are a simplified version of the FIPS documentation file, `fips.doc`, located in the FIPS directory (`/dosutils/fips20/*`). These instructions should apply in most instances. If you encounter any problems, see the documentation file.

1. From Windows:

- Do a full backup.
- Run `scandisk` to verify that the hard drive contains no bad clusters.
- Decide how to distribute the available space on the hard drive between the operating systems. Use **Windows Explorer** to see the free space on the drive. Make a note of the space (in megabytes) that each operating system will have.
- If you 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.

```
dosutils/fips20/fips.exe  
dosutils/fips20/restorrb.exe  
dosutils/fips20/errors.txt
```

```
dosutils/fips20/fips.doc
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 1-1 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 1-2 FIPS Root Partition Screen

```
Partition table:
```

Part.	bootable	Start			System	End			Start Sector	Number of Sectors	MB
		Head	Cyl.	Sector		Head	Cyl.	Sector			
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 1–3 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 1–4 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 1–5 FIPS Verification Screen

```
First Cluster: 17442
Last Cluster: 65511
```

```
Testing if empty ... OK
```

```
New partition table:
```

Part.	bootable	Start Head Cyl. Sector	System	End Head Cyl. Sector	Start Sector	Number of 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 1–4, *Partition Resizing Screen* reappears, allowing you to change the partition sizes. If you answer **c**, a confirmation screen Figure 1–6, *FIPS Confirmation Screen* appears:

Figure 1–6 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 describe a number of factors that could cause the resizing operation to fail. If all else fails, you can restore Windows with the backup you made.

The second partition (`hda2` or `sda2`) contains the space that the Red Hat Linux installation program will use. When the **Disk Druid** screen appears during installation, delete this partition (the installation manual explains how), then proceed with Linux partitioning.

2 Kickstart Installations

2.1 What are Kickstart Installations?

Many system administrators would prefer to use an automated installation method to install Red Hat Linux on their machines. To answer this need, Red Hat created the kickstart installation method. Using kickstart, 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. This installation method can support the use of a single kickstart file 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

2.2 How Do You Perform a Kickstart Installation?

Kickstart installations can be performed using a local CD-ROM, a local hard drive, or via NFS, FTP or HTTP.

To use kickstart mode, you must first create a kickstart file (`ks.cfg`), and make it available to the Red Hat Linux installation program.

2.2.1 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 us 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 must be located in the boot disk's top-level directory. Note that the Red Hat Linux boot disks are in MS-DOS format, so it is easy to copy the kickstart file under Linux using the `mcopy` command:

```
mcopy ks.cfg a:
```

Alternatively, you can use Windows to copy the file. 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, 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 they are not required to.

To perform 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 provide the client with 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 replace the value after `filename` with the name of the kickstart file (or the directory in which the kickstart file resides) and the value after `next-server` with 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 particular 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.

2.3 Starting a Kickstart Installation

To begin a kickstart installation, you must boot the system from a Red Hat Linux boot diskette or the CD-ROM and enter a special boot command at the boot prompt. If the kickstart file is located on a boot diskette that was created from the `boot.img` or `bootnet.img` image file, the correct boot command would be:

```
boot: linux ks=floppy
```

The `linux ks=floppy` command also works if the `ks.cfg` file is located on a vfat filesystem on a floppy diskette and you boot from the Red Hat Linux CD-ROM.

An alternate boot command for booting of the Red Hat Linux CD-ROM and having the kickstart file on a vfat filesystem on a floppy diskette is:

```
boot: linux ks=hd:fd0/ks.cfg
```

The Red Hat Linux installation program looks for a kickstart file if the `ks` command line argument is passed to the kernel. The command line argument can take a number of forms:

`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. For example, if your NFS server is `server.example.com` and the kickstart file is on the NFS share `/mydir/ks.cfg`, the correct boot command would be `ks=nfs:server.example.com/mydir/ks.cfg`.

`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=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 `/`, the bootfile provided by DHCP is looked for on the NFS server.
- If DHCP is specified and the bootfile begins with something other than a `/`, the bootfile provided by DHCP is looked for in the `/kickstart` directory on the NFS server.
- If DHCP 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.

2.4 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 `RH-DOCS` directory of the Red Hat Linux Documentation 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, be aware of the following issues when you are 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 the user would be prompted during a typical installation. Once the answer is given, the installation will continue unattended (unless it finds 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 mark is necessary to end the file other than the post section itself.

2.5 Kickstart Configurator

Kickstart Configurator allows you to create a kickstart file using a graphical user interface, so that you do not have to remember the correct syntax of the file. After choosing the kickstart options, click the **Save File** button, and a kickstart file is generated.

For a more detailed explanation of kickstart options, refer to Section 2.6, *Kickstart Options*.

Figure 2–1 Kickstart Configurator

The screenshot shows the Kickstart Configurator window with the following sections and options:

- Basic configuration:**
 - Language: English
 - Keyboard: us
 - Mouse: Generic - 2 Button Mouse (PS/2)
 - Time Zone: America/New_York
 - Root Password: [Empty field]
 - LILO: MBR (selected), None
 - Authentication: Use shadow passwords, Use MD5
- Installation Source:**
 - Installation Source: CD-ROM (selected), NFS, FTP, Hard drive
 - NFS Server: [Empty field]
 - NFS Directory: [Empty field]
 - FTP Server: [Empty field]
 - FTP Directory: [Empty field]
 - Hard Drive partition: [Empty field], Directory: [Empty field]
- Partition Information:**
 - Clear Master Boot Record: Yes (selected), No
 - Remove Existing Partitions: None (selected), All, Linux

Mount Point	Type	Size (M)	Growable
/boot	ext2	35	No
	Linux Swap	128	No
/	ext2	1000	Yes

Buttons: Add, Edit, Delete
- Additional Options:**
 - Networking, Authentication, Firewall, Packages
- Buttons: Save File, Exit

To use Kickstart Configurator, you must be running the X Window System. To start Kickstart Configurator, use one of the following methods:

- On the GNOME desktop, go to the **Main Menu Button** (on the Panel) => **Programs** => **System** => **Kickstart Configurator**.
- On the KDE desktop, go to the **Main Menu Button** (on the Panel) => **Red Hat** => **System** => **Kickstart Configurator**
- Type the command `ksconfig` at a shell prompt (for example, in an XTerm or GNOME terminal)

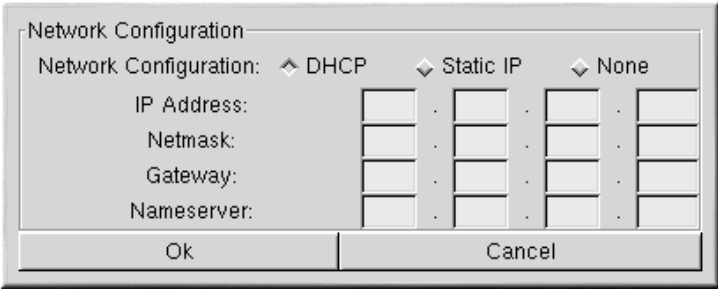
All options in the **Basic Configuration** section are required. Choose the default language, keyboard, mouse, and time zone. Enter a root password in the **Root Password** text field to set the root password for the system. Choose whether to install LILO on the master boot record (MBR) or not to install LILO. Next to **Authentication** choose whether to enable shadow passwords and MD5.

In the **Installation Source** section, choose either **CD-ROM**, **NFS**, **FTP**, or **Hard drive** as the installation source. If you choose **NFS**, enter the NFS server and NFS directory. If you choose **FTP**, enter the FTP server and FTP directory. If you choose **Hard drive**, enter the hard drive partition and directory.

In the **Partition Information** section, choose whether to clear the MBR. You can also choose not to remove the existing partitions, to remove all the existing partitions, or to remove only the existing Linux partitions. Configure your partition table and mount points (this can be compared to **Disk Druid** or **fdisk**). By default, a 35 MB `/boot` directory, a 128 MB swap partition, and a `/` (root) partition is created. Delete them if you want to configure different partitions. To add additional partitions, click the **Add** button.

In the **Additional Options** section, click the **Networking** button to configure the network settings. Choose **DHCP**, **Static IP**, or **None** as shown in Figure 2–2, *Network Configuration*. If you choose **Static IP**, enter your IP address, netmask, gateway, and primary nameserver in the given text fields.

Figure 2–2 Network Configuration



Network Configuration

Network Configuration: DHCP Static IP None

IP Address: [] . [] . [] . []

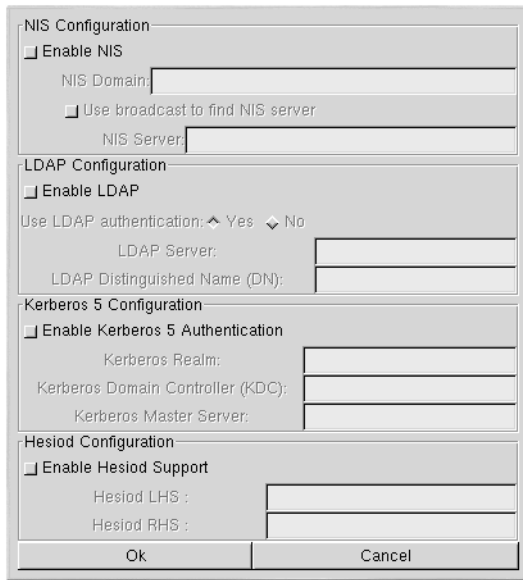
Netmask: [] . [] . [] . []

Gateway: [] . [] . [] . []

Nameserver: [] . [] . [] . []

Ok Cancel

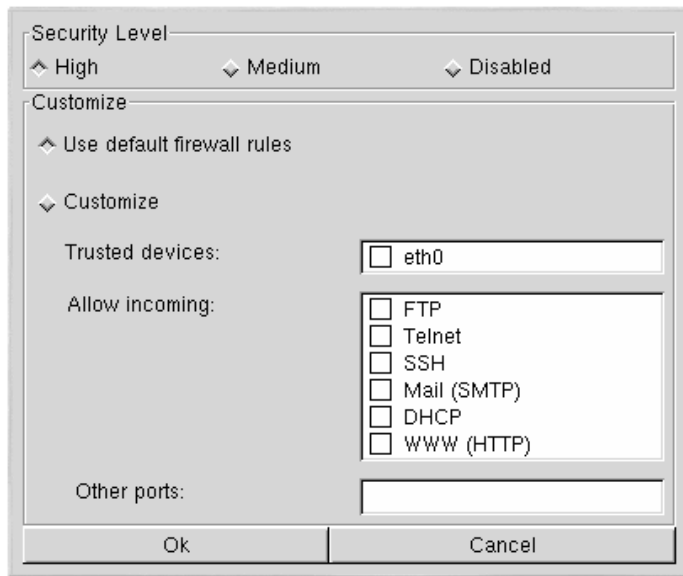
Click the **Authentication** button to enable NIS, LDAP, Kerberos 5, and Hesiod support. The window in Figure 2–3, *Authentication* will appear.

Figure 2–3 Authentication

The image shows a graphical user interface window titled "Authentication" with four main sections, each with an "Enable" checkbox and several input fields. At the bottom are "Ok" and "Cancel" buttons.

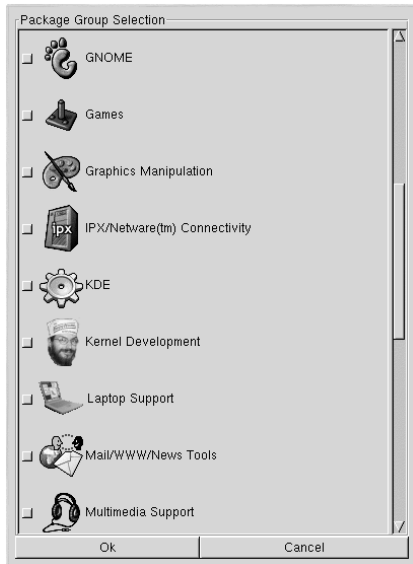
- NIS Configuration:**
 - Enable NIS
 - NIS Domain:
 - Use broadcast to find NIS server
 - NIS Server:
- LDAP Configuration:**
 - Enable LDAP
 - Use LDAP authentication: Yes No
 - LDAP Server:
 - LDAP Distinguished Name (DN):
- Kerberos 5 Configuration:**
 - Enable Kerberos 5 Authentication
 - Kerberos Realm:
 - Kerberos Domain Controller (KDC):
 - Kerberos Master Server:
- Hesiod Configuration:**
 - Enable Hesiod Support
 - Hesiod LHS:
 - Hesiod RHS:

Click the **Firewall** button to configure the firewall settings for the system. The firewall window, as shown in Figure 2–4, *Firewall Configuration*, is identical to the firewall screen that is used during the Red Hat Linux installation program. Refer to the *Official Red Hat Linux x86 Installation Guide* for further information about configuring the firewall settings.

Figure 2–4 Firewall Configuration

Click the **Packages** button to configure which packages to install. This section allows you to choose which packages you want to install. These are the same packages groups that are in the Red Hat Linux installation program.

Figure 2–5 Package Selection



After selecting your kickstart options, click the **Save File** button to display a file dialog box. Save the file as `ks.cfg`. Refer to Section 2.3, *Starting a Kickstart Installation* to start the kickstart installation.

2.6 Kickstart Options

The following options can be placed in a kickstart file.

2.6.1 `auth` — Authentication Options

`auth` (required)

Sets up the authentication options for the system. It's similar to the `authconfig` command, which 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. To use this option, you must have the `nss_ldap` package 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 changing passwords, using an LDAP directory. To use this option, you must have the `nss_ldap` package installed. You must also specify a server and a base DN.

--ldapserver=

The name of the LDAP server to use, if you 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 does not know about home directories, UIDs, or shells. So if you enable Kerberos you will need to make users' accounts known to this workstation by enabling LDAP, NIS, or Hesiod or by using the `/usr/sbin/useradd` command to make their accounts known to this workstation. If you use this option, you must have the `pam_krb5` package installed.

--krb5realm

The Kerberos 5 realm to which your workstation belongs.

--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 handles password changing and other administrative requests. This server must be run on the master KDC if you have more than one KDC.

--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 for which to search, so the LHS and RHS usually begin with periods.

2.6.2 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.

2.6.3 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, however, kickstart needs a hint to find the proper devices. The device command, which tells Anaconda to install extra modules, is in this format:

```
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 they are put in quotes. For example:

```
--opts "aic152x=0x340 io=11"
```

2.6.4 Driver Disk

driverdisk (optional)

Driver disks can be used during kickstart installations. You will need to copy the driver disk's contents to the root directory of a partition on the system's hard drive. Then you will need to use the `driverdisk` command to tell the installation program where to look for the driver disk.

```
driverdisk <partition> [--type <fstype>]
```

<partition> is the partition containing the driver disk.

--type

Filesystem type (for example, vfat or ext2).

2.6.5 firewall

firewall (optional)

Firewall options can be configured in kickstart. This configuration corresponds to the **Firewall Configuration** screen in the installation program.

```
firewall [--high | --medium | --disabled] [--trust
<device>] [--dhcp] [--ssh] [--telnet] [--smtp] [--http]
[--ftp] [--port <portspec>]
```

Levels of security

Choose one of the following levels of security:

- --high
- --medium
- --disabled

--trust <device>

Listing a device here, such as eth0, allows all traffic coming from that device to go through the firewall. To list more than one device, use `--trust eth0 --trust eth1`. Do NOT use a comma-separated format such as `--trust eth0, eth1`.

Allow incoming

Enabling these options allow the specified services to pass through the firewall.

- `--dhcp`
- `--ssh`
- `--telnet`
- `--smtp`
- `--http`
- `--ftp`

--port <portspec>

You can specify that ports be allowed through the firewall using the port:protocol format. For example, if you wanted to allow IMAP access through your firewall, you can specify `imap:tcp`. You can also specify numeric ports explicitly; for example, to allow UDP packets on port 1234 through, specify `1234:udp`. To specify multiple ports, separate them by commas.

2.6.6 install

install (optional)

Tells the system to install a fresh system rather than upgrade an existing system. This is the default mode.

2.6.7 Installation Methods

You must use one of these four commands to specify what type of kickstart installation is being performed:

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>
```

cdrom

Install from the first CD-ROM drive on the system.

For example:

```
cdrom
```

harddrive

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 --url http://<server>/<dir>
url --url ftp://<username>:<password>@<servername>;/<dir>
```

2.6.8 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.14, nl, no-latin1, no, pc110, pl, pt-latin1,
```

```
pt-old, ro, ru-cpl251, ru-ms, ru-yawerty, ru, rul, 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
```

2.6.9 language

lang (required)

Sets the default language for the installed system. The language you specify will be used during the installation and will be used 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 language codes are the following (please note that these are subject to change at any time):

```
cs_CZ, da_DK, en_US, fr_FR, de_DE, hu_HU, is_IS, it_IT,
ja_JP.eucJP, no_NO, ro_RO, sk_SK, sl_SI, sr_YU, es_ES,
ru_RU.KOI8-R, uk_UA.KOI8-U, sv_SE, tr_TR
```

2.6.10 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).

--nonlinear

Use the `nolinux` LILO option; `linear` is now used by default.

--location

Specifies where the LILO boot record is written. Valid values are the following: `mbr` (the default) or `partition` (installs the boot loader on the first sector of the partition containing the kernel). If no location is specified, LILO is not installed.

2.6.11 lilocheck

lilocheck (optional)

If `lilocheck` is present, the installation program checks for LILO on the MBR of the first hard drive, and reboots the system if it is found — in this case, no installation is performed. This can prevent kickstart from reinstalling an already installed system.

2.6.12 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, simultaneous clicks on the left and right mouse buttons will be recognized as the middle mouse button by the X Window System. This option should not be used if you have a two button mouse.

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,  
geniusnmps/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. This procedure works for most modern mice.

2.6.13 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, the Red Hat Linux installation program assumes that the installation 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` option configures networking information for kickstart installations via a network as well as for the 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 that 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 the 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
- Nameserver 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
```

If you use the static method, be aware of the following two restrictions:

- 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 nameserver here. However, you can use the kickstart file's `%post` section (described in Section 2.6.25, `%post — Post-Installation Configuration Section`) to add more name servers, if needed.

2.6.14 partition

part (required for installs, ignored for upgrades)

Creates a partition on the system. Partition requests are of the form:



```
part <mntpoint> --size <size> [--grow]
[--onpart <partc>] [--ondisk <disk>]
[--onprimary <N>] [--asprimary]
```

The *<mntpoint>* is where the partition will be mounted and must be of one of the following forms:

/<mntpoint>

For example, */*, */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>

The minimum partition size in megabytes. Specify an integer value here such as 500. Do not append the number with MB.

--grow

Tells the partition to grow to fill available space (if any), or up to the maximum size setting.

--maxsize <size>

The maximum partition size in megabytes when the partition is set to grow. Specify an integer value here, and do not append the number with MB.

--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 the primary partition *<N>* or fail. *<N>* can be 1 through 4. For example, `--onprimary=1` specifies that the partition is to be created on the first primary partition.

--asprimary

Forces automatic allocation of the partition as a primary partition or the partitioning will fail.

--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.

Note

If `--clearpart` is used in the `ks.cfg` file, then `--onpart` cannot be used on a logical partition.

Note

If partitioning fails for any reason, diagnostic messages will appear on virtual console 3.

2.6.15 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 where the RAID filesystem is mounted. If it is `/`, the RAID level must be 1 unless a boot partition (`/boot`) is present. If a boot partition is present, the `/boot` partition must 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.

The following example shows 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
```

2.6.16 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.

2.6.17 rootpw

rootpw (required)

`rootpw [--iscrypted] <password>`

Sets the system's root password to the *<password>* argument.

--iscrypted

If this is present, the password argument is assumed to already be encrypted.

2.6.18 skipx

skipx (optional)

If present, X is not configured on the installed system.

2.6.19 timezone

timezone (required)

```
timezone [--utc] <timezone>
```

Sets the system time zone to <timezone> which may be any of the time zones listed by `timeconfig`.

--utc

If present, the system assumes the hardware clock is set to UTC (Greenwich Mean) time.

2.6.20 upgrade

upgrade (optional)

Tells the system to upgrade an existing system rather than install a fresh system.

2.6.21 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 installation program tries to probe for it automatically.

--hsync <sync>

Specifies the horizontal sync frequency of the monitor.

--vsync <sync>

Specifies the vertical sync frequency of the monitor.

--defaultdesktop=GNOME or --defaultdesktop=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 on the installed system.

2.6.22 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 in the following format:

```
zerombr yes
```

No other format is effective.

2.6.23 %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:

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 used by the installation program and should not be altered.

Lines beginning with --hide

If a package name begins with `--hide`, you only need to type in the package name, without 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).

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
```

2.6.24 %pre — Pre-Installation Configuration Section

You can add 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 that 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 limitation of only one name server by adding another name server to `/etc/resolv.conf`.

Note

Note that the pre-install script is not run in the change root environment.

2.6.25 %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 limitation of one name server only by adding another name server to `/etc/resolv.conf`.

Note

Note that the post-install script is run in a chroot 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 environment.

The following example copies the file `/etc/resolv.conf` to the filesystem that was just installed.

```
%post --nochroot
cp /etc/resolv.conf /mnt/sysimage/etc/resolv.conf
```

--interpreter */usr/bin/perl*

Allows you to specify a different scripting language, such as Perl. Replace `/usr/bin/perl` with the scripting language of your choice.

The following example uses a Perl script to replace `/etc/HOSTNAME`.

```
%post --interpreter /usr/bin/perl

# replace /etc/HOSTNAME
open(HN, ">HOSTNAME");
print HN "1.2.3.4 an.ip.address\n";
```

3 Rescue Mode

When things go wrong, there are ways to fix problems. However, these methods require that you understand the system well. This chapter will describe the ways that you can boot into rescue modes, where you can use your own knowledge to repair the system.

3.1 What is Rescue Mode?

Rescue mode is the ability to boot a small Linux environment entirely from a diskette or CD, or using some other method.

As the name implies, rescue mode is provided to rescue you from something. During 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, you can access the files stored on your system's hard drive, even if you cannot actually run Linux from that hard drive.

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

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

Next, we will take a closer look at each of these scenarios.

3.1.1 Unable to Boot Linux

This problem is often caused by the installation of another operating system after you have 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 (MBR) that originally contained the LILO boot-loader. If LILO is overwritten in this manner, you will not be able to boot Red Hat Linux unless you can get into rescue mode.

3.1.2 Hardware/Software Problems

This category includes a wide variety of different situations. Two examples include failing hard drives and forgetting to run LILO after building a new kernel. In both of these situations, you may be unable to boot 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 command 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 an installation boot diskette ¹ or the Red Hat Linux CD-ROM #1.
- By booting from a network or PCMCIA boot diskette. You can only do this if your network connection is working. You will need to identify the network host and transfer type. For an explanation of how to specify this information, see *Installing over the Network* in the *Official Red Hat Linux x86 Installation Guide*.

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 many useful commands including:

anaconda	gzip	mkfs.ext2	ps
badblocks	head	mknod	python
bash	hwclock	mkraid	python1.5
cat	ifconfig	mkswap	raidstart
chattr	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
fsck.ext3	mdu	mzip	traceroute
ftp	mformat	open	umount
gnome-pty-helper	minfo	pico	uncpio
grep	mkdir	ping	uniq
gunzip	mke2fs	probe	zcat

¹ To create an installation boot diskette use the `images/boot.img` file on the Red Hat Linux CD-ROM #1 with the command `dd if=boot.img of=/dev/fd0` and a blank diskette.

However, if your root filesystem is undamaged, you can mount it and then run any standard Linux utility. For example, if your root filesystem is in `/dev/hda5`, you can mount this partition with the following command:

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

In the above command, `/foo` is a directory that you have created.

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

If you do not know the names of your Linux partitions, you can guess what they are. Mounting non-existent partitions will do no harm.

3.1.3 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 (if you are using the graphical LILO, you must press `[Ctrl]-[x]` to exit the graphical screen and go to the `boot :` prompt):

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

In single-user mode, your computer boots to runlevel 1. Your local filesystems will be mounted, but your network will not be activated. You will have 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 emergency mode 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.

Have you ever rebuilt a kernel and, eager to try out your new handiwork, rebooted before running `/sbin/lilo`? If you did not have an entry for an older kernel in `lilo.conf`, you had a problem. If you would like to know a solution to this problem, read this section.

In many cases, you can boot your Red Hat Linux system from the Red Hat Linux boot disk ¹ with your root filesystem mounted and ready to go. Here is how to do it:

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

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

Replace the `XX` in `/dev/hdXX` with the appropriate letter and number for your root partition.

What does this command do? First, it starts the boot process 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 negative side to using this technique? Unfortunately, yes. Because the kernel on the Red Hat Linux boot disk only has support for IDE built-in, if your system is SCSI-based, you will not be able to do this. In that case, you will have to access rescue mode using the **linux rescue** command mentioned above.

4 Software RAID Configuration

Read the Appendix on RAID in the *Official Red Hat Linux Reference Guide* first to tell about RAID and the differences between Hardware RAID versus Software RAID.

Software RAID can be configured during the graphical installation of Red Hat Linux or during a kickstart installation. 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.

Tip: If You Use `fdisk`

If you are using `fdisk` to create a RAID partition, remember that instead of creating a partition as type 83, which is Linux native, you must create the partition as type `fd` (Linux RAID). Also, for best performance, partitions within a given RAID array should span identical cylinders on drives.

- Create a partition. In Disk Druid, choose **Add** to create a new partition (see Figure 4–1, *Creating a New RAID Partition*).

Figure 4–1 Creating a New RAID Partition

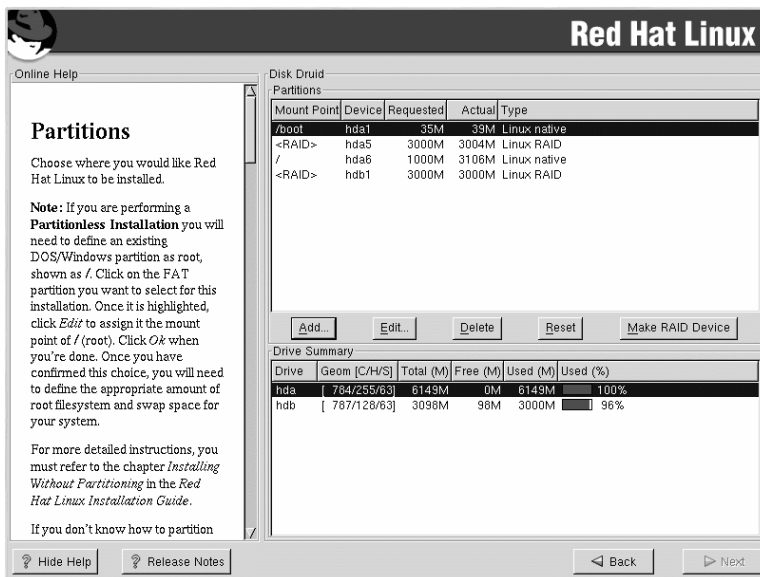


- You will not be able to enter a mount point (you will be able to do that once you have created your RAID device).
 - Enter the size that you want the partition to be.
-

- Select **Use remaining space** 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 growable, the partitions will share the available free space on the disk.
- Choose **Linux RAID** from the **Partition Type** pull-down menu.
- Finally, for **Allowable Drives**, select the drive on which RAID will be created. If you have multiple drives, all drives will be selected here and you must deselect those drives which will *not* have the RAID array on them.

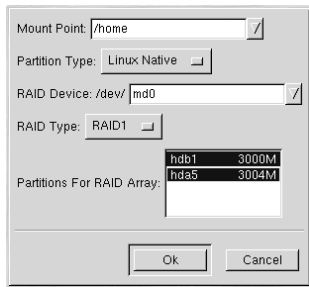
Continue these steps to create as many partitions as needed for your RAID setup. Notice that all the partitions do not have to be RAID partitions. For example, in Figure 4–2, *RAID Partitions*, only the /home partition is a software RAID device.

Figure 4–2 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 4–2, *RAID Partitions*).

Next, Figure 4–3, *Making a RAID Device* will appear, where you can make a RAID device.

Figure 4–3 Making a RAID Device

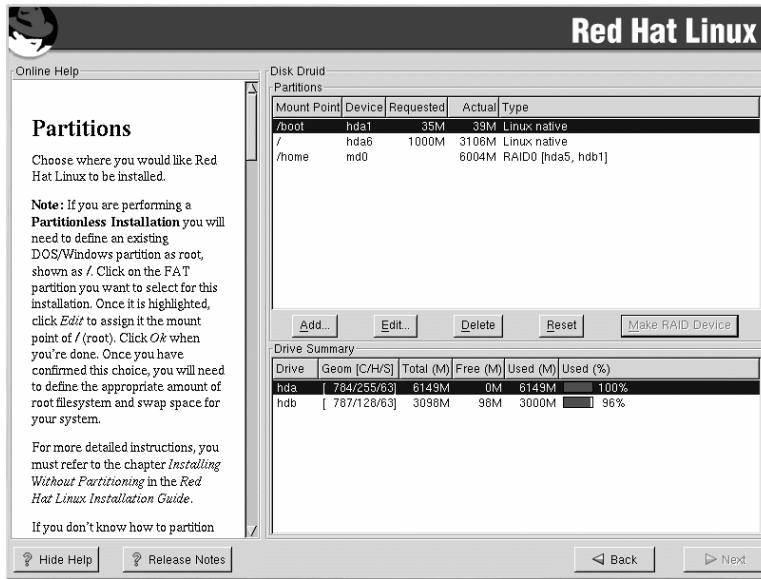
- First, enter a mount point.
- Next, choose the partition type for the partition.
- Choose your RAID device. You should choose md0 for your first device, md1 for your second device, and so on, unless you have a specific reason to make it something else. Raid devices range from md0 to md7, and each may 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 you 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 4–4, *Creating a RAID Array*) and then click **Next**.

