

Red Hat Linux 8.0

**The Official Red Hat Linux
Customization Guide**



Red Hat Linux 8.0: The Official Red Hat Linux Customization Guide

Copyright © 2002 by Red Hat, Inc.



Red Hat, Inc.

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

rhl-cg(EN)-8.0-Print-RHI (2002-08-14T17:28-0400)

Copyright © 2002 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.

Intel and Pentium are a registered trademarks of Intel Corporation. Itanium and Celeron are trademarks of Intel Corporation. AMD, AMD Athlon, AMD Duron, and AMD K6 are trademarks of Advanced Micro Devices, Inc.

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.

The GPG fingerprint of the security@redhat.com key is:

CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

Table of Contents

Introduction	vii
1. Changes to This Manual	vii
2. Document Conventions	viii
3. More to Come	xi
3.1. Send in Your Feedback	xi
4. Sign Up for Support	xi
I. File Systems	xiii
1. The ext3 File System	15
1.1. Features of ext3	15
1.2. Creating an ext3 File System	15
1.3. Converting to an ext3 File System	16
1.4. Reverting to an ext2 File System	16
2. Swap Space	19
2.1. What is Swap Space?	19
2.2. Adding Swap Space	19
2.3. Removing Swap Space	20
2.4. Moving Swap Space	21
3. Redundant Array of Independent Disks (RAID)	23
3.1. What is RAID?	23
3.2. Who Should Use RAID?	23
3.3. Hardware RAID versus Software RAID	23
3.4. RAID Levels and Linear Support	24
4. Logical Volume Manager (LVM)	27
5. Managing Disk Storage	29
5.1. Viewing the Partition Table	30
5.2. Creating a Partition	30
5.3. Removing a Partition	32
5.4. Resizing a Partition	33
II. Installation-Related Information	35
6. Kickstart Installations	37
6.1. What are Kickstart Installations?	37
6.2. How Do You Perform a Kickstart Installation?	37
6.3. Creating the Kickstart File	37
6.4. Kickstart Options	38
6.5. Package Selection	52
6.6. Pre-installation Script	53
6.7. Post-installation Script	54
6.8. Making the Kickstart File Available	55
6.9. Making the Installation Tree Available	56
6.10. Starting a Kickstart Installation	57
7. Kickstart Configurator	59
7.1. Basic Configuration	59
7.2. Installation Method	60
7.3. Boot Loader Options	61
7.4. Partition Information	62
7.5. Network Configuration	65
7.6. Authentication	66
7.7. Firewall Configuration	67
7.8. X Configuration	68
7.9. Package Selection	71
7.10. Pre-Installation Script	72
7.11. Post-Installation Script	72
7.12. Saving the File	74

8. Rescue Mode.....	75
8.1. What is Rescue Mode?.....	75
8.2. Booting Rescue Mode.....	75
8.3. Booting Single-User Mode.....	77
9. Software RAID Configuration.....	79
10. LVM Configuration.....	83
III. Network-Related Configuration.....	87
11. Network Configuration.....	89
11.1. Overview.....	89
11.2. Establishing an Ethernet Connection.....	89
11.3. Establishing an ISDN Connection.....	91
11.4. Establishing a Modem Connection.....	92
11.5. Establishing an xDSL Connection.....	94
11.6. Establishing a Token Ring Connection.....	96
11.7. Establishing a CIPE Connection.....	97
11.8. Establishing a Wireless Connection.....	97
11.9. Managing Hosts.....	99
11.10. Managing DNS Settings.....	100
11.11. Activating Devices.....	101
11.12. Working with Profiles.....	101
11.13. Device Aliases.....	103
12. Basic Firewall Configuration.....	105
12.1. Security Level Configuration Tool	105
12.2. GNOME Lokkit	108
12.3. Activating the <code>iptables</code> Service.....	111
13. Controlling Access to Services.....	113
13.1. Runlevels.....	113
13.2. TCP Wrappers.....	114
13.3. Services Configuration Tool	115
13.4. <code>ntsysv</code>	116
13.5. <code>chkconfig</code>	116
13.6. Additional Resources.....	117
14. OpenSSH.....	119
14.1. Why Use OpenSSH?.....	119
14.2. Configuring an OpenSSH Server.....	119
14.3. Configuring an OpenSSH Client.....	119
14.4. Additional Resources.....	124
15. Network File System (NFS).....	125
15.1. Why Use NFS?.....	125
15.2. Mounting NFS File Systems.....	125
15.3. Exporting NFS File Systems.....	126
15.4. Additional Resources.....	130
16. Samba.....	131
16.1. Why Use Samba?.....	131
16.2. Configuring Samba.....	131
16.3. Connecting to a Samba Share.....	133
16.4. Additional Resources.....	133
17. Dynamic Host Configuration Protocol (DHCP).....	135
17.1. Why Use DHCP?.....	135
17.2. Configuring a DHCP Server.....	135
17.3. Configuring a DHCP Client.....	139
17.4. Additional Resources.....	140
18. Apache HTTP Server Configuration.....	143
18.1. Basic Settings.....	143
18.2. Default Settings.....	145