Figure 4-4 Creating a RAID Array



- At this point, you can continue with your installation process. Refer to the *Official Red Hat Linux x86 Installation Guide* for further instructions.

Part II Network-Related References

5 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`.

The `ntsysv` utility 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 that 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/unselects 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.

The `chkconfig` command 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).

If you use `chkconfig --list` to query a service managed by `xinetd`, you'll see whether the `xinetd` service is enabled (`on`) or disabled (`off`). For example, the following command shows that `finger` is enabled as an `xinetd` service:

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

As shown above, if `xinetd` is running, `finger` is enabled.

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.

WARNING

Changes do not take effect immediately after using `ntsysv` or `chkconfig`. You must stop or start the individual service with the command `service daemon stop`. In the previous example, replace `daemon` with the name of the service you want to stop; for example, `httpd`. Replace `stop` with `start` or `restart` to start or restart the service. If you want to start or stop a service which is managed by `xinetd`, use the command `service xinetd restart`.

To control access to Internet services, you can use `xinetd`, a secure replacement for `inetd`. The `xinetd` daemon conserves system resources, provides access control and logging, and can be used to start special-purpose servers. `xinetd` can be used to provide or 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. To enable or disable a `xinetd` service, edit its configuration file in the `/etc/xinetd.d` directory. If the `disable` attribute is set to **yes**, the service is disabled. If the `disable` attribute is set to **no**, the service is enabled. If you edit any of the `xinetd` configuration files or change its enabled status using `ntsysv` or `chkconfig`, you must restart `xinetd` with the command `service xinetd restart` before the changes will take effect.

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.

5.1 Additional Resources

For more information on `xinetd`, refer to the following resources.

5.1.1 Installed Documentation

- `man xinetd` — The `xinetd` manual page.
- `man xinetd.conf` — The manual page for the `xinetd.conf` configuration file.

5.1.2 Useful Websites

- <http://www.xinetd.org> — The `xinetd` webpage. It contains the a more detailed list of features and sample configuration files.

6 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 and can be configured to 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.

7 OpenSSH

OpenSSH is a free, open source implementation of the SSH (Secure SHell) protocols. It replaces `telnet`, `ftp`, `rlogin`, `rsh`, and `rcp` with secure, encrypted network connectivity tools. OpenSSH supports versions 1.3, 1.5, and 2.0 of the SSH protocol. By default, Red Hat Linux 7.1 uses version 2.0.

7.1 Why Use OpenSSH?

If you use OpenSSH tools, you are enhancing the security of your machine. All communications using OpenSSH tools, including passwords, are encrypted. `Telnet` and `ftp` use plaintext passwords and send all information unencrypted. The information can be intercepted, the passwords can be retrieved, and then your system can be compromised by an unauthorized person logging in to your system using one of the intercepted passwords. The OpenSSH set of utilities should be used whenever possible to avoid these security problems.

Another reason to use OpenSSH is that it automatically forwards the `DISPLAY` variable to the client machine. In other words, if you are running the X Window System on your local machine, and you log in to a remote machine using the `ssh` command, when you execute a program on the remote machine that requires X, it will be displayed on your local machine. This is convenient if you prefer graphical system administration tools but do not always have physical access to your server.

7.2 Configuring an OpenSSH Server

To run an OpenSSH server, you must first make sure that you have the proper RPM packages installed. The `openssh-server` package is required and depends on the `openssh` package. Both of these packages are included in Red Hat Linux 7.1.

The OpenSSH daemon uses the configuration file `/etc/ssh/sshd_config`. The default configuration file installed with Red Hat Linux 7.1 should be sufficient for most purposes. If you want to configure the daemon in ways not provided by the default `sshd_config`, read the `sshd` manual page for a list of the keywords that can be defined in the configuration file.

To start the OpenSSH service, use the command `/sbin/service sshd start`. To stop the OpenSSH server, use the command `/sbin/service sshd stop`. If you want the daemon to start automatically at boot time, see Chapter 5, *Controlling Access to Services* for information on how to manage services.

7.3 Configuring an OpenSSH Client

To connect to an OpenSSH server from a client machine, you must have the `openssh-clients` and `openssh` packages installed on the client machine.

7.3.1 Using the `ssh` Command

The `ssh` command is a secure replacement for the `rlogin`, `rsh`, and `telnet` commands. It allows you to log in to and execute commands on a remote machine.

Logging in to a remote machine with `ssh` is similar to using `telnet`. To log in to a remote machine named `penguin.example.net`, type the following command at a shell prompt:

```
ssh penguin.example.net
```

The first time you `ssh` to a remote machine, you will see a message similar to the following:

```
The authenticity of host 'penguin.example.net' can't be established.  
DSA key fingerprint is 94:68:3a:3a:bc:f3:9a:9b:01:5d:b3:07:38:e2:11:0c.  
Are you sure you want to continue connecting (yes/no)?
```

Type **yes** to continue. This will add the server to your list of known hosts as seen in the following message:

```
Warning: Permanently added 'penguin.example.net' (DSA) to the list of known hosts.
```

Next, you'll see a prompt asking for your password for the remote machine. After entering your password, you will be at a shell prompt for the remote machine. If you use `ssh` without any command line options, the username that you are logged in as on the local client machine is passed to the remote machine. If you want to specify a different username, use the following command:

```
ssh -l username penguin.example.net
```

You can also use the syntax `ssh username@penguin.example.net`.

The `ssh` command can be used to execute a command on the remote machine without logging in to a shell prompt. The syntax is `ssh hostname command`. For example, if you want to execute the command `ls /usr/share/doc` on the remote machine `penguin.example.net`, type the following command at a shell prompt:

```
ssh penguin.example.net ls /usr/share/doc
```

After you enter the correct password, the contents of `/usr/share/doc` will be displayed, and you will return to your shell prompt.

7.3.2 Using the `scp` Command

The `scp` command can be used to transfer files between machines over a secure, encrypted connection. It is similar to `r``cp`.

The general syntax to transfer a local file to a remote system is `scp localfile username@to-hostname:/newfilename`. The *localfile* specifies the source, and the group of *username@to-hostname:/newfilename* specifies the destination.

To transfer the local file `shadowman` to your account on `penguin.example.net`, type the following at a shell prompt (replace *username* with your username):

```
scp shadowman username@penguin.example.net:/home/username
```

This will transfer the local file `shadowman` to `/home/username/shadowman` on `penguin.example.net`.

The general syntax to transfer a remote file to the local system is `scp username@to-hostname:/remotefile/newlocalfile`. The *remotefile* specifies the source, and *newlocalfile* specifies the destination.

Multiple files can be specified as the source files. For example, to transfer the contents of the directory `/downloads` to an existing directory called `uploads` on the remote machine `penguin.example.net`, type the following at a shell prompt:

```
scp /downloads/* username@penguin.example.net:/uploads/
```

7.3.3 Using the `sftp` Command

The `sftp` utility can be used to open a secure, interactive FTP session. It is similar to `ftp` except that it uses a secure, encrypted connection. The general syntax is `sftp username@hostname.com`. Once authenticated, you can use a set of commands similar to using FTP. Refer to the `sftp` manual page for a list of these commands. To read the manual page, execute the command `man sftp` at a shell prompt. The `sftp` utility is only available in OpenSSH version 2.5.0p1 and higher.

7.3.4 Generating Key Pairs

If you do not want to enter your password every time you `ssh`, `scp`, or `sftp` to a remote machine, you can generate an authorization key pair.

Separate Authorization Key Pairs

You must have separate authorization key pairs for SSH Protocol 1 (RSA) and SSH Protocol 2 (DSA).

WARNING

Keys must be generated for each user. To generate keys for a user, follow the following steps as the user who wants to connect to remote machines. If you complete the following steps as root, only root will be able to use the keys.

Generating a DSA Key Pair

Use the following steps to generate a DSA key pair. DSA is used by SSH Protocol 2 and is the default for Red Hat Linux 7.1.

1. To generate a DSA key pair to work with version 2.0 of the protocol, type the following command at a shell prompt:

```
ssh-keygen -t dsa
```

Accept the default file location of `~/.ssh/id_dsa`. Enter a passphrase different from your account password and confirm it by entering it again. ¹

What is a Passphrase?

A passphrase is a string of words and characters used to authenticate a user. Passphrases differ from passwords in that you can use spaces or tabs in the passphrase. Passphrases are generally longer than passwords because they are usually phrases instead of just a word.

2. Change the permissions of your `.ssh` directory using the command `chmod 755 ~/.ssh`.
 3. Copy the contents of `~/.ssh/id_dsa.pub` to `~/.ssh/authorized_keys2` on the machine to which you want to connect. If the file `~/.ssh/authorized_keys2` doesn't exist, you can copy the file `~/.ssh/id_dsa.pub` to the file `~/.ssh/authorized_keys2` on the other machine. ¹
-

4. If you are running GNOME, skip to *Configuring ssh-agent with GNOME* in Section 7.3.4. If you are not running the X Window System, skip to *Configuring ssh-agent* in Section 7.3.4.

Generating an RSA Key Pair for Version 2.0

Use the following steps to generate a RSA key pair for version 2.0 of the SSH protocol.

1. To generate a RSA key pair to work with version 2.0 of the protocol, type the following command at a shell prompt:

```
ssh-keygen -t rsa
```

Accept the default file location of `~/.ssh/id_rsa`. Enter a passphrase different from your account password and confirm it by entering it again. ¹

2. Change the permissions of your `.ssh` directory using the command `chmod 755 ~/.ssh`.
3. Copy the contents of `~/.ssh/id_rsa.pub` to `~/.ssh/authorized_keys2` on the machine to which you want to connect. If the file `~/.ssh/authorized_keys2` doesn't exist, you can copy the file `~/.ssh/id_rsa.pub` to the file `~/.ssh/authorized_keys2` on the other machine.¹
4. If you are running GNOME, skip to *Configuring ssh-agent with GNOME* in Section 7.3.4. If you are not running the X Window System, skip to *Configuring ssh-agent* in Section 7.3.4.

Generating an RSA Key Pair for Version 1.3 and 1.5

Use the following steps to generate an RSA key pair, which is used by version 1 of the SSH Protocol. If you are only connecting between Red Hat Linux 7.1 systems, you do not need an RSA key pair.

1. To generate an RSA (for version 1.3 and 1.5 protocol) key pair, type the following command at a shell prompt:

```
ssh-keygen
```

Accept the default file location (`~/.ssh/identity`). Enter a passphrase different from your account password. Confirm the passphrase by entering it again.

2. Change the permissions of your `.ssh` directory and your keys with the commands `chmod 755 ~/.ssh` and `chmod 644 ~/.ssh/identity.pub`.
3. Copy the contents of `~/.ssh/identity.pub` to the file `~/.ssh/authorized_keys` on the machine to which you wish to connect. If the file `~/.ssh/authorized_keys` doesn't exist, you can copy the file `~/.ssh/identity.pub` to the file `~/.ssh/authorized_keys` on the remote machine. ¹

¹ The `~` stands for the home directory of the currently logged in user. See the *Official Red Hat Linux Getting Started Guide* for more details.

4. If you are running GNOME, skip to *Configuring ssh-agent with GNOME* in Section 7.3.4. If you are not running GNOME, skip to *Configuring ssh-agent* in Section 7.3.4.

Configuring ssh-agent with GNOME

The `ssh-agent` utility can be used to save your passphrase so that you do not have to enter it each time you initiate an `ssh` or `scp` connection. If you are using GNOME, the `openssh-askpass-gnome` utility can be used to prompt you for your passphrase when you log in to GNOME and save it until you log out of GNOME. You will not have to enter your password or passphrase for any `ssh` or `scp` connection made during that GNOME session. If you are not using GNOME, refer to *Configuring ssh-agent* in Section 7.3.4.

To save your passphrase during your GNOME session, follow the following steps:

1. You'll need to have the package `openssh-askpass-gnome` installed; you can use the command `rpm -q openssh-askpass-gnome` to determine if it is installed or not. If it is not installed, install it from your Red Hat CD-ROM set, from a Red Hat FTP mirror site, or using Red Hat Network.
2. If you do not have an `~/.Xclients` file, you can run `switchdesk` to create it. In your `~/.Xclients` file, edit the following line:

```
exec $HOME/.Xclients-default
```

Change the line so that it instead reads:

```
exec /usr/bin/ssh-agent $HOME/.Xclients-default
```

3. Open the GNOME Control Center (**GNOME Main Menu Button => Programs => Settings => GNOME Control Center**) and go to **Session => Startup Programs**. Click **Add** and enter `/usr/bin/ssh-add` in the **Startup Command** text area. Set it a priority to a number higher than any existing commands to ensure that it is executed last. A good priority number for `ssh-add` is 70 or higher. The higher the priority number, the lower the priority. If you have other programs listed, this one should have the lowest priority. Click **OK** to save your settings, and exit the GNOME Control Center.
4. Log out and then log back into GNOME; in other words, restart X. After GNOME is started, a dialog box will appear prompting you for your passphrase(s). Enter the passphrase requested. If you have both DSA and RSA key pairs configured, you will be prompted for both. From this point on, you should not be prompted for a password by `ssh`, `scp`, or `sftp`.

Configuring ssh-agent

The `ssh-agent` can be used to store your passphrase so that you do not have to enter it each time you make a `ssh` or `scp` connection. If you are not running the X Window System, follow these steps from a shell prompt. If you are running GNOME but you do not want to configure it to prompt you for your passphrase when you log in (see *Configuring ssh-agent with GNOME* in Section 7.3.4), this procedure will work in a terminal window, such as an `xterm`. If you are running X but not GNOME, this procedure will work in a terminal window, such as an `xterm`. However, your passphrase will only be remembered for that terminal window; it is not a global setting.

1. At a shell prompt, type the following command:

```
exec /usr/bin/ssh-agent $SHELL
```

Then type the command

```
ssh-add
```

and enter your passphrase(s). If you have both DSA and RSA key pairs configured, you will be prompted for both.

2. When you log out, your passphrase will be forgotten. You must execute these two commands each time you log in to a virtual console or open a terminal window.

7.4 Additional Resources

The OpenSSH and OpenSSL projects are in constant development, so the most up-to-date information for them will be found on their websites. The man pages for OpenSSH and OpenSSL tools are also good sources of detailed information.

7.4.1 Installed Documentation

- The `ssh`, `scp`, `sshd`, and `ssh-keygen` commands — These man pages include information on how to use these commands as well as all the parameters that can be used with them.

7.4.2 Useful Websites

- <http://www.openssh.com> — The OpenSSH FAQ page, bug reports, mailing lists, project goals, and a more technical explanation of the security features.

- <http://www.openssl.org> — The OpenSSL FAQ page, mailing lists, and a description of the project goal.
- <http://www.freessh.org> — SSH client software for other platforms.

8 Network File System (NFS)

Network File System (NFS) is a way to share files between machines on a network as if the files were located on your local hard drive. Red Hat 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.

8.1 Why Use NFS?

NFS is useful for sharing directories of files between multiple users on the same network. For example, a group of users working on the same project can have access to the files for that project using a shared portion of the NFS filesystem (commonly known as an NFS share) mounted in the directory `/myproject`. To access the shared files, the user goes into the `/myproject` directory on his machine. There are no passwords to enter or special commands to remember. The user works as if the directory is on his local machine.

8.2 Mounting NFS Filesystems

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

```
mount shadowman:/mnt/export /mnt/local
```

Directory Must Exist

The mount point directory on local machine (`/mnt/local` in the above example) must exist.

In this command, `shadowman` is the hostname of the NFS fileserver, `/mnt/export` is the filesystem that `shadowman` 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 `shadowman`) we can enter `ls /mnt/local` and get a listing of the files in `/mnt/export` on `shadowman`.

8.2.1 Mounting NFS Filesystems using `/etc/fstab`

An alternate way to mount an NFS share from another machine is to add a line to your `/etc/fstab` file. The line must state the hostname of the NFS server, the directory on the server being exported, and the directory on the local machine where you want to mount the filesystem. You must be root to modify the `/etc/fstab` file.

The general syntax for the line in `/etc/fstab` is as follows:

```
server:/usr/local/pub /pub nfs rsize=8192,wsiz=8192,timeo=14,intr
```

The mount point `/pub` must exist on your machine. After adding this line to `/etc/fstab`, you can type the command `mount /pub` at a shell prompt, and the mount point `/pub` will be mounted from the server.

8.2.2 Mounting NFS Filesystems using autofs

A third option for mounting an NFS share is the use of autofs. Autofs uses the automount daemon to manage your mount points by only mounting them dynamically when they are accessed.

Autofs consults the master map configuration file `/etc/auto.master` to determine which mount points are defined. It then starts an automount process with the appropriate parameters for each mount point. Each line in the master map defines a mount point and a separate map file that defines the filesystems to be mounted under this mount point. For example, the `/etc/auto.mnt` file might define mount points in the `/mnt` directory; this relationship would be defined in the `/etc/auto.master` file.

Each entry in `auto.master` has three fields. The first field is the mount point. The second field is the location of the map file, and the third field is optional. The third field can contain information such as a timeout value.

For example, to mount the directory `/project52` on the remote machine `penguin.host.net` at the mount point `/mnt/myproject` on your machine, add the following line to `auto.master`:

```
/mnt /etc/auto.mnt --timeout 60
```

Add the following line to `/etc/auto.mnt`:

```
myproject -rw,soft,intr,rsize=8192,wsiz=8192 penguin.host.net:/project52
```

The first field in `/etc/auto.mnt` is the name of the `/mnt` subdirectory. This directory is created dynamically by automount. It should not actually exist on the client machine. The second field contains mount options such as `rw` for read and write access. The third field is the location of the NFS export including the hostname and directory.

The directory `/mnt` must exist on the local filesystem. There should be no subdirectories to `/mnt` on the local filesystem.

Autofs is a service. To start the service, at a shell prompt, type the following commands:

```
service autofs restart
```

To view the active mount points, type the following command at a shell prompt:


```
service autofs status
```

If you modify the `/etc/auto.master` configuration file while `autofs` is running, you must tell the automount daemon(s) to reload by typing the following command at a shell prompt:

```
service autofs reload
```

To learn how to configure `autofs` to start at boot time, refer to Chapter 5, *Controlling Access to Services* for information on managing services.

8.3 Exporting NFS Filesystems

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

```
directory      hostname(options)
```

The `(options)` are not required. For example:

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

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

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

would 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 command will also work:

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

8.4 Additional Resources

This chapter discusses the basics of using NFS. For more detailed information, refer to the following resources.

8.4.1 Installed Documentation

- `nfsd(8)`, `mountd(8)`, `exports(5)`, `auto.master(5)`, `autofs(5)`, and `autofs(8)` man pages — These man pages show the correct syntax for the NFS and autofs configuration files.

8.4.2 Related Books

- *Managing NFS and NIS Services* by Hal Stern; O'Reilly & Associates, Inc.

9 Samba

Samba uses the SMB protocol to share files and printers across a network connection. Operating systems that support this protocol include Microsoft Windows (through its Network Neighborhood), OS/2, and Linux.

9.1 Why Use Samba?

Samba is useful if you have a network of both Windows and Linux machines. Samba will allow files and printers to be shared by all the systems in your network. If you want to share files between Red Hat Linux machines only, refer to Chapter 8, *Network File System (NFS)*. If you want to share printers between Red Hat Linux machines only refer to Chapter 13, *Printer Configuration*.

9.2 Configuring Samba

Samba uses `/etc/samba/smb.conf` as its configuration file. If you change this configuration file, the changes will not take effect until you restart the Samba daemon with the command `service smb restart`.

The default configuration file (`smb.conf`) in Red Hat Linux 7.1 allows users to view their Linux home directories as a Samba share on the Windows machine after they log in using the same username and password. It also shares any printers configured for the Red Hat Linux system as Samba shared printers. In other words, you can attach a printer to your Red Hat Linux system and print to it from the Windows machines on your network.

To specify the Windows workgroup and description string, edit the following lines in your `smb.conf` file:

```
workgroup = WORKGROUPNAME
server string = BRIEF COMMENT ABOUT SERVER
```

Replace `WORKGROUPNAME` with the name of the Windows workgroup to which this machine should belong. The `BRIEF COMMENT ABOUT SERVER` is optional and will be the Windows comment about the Samba system.

To create a Samba share directory on your Linux system, add the following section to your `smb.conf` file (after modifying it to reflect your needs and your system):

```
[sharename]
comment = Insert a comment here
path = /home/share/
valid users = tfox carole
public = no
writable = yes
```

```
printable = no
create mask = 0765
```

The above example allows the users `tfox` and `carole` to read and write to the directory `/home/share`, on the Samba server, from a Samba client.

9.3 Connecting to a Samba Share

To connect to a Linux Samba share from a Microsoft Windows machine, use Network Neighborhood or Windows Explorer.

To connect to a Samba share from a Linux system, from a shell prompt, type the following command:

```
smbclient //hostname/sharename -U username
```

You will need to replace *hostname* with the hostname or IP address of the Samba server you want to connect to, *sharename* with the name of the shared directory you want to browse, and *username* with the Samba username for the system. Enter the correct password or press [Enter] if no password is required for the user.

If you see the `smb: \>` prompt, you have successfully logged in. Once you are logged in, type **help** for a list of commands. If you wish to browse the contents of your home directory, replace *sharename* with your username. If the `-U` switch is not used, the username of the current user is passed to the Samba server.

To exit `smbclient`, type **exit** at the `smb: \>` prompt.

9.4 Using Samba with Windows NT 4.0 and Windows 2000

The Microsoft SMB Protocol originally used plaintext passwords. However, Windows 2000 and Windows NT 4.0 with Service Pack 3 or higher require encrypted Samba passwords. To use Samba between a Red Hat Linux system and a system with Windows 2000 or Windows NT 4.0 Service Pack 3 or higher, you can either edit your Windows registry to use plaintext passwords or configure Samba on your Linux system to use encrypted passwords. If you choose to modify your registry, you must do so for all your Windows NT or 2000 machines — this is risky and may cause further conflicts.

To configure Samba on your Red Hat Linux system to use encrypted passwords, follow these steps:

1. Create a separate password file for Samba. To create one based on your existing `/etc/passwd` file, at a shell prompt, type the following command:

```
cat /etc/passwd | mksmbpasswd.sh > /etc/samba/smbpasswd
```

The `mksmbpasswd.sh` script is installed in your `/usr/bin` directory with the `samba` package.

2. Use the command `chmod 600 /etc/samba/smbpasswd` to change permissions on the Samba password file so that only root has read and write permissions.
3. The script does not copy user passwords to the new file. To set each Samba user's password, use the command `smbpasswd username` (replace *username* with each user's username). A Samba user account will not be active until a Samba password is set for it.
4. The next step is to enable encrypted passwords in the Samba configuration file. In the file `smb.conf`, uncomment the following lines:

```
encrypt password = yes
smb passwd file = /etc/samba/smbpasswd
```

5. To have the changes take effect, restart Samba by typing the command `service smb restart` at a shell prompt.

Additional Information

To read more about *Using Samba with Windows NT 4.0 and Windows 2000*, read `ENCRYPTION.txt`, `Win95.txt`, and `WinNT.txt` in the directory `/usr/share/doc/samba-version-number/docs/textdocs/` (replace *version-number* with the version-number of Samba that you have installed).

9.5 Additional Resources

For configuration options not covered here, please refer to the following resources.

9.5.1 Installed Documentation

- `smb.conf` man page — explains how to configure the Samba configuration file
 - `smbd` man page — describes how the Samba daemon works
 - `/usr/share/doc/samba-version-number/docs/` — HTML and text help files included with the samba package
-

9.5.2 Useful Websites

- <http://www.samba.org> — The Samba Web page contains useful documentation, information about mailing lists, and a list of GUI interfaces.

Part III System Configuration

10 Gathering System Information

Before you learn how to configure your system, you should learn how to gather essential system information. For example, you should know how to find the amount of free memory, how your hard drive is partitioned, and what processes are running. This chapter discusses how to retrieve this type of information from your Red Hat Linux system using simple commands and a few simple programs.

10.1 System Processes

The `ps aux` command displays a list of current system processes, including processes owned by other users. To display the owner of the processes along with the processes use the command `ps aux`. This list is a static list; in other words, it is a snapshot of what is running when you invoked the command. If you want a constantly updated list of running processes, use `top` as described below.

You can use the `ps` command in combination with the `grep` command to see if a process is running. For example, to determine if Netscape is still running, use the command `ps aux | grep netscape`.

The `top` command displays currently running processes and important information about them including their memory and CPU usage. The list is both real-time and interactive. An example of `top`'s output is provided as follows:

```

6:14pm up 2 days, 19:29, 5 users, load average: 0.10, 0.06, 0.07
71 processes: 68 sleeping, 2 running, 1 zombie, 0 stopped
CPU states: 2.7% user, 0.5% system, 0.0% nice, 96.6% idle
Mem: 256812K av, 252016K used, 4796K free, 97228K shrd, 43300K buff
Swap: 265032K av, 1328K used, 263704K free, 86180K cached

  PID USER      PRI  NI  SIZE  RSS SHARE STAT  %CPU  %MEM   TIME COMMAND
 15775 joe         5   0 11028  10M  3192 S    1.5   4.2   0:46 emacs
 14429 root        15   0 63620  62M  3284 R    0.5  24.7  63:33 X
 17372 joe        11   0  1056  1056   840 R    0.5   0.4   0:00 top
 17356 joe         2   0  4104  4104  3244 S    0.3   1.5   0:00 gnome-terminal
 14461 joe         1   0  3584  3584  2104 S    0.1   1.3   0:17 sawfish
    1 root         0   0    544   544   476 S    0.0   0.2   0:06 init
    2 root         0   0     0     0     0 SW    0.0   0.0   0:00 kflushd
    3 root         1   0     0     0     0 SW    0.0   0.0   0:24 kupdate
    4 root         0   0     0     0     0 SW    0.0   0.0   0:00 kpiod
    5 root         0   0     0     0     0 SW    0.0   0.0   0:29 kswapd
   347 root         0   0    556   556   460 S    0.0   0.2   0:00 syslogd
   357 root         0   0    712   712   360 S    0.0   0.2   0:00 klogd
   372 bin          0   0    692   692   584 S    0.0   0.2   0:00 portmap
   388 root         0   0     0     0     0 SW    0.0   0.0   0:00 lockd

```

```

389 root      0  0      0   0      0 SW    0.0  0.0   0:00 rpciod
414 root      0  0    436  432   372 S    0.0  0.1   0:00 apmd
476 root      0  0    592  592   496 S    0.0  0.2   0:00 automount

```

To exit `top`, press the [q] key.

Useful interactive commands that you can use with `top` include the following:

Table 10–1 Interactive `top` commands

Command	Description
[Space]	Immediately refresh the display
[h]	Display a help screen
[k]	Kill a process. You will be prompted for the process ID and the signal to send to it.
[n]	Change the number of processes displayed. You will be prompted to enter the number.
[u]	Sort by user.
[M]	Sort by memory usage.
[P]	Sort by CPU usage.

If you would like to use a graphical interface for `top`, you can use **GNOME System Monitor**. To start it, go to the **GNOME Main Menu Button => Programs => System => System Monitor** or type `gtop` at a shell prompt.

Figure 10–1 GNOME System Monitor

PID	User	Pri	Size	Resident	Stat	CPU	MEM	Time	Cmd
14429	root	9	63620	63620	S	1.3	8.7	1:03h	/etc/X11/X
14444	tfox	0	3064	3064	S	0.0	1.1	5:04m	magicdev
16154	tfox	19	4388	4388	R	4.6	1.7	3:38m	gtop
14501	tfox	0	41132	41132	S	0.0	15.9	1:15m	/usr/lib/netscape/ne
15775	tfox	0	11008	11008	S	0.0	4.2	43.97s	emacs
5	root	0	0	0	SU	0.0	0.0	29.58s	kswapd
3	root	0	0	0	SU	0.0	0.0	24.81s	kupdate
14837	tfox	0	6772	6772	S	0.0	2.6	23.69s	emacs
14461	tfox	1	3584	3584	S	0.0	1.3	16.52s	sawfish
659	root	0	364	356	S	0.0	0.1	13.39s	gpm
14466	tfox	0	6292	6292	S	0.0	2.4	13.15s	panel
1	root	0	544	544	S	0.0	0.2	6.49s	init
14475	tfox	0	4416	4416	S	0.0	1.7	6.33s	tasklist_applet
14477	tfox	0	3932	3932	S	0.0	1.5	5.89s	deskguide_applet
16151	tfox	3	21132	21132	S	1.9	8.2	5.45s	gimp
14462	tfox	1	1840	1840	S	0.6	0.7	4.59s	xscreensaver
817	xfs	0	5532	4780	S	0.0	1.8	4.52s	xfs
14483	tfox	0	4244	4244	S	0.0	1.6	2.28s	gnome-terminal
15800	tfox	0	4096	4096	S	0.0	1.5	1.46s	gnome-terminal
14442	tfox	0	2124	2124	S	0.0	0.8	1.43s	gnome-smpoxy

poofy.chinfox CPU: 8.37% user, 7.57% system 5:44pm, up 2 days loadavg: 0.08, 0.15, 0.25

10.2 Memory Usage

The `free` command displays the total amount of physical memory and swap space for the system as well as the amount of memory that are used, free, shared, in kernel buffers, and cached.

```

                total          used          free          shared    buffers     cached
Mem:            256812        240668        16144         105176     50520      81848
-/+ buffers/cache: 108300        148512
Swap:           265032           780         264252

```

The command `free -m` shows the same information in megabytes, which are easier to read.

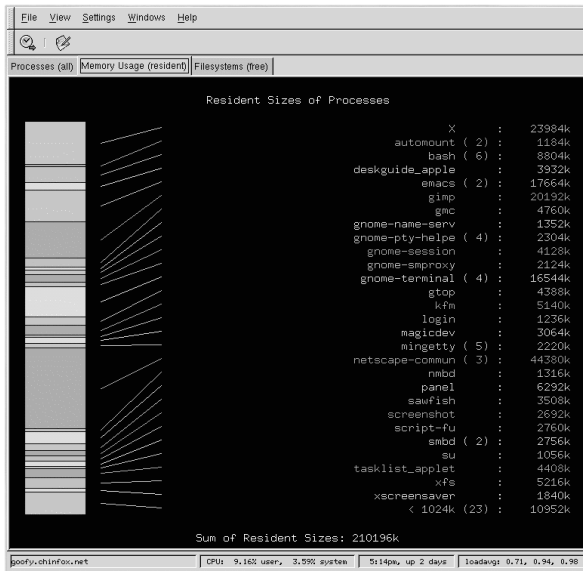
```

                total          used          free          shared    buffers     cached
Mem:              250           235           15           102           49           79
-/+ buffers/cache: 105           145
Swap:             258             0           258

```

If you would like to use a graphical interface with `free`, you can use GNOME System Monitor. To start it, go to the **GNOME Main Menu Button** => **Programs** => **System** => **System Monitor** or type `gtop` at a shell prompt. Then choose the **Memory Usage** tab.

Figure 10–2 GNOME System Monitor



10.3 Filesystems

The `df` command reports the system's disk space usage. If you type the command `df` at a shell prompt, the output looks similar to the following:

Filesystem	1k-blocks	Used	Available	Use%	Mounted on
/dev/hda2	10325716	2902060	6899140	30%	/
/dev/hda1	15554	8656	6095	59%	/boot
/dev/hda3	20722644	2664256	17005732	14%	/home

By default, this utility shows the partition size in 1 kilobyte blocks and the amount of used and available disk space in kilobytes. To view the information in megabytes and gigabytes, use the command `df -h`. The `-h` argument stands for human-readable format. The output looks similar to the following:

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/hda2	9.8G	2.8G	6.5G	30%	/
/dev/hda1	15M	8.5M	5.9M	59%	/boot
/dev/hda3	20G	2.6G	16G	14%	/home

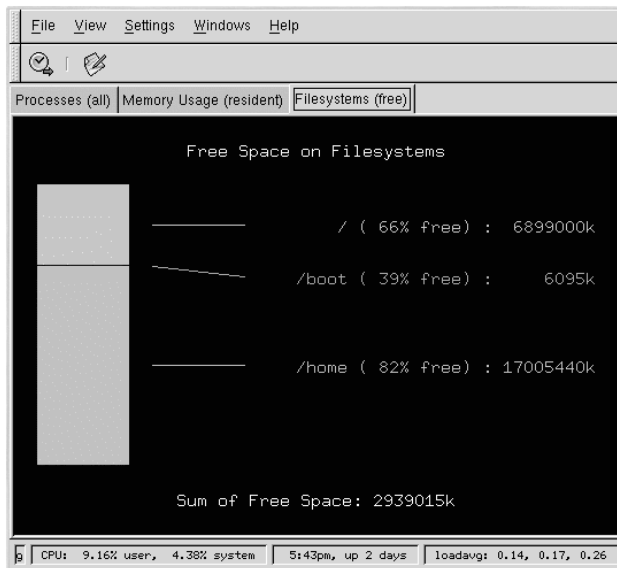
To view the system's disk space usage in a graphical format, use **GDiskFree**. To start it, go to the **GNOME Main Menu Button => Programs => System => GNOME Disk Free** or type the command `gdiskfree` at a shell prompt. This utility displays all mounted file systems and their disk usage using a dial diagram.

Figure 10–3 GDiskFree



You can also choose the **Filesystems** tab in the **GNOME System Monitor**. To start it, go to the **GNOME Main Menu Button => Programs => System => System Monitor** or type `gtop` at a shell prompt. Then choose the **Filesystems** tab.

Figure 10–4 GNOME System Monitor



The `du` command displays the estimated amount of space being used by files in a directory. If you type `du` at a shell prompt, the disk usage for each of the subdirectories will be displayed in a list. The grand total for the current directory and subdirectories will also be shown, as the last line in the list. If you do not want to see all the subdirectories, use the command `du -hs` to see only the grand total for the directory in human-readable format. Use the `du --help` command to see more options.

10.4 Sysreport

`Sysreport` is a system utility created to collect important system data, in order to assist the Red Hat Technical Support and Development Teams in solving customer problems. `Sysreport` gathers as much system information as is possible, while avoiding certain actions: the creation of a very large file; the invasion of the user's privacy; and the collection of information that could be detrimental to the integrity of the system.

To start `Sysreport`, you must be logged in as root. As root, at a shell prompt type the command `sysreport`.

You will then see the following message:

```
This utility will go through and collect some detailed information
about the hardware and setup of your Red Hat Linux system.
This information will be used to diagnose problems with your system,
and will be considered confidential information. Red Hat will use
this information for diagnostic purposes ONLY.
```

```
Please wait while we collect information about your system.
```

```
This process may take awhile to complete...
No changes will be made to your system during this process.
```

```
NOTE: You can safely ignore a failed message.This only means a file
we were checking for did not exist.
```

```
Press ENTER to continue, or CTRL-C to quit.
```

As the message says, ignore any failed messages. `Sysreport` checks for all possible Red Hat Linux system packages. If you do not have every Red Hat Linux package installed, you will see failed messages.

After pressing [Enter], `Sysreport` will gather information about your system's configuration. When it is finished, you will see the following message:

```
Enter your first initial and last name with no spaces (example: jsmith):
```

Enter the information requested, and press [Enter]. `Sysreport` will place a compressed TAR file in the `/tmp` directory beginning with the initial and last name you just entered. You will see a message

telling you to email this file to the Red Hat support team. However, even if you do not need support, you can use this information to back up most of your system's configuration.

Use the command `tar ztvf filename` with the name of the compressed TAR file that you created to display a list of its contents.

10.5 Additional Resources

To learn more about gathering system information, refer to the following resources.

10.5.1 Installed Documentation

- `ps --help` — The `ps --help` displays a list of options that can be used with `ps`.
- `top` manual page — Type `man top` to learn more about `top` and its many options.
- `free` manual page — type `man free` to learn more about `free` and its many options.
- `df` manual page — Type `man df` to learn more about the `df` command and its many options.
- `du` manual page — Type `man du` to learn more about the `du` command and its many options.
- `/proc` — The contents of the `/proc` directory can also be used to gather more detailed system information. Refer to the *Official Red Hat Linux Reference Guide* for additional information about the `/proc` directory.

10.5.2 Useful Websites

- <http://www.ibiblio.org/shadow/sysreport/> — The Sysreport Web page provides the latest version and instructions.
-

11 Apache Configuration

Apache Configuration Tool requires the X Window System and root access. To start Apache Configuration Tool, use one of the following methods:

- On the GNOME desktop, go to the **Main Menu Button** (on the Panel) => **Programs** => **System** => **Apache Configuration**.
- On the KDE desktop, go to the **Main Menu Button** (on the Panel) => **Red Hat** => **System** => **Apache Configuration**.
- Type the command `apacheconf` at a shell prompt (for example, in an XTerm or GNOME-terminal).

Do Not Edit `httpd.conf`

Do not edit the `/etc/httpd/conf/httpd.conf` Apache configuration file if you wish to use this tool. Apache Configuration Tool generates this file after you save your changes and exit the program. If you want to add additional modules or configuration options that are not available in Apache Configuration Tool, you cannot use this tool.

Apache Configuration Tool allows you to configure the `/etc/httpd/conf/httpd.conf` configuration file for your Apache Web server. It does not use the old `srml.conf` or `access.conf` configuration files; leave them empty. Through the graphical interface, you can configure Apache directives such as virtual hosts, logging attributes, and maximum number of connections.

Only modules that are shipped with Red Hat Linux can be configured with Apache Configuration Tool. If additional modules are installed, they can not be configured using this tool.

The general steps for configuring the Apache Web Server using the Apache Configuration Tool are as following:


1. Configure the basic settings under the **Main** tab.
 2. Click on the **Virtual Hosts** tab and configure the default settings.
 3. Under the **Virtual Hosts** tab, configure the Default Virtual Host.
 4. If you want to serve more than one URL or virtual host, add the additional virtual hosts.
 5. Configure the server settings under the **Server** tab.
 6. Configure the connections settings under the **Performance Tuning** tab.
-

7. Copy all necessary files to the DocumentRoot and cgi-bin directories, and save your settings in the Apache Configuration Tool.

11.1 Basic Settings

Use the **Main** tab to configure the basic server settings.

Figure 11–1 Basic Settings




Enter a fully qualified domain name that you have the right to use in the **Server Name** text area. This option corresponds to the `ServerName` directive in `httpd.conf`. The `ServerName` directive sets the hostname of the Web server. It is used when creating redirection URLs. If you do not define a `Server Name`, Apache attempts to resolve it from the IP address of the system. The `Server Name` does not have to be the domain name resolved from the IP address of the server. For example, you might want to set the `Server Name` to `www.your_domain.com` when your server's real DNS name is actually `foo.your_domain.com`.

Enter the email address of the person who maintains the Web server in the **Webmaster email address** text area. This option corresponds to the `ServerAdmin` directive in `httpd.conf`. If you configure the server's error pages to contain an email address, this email address will be used so that users can report a problem by sending email to the server's administrator. The default value is `root@localhost`.

Use the **Available Addresses** area to define the ports on which Apache will accept incoming requests. This option corresponds to the `Listen` directive in `httpd.conf`. By default, Red Hat configures Apache to listen to ports 80 and 8080 for non-secure Web communications. Click the **Add** button to define additional ports on which to accept requests. A window as shown in Figure 11–2, *Available Addresses* will appear. Either choose the **Listen to all addresses** option to listen to all IP addresses on the defined port or specify a particular IP address over which the server will accept connections in the **Address** field. Only specify one IP address per port number. If you want to specify more than one IP address with the same port number, create an entry for each IP address. If at all possible, use an IP address instead of a domain name to prevent a DNS lookup failure. Refer to <http://httpd.apache.org/docs/dns-caveats.html> for more information about *Issues Regarding DNS and Apache*. Entering an asterisk (*) in the **Address** field is the same as choosing **Listen to all addresses**. Clicking the **Edit** button shows the same window as the **Add** button except with the fields populated for the selected entry. To delete an entry, select it and click the **Delete** button.

Figure 11–2 Available Addresses



Tip

If you set Apache to listen to a port under 1024, you must be root to start it. For port 1024 and above, `httpd` can be started as a regular user.

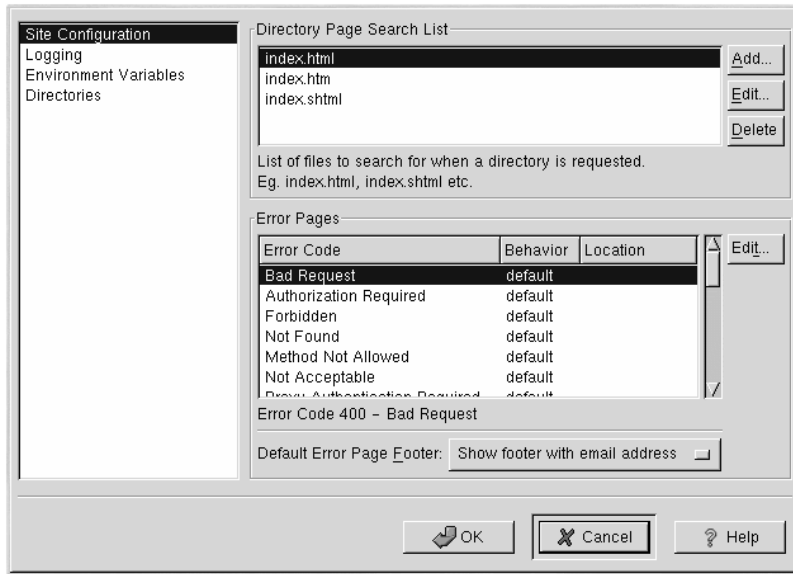
11.2 Default Settings

After defining the Server Name, Webmaster email address, and Available Addresses, click the **Virtual Hosts** tab and click the **Edit Default Settings** button. The window shown in Figure 11–3, *Site Configuration* will appear. Configure the default settings for your Web server in this window. If you add a virtual host, the settings you configure for the virtual host take precedence for that virtual host. For a directive not defined within the virtual host settings, the default value is used.

11.2.1 Site Configuration

The default values for the **Directory Page Search List** and **Error Pages** will work for most servers. If you are unsure of these settings, do not modify them.

Figure 11–3 Site Configuration



The entries listed in the **Directory Page Search List** define the `DirectoryIndex` directive. 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 list. The server will try to find one of the files listed in the `DirectoryIndex` directive 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 list, in HTML format, of the subdirectories and files in the directory.

Use the **Error Code** section to configure Apache to redirect the client to a local or external URL if the event of a problem or error. This option corresponds to the `ErrorDocument` directive. If a problem or error occurs when a client tries to connect to the Apache Web server, the default action is to display the short error message shown in the **Error Code** column. To override this default configuration, select the error code and click the **Edit** button. Choose **Default** to display the default short error message. Choose **URL** to redirect the client to an external URL and enter a complete URL including the `http://`

in the **Location** field. Choose **File** to redirect the client to an internal URL and enter a file under the Document Root for the Web server. The location must begin the a slash (/) and be relative to the Document Root.

For example, to redirect a 404 Not Found error code to a Web page that you created in a file called 404.html, copy 404.html to *DocumentRoot/errors/404.html*. In this case, *DocumentRoot* is the Document Root directory that you have defined (the default is */var/www/html*). Then, choose **File** as the Behavior for **404 - Not Found** error code and enter */errors/404.html* as the **Location**.


From the **Default Error Page Footer** menu, you can choose one of the following options:

- **Show footer with email address** — Display the default Apache footer at the bottom of all error pages along with the email address of the website maintainer specified by the `ServerAdmin` directive. Refer to *General Options* in Section 11.3.1 for information about configuring the `ServerAdmin` directive.
- **Show footer** — Display just the default Apache footer at the bottom of error pages.
- **No footer** — Do not display a footer at the bottom of error pages.

11.2.2 Logging

By default, Apache writes the transfer log to the file */var/log/httpd/access_log* and the error log to the file */var/log/httpd/error_log*.

Figure 11–4 Logging



The transfer log contains a list of all attempts to access the Web server. It records the IP address of the client that is attempting to connect, the date and time of the attempt, and the file on the Web server that it is trying to retrieve. Enter the name of the path and file in which to store this information. If the path and filename does not start with a slash (/), the path is relative to the server root directory as configured. This option corresponds to the `TransferLog` directive.

You can configure a custom log format by checking **Use custom logging facilities** and entering a custom log string in the **Custom Log String** field. This configures the `LogFormat` directive. Refer to http://httpd.apache.org/docs/mod/mod_log_config.html#formats for details on the format of this directive.

The error log contains a list of any server errors that occur. Enter the name of the path and file in which to store this information. If the path and filename does not start with a slash (/), the path is relative to the server root directory as configured. This option corresponds to the `ErrorLog` directive.

Use the **Log Level** menu to set how verbose the error messages in the error logs will be. It can be set (from least verbose to most verbose) to `emerg`, `alert`, `crit`, `error`, `warn`, `notice`, `info` or `debug`. This option corresponds to the `LogLevel` directive.

The value chosen with the **Reverse DNS Lookup** menu defines the `HostnameLookups` directive. Choosing **No Reverse Lookup** sets the value to off. Choosing **Reverse Lookup** sets the value to on. Choosing **Double Reverse Lookup** sets the value to double.

If you choose **Reverse Lookup**, 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.

If you choose **Double Reverse Lookup**, your server will perform a double-reverse DNS. In other words, after a reverse lookup is performed, a forward lookup is performed on the result. At least one of the IP addresses in the forward lookup must match the address from the first reverse lookup.


Generally, you should leave this option set to **No Reverse Lookup**, because the DNS requests add a load to your server and may slow it down. If your server is busy, the effects of trying to perform these reverse lookups or double reverse lookups may be quite noticeable.

Reverse lookups and double reverse lookups 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 Internet's benefit, you should leave this option set to **No Reverse Lookup**.

11.2.3 Environment Variables

Apache can use the `mod_env` module to configure the environment variables which are passed to CGI scripts and SSI pages. Use the **Environment Variables** page to configure the directives for this Apache module.

Figure 11–5 Environment Variables



Use the **Set for CGI Scripts** section to set an environment variable that is passed to CGI scripts and SSI pages. For example, to set the environment variable `MAXNUM` to 50, click the **Add** button inside the **Set for CGI Script** section as shown in Section 11.2.3, *Environment Variables* and type **MAXNUM** in the **Environment Variable** text field and **50** in the **Value to set** text field. Click **OK**. The **Set for CGI Scripts** section configures the `SetEnv` directive.

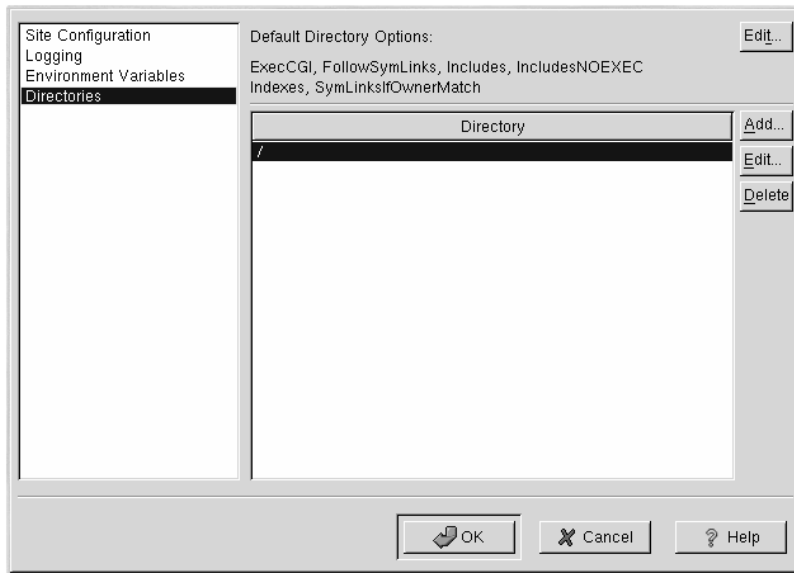
Use the **Pass to CGI Scripts** section to pass the value of an environment variable when Apache was first started to CGI scripts. To see this environment variable, type the command `env` at a shell prompt. Click the **Add** button inside the **Pass to CGI Scripts** section and enter the name of the environment variable in the resulting dialog box. Click **OK**. The **Pass to CGI Scripts** section configures the `PassEnv` directive.

If you want to remove an environment variable so that the value is not passed to CGI scripts and SSI pages, use the **Unset for CGI Scripts** section. Click **Add** in the **Unset for CGI Scripts** section, and enter the name of the environment variable to unset. This corresponds to the `UnsetEnv` directive.

11.2.4 Directories

Use the **Directories** page to configure options for specific directories. This corresponds to the `<Directory>` directive.

Figure 11–6 Directories



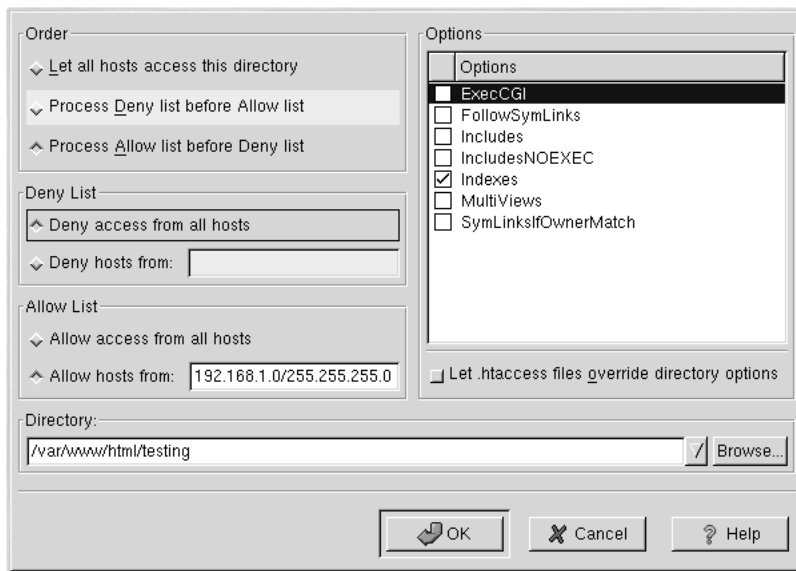
Click the **Edit** button in the top right-hand corner to configure the **Default Directory Options** for all directories that are not specified in the **Directory** list below it. The options that you choose are listed as the `Options` directive within the `<Directory>` directive. You can configure the following options:

- **ExecCGI** — Allow execution of CGI scripts. CGI scripts are not executed if this option is not chosen.
- **FollowSymLinks** — Allow symbolic links to be followed.
- **Includes** — Allow server-side includes.
- **IncludesNOEXEC** — Allow server-side includes, but disable the `#exec` and `#include` commands in CGI scripts.
- **Indexes** — Display a formatted list of the directory's contents, if no `DirectoryIndex` (such as `index.html`) exists in the requested directory.
- **Multiview** — Support content-negotiated multiviews; this option is disabled by default.
- **SymLinksIfOwnerMatch** — Only follow symbolic links if the target file or directory has the same owner as the link.

To specify options for specific directories, click the **Add** button beside the **Directory** list box. The window shown in Figure 11–7, *Directory Settings* appears. Enter the directory to configure in the **Directory** text field at the bottom of the window. Select the options in the right-hand list, and configure the **Order** directive with the left-hand side options. The **Order** directive controls the order in which allow and deny directives are evaluated. In the **Allow hosts from** and **Deny hosts from** text field, you can specify one of the following:

- Allow all hosts — Type **a11** to allow access to all hosts.
- Partial domain name — Allow all hosts whose names match or end with the specified string.
- Full IP address — Allow access to a specific IP address.
- A subnet — Such as **192.168.1.0/255.255.255.0**
- A network CIDR specification — such as **10.3.0.0/16**

Figure 11–7 Directory Settings



If you check the **Let .htaccess files override directory options**, the configuration directives in the `.htaccess` file take precedence.


11.3 Virtual Hosts Settings

You can use Apache Configuration Tool to configure virtual hosts. Virtual hosts allow you to run different servers for different IP addresses, different host names, or different ports on the same machine. For example, you can run the website for `http://www.your_domain.com` and `http://www.your_second_domain.com` on the same Apache server using virtual hosts. This option corresponds to the `<VirtualHost>` directive for the default virtual host and IP based virtual hosts. It corresponds to the `<NameVirtualHost>` directive for a name based virtual host.

The Apache directives set for a virtual host only apply to that particular virtual host. If a directive is set server-wide using the **Edit Default Settings** button and not defined within the virtual host settings, the default setting is used. For example, you can define a **Webmaster email address** in the **Main** tab and not define individual email addresses for each virtual host.

Apache Configuration Tool includes a default virtual host as shown in Figure 11–8, *Virtual Hosts*. Refer to **Default Virtual Host** in Section 11.3.1 for details about the default virtual host.

Figure 11–8 Virtual Hosts

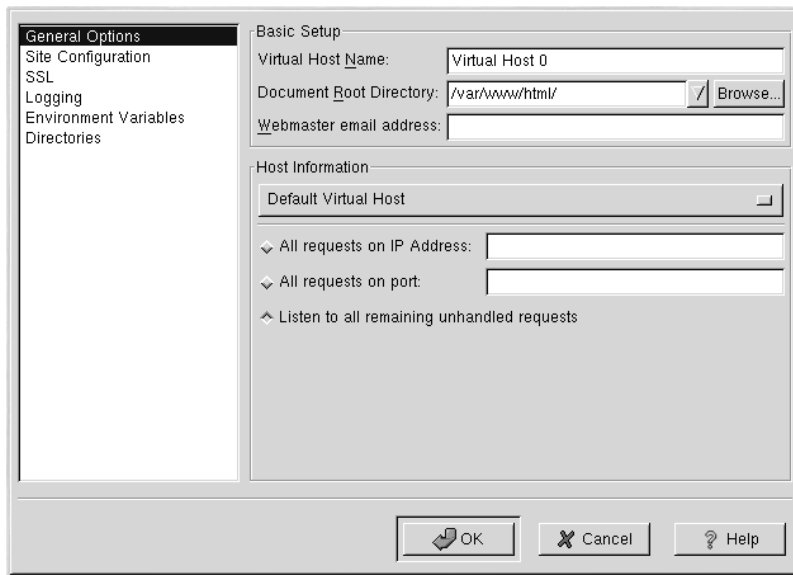


The Apache documentation on your machine or on the Web at <http://www.apache.org/docs/vhosts/> provides more information about virtual hosts.

11.3.1 Adding and Editing a Virtual Host

To add a virtual host, click the **Virtual Hosts** tab and then click the **Add** button. The window as shown in Figure 11–9, *Virtual Hosts Configuration* appears. You can also edit a virtual host by selecting it in the list and clicking the **Edit** button.

Figure 11–9 Virtual Hosts Configuration



General Options

The **General Options** settings only apply to the virtual host that you are configuring. Set the name of the Virtual Host in the **Virtual Host Name** text area. This name is used by Apache Configuration Tool to distinguish between virtual hosts.

Set the **Document Root Directory** value to the directory that contains the root document (such as index.html) for the virtual host. This option corresponds to the `DocumentRoot` directive within the `VirtualHost` directive.

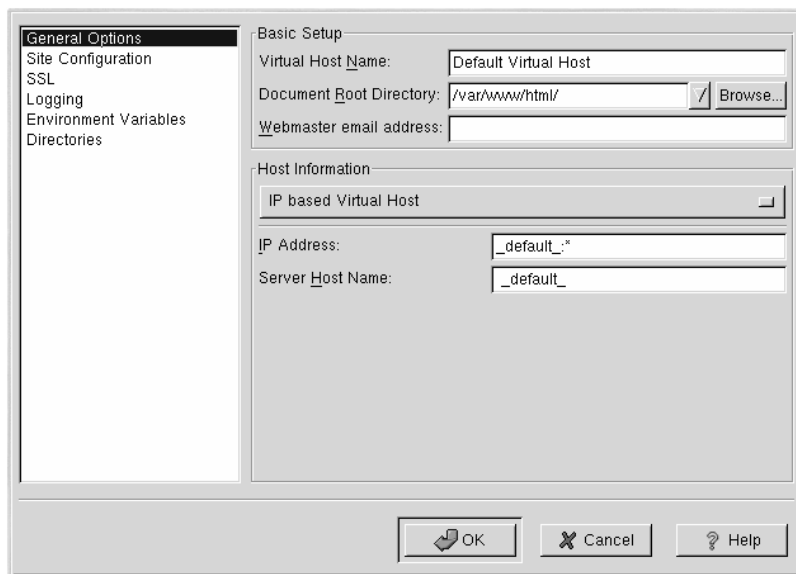
The **Webmaster email address** corresponds to the `ServerAdmin` directive within the `VirtualHost` directive. This email address is used in the footer of error pages if you choose to show a footer with an email address on the error pages.

In the **Host Information** section, choose **Default Virtual Host**, **IP based Virtual Host**, or **Name based Virtual Host**.

Default Virtual Host

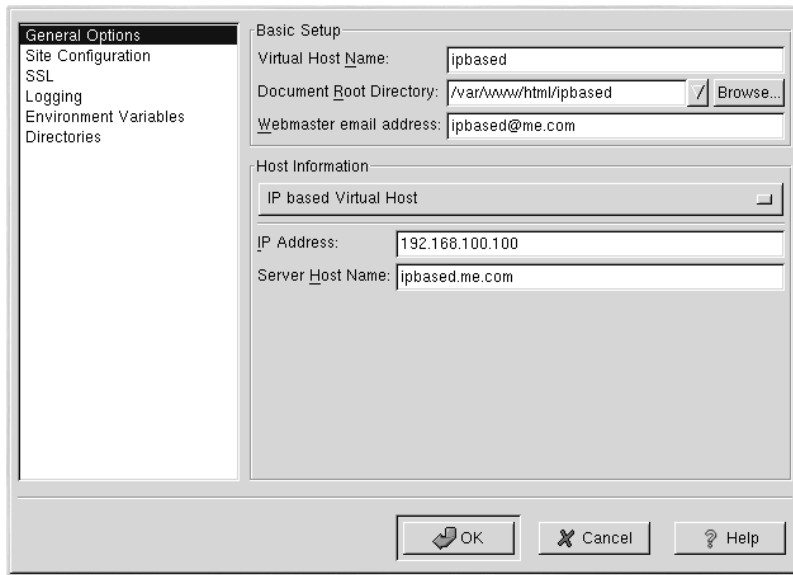
If you choose **Default Virtual Host**, Figure 11–10, *Default Virtual Hosts* appears. You should only configure one default virtual host. The default virtual host settings are used when the requested IP address is not explicitly listed in another virtual host. If there is no default virtual host defined, the main server settings are used.