18.3. Virtual Hosts Settings	150
18.4. Server Settings	153
18.5. Performance Tuning	154
18.6. Saving Your Settings	155
18.7. Additional Resources	156
19. Apache HTTP Secure Server Configuration	157
19.1. Introduction	157
19.2. An Overview of Security-Related Packages	157
19.3. An Overview of Certificates and Security	159
19.4. Using Pre-Existing Keys and Certificates	160
19.5. Types of Certificates	160
19.6. Generating a Key	161
19.7. Generating a Certificate Request to Send to a CA	163
19.8. Creating a Self-Signed Certificate	164
19.9. Testing Your Certificate	165
19.10. Accessing Your Secure Server	166
19.11. Additional Resources	167
20. BIND Configuration	169
20.1. Adding a Forward Master Zone	169
20.2. Adding a Reverse Master Zone	171
20.3. Adding a Slave Zone	173
21. Mail Transport Agent (MTA) Configuration	175
IV. System Configuration	177
22. Console Access	179
22.1. Disabling Shutdown Via Ctrl-Alt-Del	179
22.2. Disabling Console Program Access	179
22.3. Disabling All Console Access	180
22.4. Defining the Console	180
22.5. Making Files Accessible From the Console	180
22.6. Enabling Console Access for Other Applications	181
22.7. The floppy Group	182
23. Time and Date Configuration	183
23.1. Time and Date Properties	183
23.2. Time Zone Configuration	184
24. User and Group Configuration	185
24.1. Adding a New User	185
24.2. Modifying User Properties	186
24.3. Adding a New Group	187
24.4. Modifying Group Properties	187
25. Gathering System Information	189
25.1. System Processes	189
25.2. Memory Usage	191
25.3. File Systems	192
25.4. Hardware	193
25.5. Additional Resources	194
26. Printer Configuration	197
26.1. Adding a Local Printer	198
26.2. Adding a Remote UNIX Printer	200
26.3. Adding a Samba (SMB) Printer	202
26.4. Adding a Novell NetWare (NCP) Printer	203
26.5. Adding a JetDirect Printer	205
26.6. Selecting the Print Driver and Finishing	206
26.7. Printing a Test Page	207
26.8. Modifying Existing Printers	208
26.9. Saving the Configuration File	209

26.10. Command Line Configuration	210
26.11. Managing Your Print Jobs.....	211
26.12. Configuring the CUPS Printing System.....	212
26.13. Additional Resources	214
27. Automated Tasks.....	215
27.1. Cron.....	215
27.2. Anacron.....	217
27.3. At and Batch	218
27.4. Additional Resources	220
28. Log Files	221
28.1. Locating Log Files	221
28.2. Viewing Log Files	221
28.3. Examining Log Files.....	222
29. Upgrading the Kernel.....	225
29.1. The 2.4 Kernel	225
29.2. Preparing to Upgrade	225
29.3. Downloading the Upgraded Kernel	226
29.4. Performing the Upgrade.....	227
29.5. Configuring the Boot Loader	228
30. Kernel Modules.....	231
30.1. Kernel Module Utilities	231
30.2. Additional Resources	232
V. Package Management	233
31. Package Management with RPM.....	235
31.1. RPM Design Goals	235
31.2. Using RPM	236
31.3. Checking a Package's Signature	241
31.4. Impressing Your Friends with RPM	242
31.5. Additional Resources	244
32. Package Management Tool	247
32.1. Installing Packages	247
32.2. Removing Packages	249
33. Red Hat Network	251
VI. Appendixes	253
A. Building a Custom Kernel.....	255
A.1. Building a Modularized Kernel.....	255
A.2. Making an initrd Image.....	257
A.3. Configuring the Boot Loader	257
A.4. Building a Monolithic Kernel	259
A.5. Additional Resources	260
B. Getting Started with Gnu Privacy Guard.....	261
B.1. An Introduction to GnuPG	261
B.2. Warning Messages.....	261
B.3. Generating a Keypair.....	262
B.4. Generating a Revocation Certificate.....	263
B.5. Exporting your Public Key	264
B.6. Importing a Public Key	266
B.7. What Are Digital Signatures?	267
B.8. Additional Resources	267
Index.....	269
Colophon.....	277



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 a step-by-step, task-oriented guide for configuring and customizing your system, this is the manual for you. This manual discusses many intermediate topics such as the following:

- Setting up a network interface card (NIC)
- Performing a Kickstart installation
- Configuring Samba shares
- Managing your software with RPM
- Determining information about your system
- Upgrading your kernel

This manual is divided into the following main categories:

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

This guide assumes you have a basic understanding of your Red Hat Linux system. If you need reference material which covers more basic issues such as configuring your desktop or playing audio CD-ROMs, please refer to the *Official Red Hat Linux Getting Started Guide*. If you need more advanced documentation such as an overview of the Red Hat Linux filesystem, please refer to the *Official Red Hat Linux Reference Guide*.