Figure 11–10 Default Virtual Hosts



IP based Virtual Host

If you choose **IP based Virtual Host**, Figure 11–11, *IP Based Virtual Hosts* appears to configure the `<VirtualHost>` directive based on the IP address of the server. Specify this IP address in the **IP address** field. To specify more than one IP address, separate each IP address with spaces. To specify a port, use the syntax *IP Address:Port*. Use `:*` to configure all ports for the IP address. Specify the host name for the virtual host in the **Server Host Name** field.

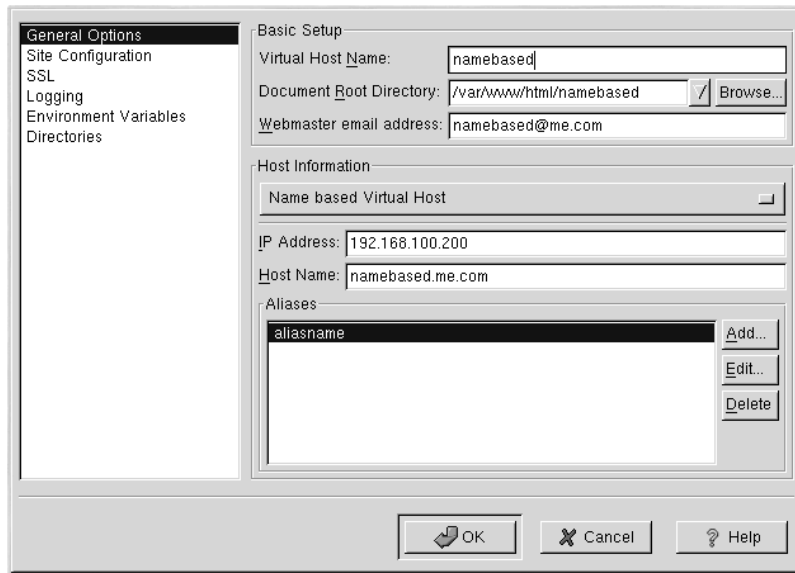
Figure 11–11 IP Based Virtual Hosts

The screenshot shows a configuration window with a sidebar on the left containing a tree view with the following items: General Options (selected), Site Configuration, SSL, Logging, Environment Variables, and Directories. The main area is divided into two sections: 'Basic Setup' and 'Host Information'. In the 'Basic Setup' section, there are three input fields: 'Virtual Host Name' with the value 'ipbased', 'Document Root Directory' with the value '/var/www/html/ipbased' and a 'Browse...' button, and 'Webmaster email address' with the value 'ipbased@me.com'. The 'Host Information' section contains a dropdown menu with 'IP based Virtual Host' selected, and two input fields: 'IP Address' with the value '192.168.100.100' and 'Server Host Name' with the value 'ipbased.me.com'. At the bottom of the window are three buttons: 'OK', 'Cancel', and 'Help'.

Name based Virtual Host

If you choose **Name based Virtual Host**, Figure 11–12, *Name Based Virtual Hosts* appears to configure the `NameVirtualHost` Directive based on the host name of the server. Specify the IP address in the **IP address** field. To specify more than one IP address, separate each IP address with spaces. To specify a port, use the syntax `IP Address:Port`. Use `*` to configure all ports for the IP address. Specify the host name for the virtual host in the **Server Host Name** field. In the **Aliases** section, click **Add** to add a host name alias. Adding an alias here adds a `ServerAlias` directive within the `NameVirtualHost` Directive.

Figure 11–12 Name Based Virtual Hosts



SSL

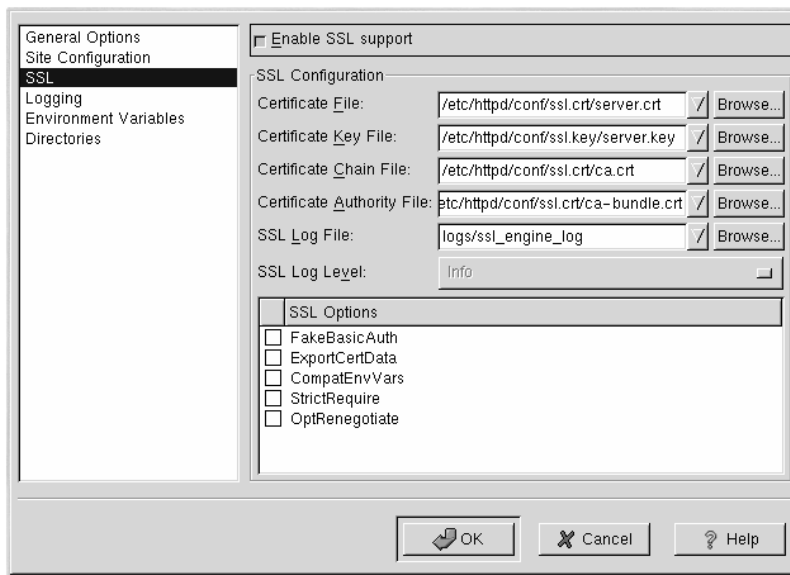
Note

You can't use name based virtual hosts with SSL, 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.

If an Apache server is not configured with SSL support, communication between an Apache server and its clients are not encrypted. This is appropriate for websites without personal or confidential information. For example, an open source website that distributes open source software and documentation has no need for secure communications. However, an ecommerce website that requires credit card information should use the Apache SSL support to encrypt its communications. Enabling Apache SSL support enables the use of the `mod_ssl` security module. To enable it through Apache Configuration Tool you must allow access through port 443 under the **Main** tab => Available Addresses. Refer to Section 11.1, *Basic Settings* for details. Then, select the virtual host name in the **Virtual Hosts** tab,

click the **Edit** button, choose **SSL** from the left-hand menu, and check the **Enable SSL Support** option as shown in Figure 11–13, *SSL Support*. The **SSL Configuration** section is pre-configured with the dummy digital certificate. The digital certificate provides authentication for your secure Web server and identifies the secure server to client Web browsers. You must purchase your own digital certificate. Do not use the dummy one provided in Red Hat Linux for your website. For details on purchasing a CA-approved digital certificate, refer to the *Apache-Related Reference* portion of the *Official Red Hat Linux Reference Guide*.

Figure 11–13 SSL Support



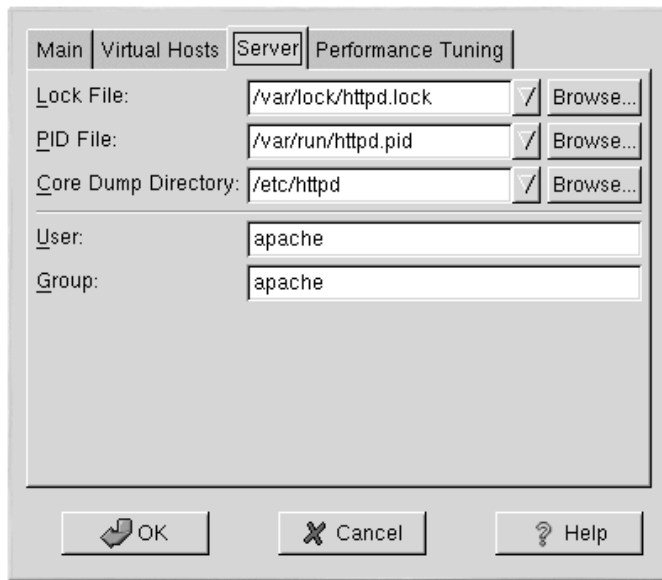
Additional Virtual Host Options

The **Site Configuration**, **Environment Variables**, and **Directories** options for the virtual hosts are the same directives that you set when you clicked the **Edit Default Settings** button, except the options set here are for the individual virtual hosts that you are configuring. Refer to Section 11.2, *Default Settings* for details on these options.

11.4 Server Settings

The **Server** tab allows you to configure basic server settings. The default settings for these options are appropriate for most situations.

Figure 11–14 Server Configuration



The **Lock File** value corresponds to the `LockFile` directive. This directive sets the path to the lockfile used when Apache is compiled with either `USE_FCNTL_SERIALIZED_ACCEPT` or `USE_FLOCK_SERIALIZED_ACCEPT`. It must be stored on the local disk. IT should be left to the default value unless the `logs` directory is located on an NFS share. If this is the case, the default value should be changed to a location on the local disk and to a directory that is readable only by root.

The **PID File** value corresponds to the `PidFile` directive. This directive sets the file in which the server records its process ID (pid). This file should only be readable by root. In most cases, it should be left to the default value.

The **Core Dump Directory** value corresponds to the `CoreDumpDirectory` directive. Apache tries to switch to this directory before dumping core. The default value is the `ServerRoot`. However, if the user that the server runs as can not write to this directory, the core dump can not be written. Change this value to a directory writable by the user the server runs as, if you want to write the core dumps to disk for debugging purposes.

The **User** value corresponds to the `User` directive. It sets the userid used by the server to answer requests. This user's settings determine 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.

WARNING

Unless you know exactly what you're doing, don't set the User to root. Using root as the User will create large security holes for your Web server.

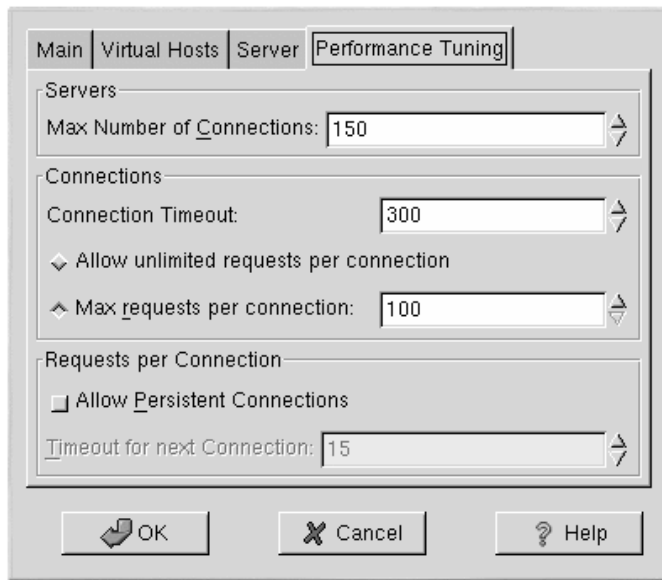
The parent `httpd` process first runs as root during normal operations, but is then immediately handed off to the apache user. The server must start as root because it needs to bind to a port below 1024. Ports below 1024 are reserved for system use, so they can't be used by anyone but root. Once the server has attached itself to its port, however, it hands the process off to the apache user before it accepts any connection requests.

The **Group** value corresponds to the `Group` directive. 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`.

11.5 Performance Tuning

Click on the **Performance Tuning** tab to configure the maximum number of child server processes you want and to configure the Apache options for client connections. The default settings for these options are appropriate for most situations. Altering these settings may affect the overall performance of your Web server.

Figure 11–15 Performance Tuning



Set **Max Number of Connections** to the maximum number of simultaneous client requests that the server will handle. For each connection, a child `httpd` process is created. After this maximum number of process is reached, no one else will be able to connect to the Web server until a child server process is freed. You can not set this value to higher than 256 without recompiling Apache. This option corresponds to the `MaxClients` directive.

Connection Timeout defines, in seconds, the amount of time that your server will wait for receipts and transmissions during communications. Specifically, `Connection 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. By default, `Connection Timeout` is set to 300 seconds, which is appropriate for most situations. This option corresponds to the `TimeOut` directive.

Set the **Max requests per connection** to the maximum number of requests allowed per persistent connection. The default value is 100, which should be appropriate for most situations. This option corresponds to the `MaxRequestsPerChild` directive.

If you check the **Allow unlimited requests per connection** option, the `MaxKeepAliveRequests` directive to 0, and unlimited requests are allowed.

If you uncheck the **Allow Persistent Connections** option, the `KeepAlive` directive is set to `false`. If you check it, the `KeepAlive` directive is set to `true`, and the `KeepAliveTimeout` directive is set to the number that is selected as the **Timeout for next Connection** value. This directive 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 **Connection Timeout** value applies instead.


Setting the **Persistent Connections** to a high value may cause a server to slow down, depending on how many users are trying to connect to it. The higher the number, the more server processes waiting for another connection from the last client that connected to it.

11.6 Saving Your Settings


If you do not want to save your Apache configuration settings, click the **Cancel** button in the bottom right corner of the Apache Configuration Tool window. You will be prompted to confirm this decision. If you click **Yes** to confirm this choice, your settings will not be saved.

If you want to save your Apache configuration settings, click the **OK** button in the bottom right corner of the Apache Configuration Tool window. The dialog window shown in Figure 11–16, *Save and Exit* will appear. If you answer **Yes**, your settings will be saved in `/etc/httpd/conf/httpd.conf`. Remember that your original configuration file will be overwritten.

Figure 11–16 Save and Exit



If this is the first time that you have used Apache Configuration Tool, you will see the dialog window shown in Figure 11–17, *Configuration File Manually Modified*, warning you that the configuration file has been manually modified. If Apache Configuration Tool detects that the `httpd.conf` configuration file has been manually modified, it will save the manually modified file as `/etc/httpd/conf/httpd.conf.bak`.

Figure 11–17 Configuration File Manually Modified

Restart Daemon

After saving your settings, you must restart the Apache daemon with the command `service httpd restart`. You must be logged in as root to execute this command.

11.7 Additional Resources

To learn more about Apache, refer to the following resources.

11.7.1 Installed Documentation

- Apache documentation — If you have the `apache-manual` package installed and the Apache Web server daemon (`httpd`) running, you can view the Apache documentation. Open a Web browser, and go to the URL `http://localhost` on the server that is running Apache. Then, click the **Documentation** link.

11.7.2 Useful Websites

- <http://www.apache.org> — *The Apache Software Foundation*
 - <http://httpd.apache.org/docs/> — *Apache HTTP Server Version 1.3 User's Guide*
 - <http://localhost/manual/index.html> — After starting the Apache server on your local system, you can view the *Apache HTTP Server Version 1.3 User's Guide* using this URL.
-

11.7.3 Related Books

- *Apache: The Definitive Guide* by Ben Laurie and Peter Laurie; O'Reilly & Associates, Inc.

12 BIND Configuration

This chapter assumes that you have a basic understanding of BIND and DNS; it does not attempt to explain the concepts of BIND and DNS. This chapter does explain how to use BIND Configuration Tool (`bindconf`) to configure basic BIND server zones for BIND version 8. BIND Configuration Tool creates the `/etc/named.conf` configuration file and the zone configuration files in the `/var/named` directory each time you apply your changes.


If you require more functionality than this tool provides, you can create the `/etc/named.conf` configuration file using BIND Configuration Tool and then add your customized settings. However, once you manually modify the configuration file, you cannot use BIND Configuration Tool to edit the custom configuration settings that were added manually.

Do Not Edit `/etc/named.conf`

Do not edit the `/etc/named.conf` configuration file. BIND Configuration Tool generates this file after you apply your changes. If you want to configure settings that are not configurable using BIND Configuration Tool, then do not use it.

BIND Configuration Tool requires the X Window System and root access. To start BIND Configuration Tool, use one of the following methods:

- On the GNOME desktop, go to the **Main Menu Button** (on the Panel) => **Programs** => **System** => **bindconf**.
 - On the KDE desktop, go to the **Main Menu Button** (on the Panel) => **Red Hat** => **System** => **bindconf**.
 - Type the command `bindconf` at a shell prompt (for example, in an XTerm or GNOME-terminal).
-

Figure 12–1 bindconf

BIND Configuration Tool configures the default zone directory to be `/var/named`. All zone files specified are relative to this directory. BIND Configuration Tool also includes basic syntax checking when values are entered. For example, if a valid entry is an IP address, you are only allowed to type numbers and the dot (.) character into the text area.

BIND Configuration Tool allows you to add a forward master zone, a reverse master zone, and a slave zone. After adding the zones, you can edit or delete them from the main window as shown in Figure 12–1, *bindconf*.

After adding, editing, or deleting a zone, you must choose **File => Apply** to write the `/etc/named.conf` configuration file and all the individual zone files in the `/var/named` directory. Applying your changes will also have the named service reload the configuration files. You can also choose **File => Exit** and click **Yes** to **Do you want to apply your changes before exiting?**

12.1 Adding a Forward Master Zone

To add a forward master zone (also known as a primary master), click the **Add** button, select **Forward Master Zone**, and enter the domain name for the master zone in the **Domain name** text area.

A new window as shown in Figure 12–2, *Adding a Forward Master Zone* will appear with the following options:

- **Name** — Domain name that was just entered in the previous window.
- **File Name** — File name of the DNS database file, relative to `/var/named`.
- **Contact** — Email address of the main contact for the master zone.
- **Primary Name Server (SOA)** — State of authority (SOA) record. This specifies the name server that is the best resource of information for this domain. The default value is `@`, which means that the SOA is the same as the domain name entered in the **Name** field above.
- **Serial Number** — The serial number of the DNS database file. This number must be incremented each time the file is changed, so that the slave name servers for the zone will retrieve the latest data. BIND Configuration Tool increments this number each time the configuration changes. It can also be incremented manually by clicking the **Set** button next to the **Serial Number** value.
- **Time Settings** — The **Refresh**, **Retry**, **Expire**, and **Minimum TTL** (Time to Live) values that are stored in the DNS database file.
- **Records** — Add, edit, and delete record resources of type **Host**, **Alias**, and **Name server**.

Figure 12–2 Adding a Forward Master Zone

The screenshot shows a dialog box titled "Master Zone" with the following fields and controls:

- Name:** forward.example.com
- File Name:** forward.example.com.zone
- Contact:** root@localhost
- Primary Name Server (SOA):** @
- Serial Number:** 1 (with a "Set..." button next to it)
- Time Settings...** button
- Records:** A list box containing "forward.example.com" with "Add...", "Edit...", and "Delete" buttons to its right.
- OK** and **Cancel** buttons at the bottom.

The configuration shown in Figure 12–2, *Adding a Forward Master Zone* creates the following entry in `/etc/named.conf`:

```
zone "forward.example.com" {
    type master;
    file "forward.example.com.zone";
};
```

It also creates the file `/var/named/forward.example.com.zone` with the following information:

```
$TTL 86400
@ IN SOA @ root.localhost (
    1 ; serial
    28800 ; refresh
    7200 ; retry
    604800 ; expire
    86400 ; ttl
)
```

After configuring the Forward Master Zone, click **OK** to return to the main window as shown in Figure 12–1, *bindconf*. From the pulldown menu, choose **File => Apply** to write the `/etc/named.conf` configuration file, write all the individual zone files in the `/var/named` directory, and have the daemon reload the configuration files.

12.2 Adding a Reverse Master Zone

To add a reverse master zone, click the **Add** button and select **Reverse Master Zone**. Enter the first three octets of the IP address range that you want to configure. For example, if you are configuring the IP address range `192.168.10.0/255.255.255.0`, enter `192.168.10` in the **IP Address (first 3 Octets)** text area.

A new window will appear, as shown in Figure 12–3, *Adding a Reverse Master Zone*, with the following options:

1. **IP Address** — The first three octets that you just entered in the previous window.
2. **Reverse IP Address** — Non-editable. Pre-populated based on the IP Address entered.
3. **File Name** — File name of DNS database file in the `/var/named` directory.
4. **Primary Name Server (SOA)** — State of authority (SOA) record. This specifies the name server that is the best resource of information for this domain. The default value is `@`, which means that the SOA is the same as the domain name entered in the **Name** field above.
5. **Time Settings** — The **Refresh**, **Retry**, **Expire**, and **Minimum TTL** (Time to Live) values that are stored in the DNS database file.

6. **Name Servers** — Add, edit, and delete name servers for for the reverse master zone. At least one name server is required.
7. **Reverse Address Table** — List of IP addresses within the reverse master zone and their hostnames. For example, for the reverse master zone 1.2.3, you can add 1.2.3.100 in the **Reverse Address Table** with the hostname foo.example.com. The hostname must end with a period (.) to specify that it is a full hostname.

Figure 12–3 Adding a Reverse Master Zone

The screenshot shows a window titled "Reverse Master Zone". It contains the following fields and sections:

- Reverse Master Zone** section:
 - IP Address: 1.2.3
 - Reverse IP Address: 3.2.1.in-addr.arpa
 - File Name: 3.2.1.in-addr.arpa.zone
 - Primary Name Server (SOA): @
 - Time Settings... button
- Name Servers** section:
 - List: ns.example.com.
 - Buttons: Add..., Edit..., Delete
- Reverse Address Table** section:

Address	Host or Domain:	Buttons
1.2.3.1	one.example.com.	Edit...
1.2.3.2	two.example.com.	Delete

At the bottom of the window are OK and Cancel buttons.

The configuration shown in Figure 12–3, *Adding a Reverse Master Zone* creates the following entry in `/etc/named.conf`:

```
zone "3.2.1.in-addr.arpa" {
    type master;
    file "3.2.1.in-addr.arpa.zone";
};
```

It also creates the file `/var/named/3.2.1.in-addr.arpa.zone` with the following information:

```
$TTL 86400
@ IN SOA @ root.localhost (
    2 ; serial
```

```

28800 ; refresh
7200 ; retry
604800 ; expire
86400 ; ttl
)

@ IN NS ns.example.com.

1 IN PTR one.example.com.
2 IN PTR two.example.com.

```

After configuring the Reverse Master Zone, click **OK** to return to the main window, as shown in Figure 12-1, *bindconf*. From the pulldown menu, choose **File => Apply** to write the `/etc/named.conf` configuration file, write all the individual zone files in the `/var/named` directory, and have the daemon reload the configuration files.

12.3 Adding a Slave Zone

To add a slave zone (also known as a secondary master), click the **Add** button and select **Slave Zone**. Enter the domain name for the slave zone in the **Domain name** text area.

A new window will appear, as shown in Figure 12-4, *Adding a Slave Zone*, with the following options:

- **Name** — The domain name that was entered in the previous window.
- **Masters List** — The name server from which the slave zone retrieves its data. This value must be a valid IP address. You can only enter numbers and dots (.) in the text area.
- **File Name** — File name of the DNS database file in `/var/named`.

Figure 12-4 Adding a Slave Zone



The configuration shown in Figure 12-4, *Adding a Slave Zone* creates the following entry in `/etc/named.conf`:

```
zone "slave.example.com" {
    type slave;
    file "slave.example.com.zone";
    masters {
        1.2.3.4;
    };
};
```

The configuration file `/var/named/slave.example.com.zone` is created by the `named` service when it downloads the zone data from the master server(s).

After configuring the slave zone, click **OK** to return to the main window as shown in Figure 12-1, *bindconf*. From the pulldown menu, choose **File => Apply** to write the `/etc/named.conf` configuration file and have the daemon reload the configuration files.

13 Printer Configuration

Red Hat Linux no longer includes `printtool`. The `printconf` utility has replaced `printtool`. The `printconf` utility maintains the `/etc/printcap` configuration file, print spool directories, and print filters.

To use `printconf`, you must be running the X Window System and have root privileges. To start `printconf`, use one of the following methods:

- On the GNOME desktop, go to the **Main Menu Button** (on the Panel) => **Programs** => **System** => **Printer Configuration**
- On the KDE desktop, go to the **Main Menu Button** (on the Panel) => **Red Hat** => **System** => **Printer Configuration**.
- Type the command `printconf-gui` at a shell prompt (for example, in an XTerm or a GNOME terminal).¹

Do Not Edit `/etc/printcap`

Do not edit the `/etc/printcap` file. Each time the printer daemon (`lpd`) is started or restarted, a new `/etc/printcap` file is dynamically created.

If you want to add a printer without using `printconf`, edit the `/etc/printcap.local` file. The entries in `/etc/printcap.local` are not displayed in `printconf` but are read by the printer daemon. If you upgrade your system from a previous version of Red Hat Linux, your existing configuration file is converted to the new format used by `printconf`. Each time a new configuration file is generated by `printconf`, the old file is saved as `/etc/printcap.old`.

¹ If you type `printtool` at a shell prompt, `printconf` will start.

Figure 13–1 `printconf`

The screenshot shows a window titled 'printconf' with a menu bar (File, Test, Help) and a toolbar with icons for New, Edit, Delete, Default, and Apply. Below the toolbar is a table with the following data:

	Queue	Alias List	Queue Type	Details
<input checked="" type="checkbox"/>	test		LOCAL	PostScript queue on local device /dev/lp0
<input type="checkbox"/>	test2		LPD	HP Color LaserJet 5 lpd queue lp@servername
<input type="checkbox"/>	test3		SMB	PostScript SMB queue on share //machinename/printer
<input type="checkbox"/>	test4		NCP	Canon BJ-10e Novell queue queue on server servername
<input type="checkbox"/>	test5		JETDIRECT	HP Color LaserJet 5000 JetDirect queue 192.168.1.10:9100

Five types of print queues can be configured with `printconf`:

- **Local Printer** — a printer attached directly to your computer through a parallel or USB port. In the main printer list as shown in Figure 13–1, *printconf*, the **Queue Type** for a local printer is set to **LOCAL**.
- **Unix Printer (lpd Spool)** — a printer attached to a different UNIX system that can be accessed over a TCP/IP network (or example, a printer attached to another Red Hat Linux system on your network). In the main printer list as shown in Figure 13–1, *printconf*, the **Queue Type** for a remote UNIX printer is set to **LPD**.
- **Windows Printer (SMB Share)** — a printer attached to a different system which is sharing a printer over a SMB network (for example, a printer attached to a Microsoft Windows machine). In the main printer list as shown in Figure 13–1, *printconf*, the **Queue Type** for a remote Windows printer is set to **SMB**.
- **Novell Printer (NCP Queue)** — a printer attached to a different system which uses Novell’s NetWare network technology. In the main printer list as shown in Figure 13–1, *printconf*, the **Queue Type** for a remote Novell printer is set to **NCP**.
- **JetDirect Printer** — a printer connected directly to the network instead of to a computer. In the main printer list as shown in Figure 13–1, *printconf*, the **Queue Type** for a JetDirect printer is set to **JETDIRECT**.

Important

If you add a new print queue or modify an existing one, you need to restart the printer daemon (lpd) for the changes to take effect.

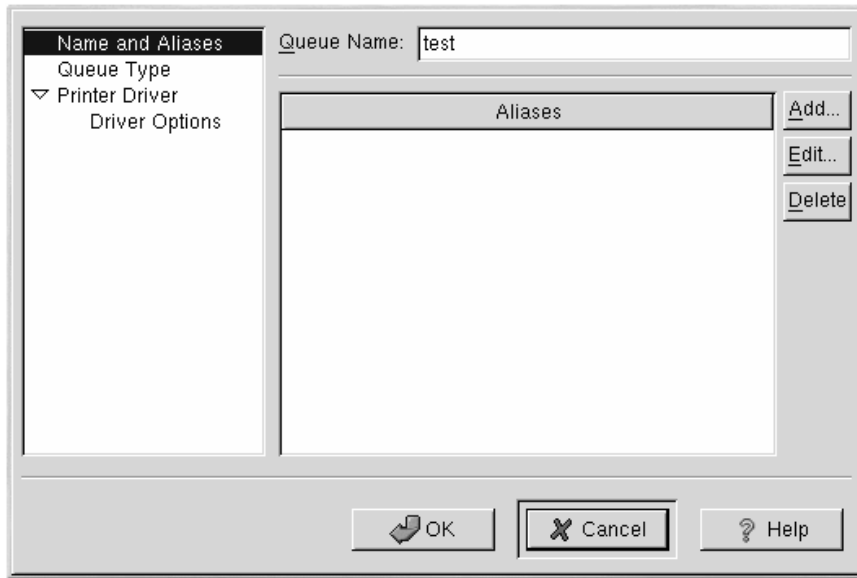
Clicking the **Apply** button saves any changes that you have made and restarts the printer daemon.² Alternatively, you can choose **File => Save Changes** and then choose **File => Restart lpd** to save your changes and then restart the printer daemon.

If a printer appears in the main printer list with the **Queue Type** set to **INVALID**, the printer configuration is missing options that are required for the printer to function properly. To remove this printer from the list, select it from the list and click the **Delete** button.

13.1 Adding a Local Printer

To add a local printer such as one attached to the parallel port or USB port of your computer, click the **Add** button in the main printconf window. The window shown in Figure 13–2, *Adding a Printer* will appear.

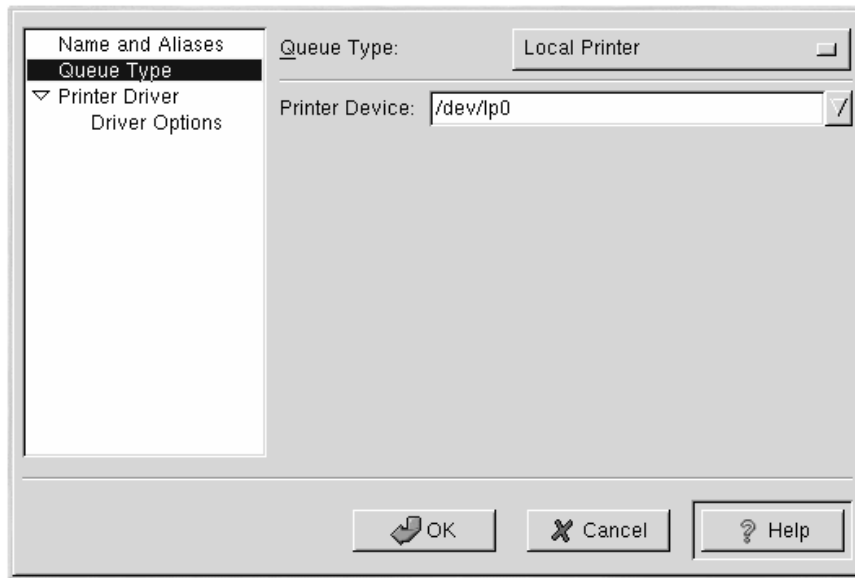
² The changes are not written to the `/etc/printcap` configuration file until the printer daemon (lpd) is restarted.

Figure 13–2 Adding a Printer

Enter a unique name for the printer in the **Queue Name** text field. This can be any descriptive name for your printer. You can also create alias names for the printer by clicking the **Add** button beside the **Aliases** list. Refer to Section 13.7, *Creating Printer Aliases* for more information about aliases. The printer name and aliases cannot contain spaces and must begin with a letter a through z or A through Z. The valid characters are a through z, A through Z, 0 through 9, -, and _.

Click **Queue Type** from the left side menu and choose **Local Printer** from the **Queue Type** menu. Also enter the printer device in the **Printer Device** text field or choose it from the pulldown menu as shown in Figure 13–3, *Adding a Local Printer*.

Figure 13–3 Adding a Local Printer



Next, select the type of printer that is connected to the system by clicking **Printer Driver** from the left side menu. After choosing the manufacturer and model number of the printer, a list of drivers will appear. If there is more than one driver for the printer, choose the preferred driver in the **Printer Driver** list. If you are not sure which one to use, do not change this value. Click the **Printer Notes** button to view notes about the printer driver from the Linux Printing Database.

Click **Driver Options** from the left side menu after selecting a printer driver. These options will vary depending on the printer driver that you selected. Typical options include paper size, print quality, and printer resolution.

Click the **OK** button. The new printer will appear in the printer list in the main window. Click the **Apply** button in the main window to save your changes to the `/etc/printcap` configuration file and restart the printer daemon (`lpd`). After applying the changes, print a test page to ensure the configuration is correct. Refer to Section 13.6, *Printing a Test Page* for details.

13.2 Adding a Remote UNIX Printer

To add a remote UNIX printer, such as one attached to a different Linux system on the same network, click the **Add** button in the main `printconf` window. The window shown in Figure 13–2, *Adding a Printer* will appear. Enter a unique name for the printer in the **Queue Name** text field. This can be

any descriptive name for your printer. You can also create alias names for the printer by clicking the **Add** button beside the **Aliases** list. Refer to Section 13.7, *Creating Printer Aliases* for more information about aliases. The printer name and aliases can not contain spaces and must begin with a letter a through z or A through Z. The valid characters are a through z, A through Z, 0 through 9, -, and _.


Click **Queue Type** from the left side menu and choose **Unix Printer (lpd Spool)** from the **Queue Type** menu.

Text fields for the following options appear below the **Queue Type** menu as shown in Figure 13–4, *Adding a Remote Printer*:

- **Server** — The hostname or IP address of the remote machine to which the printer is attached.
- **Queue** — The remote printer queue. The default printer queue is usually `lp`.

By default, the **Strict RFC1179 Compliance** option is not chosen. If you are having problems printing to a non-Linux `lpd` queue, choose this option to disable enhanced LPRng printing features.

Figure 13–4 Adding a Remote Printer



Next, select the type of printer that is connected to the remote system from the left side list by clicking **Printer Driver** from the left side menu. After choosing the manufacturer and model number of the

printer, a list of drivers will appear. If there is more than one driver for the printer, choose the preferred driver in the **Printer Driver** list. If you are not sure which one to use, do not change this value. Click the **Printer Notes** button to view notes about the printer driver from the Linux Printing Database. Click **OK**.

Click **Driver Options** from the left side menu after selecting a printer driver. These options will vary depending on the printer driver that you selected. Typical options include paper size, print quality, and printer resolution.

Click the **Apply** button in the main window to save your changes to the `/etc/printcap` configuration file and restart the printer daemon (`lpd`). After applying the changes, print a test page to ensure the configuration is correct. Refer to Section 13.6, *Printing a Test Page* for details.

Important

The remote machine must be configured to allow the local machine to print on the desired queue. As root, create the file `/etc/hosts.lpd` on the remote machine to which the printer is attached. On separate lines in the file, add the IP address or hostname of each machine which should have printing privileges.

13.3 Adding a Samba (SMB) Printer

To add printer which is accessed using the SMB protocol, click the **Add** button in the main `printconf` window. The window shown in Figure 13–2, *Adding a Printer* will appear. Enter a unique name for the printer in the **Queue Name** text field. This can be any descriptive name for your printer. You can also create any alias names for the printer by clicking the **Add** button beside the **Aliases** list. Refer to Section 13.7, *Creating Printer Aliases* for more information about aliases. The printer name and aliases cannot contain spaces and must begin with a letter a through z or A through Z. The valid characters are a through z, A through Z, 0 through 9, -, and _.

Click **Queue Type** from the left side menu and choose **Windows Printer (SMB Share)** from the **Queue Type** menu. If the printer is attached to a Microsoft Windows system, choose this queue type.

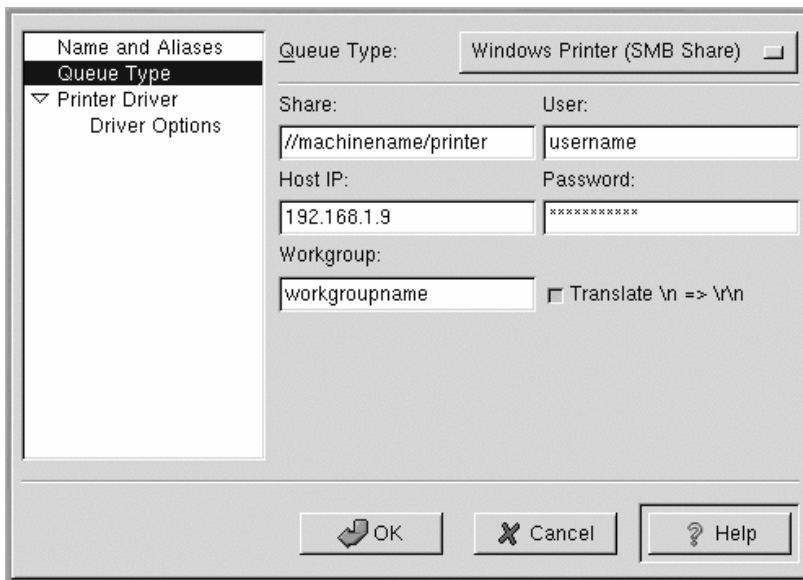
Text fields for the following options appear below the **Queue Type** menu as shown in Figure 13–5, *Adding a SMB Printer*:

- **Share** — The name of the shared printer on which you want to print. This name must be the same name defined as the Samba printer on the remote Windows machine. Notice the syntax of `//machinename/sharename`.
-

- **User** — The name of the user you must login as to access the printer. This user must exist on the Windows system, and the user must have permission to access the printer. The user name is typically **guest** for Windows servers, or **nobody** for Samba servers.
- **Host IP** — The hostname or IP address of the remote system that is sharing the SMB printer.
- **Password** — The password (if required) for the user specified in the **User** field.
- **Workgroup** — The name of the workgroup on the machine running Samba.

Click the **Translate \n => \r\n** button to translate the end of line characters to a form that is readable by a Microsoft Windows system.

Figure 13–5 Adding a SMB Printer



Next, select the type of printer that is connected to the remote SMB system from the left side list by clicking **Printer Driver** from the left side menu. After choosing the manufacturer and model number of the printer, a list of drivers will appear. If there is more than one driver for the printer, choose the preferred driver in the **Printer Driver** list. If you are not sure which one to use, do not change this value. Click the **Printer Notes** button to view notes about the printer driver from the Linux Printing Database. Click **OK**.

Click **Driver Options** from the left side menu after selecting a printer driver. These options will vary depending on the printer driver that you selected. Typical options include paper size, print quality, and printer resolution.

Click the **Apply** button in the main window to save your changes to the `/etc/printcap` configuration file and restart the printer daemon (`lpd`). After applying the changes, print a test page to ensure the configuration is correct. Refer to Section 13.6, *Printing a Test Page* for details.

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. To avoid this, the username and password to use the printer should be different from the username and password used for the user's account on the local Red Hat Linux system. If they are different, then 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.

13.4 Adding a Novell NetWare (NCP) Printer

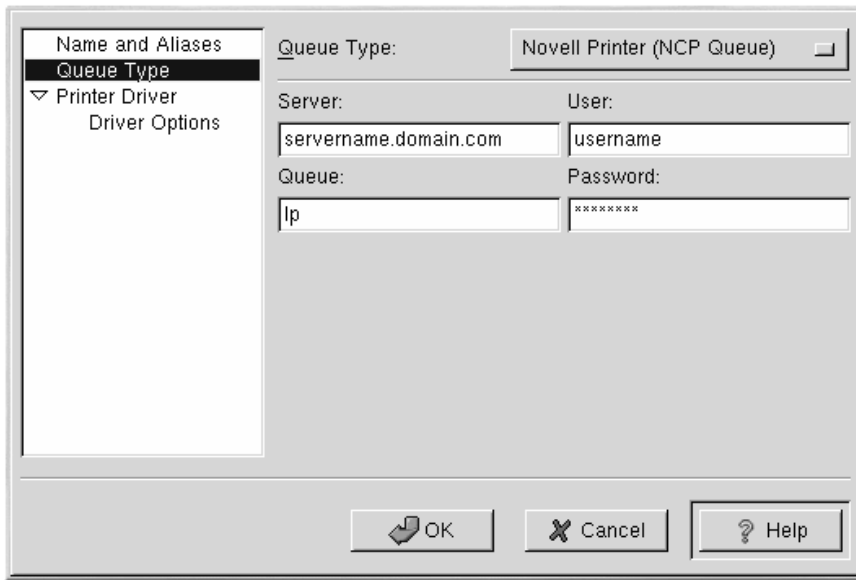
To add a NetWare (NCP) printer, click the **Add** button in the main `printconf` window. The window shown in Figure 13–1, *printconf* will appear. Enter a unique name for the printer in the **Queue Name** text field. This can be any descriptive name for your printer. You can also create any alias names for the printer by click the **Add** button beside the **Aliases** list. Refer to Section 13.7, *Creating Printer Aliases* for more information about aliases. The printer name and aliases cannot contain spaces and must begin with a letter a through z or A through Z. The valid characters are a through z, A through Z, 0 through 9, -, and _.

Click **Queue Type** from the left side menu and choose **Novell Printer (NCP) Queue** from the **Queue Type** menu.

Text fields for the following options appear below the **Queue Type** menu as shown in Figure 13–6, *Adding an NCP Printer*:

- **Server** — The hostname or IP address of the NCP system to which the printer is attached.
 - **Queue** — The remote queue for the printer on the NCP system.
 - **User** — The name of the user you must login as to access the printer.
 - **Password** — The password for the user specified in the **User** field above.
-

Figure 13–6 Adding an NCP Printer



Next, select the type of printer that is connected to the remote NCP system by clicking **Printer Driver** from the left side menu. After choosing the manufacturer and model number of the printer, a list of drivers will appear. If there is more than one driver for the printer, choose the preferred driver in the **Printer Driver** list. If you are not sure which one to use, do not change this value. Click the **Printer Notes** button to view notes about the printer driver from the Linux Printing Database. Click **OK**.

Click **Driver Options** from the left side menu after selecting a printer driver. These options will vary depending on the printer driver that you selected. Typical options include paper size, print quality, and printer resolution.

Click the **Apply** button in the main window to save your changes to the `/etc/printcap` configuration file and restart the printer daemon (`lpd`). After applying the changes, print a test page to ensure the configuration is correct. Refer to Section 13.6, *Printing a Test Page* for details.

13.5 Adding a JetDirect Printer

To add a JetDirect printer, click the **Add** button in the main `printconf` window. The window shown in Figure 13–1, *printconf* will appear. Enter a unique name for the printer in the **Queue Name** text field. This can be any descriptive name for your printer. You can also create any alias names for the printer by click the **Add** button beside the **Aliases** list. Refer to Section 13.7, *Creating Printer Aliases* for

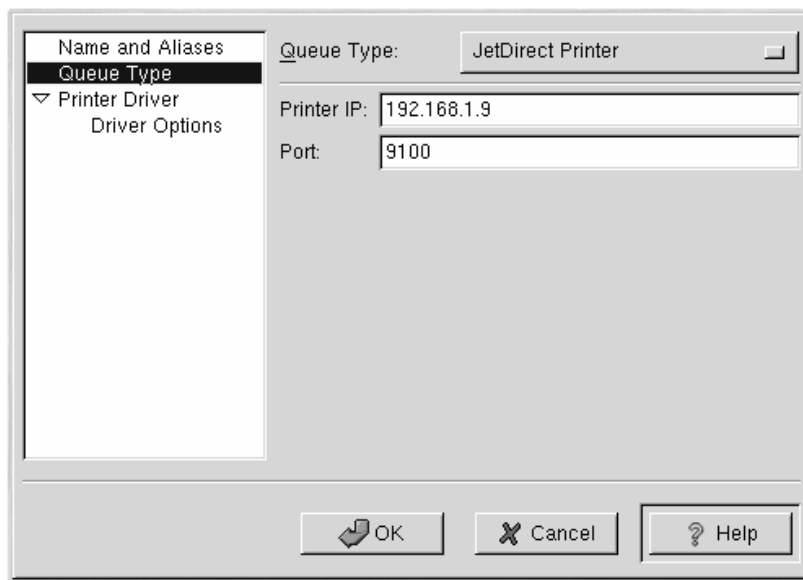
more information about aliases. The printer name and aliases cannot contain spaces and must begin with a letter a through z or A through Z. The valid characters are a through z, A through Z, 0 through 9, -, and _.

Click **Queue Type** from the left side menu and choose **JetDirect Printer** from the **Queue Type** menu.

Text fields for the following options appear below the **Queue Type** menu as shown in Figure 13–7, *Adding a JetDirect Printer*:

- **Printer IP** — The hostname or IP address of the JetDirect printer.
- **Port** — The port on the JetDirect printer that is listening for print jobs.

Figure 13–7 Adding a JetDirect Printer



Next, select the type of printer by clicking **Printer Driver** from the left side menu. After choosing the manufacturer and model number of the printer, a list of drivers will appear. If there is more than one driver for the printer, choose the preferred driver in the **Printer Driver** list. If you are not sure which one to use, do not change this value. Click the **Printer Notes** button to view notes about the printer driver from the Linux Printing Database. Click **OK**.

Click **Driver Options** from the left side menu after selecting a printer driver. These options will vary depending on the printer driver that you selected. Typical options include paper size, print quality, and printer resolution.

Click the **Apply** button in the main window to save your changes to the `/etc/printcap` configuration file and restart the printer daemon (`lpd`). After applying the changes, print a test page to ensure the configuration is correct. Refer to Section 13.6, *Printing a Test Page* for details.

13.6 Printing a Test Page

After you have configured your printer, you should print a test page to make sure the printer is functioning properly. To print a test page, select the printer that you want to test from the printer list, and choose **Test => Print Postscript Test Page, Print A4 Postscript Test Page, or Print ASCII Test Page** from the pulldown menu. Do not choose **Print Postscript Test Page** if the printer can not print PostScript.

13.7 Creating Printer Aliases

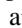
A printer alias is an alternate name for a printer. To add an alias for an existing printer, select the printer from the printer list, and click the **Alias** button on the toolbar. In the new dialog window that appears add new aliases for the printer or delete existing alias names. A printer can have more than one alias. All aliases for each printer are listed in the **Alias List** column in the printer list. Click **Apply** to save the aliases and restart the printer daemon.

13.8 Modifying Existing Printers

After adding your printer(s), you can edit settings by selecting the printer from the printer list and clicking the **Edit** button. The same window that is used for adding a printer appears, as shown in Figure 13-2, *Adding a Printer*. The window contains the current values for the printer that you selected to edit. Make any changes, and click **OK**. Click **Apply** to save the changes and restart the printer daemon.

If you want to rename a printer, select the printer from the printer list, and click the **Rename** button on the toolbar. A dialog box will appear with the current name of the printer. Rename the printer, and click the **OK** button. The name of the printer should change in the printer list. Click **Apply** to save the change and restart the printer daemon.

To delete an existing printer, select the printer and click the **Delete** button on the toolbar. The printer will be removed from the printer list. Click **Apply** to save the changes and restart the printer daemon.

To set the default printer, select the printer from the printer list and click the **Default** button on the toolbar. The default printer icon  appears in the first column of the printer list beside the default printer.

If you want to modify an imported printer's settings, you cannot modify its settings directly. You must override the printer. You can only override an imported printer that has been imported using the alchemist libraries. Imported printers have the `*` symbol beside them in the first column of the printer list.

To override the printer, select the printer, and choose **File => Override Queue** from the pulldown menu. After overriding a printer, the original imported printer will have the `*` symbol beside it in the first column of the printer list.

13.9 Additional Resources

To learn more about printing on Red Hat Linux, refer to the following resources.

13.9.1 Installed Documentation

- `man printcap` — The manual page for the `/etc/printcap` printer configuration file.

13.9.2 Useful Websites

- <http://www.linuxprinting.org> — *GNU/Linux Printing* contains a large amount information about printing in Linux.

14 Linuxconf

Linuxconf allows you to configure and control various aspects of your system. 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.

After configuring your systems settings through Linuxconf, the changes are not activated immediately. You must activate the changes by choosing **File => Act/Changes** from the pulldown menu in the GUI version of Linuxconf, clicking on an **Accept** button in Web-based Linuxconf, or selecting the **Accept** button in text-mode Linuxconf.

14.1 Starting Linuxconf

You'll need to be root to run Linuxconf. If you are in your user account, type `su -` at a shell prompt to become root and then type the command `linuxconf`. If the directory `/sbin` is not in your path, type the command with the full path: `/sbin/linuxconf`. If you want to run the GUI version of Linuxconf, you need to have the X Window System installed as well as GNOME.

14.2 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 pulldown menu exists on that line. The [Ctrl]-[X] key combination will make pulldown 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 easy; 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 14.4, *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.

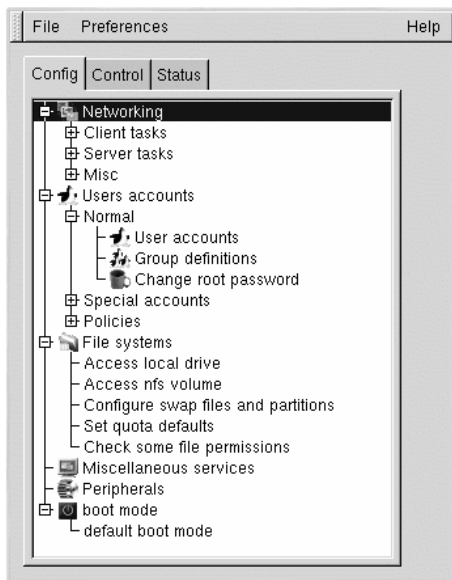
Linuxconf will start in either character-cell or X mode, depending on your **DISPLAY** environment variable. The first time you run Linuxconf, an introductory message will be shown; although it is only displayed once, accessing help from the main screen will give you the same basic information.

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.

14.3 Gnome-Linuxconf Interface

Gnome-Linuxconf makes it easy to navigate the hierarchical structure of Linuxconf.

Figure 14–1 Linuxconf Menu View



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** checkbox.
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 panel on the right side 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 the changes you have made.

14.4 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 checkbox.
3. Select the **Accept** button.
4. Also, verify that the disable line in the `/etc/xinetd.d/linuxconf-web` file reads

```
disable=no
```

Then run the command `/sbin/service xinetd reload` from a shell prompt.

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.) 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.

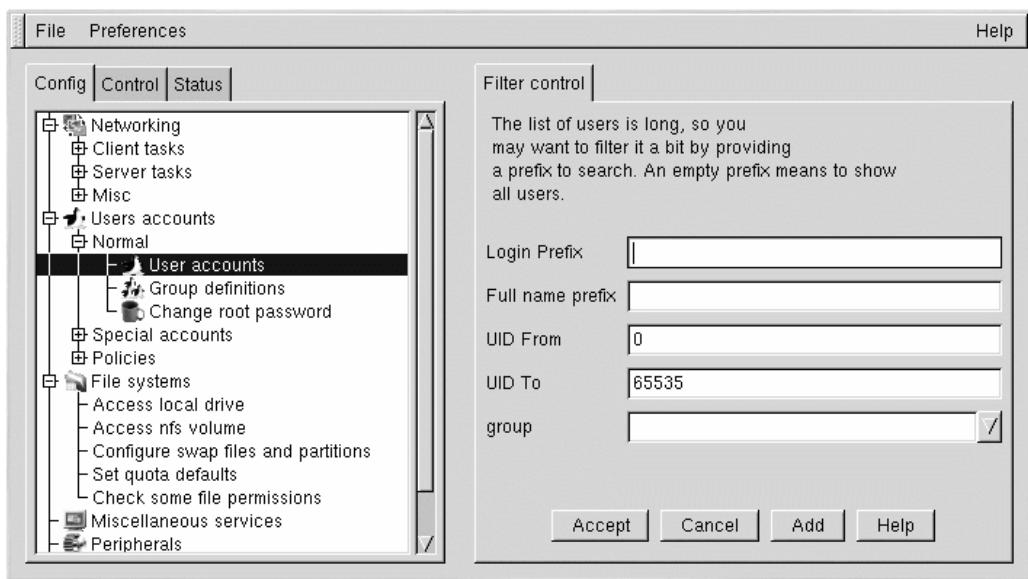
You can also enable network-wide access to Linuxconf by following the same steps and entering a network name instead of a hostname.

14.5 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:

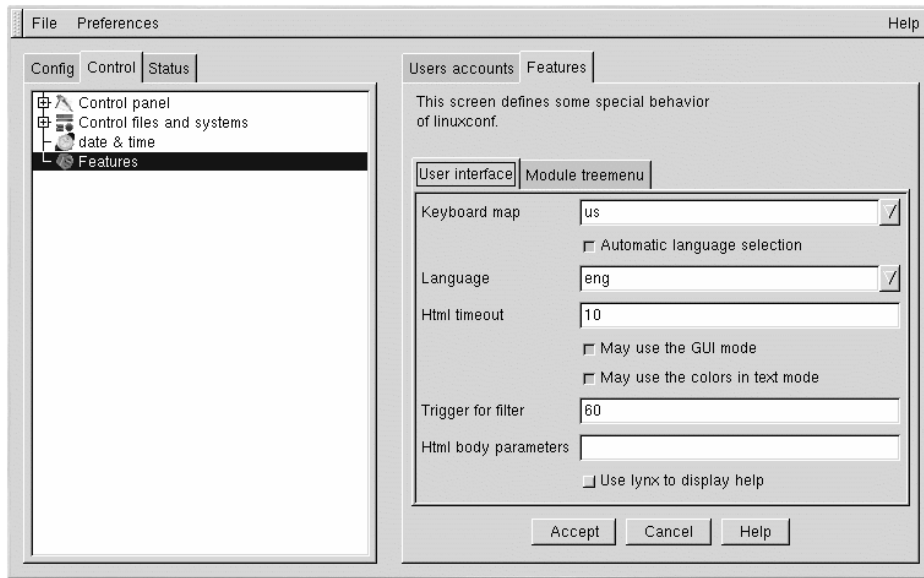
- Open **Config => Users accounts => Normal => User accounts**. Linuxconf may show you a filter screen (see Figure 14–2, *Filter Control Screen*).

Figure 14–2 Filter Control Screen



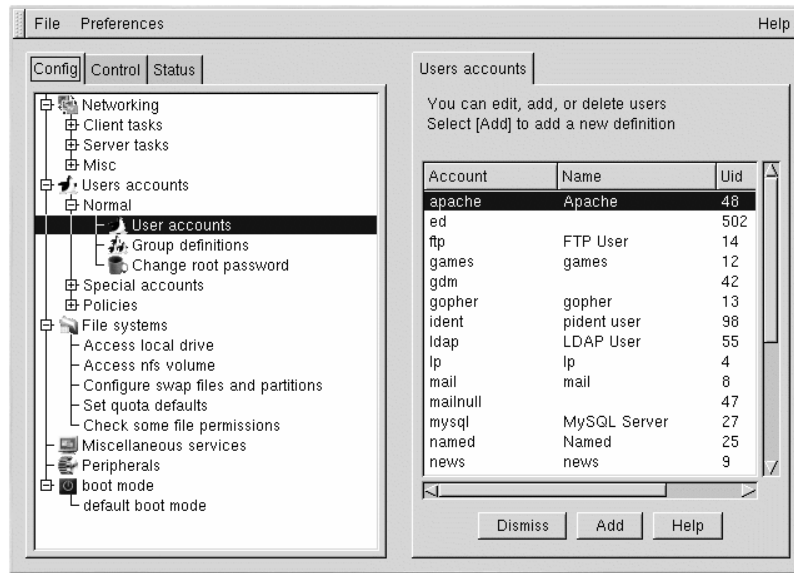
You can use the filter screen to select a smaller range of accounts than the full list. To get the full list, select **Accept** without changing any of the parameters. For detailed information on the various filters, select the **Help** button on the **Filter control** screen. Once you've applied or bypassed the filter, you'll see the **Users accounts** tab (see Figure 14–4, *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 14–3, *Setting the Trigger for Filter*.

Figure 14–3 Setting the Trigger for Filter

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

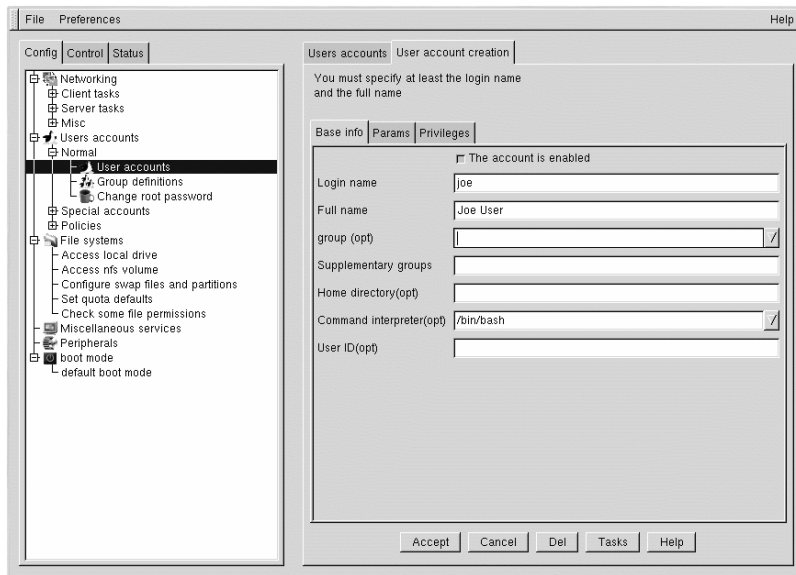
Figure 14–4 Users Accounts Screen



- Select **Add**. This will open the **User account creation** tab (see Figure 14–5, *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 14–5 User Account Creation



14.5.1 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 the *Official Red Hat Linux Reference Guide*.

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.

14.5.2 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 an **Expiration date**.

14.5.3 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 choosing **File => Act/Changes** from the pulldown menu in the GUI version of 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. 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. A password 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.

14.6 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 14.5, *Adding a User Account* if you need guidance for how to fill in the user accounts fields.

14.7 Changing a User's Password

- Open **Config => Users accounts => Normal => User accounts**. This will open the **Users accounts** tab (see Figure 14–4, *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 14.5, *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**.

14.8 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 the current root password, you will be prompted for a new password. In the **Confirmation** field, type the password again. This is to prevent you from mistyping the password. See Section 14.5, *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**.

14.9 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.
 - Unselect the checkbox 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.

14.10 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.

14.11 Deleting a User Account

Note

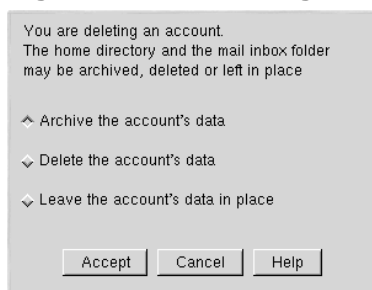
There are ways to retain the files associated with an account, but any files that are deleted are gone and effectively unrecoverable. Take care when using this option!

To delete an account:

- Open **Config => Users accounts => Normal => User accounts**.
- On the **User accounts** screen (see Figure 14–4, *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 14–6 Deleting Account Screen



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

1. It removes the user from the user accounts list.
2. It 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, created automatically the first time you remove a user account this way, is put in the same place as all of your user directories.

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 14-6, *Deleting Account Screen*) will:

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

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 14-6, *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.
-

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.

14.12 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.

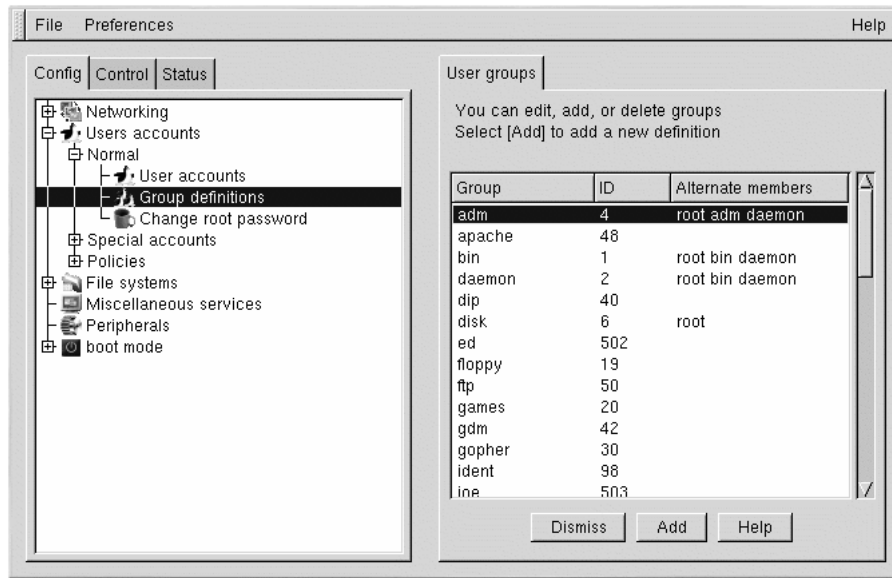
14.12.1 Creating a Group

To create a new group:

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

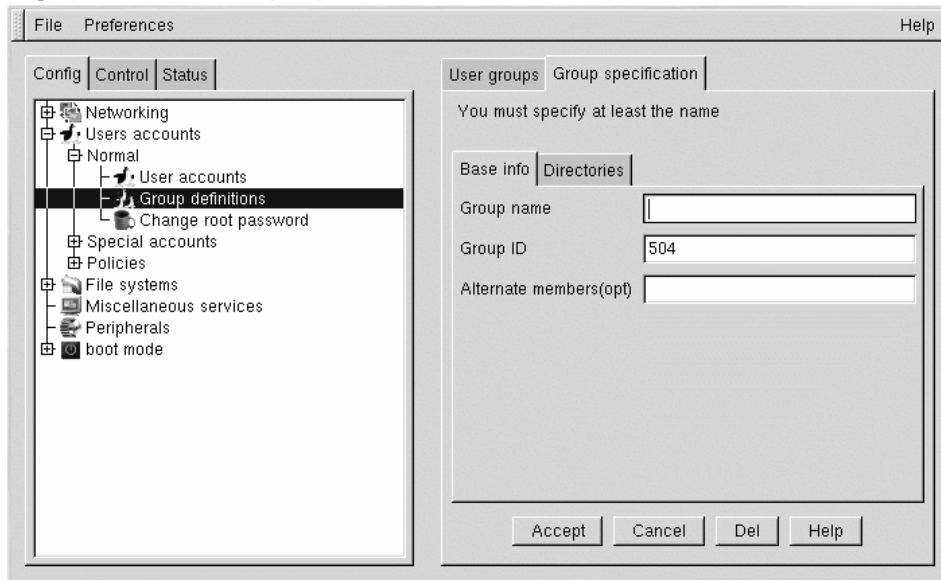
You may see a filter screen, depending upon the settings in **Control** => **Filters**. Either provide a filter, or select **Accept** to bypass the filter.

Figure 14–7 User Groups Screen



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

Figure 14–8 Group Specification Screen



Enter a group name. You may also wish to specify members of the group; you 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.

14.12.2 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 14–7, *User Groups Screen*), select the group you wish to delete.
- You'll be presented with the **Group specification** screen (see Figure 14–8, *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. Linuxconf does not recycle used group numbers or used user IDs, so it won't happen by accident.

14.12.3 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 be simpler.

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 14–8, *Group Specification Screen*).
- Add or remove each user from the **Alternate members** field. Make sure that all of the usernames are separated by a space character.
- Select **Accept** 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 14–4, *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.
- 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.

14.13 Filesystems

A filesystem is composed of files and directories, all under a single root directory. The root directory and the directories below it may contain any number of files and subdirectories. A 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 on 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 called **mounting**. When a device is mounted, it is then accessible to the system's users. The directory under which a mounted device's filesystem is 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) on 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:
 - `/` — the system root directory
 - `/foo` — the mount point for the CD-ROM

- A CD-ROM:
 - / — the 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):
 - / — the system root directory
 - /foo — the 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.

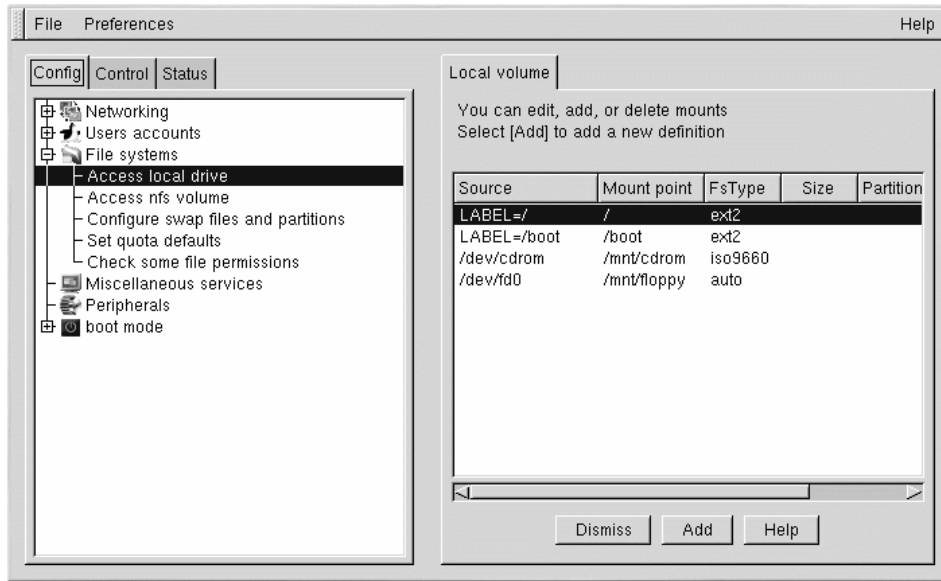
Using the information in `/etc/fstab`, you can type `mount` and then either the mount point or the device. The `mount` command will look for the rest of the information in `/etc/fstab`. It's possible to modify this file by hand, or by using `Linuxconf`.

14.13.1 Reviewing Your Current Filesystem

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

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

Figure 14–9 Local Volume Screen



The fields, as shown in Figure 14–9, *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; for diskette drives, the number refers to the actual unit.
- **Mount point:** The place in the filesystem from which the drive (or other device) is accessible after it is 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.

Note

Red Hat Linux 7.1 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, but it 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 one small directory to entire volumes. No information on **Size** or **Partition type** is available for these partitions. Additional information on NFS filesystems (if any are available) will be contained under:

Config => File systems => Access nfs volume

The screen is similar to the **Local Volume** screen (see Figure 14–9, *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, you might see a value of `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."

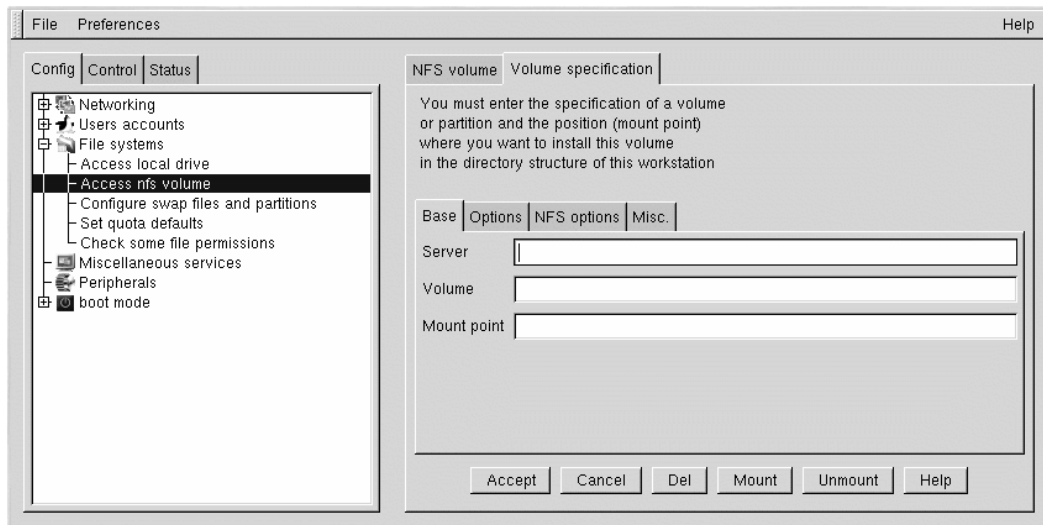
14.13.2 Adding NFS Mounts

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 14–10 Volume Specification Screen



You will need to fill in the three fields on the **Base** tab next (see Figure 14–10, *Volume Specification Screen*).

- **Server:** The hostname of the machine on which the desired filesystem is located. For example, `foo.bar.com`.
- **Volume:** The filesystem you wish to add. For example, `/var/spool/mail`.
- **Mount point:** The directory in your system from which you want the remote file system to be accessible. 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**.

14.14 Network Configuration with Linuxconf

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, make sure your modem is installed and that the cables

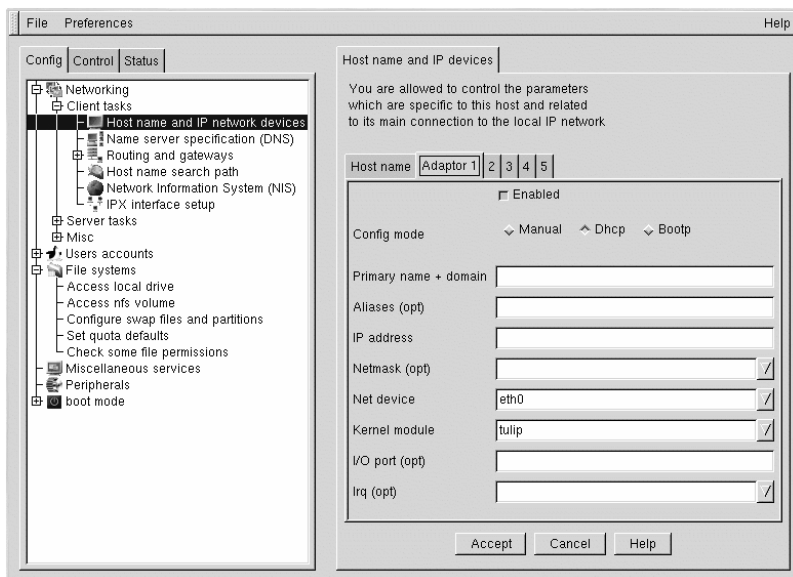
are attached correctly. If you're using a network card, make sure it is installed properly and that the cables are correctly connected. In other words, no matter what network configuration you specify, if every phone line or cable is not in place, you'll never get connected.

14.14.1 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 => Host names and IP network devices**. The **Host name** tab will request a hostname, 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 14–11 Adaptor 1



The first item on this screen is a checkbox to indicate whether this adaptor is enabled or not. If this is the one you intend to use, it should be checked. 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; refer to the *Official Red Hat Linux Reference Guide* for a list of kernel modules.

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 in which 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 *Official Red Hat Linux Reference Guide*. For the most up-to-date list, please see our website at:

<http://hardware.redhat.com/>

The netmask information may be already be set for you. However, depending upon the type of network you are joining or setting up, you may need to fill in this field. The most common value for this field is **255.255.255.0**.

Required fields for manual configuration:

- **Primary name + domain** — The primary name is the name of your computer; 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 will follow this pattern: 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 network 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.

If you're setting up a private network that will not *ever* be connected to the Internet, then you can choose any **Primary name + domain name** you want and you have several choices for **IP addresses** (see Table 14–1, *Addresses and Examples*).

Table 14–1 Addresses and Examples

Available addresses	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 ¹	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.

14.14.2 Name Server Specification

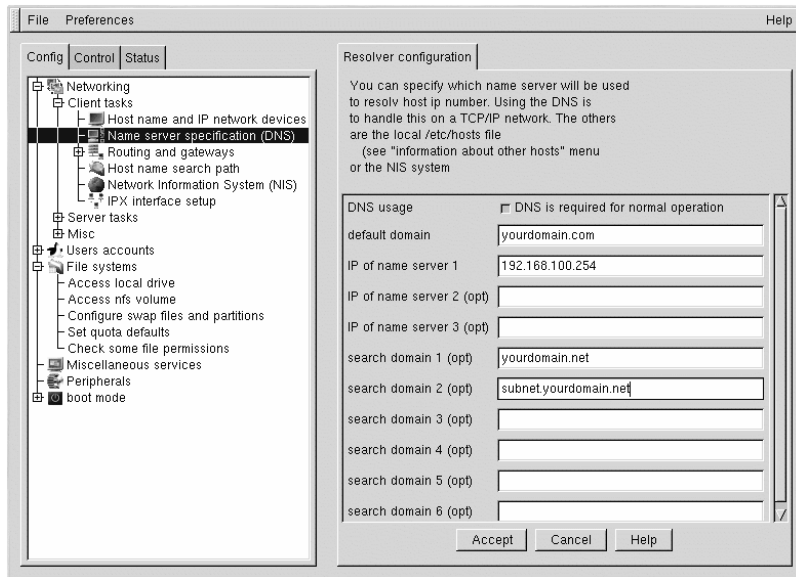
A nameserver and a default domain are also needed to establish a network connection. The nameserver is used to translate hostnames, such as `private.network.com`, to corresponding IP addresses, such as `192.168.7.3`.

The default domain tells the computer where to look if a fully qualified hostname is not 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 was not.

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

¹ You should not use these IP addresses if you connect to the Internet, since `192.168.0.*` and `192.168.255.*` cannot be 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.

Figure 14–12 Resolver Configuration Screen

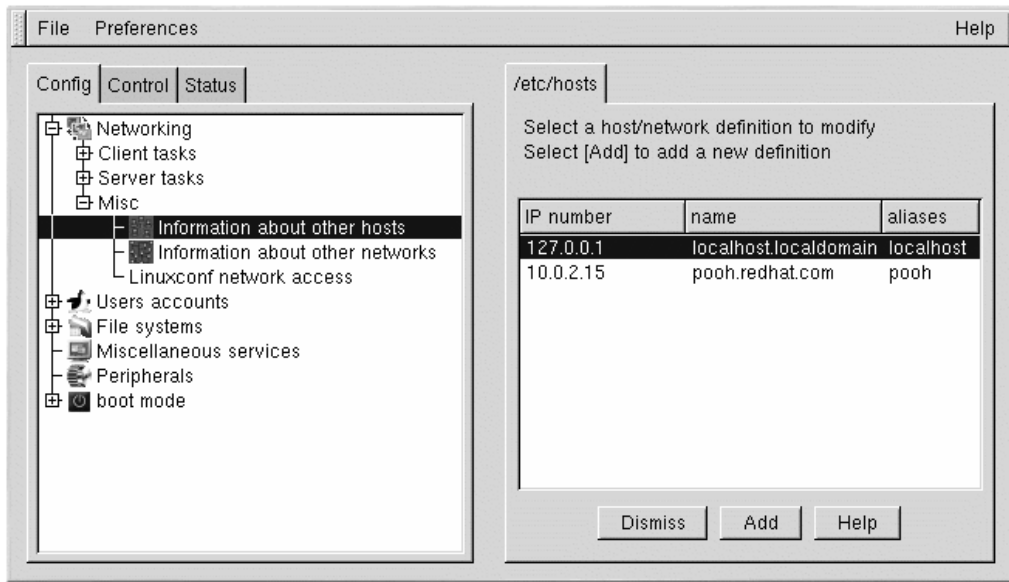


Nameservers are ranked according to the order in which they are accessed. It is not unusual to see nameservers referred to as primary, secondary, tertiary and so on down the list if more than one is specified. Each of these must be an IP address and not a name, since the computer has no way to resolve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

In addition to a default domain, you can also specify up to six search domains. Search domains are useful if you want to specify many hostnames with the same domain but do not want to type the domain name. For example, if your search domain is redhat.com, the hostname falcon would resolve to falcon.redhat.com. Search domains take precedence over the default domain.

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

Figure 14–13 /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.

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 is specified. For example, in `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 set `foo` as an alias.
- **Comment** — A comment about the machine. For example, "The remote nameserver."

Once finished, select `Accept`.

14.15 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 14–2 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 <code>/etc/hosts</code>	Config => Networking => Misc => Information about other hosts
Modify group membership	Config => Users accounts => Normal => Group definitions or Config => Users accounts => Normal => User accounts
Set filter parameters	Control => Features

What do you want to do?	Where to find it in Linuxconf
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

14.16 Additional Resources

For more in-depth information about Linuxconf, refer to the following resources.

14.16.1 Useful Websites

- <http://www.solucorp.qc.ca/linuxconf/> — More information about Linuxconf, including its most recent release, can be found at the Linuxconf website.
- <http://www.xc.org/jonathan/linuxconf-faq.html> — The Linuxconf FAQ website.
- <http://hub.xc.org/scripts/lyris.pl?visit=linuxconf> — The archives of the Linuxconf mailing list.
- 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 the Linuxconf list is for questions pertaining to Linuxconf, and is not intended for questions about Linux.

15 Control Panel

Note

Most of what can be done with the Control Panel applications can also be done using `linuxconf`.

The Control Panel is a launching pad for a number of different system administration tools (see Figure 15–1, *The Control Panel*). These tools make system administration easier because you can configure your system without having to remember configuration file formats and awkward command line options.

Figure 15–1 The Control Panel



To start the Control Panel, type `control-panel` at a shell prompt. You will need to be root to run the Control Panel tools and must be running the X Window System because it is a graphical utility. 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 can type `su -` followed by the root password when prompted.

Note

If you are not running X as root, you might see the following error message:

```
Xlib: connection to "server.domain.net:0.0" refused by server
Xlib: Client is not authorized to connect to Server
kmail: cannot connect to X server server.domain.net:0
```

If you see this, you 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 you have started the Control Panel, click on an icon to start 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.

WARNING

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.

15.1 Network Configurator

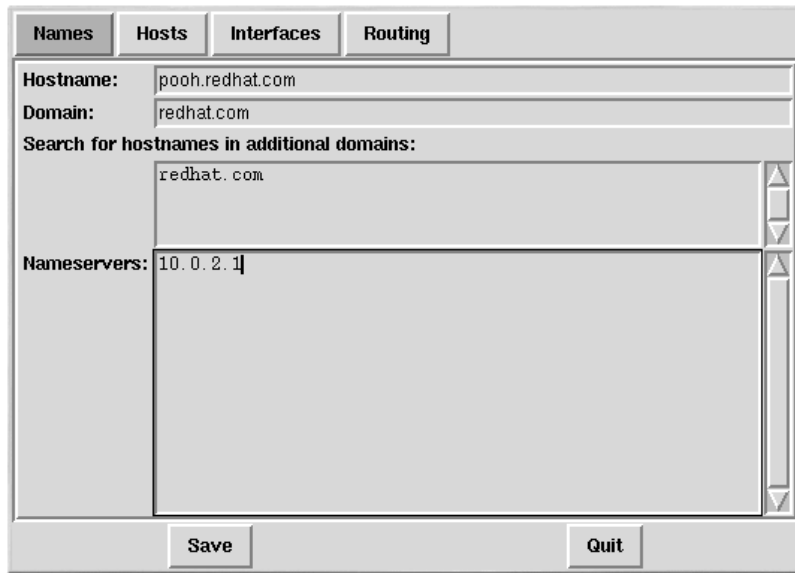
Note

Documentation on network configuration using `linuxconf` can be found in Section 14.14, *Network Configuration with Linuxconf*.

Network Configurator, shown in Figure 15–2, *Network Configurator*, allows you to easily manipulate parameters such as IP addresses, gateway addresses, and network addresses, as well as name servers and `/etc/hosts`.

To start Network Configurator start the Control Panel and click on the Network Configurator icon or type `netcfg` from a shell prompt.

Figure 15–2 Network Configurator



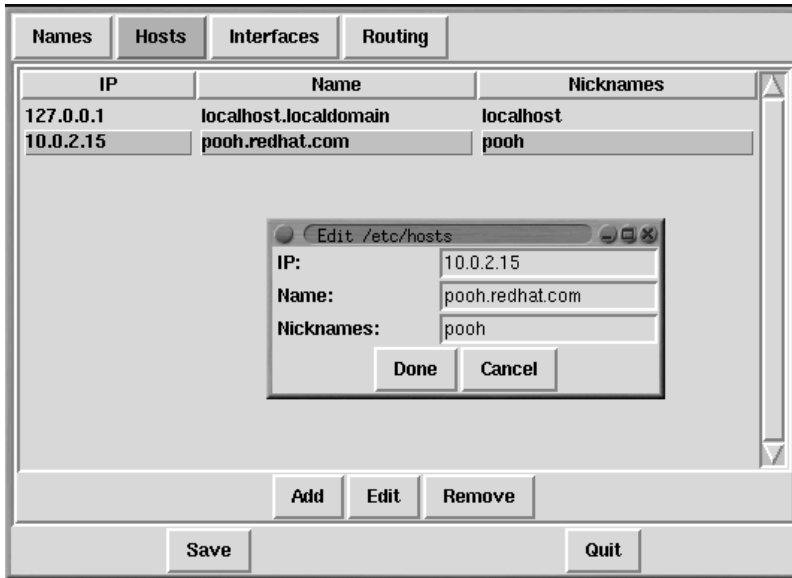
Network devices can be added, removed, configured, activated, deactivated and aliased. Ethernet, ARCnet, Token Ring, pocket (ATP), SLIP, PLIP and loopback devices are supported. SLIP/PLIP support works well on most hardware, but some hardware setups may exhibit unpredictable behavior.

When using Network Configurator, click **Save** to write your changes to disk. To quit without making any changes, select **Quit**.

15.1.1 Managing Names

Network Configurator's **Names** panel 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 **Names** panel can not configure a machine to be a name server. 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 15–3 Adding/Editing Hosts



15.1.2 Managing Hosts

In the **Hosts** management panel, you can 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 15–3, *Adding/Editing Hosts* for an example.

15.1.3 Adding a Network Interface

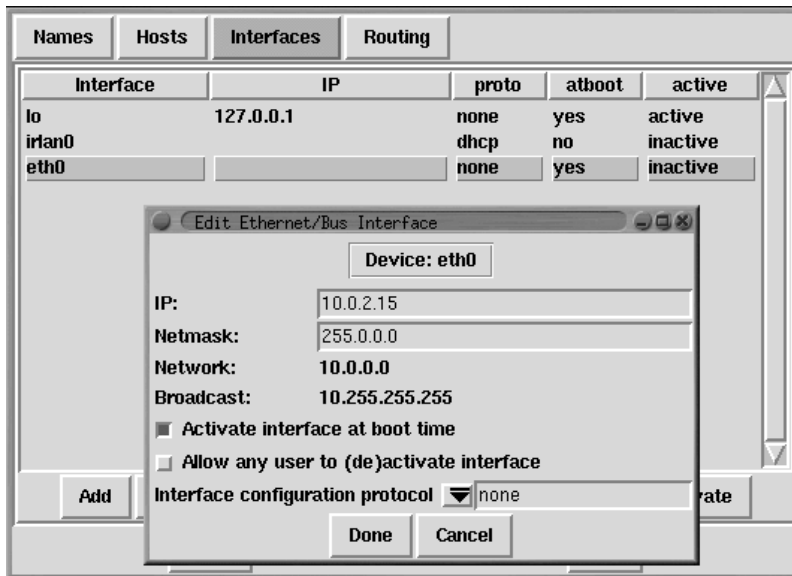
If you have added a network interface to your machine since installing Red Hat Linux, or if you didn't configure your Ethernet card at install time, you can configure it here with a few clicks of a mouse.

Note

You may need to configure `kmod` to load a driver for the network interface you are adding (e.g., `eth0`); see Section 16.5, *Loading Kernel Modules* for more information.

Begin adding an interface by clicking on the **Interfaces** in the main panel. A list of configured devices and a row of available options will be displayed; see Figure 15–4, *Configured Interfaces*.

Figure 15–4 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.

Note

A clone button is now available in Network Configurator. This button can be used to create a "clone" of an already-existing interface. Using clone interfaces, a laptop can have one Ethernet interface defined for a work LAN, and a clone Ethernet device defined for a home LAN.

SLIP Interface

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 whether you want to activate the interface at boot time.

Ethernet, ARCnet, Token Ring and Pocket Adapter 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 option.
- **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**.

15.1.4 Managing Routes

In the **Routing** 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 15-5, *Adding/Editing Routes* for an example.

Figure 15–5 Adding/Editing Routes



15.2 Time and Date

Time Tool 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.

Note

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

16 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 already know why you need to build your own kernel, you probably do not need to."

In the past, you had to recompile the kernel if you added new hardware on your system. In other words, the kernel was **static**. Improvements in the Linux 2.0.x kernels allowed for many hardware drivers to be **modularized** into components that are loaded on demand. However, major problems existed when users had multiple kernels that had been compiled for different configuration options on their system; for example, SMP versus UP kernels. Further Linux 2.4.x kernel modularization advancements allow for multiple kernels to co-exist more easily, but they can not share modules.

For information on handling kernel modules see Section 16.5, *Loading Kernel Modules*. Unless you are recompiling a customized kernel for your system, you will not see many changes in how kernel modules are handled.

16.1 The 2.4 Kernel

Red Hat Linux now ships with the 2.4 kernel. Here are the highlights of the 2.4 kernel as shipped with Red Hat Linux:

- The directory for the kernel source is now `/usr/src/linux-2.4` instead of `/usr/src/linux`.
- Better SMP support.
- Support for up to 64 gigabytes of physical RAM — the enterprise kernel installed with Red Hat Linux 7.1 is compiled to support 64 gigabytes of physical memory.
- Better multimedia support including the `maestro3` module for the ESS Allegro sound card.
- Better USB support.
- Now support for IEEE 1394, also referred to as FireWire™, devices.

16.2 Building a Modularized Kernel

The instructions in this section apply to building a modularized kernel. If you are interested in building a monolithic kernel instead, see Section 16.4, *Building a Monolithic Kernel* for an explanation of the different aspects of building and installing a monolithic kernel.

The following steps will guide you through building a custom kernel for the x86 architecture:

Note

This example uses 2.4.2-0.1.21 as the kernel version. Your kernel version might differ. To determine your kernel version, type the command `uname -r`. Replace 2.4.2-0.1.21 with your kernel version.

1. The most important step is to make sure that you have a working emergency boot disk in case you make a mistake. If you didn't make a boot disk during the installation, use the `mkbootdisk` command to make one now. The standard command is similar to `mkbootdisk --device /dev/fd0 2.4.x`, where 2.4.x is the full version of your kernel (such as 2.4.2-0.1.21). Once done, test the boot disk to make sure that it will boot the system.
2. You must have both the `kernel-headers` and `kernel-source` packages installed. Issue the commands `rpm -q kernel-headers` and `rpm -q kernel-source` to determine their versions, if they are installed. If they are not installed, install them from the Red Hat Linux CD 1 or the Red Hat FTP site available at <ftp://ftp.redhat.com> (a list of mirrors is available at <http://www.redhat.com/mirrors.html>). Refer to Chapter 17, *Package Management with RPM* for information on installing RPM packages.
3. At a shell prompt and change to the directory `/usr/src/linux-2.4`. All commands from this point forward must be issued from this directory.
4. It is important that you 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. If you already have a working configuration file (`/usr/src/linux-2.4/.config`) that you want to use, back it up to a different directory before running this command and copy it back after running the command. If you use an existing configuration file, skip the next step.
5. Now you must create a configuration file that will determine which components to include in your new kernel.

If you are running the X Window System, the recommended method is to use the command `make xconfig`. Components are listed in different levels of menus and are selected using a mouse. You can select **Y** (yes), **N** (no), or **M** (module). After choosing your components, click the **Save and Exit button** to create the configuration file `/usr/src/linux-2.4/.config` and exit the Linux Kernel Configuration program.

Other available methods for kernel configuration are listed below:

- `make config` — An interactive text program. Components are presented in a linear format and you answer them one at a time. This method does not require the X Window System and does not allow you to change your answers to previous questions.
- `make menuconfig` — A text-mode, menu driven program. Components are presented in a menu of categories; you select the desired components in the same manner used in the text-mode Red Hat Linux installation program. Toggle the tag corresponding to the item you want included: **[*]** (built-in), **[]** (exclude), **<M>** (module), or **< >** (module capable). This method does not require the X Window System.
- `make oldconfig` — This is a non-interactive script that will set up your configuration file to contain the default settings. If you're using the default Red Hat kernel, it will create a configuration file for the kernel that shipped with Red Hat Linux for your architecture. This is useful for setting up your kernel to known working defaults and then turning off features that you don't want.

Note

To use `kmod` (see Section 16.5, *Loading Kernel Modules* for details) and kernel modules you must answer **Yes** to `kmod support` and `module version (CONFIG_MODVERSIONS) support` during the configuration.

6. After creating a `/usr/src/linux-2.4/.config` file, use the command `make dep` to set up all the dependencies correctly.
 7. Use the command `make clean` to prepare the source tree for the build.
 8. The next step in making a modularized kernel is to edit `/usr/src/linux-2.4/Makefile` so that you do not overwrite your existing kernel. 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 at <http://www.redhat.com/mirrors/LDP/HOWTO/Kernel-HOWTO.html> or in the `Makefile` in `/usr/src/linux-2.4` on your Linux system.

Edit `/usr/src/linux-2.4/Makefile` and modify the line beginning with `EXTRAVERSION =` to match a "unique" name by appending the date to the end of the string. For example, if you are compiling kernel version 2.4.2-0.1.21 you can append the flag to look similar to `EXTRAVERSION = -0.1.21-feb2001`). This will allow you to have the old working kernel and the new kernel, version 2.4.2-0.1.21-12feb2001, on your system at the same time.
 9. Build the kernel with `make bzImage`.
 10. Build any modules you configured with `make modules`.
-

11. Install the kernel modules (even if you didn't build any) with `make modules_install`. Make sure that you type the underscore (`_`). 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/2.4.2-0.1.21-12feb2001/`.
12. If you have a SCSI adapter and you made your SCSI driver modular, build a new `initrd` image (see Section 16.3, *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`.
13. Use `make install` to copy your new kernel and its associated files to the proper directories.
14. 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. This can be accomplished by updating the `/etc/lilo.conf` file and running `/sbin/lilo`.

The default `/etc/lilo.conf` file looks similar to the following:

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

image=/boot/vmlinuz-2.4.2-0.1.21
label=linux
    initrd=initrd-2.4.2-0.1.21.img
read-only
root=/dev/hda5
```

To add your new kernel to LILO, copy the existing section to a new one and modify it to boot your new kernel image (and `initrd` image if you have any SCSI devices and created an `initrd` image). Also, rename the label of the old kernel to something such as **linux-old**. Your `/etc/lilo.conf` should look similar to the following:

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

```
image=/boot/vmlinuz-2.4.2-0.1.21-12feb2001
label=linux
initrd=initrd-2.4.2-0.1.21-12feb2001.img
read-only
root=/dev/hda5

image=/boot/vmlinuz-2.4.2-0.1.21
label=linux-old
initrd=initrd-2.4.2-0.1.21.img
read-only
root=/dev/hda5
```

15. To activate your changes, run the command `/sbin/lilo`. If all goes well, you will see output similar to the following:

```
Added linux *
Added linux-old
```

The `*` after `linux` means that the section labeled `linux` is the default kernel that LILO will boot.

16. From now on, when the system boots you will see `linux` and `linux-old` as LILO boot options.

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

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

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

- The loopback block device is available.
- The `/etc/modules.conf` 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/initrd-2.4.2-0.1.21-12feb2001.img 2.4.2-0.1.21-12feb2001
```

In the above example, `/boot/initrd-2.4.2-0.1.21-12feb2001.img` is the filename of the new `initrd` image. `2.4.2-0.1.21-12feb2001` is the kernel whose modules (from `/lib/modules`) should be used in the `initrd` image. This is not necessarily the same as the version number of the currently running kernel.

16.4 Building a Monolithic Kernel

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

- When configuring the kernel, do not compile anything as a module. In other words, only answer **Yes** or **No** to the questions. Also, you should answer **No** to `kmod support` and `module version (CONFIG_MODVERSIONS) support`.
- Omit the following steps:

```
make modules
make modules_install
```

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

16.5 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.

The mechanism that supports dynamic loading of modules is a kernel thread called `kmod`. Modules are not loaded unless they are needed. When the kernel requests a module, the module is loaded along with all its module dependencies.

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 information you provide, the installation program decides which features to compile into the resident kernel and which to put in loadable modules. The installation program sets up the dynamic loading mechanism to work transparently. If you build your own custom kernel, you can make all of these decisions for yourself.

If you add new hardware after installation and the hardware requires a kernel module, you need to set up the dynamic loading mechanism. `Kudzu` usually detects new hardware. You can also add the new driver by editing the module configuration file, `/etc/modules.conf`.

For example, if your system included a model SMC EtherPower 10 PCI network adapter at the time of installation, the module configuration file will contain this following line:

```
alias eth0 tulip
```

After installation, if you install a second identical network adapter to your system, add the following line to `/etc/modules.conf`:

```
alias eth1 tulip
```

See the *Official Red Hat Linux Reference Guide* for an alphabetical list of kernel modules and the hardware supported by the modules.

Part IV Package Management

17 Package Management with RPM

The Red Hat Package Manager (RPM) is an open packaging system, available for anyone to use, which runs on Red Hat Linux as well as other Linux and UNIX systems. Red Hat, Inc. encourages other vendors to use RPM for their own products. RPM is distributable under the terms of the GPL.

For the end user, RPM makes system updates easy. Installing, uninstalling, and upgrading RPM packages can be accomplished with short commands. RPM maintains a database of installed packages and their files, so you can invoke powerful queries and verifications on your system. If you prefer a graphical interface, you can use Gnome-RPM to perform many RPM commands.

During upgrades, RPM handles configuration files carefully, so that you never lose your customizations — something that you will not accomplish with regular `.tar.gz` files.

For the developer, RPM allows you to take software source code 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.

Run RPM Commands as Root

Because RPM makes changes to your system, you must be root in order to install, remove, or upgrade an RPM package.

17.1 RPM Design Goals

In order to understand how to use RPM, it can be helpful to understand RPM's design goals:

Upgradability

Using 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. There are no special upgrade files need to upgrade a package because the same RPM file is used to install and upgrade the package on your system.

Powerful Querying

RPM is designed to provide powerful querying options. You can do searches through your entire database for packages or just for certain files. You can also easily find out what package

a file belongs to and from where the package came. 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 an important advantage for several reasons. For instance, if a new version of a program comes out, you do not 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 using this technique.

The goal of keeping sources pristine 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.

17.2 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 17.5, *Additional Resources* for more information on RPM.

17.2.1 Finding RPMs

Before using an RPM, you must know where to find them. An Internet search will return many RPM repositories, but if you are looking for RPM packages built by Red Hat, they can be found at the following locations:

- The official Red Hat Linux CD-ROMs
 - The Red Hat Errata Page available at <http://www.redhat.com/support/errata>
 - A Red Hat FTP Mirror Site available at <http://www.redhat.com/mirrors.html>
 - Red Hat Network — See Chapter 19, *Red Hat Network* for more details on Red Hat Network
-

17.2.2 Installing

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

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

As you can see, RPM prints out the name of the package and then prints a succession of hash marks as the package is installed as a progress meter.

Note

Although a command like `rpm -ivh foo-1.0-1.i386.rpm` is commonly used to install an RPM package, 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 17.2.4, *Upgrading* for more information about using the `-U` RPM option.

Installing packages is designed to be simple, but you may sometimes see errors:

Package Already Installed

If the package of the same version is already installed, you will see:

```
# rpm -ivh foo-1.0-1.i386.rpm
foo package foo-1.0-1 is already installed
#
```

If you want to install the package anyway and the same version you are trying to install is already installed, you can use the `--replacepkgs` option, which tells RPM to ignore the error:

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

This option is helpful if files installed from the RPM were deleted or if you want the original configuration files from the RPM to be installed.

Conflicting Files

If you attempt to install a package that contains a file which has already been installed by another package or an earlier version of the same package, you'll see:

```
# rpm -ivh foo-1.0-1.i386.rpm
foo          /usr/bin/foo conflicts with file from bar-1.0-1
#
```

To make RPM ignore this error, use the `--replacefiles` option:

```
# 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 which has an unresolved dependency, you'll see:

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

17.2.3 Uninstalling

Uninstalling a package is just as simple as installing one. Type the following command at a shell prompt:

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

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 package name of the original package.

You can encounter a dependency error when uninstalling a package if another 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 this 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 the `--nodeps` option.

17.2.4 Upgrading

Upgrading a package is similar to installing one. Type the following command at a shell prompt:

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

What you do not 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 will work 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 the following:

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

This message 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 the differences between the two configuration files and resolve them as soon as possible, to ensure that your system continues to function properly.

Upgrading is really a combination of uninstalling and installing, so during an RPM upgrade you can encounter uninstalling and installing errors, 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
#
```

To cause RPM to "upgrade" anyway, use the `--oldpackage` option:

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

17.2.5 Freshening

Freshening a package is similar to upgrading one. Type the following command at a shell prompt:

```
# 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 for single packages or a group of packages. If you have just downloaded a large number of different packages, and you only want to upgrade those packages that are already installed on your system, freshening will do the job. If you use freshening, you will not have to deleting any unwanted packages from the group that you downloaded before using RPM.

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

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

RPM will automatically upgrade only those packages that are already installed.

17.2.6 Querying

Use the `rpm -q` command to query the database of installed packages. The `rpm -q foo` command will print the package name, version, and release number of the installed package `foo`:

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

Note

Notice that we used the package *name* `foo`. To query a package, you will need to replace `foo` with the actual package name.

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 which owns `<file>`. When specifying a file, you must specify the full path of the file (for example, `/usr/bin/lis`).
- `-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, READMEs, 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 (for example, `sendmail.cf`, `passwd`, `inittab`, etc.).

For the options that display lists of files, you can add `-v` to the command to display the lists in a familiar `ls -l` format.

17.2.7 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 of verifying 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 a 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 command 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 eight characters (a `c` denotes a configuration file) and then the file name. Each of the eight 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 . (a 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 fix the problem in another way.

17.3 Checking a Package's Signature

If you wish to verify that a package has not been corrupted or tampered with, examine only the md5sum by typing the following command at a shell prompt (replace coolapp with the filename of your RPM package):

```
rpm --checksig --nogpg coolapp-1.1-1.rpm
```

You'll see the message coolapp-1.1-1.rpm: md5 OK. This brief message means that the file was not corrupted by the download.

On the other hand, how trustworthy is the developer who created the package? If the package is **signed** with the developer's **GnuPG key**, you'll know that the developer really is who they say they are.

An RPM package can be signed using **Gnu Privacy Guard** (or **GnuPG**), to help you make certain your downloaded package is trustworthy.

GnuPG is a tool for secure communication; it is a complete and free replacement for the encryption technology of **PGP**, an electronic privacy program. With **GnuPG**, you can authenticate the validity of documents, and encrypt/decrypt data to and from other recipients. **GnuPG** is capable of decrypting and verifying **PGP 5.x** files, as well.

During the installation of Red Hat Linux, **GnuPG** is installed by default. That way you can immediately start using **GnuPG** to verify any packages that you receive from Red Hat. First, you will need to import Red Hat's public key.

17.3.1 Importing Keys

When you import a public key, you add that key to your **keyring** (a file in which public and secret keys are kept). Then, when you download a document or file from that entity, you can check the validity of that document against the key you added to your keyring.

To import a key, use the `--import` option. To demonstrate, download and import Red Hat's public key. That way, any time you want to validate a package from Red Hat, you will be able to check it against the key you retrieved.

You can find Red Hat's key at <http://www.redhat.com/about/contact.html>. Using your browser, download the key by pressing the [Shift] key while you click on the download link, then click the **OK** button to save the file (for example `redhat2.asc`). Then, at the shell prompt, import the key with the following command:

```
gpg --import redhat2.asc
```

The resulting message tells you that the key was processed. To check that the key was added, type `gpg --list-keys`. You'll see the key you just downloaded from Red Hat, as well as your own keys.

```
[newuser@localhost newuser]$ gpg --list-keys
/home/newuser/.gnupg/pubring.gpg
-----
pub  1024D/DB42A60E 1999-09-23 Red Hat, Inc <security@redhat.com>
sub  2048g/961630A2 1999-09-23
```

Keys Do Not Have to be Links

Sometimes, you will not be able to download a key from a link. Keys are text files, so they can be moved to your machine in any way a regular text file can be saved. As long as you know the name and location of the file you saved, you can import it to your keyring.

17.3.2 Verifying Packages

To check the GnuPG signature of an RPM file after importing the builder's GnuPG key, use the following command (replace `coolapp` with the filename of your RPM package):

```
rpm --checksig coolapp-1.1-1.rpm
```

If all goes well, you will see the message: `md5 gpg OK`. That means that the package is not corrupt.

17.3.3 More about GnuPG

For more information about GnuPG, see Appendix A, *Getting Started with Gnu Privacy Guard*.

17.4 Impressing Your Friends with RPM

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

- Perhaps you have deleted some files by accident, but you are not sure what you deleted. If you want to verify your entire system and see what might be missing, you could try the following command:

```
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.

- At some point, you might see a file that you do not recognize. To find out which package owns it, you would enter:

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

The output would look like the following:

```
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 do not know which package owns `paste`. Simply enter the following command:

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

and the appropriate package will be verified.

- Do you want to find out more information about a particular program? You can try the following command to locate the documentation which came with the package that owns that program:

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

The output would be like the following:

```
/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

```

- You may find a new RPM, but you don't know what it does. To find information about it, use the following command:

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

The output would look like the following:

```

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.

```

- Perhaps you now want to see what files the `sndconfig` RPM installs. You would enter the following:

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

The output will look like the following:

```
/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 a few examples. As you use it, you will find many more uses for RPM.

17.5 Additional Resources

RPM is an extremely complex utility with many options and methods for querying, installing, upgrading, and removing packages. Refer to the following resources to learn more about RPM.

17.5.1 Installed Documentation

- `rpm --help` — This command displays a quick reference of RPM parameters.
- `man rpm` — The RPM man page will give you more detail about RPM parameters than the `rpm --help` command.

17.5.2 Useful Websites

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

17.5.3 Related Books

- *Maximum RPM* by Ed Bailey; Red Hat Press — An online version of the book is available at <http://www.rpm.org/> and <http://www.redhat.com/support/books/>.
-

18 Gnome-RPM

Gnome-RPM provides a GUI interface for the Red Hat Package Manager (RPM). To learn more about RPM technology, turn to Chapter 17, *Package Management with RPM*.

If you do not want to use the command-line version of RPM, you can use 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 users to easily work with RPM technology and features a friendly interface.

Gnome-RPM is "GNOME-compliant," meaning that it seamlessly integrates into GNOME, a graphical X Window System desktop environment provided with Red Hat Linux.

Using Gnome-RPM, you can easily accomplish the following tasks:

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

The Gnome-RPM interface provides a menu, a toolbar, a tree, and a window which displays currently installed packages.

To perform a Gnome-RPM task, you usually find and select packages, then choose the type of operation to perform using either a button on the toolbar, from 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 components of the package except for configuration files you have modified.
- Upgrading a package installs the new version and uninstalls all other versions that were previously installed.

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 18.4, *Configuration* for more information about this feature.

Note

Be careful when using **Web find**, 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. Refer to Section 18.5.2, *Verifying Packages* to learn more about verifying packages.

Using **Gnome-RPM** to perform all of these and many other operations is the same as using **RPM** commands from the shell prompt. However, the graphical nature of **Gnome-RPM** may make these operations easier to perform. **Gnome-RPM** can display packages in a variety of different ways. Refer to Section 18.3, *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. Since **Gnome-RPM** is integrated with **GNOME**, you can also perform installation, query and verification on packages from within the **GNOME** File Manager.

18.1 Starting Gnome-RPM

To start **Gnome-RPM**, use one of the following methods:

- On the **GNOME** desktop, go to **Main Menu Button** (on the panel) => **Programs** => **System** => **GnoRPM**
- On the **KDE** desktop, go to **Main Menu Button** (on the panel) => **Programs** => **System** => **GnoRPM**
- At a shell prompt, type `gnorpm &`

You will see the main **Gnome-RPM** window (as shown in Figure 18–1, *Main Gnome-RPM Window*).

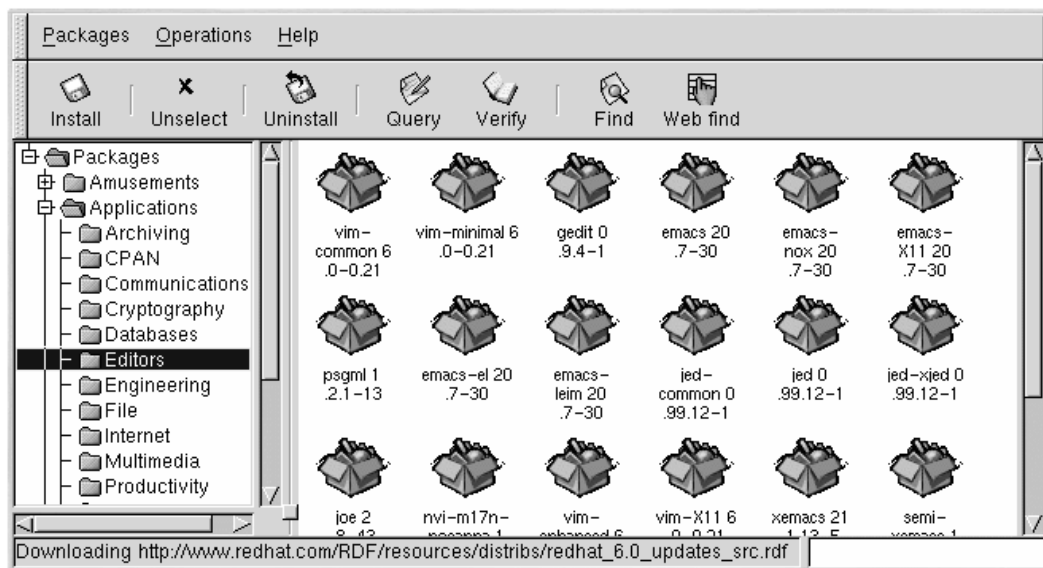
Note

If you would like to install, upgrade or uninstall packages, you must be root. The easiest way to become root is to type the `su` command and [Enter] at a shell prompt. Then type the root password. However, you do not have to be root to query and verify packages.

The Gnome-RPM interface consists of the following:

- Package Display — 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 18–1 Main Gnome-RPM Window



18.2 The Package Display

Each folder icon in the tree view at left represents a group of packages. Each group can contain subgroups. 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 can be expanded and collapsed, so you can 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 18.4, *Configuration* for more information about customizing Gnome-RPM settings.

18.2.1 Selecting Packages

To select a single package, click on it with the left mouse button. When a package is selected, its title will be highlighted as shown in Figure 18–2, *Selecting Packages in Gnome-RPM*. To unselect a package, 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 18–2 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, hold down the [Ctrl] key and left-click on packages; you'll see highlighting around each selection.

To select a group of packages within a folder, left-click on one package. Hold down the [Shift] key and left-click on the final package you wish to select. You'll see that all of the packages between your starting and ending selections will be highlighted for selection.

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

18.3 Installing New Packages

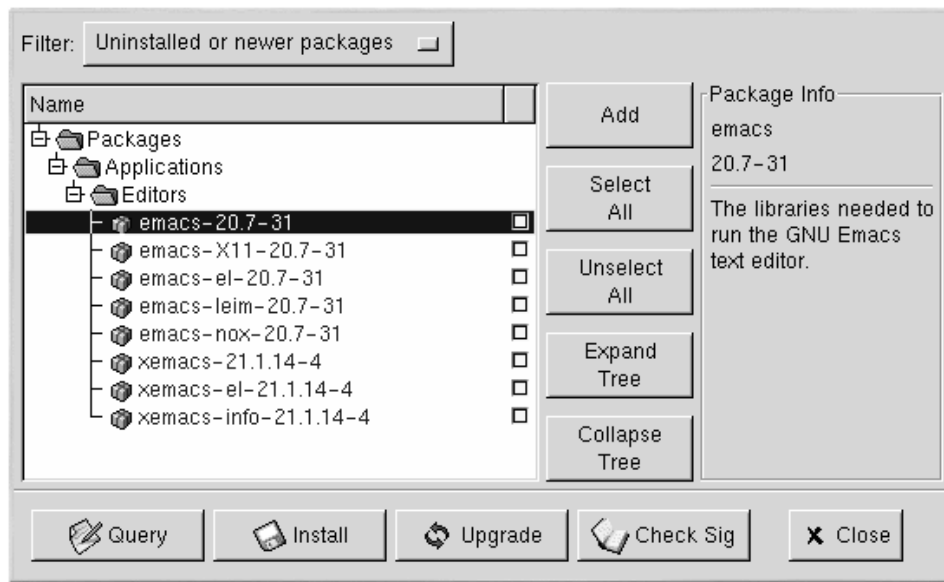
To install new packages, choose **Install** from the toolbar. In the **Install** window, your view will depend upon what you have selected under **Filter**.

Filter can be used to narrow your choices for viewing packages. Available filters include the following:

- All packages
- All packages except for those that are already installed

- Only uninstalled packages
- Only newer packages
- Uninstalled or newer packages

Figure 18–3 The Install Window



Click on 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 change the default path in the **Install Window** tab of the **Operations => Preferences** dialog. See Section 18.4, *Configuration* for more information.)

If no packages are available in the default path, you will see an **Add Packages** window. You can select the location of your new package using the standard file dialog window.

If you click on a package, you'll see 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, select the checkbox next to the package, then click on the **Install** button. You can also query a selected package. On the **Package Info** window, you can also perform the installation

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.

In addition to installing the packages from within the **Install** window, you can install a package 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 package you've selected to install, including the origination of the package, the date it was built, its size and more.

If the package already exists on your system and you're querying a newer version, the **Package Info** window provides an **Upgrade** button, which will upgrade the package to the newer version.

You can also drag and drop packages from **GNOME File Manager**. Within the File Manager, left-click on the selected package. 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 is displayed 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.

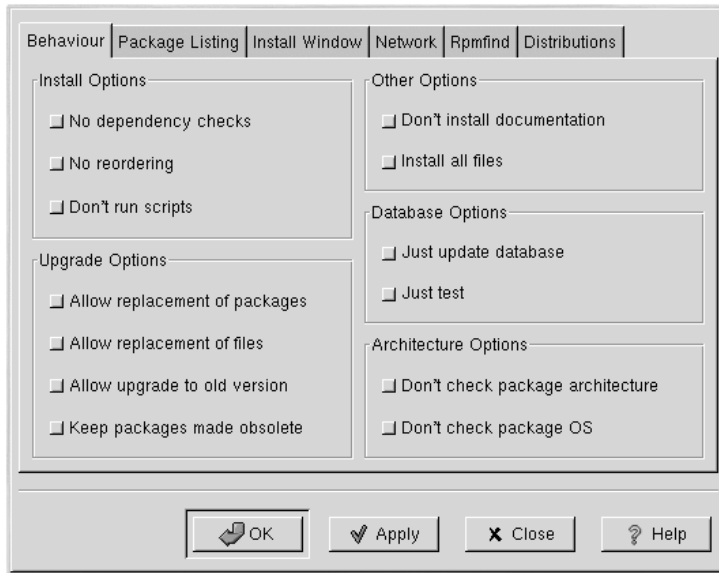
To install the package, select the **Install** button. You'll see a progress indicator as your package is being installed.

18.4 Configuration

Gnome-RPM offers a wide selection of choices for installing and uninstalling packages, documentation and other features. You can customize **Gnome-RPM** using the **Preferences** dialog, which you can access from **Operations => Preferences** on the menu. To make selections in the **Preferences** dialog, select the boxes next to the options.

Under the **Behaviour** tab, you'll find a number of options for configuring the way **Gnome-RPM** installs, uninstalls and upgrades packages. The **Behaviour** tab is divided into five sections: **Install Options**, **Upgrade Options**, **Other Options**, **Database Options** and **Architecture Options**. Note that by default these boxes are not selected (see Figure 18–4, *Behaviour Tab in Preferences*).

Figure 18–4 Behaviour 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 files that the program may depend on in order to work. Unless you know what you're doing, we strongly suggest that you not use this option as some packages may depend on other packages in order 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. Selecting this option is similar to the `--no-scripts` option when installing packages from the shell prompt.

Under **Upgrade Options**, you can select the following:

- **Allow replacement of packages** — Replaces a package with a new copy of the same package. Similar to the `--replacepks` option from the shell prompt. This option can be useful if an installed package has become damaged or requires 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 be useful when two packages include files that are named the same but contain different contents.
- **Allow upgrade to old version** — Like the shell prompt RPM command equivalent `--old-package`, this option allows you to "upgrade" to an earlier package. It can be useful if the latest version of a package doesn't function correctly on 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 where Gnome-RPM can find new RPMs on your system. Refer to Figure 18–5, *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
```

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 18–5 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 the install path and closing the dialog box, you can use the **Install** button to view the packages available in the new location.

(If the path for your RPMs doesn't match the default path in your preferences, you'll be presented with a window for browsing through your filesystem, which will allow you to select the correct path for your new RPMs.)


Under **Package Colours**, 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** field contains a list of default locations where Gnome-RPM will search for packages.

In the **Network** tab, you have the ability to specify proxies for use with HTTP and FTP transfers, as well as user and password names (see Figure 18–6, *Network Settings*). Note, however, that the password will not be stored securely.

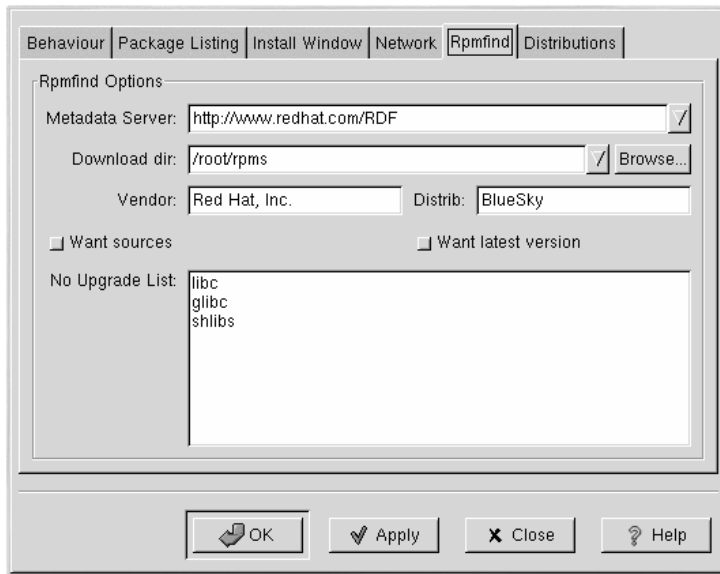
In the **Cache expire** field, you can set the length of time before data from the rpmfind database is considered to be out of date.

Figure 18–6 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 18–7, *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 18–7 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 18–8 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 18–8, *Distribution Settings in Preferences*), the higher the priority it will receive; a lower rating (such as "-1") will specify that packages not be recommended.


18.5 Package Manipulation

18.5.1 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 and then press the **Query** button on the menu.

You'll see a window like the one shown in Figure 18–9, *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 18–9 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 its group.

In the right column, you'll find the date that the package was installed 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."

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 see if any files are missing from the package (this probably means that there is a problem with the package).

If you're querying a package that's already installed, you'll also find two additional buttons 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.

18.5.2 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 18–10, *Verify Window*.

Figure 18–10 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.

18.5.3 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. Configuration files that have been modified are copied to `<filename>.rpmsave` so you can reuse them later.


Note

You must be root to uninstall packages.

If uninstalling a package would break "dependencies" (which could interfere with the operation of applications that require one or more of the removed files in the package), a dialog will pop up, asking you to confirm the deletion.

You can uninstall a selected package in a variety of ways: 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 more than one package in the same way as you would when installing, querying or verifying. The total number of selections will be displayed in the status bar on the bottom of the main window.

Figure 18–11 Uninstall Window



Once you've begun to uninstall packages, Gnome-RPM asks for confirmation, showing a window like the one in Figure 18–11, *Uninstall Window*. All of the packages that are about to be uninstalled are listed. You should carefully check the list to make sure that you're not about to remove something you want to keep. Clicking the **Yes** button will start the uninstallation process. After it is completed, the packages and groups that have been removed will disappear from any open windows.

Upgrading Packages

When a new version of a package is released, it is easy to install it on your system. Select the package from the window of available packages in the same way you select packages for installation. You can begin the upgrade process in two ways: either the **Upgrade** button on the toolbar or using **Operations => Upgrade** on the menu. You simply **Add** packages in the same manner as you would during a new package installation.

During the upgrade, you'll see a progress indicator like the one for shown when you are installing packages. When it's finished, any old versions of the packages will be removed, unless you specify otherwise (refer to Section 18.4, *Configuration* for more information).

In most cases, you should upgrade packages rather than uninstall the old versions of a package and then install the new ones. If you use upgrade, any changes you made to package configuration files are preserved properly. If you uninstall an old version of a package and then install a new package, your changes could 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.

19 Red Hat Network

Red Hat Network is an Internet solution for managing a Red Hat Linux system or a network of Red Hat Linux systems. All Security Alerts, Bug Fix Alerts, and Enhancement Alerts (collective known as Errata Alerts) can be downloaded directly from Red Hat using the Red Hat Update Agent standalone application or through a Web browser at <http://www.redhat.com/network/>.

After you register a system with Red Hat Network, your System Profile is used to deliver software packages, as soon as they are released by Red Hat, to you. Red Hat Network only informs you about Errata Alerts that are relevant to your registered system. The status of your system can be viewed through the Red Hat Update Agent or at <http://www.redhat.com/network/>.

Everyone receives a free Red Hat Network Software Manager subscription for one system. Additional subscriptions are \$19.95/month for each system. Red Hat is offering a special introductory \$9.95/month rate for systems subscribed before April 6, 2001.

If you have more than one system registered with Red Hat Network, you can view them all from one Web interface as shown in Figure 19–1, *Your Network*. You can download all the packages for all your registered systems at the same time. After downloading all the packages, you can then update them on the individual systems in your network group.

Figure 19–1 Your Network

The screenshot shows a web browser window displaying the Red Hat Network interface. The browser's address bar shows the URL <http://www.redhat.com/network/yn/>. The page header includes the Red Hat Network logo, the user name "User: William", and a "LOG OUT" button. A navigation menu contains "Main", "Your Network", "Search Errata Alerts", "Preferences", and "Help Desk". The main content area is titled "Your Network" and features a "Page Help" link. Below this, there are three tabs: "Systems Overview", "Errata Alerts", and "Assign Service Levels". The "Errata Alerts" tab is active, showing a table of alerts for various systems. The table has columns for "Alerts", "System Name", "Service Level", "Description", and "New Packages". A legend indicates that a square icon represents a Security Alert, a square with an 'X' represents a Bug Fix Alert, and a square with a plus sign represents an Enhancement Alert. The table lists several systems, including "falcon.test.redhat.com", "development.system", "rus.test.redhat.com", "hacup.system", "test3.test.redhat.com", "slot.devel.redhat.com", "luc.devel.redhat.com", "test5.system", "database.server", "web.server", and "Not Configured". Each system row includes a "Remove Profile" link. The footer of the page contains the copyright notice: "Copyright © 2001 Red Hat, Inc. All rights reserved. Security and Privacy Policy."

Alerts	System Name	Service Level	Description	New Packages
	falcon.test.redhat.com	Software Manager	7.1 running on i686	0
	development.system	Software Manager	7.1 running on i586	0
	rus.test.redhat.com	Software Manager	7.1 running on i686	0
	hacup.system	Software Manager	7.1 running on i686	0
	test3.test.redhat.com	no service [upgrade]	7.1 running on i586	0
<input checked="" type="checkbox"/>	slot.devel.redhat.com	Software Manager	7.0 running on i686	1
	luc.devel.redhat.com	Software Manager	7.1 running on i686	0
	test5.system	Software Manager	7.1 running on i686	0
	database.server	Software Manager	7.1 running on i386	0
	web.server	Software Manager	7.1 running on i386	0
	Not Configured	no service		

For more information about Red Hat Network, read the *Red Hat Network User Reference Guide* available at <http://www.redhat.com/support/manuals/RHNetwork/ref-guide/>.

Part V Appendixes

A Getting Started with Gnu Privacy Guard

A.1 An Introduction to GnuPG

Have you ever wondered if your email can be read during its transmission from you to other people, or from other people to you? Unfortunately, complete strangers could conceivably intercept or even tamper with your email.

In traditional (also known as "snail") mail, letters are usually sealed within envelopes, stamped and delivered from post office branch to branch until they reach their destination. But sending mail through the Internet is much less secure; email is usually transmitted as unencrypted text from server to server. No special steps are taken to protect your correspondence from being seen or tampered with by other people.

To help you protect your privacy, Red Hat Linux 7.1 includes GnuPG, the GNU Privacy Guard, which is installed by default during a typical Red Hat Linux installation. It is also referred to as GPG.

GnuPG is a tool for secure communication; it is a complete and free replacement for the encryption technology of PGP (Pretty Good Privacy, a widely popular encryption application). Using GnuPG, you can encrypt your data and correspondence, and authenticate your correspondence by **digitally signing** your work. GnuPG is also capable of decrypting and verifying PGP 5.x.

Because GnuPG is compatible with other encryption standards, your secure correspondence will probably be compatible with email applications on other operating systems, such as Windows and Macintosh.

GnuPG uses **public key cryptography** to provide users with a secure exchange of data. In a public key cryptography scheme, you generate two keys: a public key and a private key. You exchange your public key with correspondents or with a keyserver; you should never reveal your private key.

Encryption depends upon the use of keys. 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 key 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.

Do Not Reveal Your Private Key

Remember that your public key can be given to anyone with whom you want to communicate securely, but you must never give away your private key.

For the most part, cryptography is beyond the scope of this publication; volumes have been written about the subject. In this chapter, however, we hope you'll gain enough understanding about GnuPG to begin using cryptography in your own correspondence. For more information about GnuPG, including an online users guide, visit <http://www.gnupg.org/>. If you want to learn more about GnuPG, PGP and encryption technology, see Section A.7, *Additional Resources*.

More Information From the Shell Prompt

Like most system tools for Red Hat Linux, you'll find documentation on GnuPG in the man pages and info pages. At a shell prompt, just type `man gpg` or `info gpg` for a quick reference of GnuPG commands and options.

A.2 Generating a Keypair

To begin using GnuPG, you must first generate a new keypair: a public key and a private key.

To generate a keypair, at a shell prompt, type the following command:

```
gpg --gen-key
```

Since you work with your user account most frequently, you should perform this action while logged in to your user account (and not as root).

You will see an introductory screen, with key options, including one recommended option (the default), similar to the following:

```
gpg (GnuPG) 1.0.1; Copyright (C) 1999 Free Software Foundation, Inc.  
This program comes with ABSOLUTELY NO WARRANTY.  
This is free software, and you are welcome to redistribute it  
under certain conditions. See the file COPYING for details.
```

```
Please select what kind of key you want:  
(1) DSA and ElGamal (default)  
(2) DSA (sign only)  
(4) ElGamal (sign and encrypt)  
Your selection?
```

```
sub 1024g/E12AF9C4 2000-04-18
```

A.3 Generating a Revocation Certificate

Once you have created your keypair, you should create a revocation certificate for your public key. If you forget your passphrase, or if it has been compromised, you can publish this certificate to inform users that your public key should no longer be used.

Why Revoke a Key You Just Created?

When you generate a revocation certificate, you are not revoking the key you just created. Instead, you're giving yourself a safe way to revoke your key from public use. Let's say you create a key, then you forget your passphrase, switch ISPs (addresses), or suffer a hard drive crash. The revocation certificate can then be used to disqualify your public key.

Your signature will be valid to others who read your correspondence before your key is revoked, and you will be able to decrypt messages received prior to its revocation. To generate a revocation certificate, use the `--gen-revoke` option.

```
[newuser@localhost newuser]$ gpg --output revoke.asc
--gen-revoke <you@yourisp.net>
```

Note that if you omit the `--output revoke.asc` option from the above, your revocation certificate will be returned to the standard output, which is your monitor screen. While you can copy and paste the contents of the output into a file of your choice using a text editor, such as Pico, it is probably easier to send the output to a file in your login directory. That way, you can keep the certificate for use later, or move it to a floppy disk and store it someplace safe.

The creation of a revocation certificate will look like the following:

```
[newuser@localhost newuser]$ gpg --output revoke.asc
--gen-revoke <you@yourisp.net>

sec 1024D/823D25A9 2000-04-26 Your Name <you@yourisp.net>

Create a revocation certificate for this key? y

You need a passphrase to unlock the secret key for
user: "Your Name <you@yourisp.net>"
1024-bit DSA key, ID 823D25A9, created 2000-04-26

ASCII armored output forced.
Revocation certificate created.
```

Once your revocation certificate has been created (`revoke.asc`), it will be located in your login directory. You should copy the certificate to a floppy diskette and store it in a secure place. (If you don't know how to copy a file to a diskette in Red Hat Linux, see the *Official Red Hat Linux Getting Started Guide*.)

A.4 Exporting your Public Key

Before you can use public key cryptography, other people must have a copy of your public key. To send your key to correspondents or to a keyserver, you must **export** the key.

To export your key, so you can display it on a Web page or paste it in email, type the following:

```
[newuser@localhost newuser]$ gpg --armor --export  
<you@yourisp.net> > mykey.asc
```

You will not see any output, because not only did you export your public key, you redirected the output to a file called, for example, `mykey.asc`. (Without the addition of `> mykey.asc`, the key would have been displayed as the standard output on the monitor screen.)

Now, the file `mykey.asc` can be inserted into email or exported to a keyserver. To see the key, type `less mykey.asc` to open the file in a pager (type `[q]` to quit the pager). It should look like the following:

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
```

```
Version: GnuPG v1.0.1 (GNU/Linux)
```

```
Comment: For info see http://www.gnupg.org
```

```
mQGIBDkHP3URBACKWGsYh43pkXU9wj/X1G67K8/DSr185r7dNtHNFLL/ewill10k2  
q8saWJn26QZPsDVqdUJModHfJ6kQTAt9NzQbgcVrxLYNfgeBsvkHF/PotnYcZRgL  
tZ6syBBWs8JB4xt5V09iJSGAMPUE8Jpdn2aRXPapdoDwl79LM8Rq6r+gwCg5ZZa  
pGNlkgFu24WM5wC1zg4QTbMD/3MJCSxfL99Ek5HXcB3yhj+o0LmIrGAVBgoWdrRd  
BIGjQQFhV1NSwC8YhN/4nGHWpaTxgEtnb4CI1wI/G3DK9o1YMyRJinkGJ6XYfP3b  
cCQmqATDF5ugIAmdditnw7deXqn/eavaMxRXJM/RQSgJjYVpbAO20qKe6L6Inb5H  
kjcZA/9obTm499dDMRQ/CNR92fA5pr0zriy/ziLUow+cqI59nt+beB9nYlmfmUN6  
SW0jCH+pIQH5lerV+EookyOyq3ocUdjerYF/d2jl9xmeSyL2H3tDvnuE6vgqFU/N  
sdvby4B2Iku7S/h06W6GPQAe+pzdyX9vS+Pnf8osu7W3j60WprQkUGF1bCBHYWxs  
YWdoZXIgpPHbhdWxnYWxsQHJlZGhhdC5jb20+iFYEExECABYFAjkHP3UECwoEAwMV  
AwIDfGIBAheAAAoJEJECmvGCPSWpMjQAoNF2zvRgdR/8or9pBhu95zeSnb7AKCm  
/uXVS0a5KoN7J6l/1vEwx11poLkBDQQ5Bz+MEAQA8ztcWRJjW8cHCgLaE402jyqQ  
37gDT/n4VS66nU+YItzDFScVmgMuFRzhibLb1fO9TpZzxEBsf3T6p9hLLnHCQ1bD  
HRsKFh0eJYMMqB3+HyUpNeqCMEEd9AnWD9P4rQtO7Pes38sV01X0OSvsTyMG9wEB  
vSNZk+r1+phA55r1s8cAAwUEAJjqazvk0bgFrw1OPG9m7fEeDlvPSV6HSA0fvz4w  
c7ckfpuxg/URQnf3TJA00Acprk8Gg8J2CtebAyR/sp5IsrK511luGdk+l0M85FpT  
/cen2OdJtToAF/6fGnIkeCeP105awTbDgdAUHBrykpdWU3GJ7NS6923fVg5khQWg  
uwrAiEYEGBECAAYFAjkHP4wAcGkQkQKa8YI9JamlwCfXox/HjlorMKNQRJkeBcZ  
iLyPH1QAoI33Ft/0HBqLqtqdtP4vWYQRb1bjw
```

```
=BMEc
-----END PGP PUBLIC KEY BLOCK-----
```

A.4.1 Exporting to a Keyserver

If you are only writing to a few correspondents, you can export your public key and send it to them personally. If you correspond with many people, however, distribution of your key can be time consuming. Instead, you can use a keyserver.

Figure A-1 The Home Page of Keyserver.Net



A keyserver is a repository on the Internet which can store and distribute your public key to anyone who requests it. Many keystores are available, and most try to remain synchronized with each other; sending your key to one keyserver is like distributing it to them all. A correspondent can request your

public key from a from a keyserver, import that key to their keyring, and they are ready for secure correspondence with you.

Which Keyserver Should You Use?

Because most keyservers are synchronized, sending your public key to one keyserver is usually as good as sending it to them all. You can, however, locate different keyservers. One place to begin your search for keyservers and more information is *Keyserver.Net*, at <http://www.keyserver.net>; another location is *Robert's Crypto & PGP Links: Keyservers*, at <http://crypto.yashy.com/www/Keyservers/>.

You can send your public key from either the shell prompt or from a browser (as in Figure A-1, *The Home Page of Keyserver.Net*); of course, you must be online to send or receive keys from a keyserver.

- From the shell prompt, type the following:

```
gpg --keyserver search.keyserver.net --send-key you@yourisp.net
```

- From your browser, go to Keyserver.Net (<http://www.keyserver.net>) and select the option to add your own PGP public key.

Your next task is to copy and paste your public key into the appropriate area on the Web page. If you need instructions on how to do that, use the following:

- Open your exported public key file (such as *mykey.asc*, which was created in Section A.4, *Exporting your Public Key*) with a pager — for example, use the `less mykey.asc` command.
- Using your mouse, copy the file by highlighting all the lines from the `BEGIN PGP` to `END PGP` notations (see Figure A-2, *Copying Your Public Key*).
- Paste the contents of the file *mykey.asc* into the appropriate area of the page on Keyserver.Net by middle-clicking with your mouse (or left- and right-clicking if you're using a two-button mouse). Then select the **Submit** button on the keyserver page. (If you make a mistake, press the **Reset** button on the page to clear your pasted key.)

Figure A–2 Copying Your Public Key



```

File Edit Settings Help
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: GnuPG v1.0.1 (GNU/Linux)
Comment: For info see http://www.gnupg.org

mQGIBDKHP3URBACKwGsYh43pkXU9wj/X1G67K8/DSr185r7dNtHNFLL/ewil10k2
q8salwJn26QZPsDvqduJM0dHFJ6kQTAt9NzQbqcVrxLYNfgeBsvkHF/POtnYcZRgL
tZ6syBBWls8JB4xt5V09iJSGAMPUQE8Jpdn2aRXPapdoDw179LM8Rq6r+gwCg5Zza
pGN1kgFu24wM5wC1zg4QTbMD/3MJCSxFL99EK5HXcB3yhj+o0LmIrGAVBgolwdrRd
BIGjQQFhV1NSwC8YhN/4nGHwpaTxgEtnb4CI1wI/G3DK9o1YMyRJinkGJ6XYFP3b
cCQmqATDF5ugIAmdditnw7deXqn/eavaHxRXJM/RQ5gJJyVpbA020qKe6L6Inb5H
k,jcZA/9obTm499dDMRQ/CNR92fA5pr0zriy/ziLUow+cqI59nt+bEb9nY1mfMUN6
SW0jCH+pIQH51erV+EookyOyq3ocUdjeRYF/d2j19xmeSyL2H3tDvnuE6vgqFU/N
sdvby4B2Iku7S/h06W6GPQAE+pzdYX9vS+Pnf8osu7w3j60WprQkUGF1bCBHYlxs
YwdoZXIghPHBhdwXnYlxsQHJ1ZGhhdC5jb20+iFYEExECABYFAjKHP3UECwoEAWMV
AwIDFgIBaheAAA0JEJECmvGCPSWpMjQAoNF2zvRgdR/8or9pBhu95zeSnkb7AKCm
/uXVS0a5KoN7J6L/1vEwx11poLkBDQ05Bz+MEAQA8ztclWRJjw8cHCgLaE402jyqQ
37gDT/n4VS66nU+YItzDFScVmgMuFRzhibLb1f09TpZzxEbSF3T6p9hLLnHCQ1bD
HRsKfh0eJYMMqB3+HyUpNeqCMEEd9AnwD9P4rQt07Pes38sV01X00SvsTyMG9wEB
vSNZk+R1+phA55r1s8cAAwUEAJjqazvk0bgFru10PG9m7fEeD1vPSV6HSAOfvz4w
c7ckfpuxg/URQNF3TJA00Acprk8Gg8J2CtebAyR/sP5IsrK511luGdk+10M85FpT
/cen20dJtToAF/6fGnIkeCeP105awTbDgdAUHBRykpdlU3GJ7NS6923fVg5khQWg
uwrAiEYEGBECAAYFAjKHP4wACgkQkQKa8YI9JamlIwCfXox/HjlorMKnQRJkeBcZ
iLyPH1QAoI33Ft/OHBqLqdtP4wWYQRbIbjw
=BMEc
-----END PGP PUBLIC KEY BLOCK-----
mykey.asc (END)

```

Note that if you are submitting your key to another Web-based keyserver, the above transaction will be essentially the same.

That is all you need to do. Regardless of whether you use the shell prompt or the Web, you will see a message that your key was successfully submitted — either at the shell prompt or at the keyserver’s website. From now on, users who want to communicate securely with you can import your public key and add it to their keyring.

A.5 Importing a Public Key

The other end of key exchange — importing other people’s public keys to your keyring — is just as simple as exporting keys. When you import someone’s public key, you can decrypt their mail and check their digital signature against their public key on your keyring.

One of the easiest ways to import a key is to download the key or save it from a website. To learn how to import Red Hat's key, refer to Section 17.3.1, *Importing Keys*.

After downloading a key, use the command `gpg --import key.asc` to add it to your keyring.

Another way to save a key is to use a browser's **Save As** feature. If you are using a browser such as Navigator, and you locate a key at a keyserver, you can save the page as a text file (go to **File => Save As**). In the drop-down box next to **Format for saved document**, choose **Text**. Then, you can import the key — but remember the name of the file you saved. For example, if you saved a key as a text file called *newkey.txt*, to import the file, at a shell prompt, type:

```
[newuser@localhost newuser]$ gpg --import newkey.txt
gpg: key F78FFE84: public key imported
gpg: Total number processed: 1
gpg:          imported: 1
```

To check that the process was successful, use the `gpg --list-keys` command; you should see your newly imported key listed on your keyring.

A.6 What Are Digital Signatures?

Digital signatures can be compared to your written signature. Unlike traditional correspondence, in which it might be possible to tamper with your written signature, digital signatures can not be forged. That is because the signature is created with your unique secret key, and can be verified by your recipient using your public key.

A digital signature timestamps a document; essentially, that means that the time you signed the document is part of that signature. So if anyone tries to modify the document, the verification of the signature will fail. Some email applications, such as Exmh or KDE's KMail, include the ability to sign documents with GnuPG within the application's interface.

Two useful types of digital signatures are **clearsigned** documents and **detached signatures**. Both types of signatures incorporate the same security of authenticity, without requiring your recipient to decrypt your entire message.

In a clearsinged message, your signature appears as a text block within the context of your letter; a detached signature is sent as a separate file with your correspondence.

A.7 Additional Resources

There is more to encryption technology than can be covered in one slim introduction to GnuPG. Here are some resources where you can learn more.

A.7.1 Useful Websites

- <http://www.gnupg.org> — The GnuPG website with links to the latest GnuPG releases, a comprehensive user's guide, and other cryptography resources.
- <http://hotwired.lycos.com/webmonkey/backend/security/tutorials/tutorial1.html> — Visit the *Encryption Tutorial* from Webmonkey to learn more about encryption and how to apply encryption techniques.
- <http://www.eff.org/pub/Privacy> — The Electronic Frontier Foundation, "Privacy, Security, Crypto, & Surveillance" Archive.

A.7.2 Related Books

- *The Official PGP User's Guide* by Philip R. Zimmerman; MIT Press
 - *PGP: Pretty Good Privacy* by Simson Garfinkel; O'Reilly & Associates, Inc.
 - *E-Mail Security: How to Keep Your Electronic Messages Private* by Bruce Schneier; John Wiley & Sons
-

Index

A

accounts
 deleting with Linuxconf 151
 disabling with Linuxconf 150
 management 144
 modifying 149
 anonymous FTP 69
 Apache 97
 (See also Apache Configuration Tool)
 additional resources 117
 related books 118
 Apache Configuration Tool
 Directives
 (See Apache directives)
 error log 101
 Modules 97
 transfer log 101
 Apache directives
 DirectoryIndex 100
 ErrorDocument 100
 ErrorLog 102
 Group 114
 HostnameLookups 102
 KeepAlive 115
 KeepAliveTimeout 115
 Listen 98
 LogFormat 102
 LogLevel 102
 MaxClients 115
 MaxKeepAliveRequests 115
 Options 100
 ServerAdmin 98
 ServerName 98
 Timeout 115
 TransferLog 102
 User 113
 apacheconf 97
 autofs 80

/etc/auto.master 80

B

BIND configuration 119
 adding a forward master zone 120
 adding a reverse master zone 122
 adding a slave zone 124
 applying changes 120
 default directory 120
 bindconf 119
 (See also BIND configuration)
 booting
 emergency mode 57
 rescue mode 55
 single-user mode 57

C

chkconfig 65
 configuration
 anonymous FTP 69
 Ethernet 174
 Gnome-RPM 207
 hosts 172
 network device
 adding 172
 network routes 174
 NFS 79
 PLIP 174
 pocket network adapters 174
 selecting name servers 171
 SLIP 174
 system 141
 Token Ring 174
 Control Panel 169

D

date
 setting 175

decryption
 with GnuPG 223
 devices
 network
 clone..... 173
 df 92
 DSA keys..... 74
 du 93
 dual-boot
 FIPS partitioning tool..... 21
 LILO warning 18
 making room
 using FIPS to partition..... 21
 making room for
 adding a new hard drive 20
 creating new partitions..... 21
 using current partitions or hard drive . 20
 options
 booting Red Hat Linux or Windows.. 17
 partitionless installation 17
 Red Hat Linux as the only OS 18
 Windows NT warning 18
 OS/2..... 19
 setting up 19
 when using X..... 17

E

enabling accounts 151
 encryption
 with GnuPG 223
 /etc/auto.master 80
 /etc/fstab 79
 /etc/hosts
 managing 172
 /etc/hosts.lpd..... 133
 /etc/httpd/conf/httpd.conf 97
 /etc/printcap 127
 /etc/printcap.local..... 127
 Ethernet..... 174
 exporting NFS filesystems 81

F

FAT32 filesystems
 accessing..... 160
 filesystem
 NFS
 (See NFS)
 overview of 157
 viewing filesystem with Linuxconf 158
 free..... 91
 ftp 71
 FTP
 anonymous..... 69
 ftpassess 69
 ftphosts 69
 ftpusers 69

G

GDiskFree 92
 gnome-lokkit 67
 Gnome-RPM 201
 configuration 207
 installing packages 205
 package display 204
 package manipulation 213
 querying packages 213
 removing packages with..... 215
 selecting packages 204
 starting 202
 uninstalling packages with 215
 upgrading packages with 216
 verifying packages 215
 Gnu Privacy Guard 194
 Red Hat key 195
 using 223
 GnuPG
 (See Gnu Privacy Guard)
 GPG
 (See Gnu Privacy Guard)
 groups

- creating 153
 - deleting 155
 - management 153
 - modifying 156
- H**
-
- hostname..... 171
 - hosts
 - managing 172
- I**
-
- information
 - about your system..... 89
 - initrd 181
 - installation
 - kickstart
 - (See kickstart installations)
- K**
-
- kernel
 - building 177, 182
 - custom 177, 182
 - initrd image for..... 181
 - modular 177
 - module loader (kmod)..... 182
 - monolithic 182
 - kickstart
 - how the file is found 29
 - Kickstart Configurator
 - Basic Configuration** 32
 - Installation Source**..... 32
 - Partition Information** 32
 - kickstart file
 - auth..... 36
 - clearpart 39
 - device 39–40
 - diskette-based 28
 - driver disk 40
 - firewall 40
 - format of..... 30
 - install..... 41
 - installation methods..... 41
 - keyboard 42
 - lang 43
 - lilo 43
 - lilocheck 44
 - mouse 44
 - network..... 45–46
 - network-based 28
 - options 36
 - package selection specification 51
 - post-installation configuration 53
 - pre-installation configuration 53
 - raid 48
 - reboot 49
 - rootpw 49
 - skipx 50
 - timezone 50
 - upgrade..... 50
 - what it looks like..... 30
 - xconfig..... 50
 - zerombr..... 51
 - kickstart installations 27
 - diskette-based 28
 - file format..... 30
 - file locations 27
 - network-based 28
 - starting 29
 - ksconfig..... 32
- L**
-
- LILO
 - /etc/lilo.conf 180
 - Linuxconf 141
 - account management with..... 144
 - account modification 149
 - changing root password with..... 150
 - changing user's passwords 150
 - configuring network connections with . 162

deleting an account with 151
 deleting groups with 155
 disabling account with 150
 enabling accounts with 151
 Gnome-Linuxconf 142
 group creation with 153
 group management with 153
 group modification with 156
 nameserver specification with 164
 network configuration with 161
 NFS mount addition with 160
 quick reference 167
 reviewing filesystem 158
 user interfaces 141
 Web access 143
 loading kernel modules 182
 lpd 129

M

Maximum RPM 199
 mounting
 NFS filesystems 79

N

name servers
 selecting 171
 named.conf 119
 nameservers
 specifying
 using Linuxconf 164
 netcfg 170
 network
 adapters, pocket 174
 configuration 170
 adding device 172
 with Linuxconf 161–162
 devices
 clone 173
 interface
 aliasing 171

 routes
 managing 174
 Network Configurator 170
 Network File System
 (See NFS)
 NFS
 additional resources 81
 autofs
 (See autofs)
 configuration 79
 /etc/fstab 79
 exporting 81
 mounting 79
 with Linuxconf 160
 ntsysv 65

O

O'Reilly & Associates, Inc. 82, 118, 232
 OpenSSH 71
 additional resources 77
 authorization key pairs 73
 client 72
 /etc/ssh/sshd_config 71
 scp 71
 server 71
 starting and stopping 71
 sftp 73
 ssh 72
 ssh-keygen
 DSA 74
 RSA 75
 OpenSSL
 additional resources 77

P

packages
 dependencies 190
 determining file ownership with 196
 finding deleted files from 196

- freshening with RPM..... 191
 - installing 189
 - with Gnome-RPM 205
 - locating documentation for..... 196
 - obtaining list of files 198
 - preserving configuration files 191
 - querying 192
 - querying uninstalled 197
 - querying with Gnome-RPM 213
 - removing..... 190
 - selecting
 - with Gnome-RPM 204
 - tips..... 196
 - uninstalling with Gnome-RPM 215
 - upgrading 191
 - upgrading with Gnome-RPM..... 216
 - verifying 193
 - verifying with Gnome-RPM 213, 215
 - password
 - changing 150
 - PLIP
 - interface..... 174
 - pocket network adapters..... 174
 - printconf
 - (See printer configuration)
 - printer configuration 127
 - default printer 138
 - edit existing printer 138
 - local printer..... 129
 - modifying existing printers..... 138
 - Novell NetWare (NCP) printer..... 135
 - overriding a printer..... 139
 - printer aliases..... 138
 - remote UNIX printer..... 131
 - rename existing printer 138
 - Samba (SMB) printer 133
 - strict RFC1179 compliance 132
 - test page..... 138
 - printtool
 - (See printer configuration)
 - /proc directory 95
 - processes..... 89
 - currently running..... 89
 - ps 89
- ## R
-
- RAID
 - software..... 59
 - Red Hat Network 219
 - Red Hat Package Manager
 - (See RPM)
 - rescue mode 55
 - definition of 55
 - from CD, diskette, network, PCMCIA .. 55
 - using 55
 - utilities available 56
 - root password
 - changing 150
 - routes
 - managing 174
 - RPM 187
 - additional resources..... 198
 - book about..... 199
 - checking package signatures 194
 - dependencies 190
 - design goals 187
 - determining file ownership with 196
 - documentation with..... 196
 - file conflicts
 - resolving 189
 - finding deleted files with 196
 - freshen 191
 - freshening packages 191
 - GnuPG 194
 - installing 189
 - md5sum..... 194
 - preserving configuration files 191
 - querying 192
 - querying for file list..... 198
 - querying uninstalled packages 197
 - tips..... 196

- uninstalling 190
 - upgrading 191
 - using 188
 - verifying 193
 - website 199
- S**
-
- Samba 83
 - additional resources 85
 - configuration 83
 - smb.conf 83
 - reasons for using 83
 - share
 - connecting to 84
 - with Windows 2000 84
 - with Windows NT 4.0 84
 - scp
 - (See OpenSSH)
 - security 65
 - services
 - controlling access to 65
 - sftp
 - (See OpenSSH)
 - SLIP
 - interface 173
 - SMB protocol 83
 - smb.conf 83
 - Software RAID
 - creating partitions 59
 - ssh
 - (See OpenSSH)
 - Sysreport 94
 - system
 - configuration
 - with Linuxconf 141
 - system information
 - gathering 89
- T**
-
- TCP wrappers 67
 - telnet 71
 - time
 - setting 175
 - Token Ring 174
 - top 89
- U**
-
- upgrading
 - packages with Gnome-RPM 216
 - users
 - adding 144
- W**
-
- Windows
 - file and print sharing 83
 - Windows 2000
 - connecting to shares using Samba 84
 - Windows NT 4.0
 - connecting to shares using Samba 84
- X**
-
- xinetd 66