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



Note

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

<http://www.redhat.com/docs/manuals/linux>

1. Changes to This Manual

This manual has been expanded to include new features in Red Hat Linux 8.0 as well as topics requested by our readers. Significant changes to this manual include:

File Systems

New to this version is a file systems part. It discusses ext3, swap space, RAID, LVM, and managing disk storage with `parted`.

Kickstart

The kickstart options have been updated to include the new options in Red Hat Linux 8.0, and the **Kickstart Configurator** chapter has been updated to include many new features.

LVM Configuration

This new chapter discusses how to configure LVM during installation.

Network Configuration

The network configuration chapter now includes information on network profiles and network aliases.

Basic Firewall Configuration

This chapter includes the new **Security Level Configuration Tool**.

Network File System (NFS)

This chapter now includes how to export directories using the **NFS Server Configuration Tool**.

Apache HTTP Server

Both Apache HTTP Server chapters have been updated for version 2.0.

Console Access

How to use the `pam_timestamp` module has been added.

Log Files

This new chapter contains information on how to examine log files with the **Log Viewer**.

Package Management Tool

The **Gnome-RPM** application has been replaced with **Package Management Tool**, which this new chapter discusses.

2. Document Conventions

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

`command`

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

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

`filename`

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

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

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

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

application

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

Use **Mozilla** to browse the Web.

[key]

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

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

[key]-[combination]

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

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

text found on a GUI interface

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

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

top level of a menu on a GUI screen or window

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

Under **File** on a GNOME terminal, you will see the **New Tab** option that allows you to open multiple shell prompts in the same window.

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

Go to **Main Menu Button** (on the Panel) => **Programming** => **Emacs** to start the **Emacs** text editor.

button on a GUI screen or window

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

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

computer output

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

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

```
$ ls
Desktop          about.html      logs           paulwesterberg.png
Mail             backupfiles    mail           reports
```

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

prompt

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

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

user input

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

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

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



Note

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



Tip

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



Important

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



Caution

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

**Warning**

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

3. 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.

3.1. Send in Your Feedback

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

Be sure to mention the manual's identifier:

```
rhl-cg(EN)-8.0-Print-RHI (2002-08-14T17:28-0400)
```

By mentioning this manual's identifier, we will know exactly which version of the guide you have.

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

4. Sign Up for Support

If you have an official edition of Red Hat Linux 8.0, please remember to sign up for the benefits you are entitled to as a Red Hat customer.

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

- Official Red Hat support — Get help with your installation questions from Red Hat, Inc.'s support team.
- Red Hat Network — Easily update your packages and receive security notices that are customized for your system. Go to <http://rhn.redhat.com> for more details.
- *Under the Brim: The Official Red Hat E-Newsletter* — Every month, get the latest news and product information directly from Red Hat.

To sign up, go to <http://www.redhat.com/apps/activate/>. You will find your Product ID on 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 Installation Guide*.

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

The Red Hat Documentation Team

File Systems



Chapter 1.

The ext3 File System

Beginning with the release of Red Hat Linux 7.2, the default file system changed from the venerable ext2 format to the journaling *ext3* file system.

1.1. Features of ext3

The ext3 file system is essentially an enhanced version of the ext2 file system. These improvements provide the following advantages:

Availability

After an unexpected power failure or system crash (also called an *unclean system shutdown*), each mounted ext2 file system on the machine must be checked for consistency by the `e2fsck` program. This is a time-consuming process that can delay system boot time significantly, especially with large volumes containing a large number of files. During this time, any data on the volumes is unreachable.

The journaling provided by the ext3 file system means that this sort of file system check is no longer necessary after an unclean system shutdown. The only time a consistency check occurs using ext3 is in certain rare hardware failure cases, such as hard drive failures. The time to recover an ext3 file system after an unclean system shutdown does not depend on the size of the file system or the number of files; rather, it depends on the size of the *journal* used to maintain consistency. The default journal size takes about a second to recover, depending on the speed of the hardware.

Data Integrity

The ext3 file system provides stronger data integrity in the event that an unclean system shutdown occurs. The ext3 file system allows you to choose the type and level of protection that your data receives. By default, Red Hat Linux 8.0 configures ext3 volumes to keep a high level of data consistency with regard to the state of the file system.

Speed

Despite writing some data more than once, ext3 has a higher throughput in most cases than ext2 because ext3's journaling optimizes hard drive head motion. You can choose from three journaling modes to optimize speed, but doing so means trade offs in regards to data integrity.

Easy Transition

It is easy to change from ext2 to ext3 and gain the benefits of a robust journaling file system without reformatting. See Section 1.3 for more on how to perform this task.

If you perform a fresh installation of Red Hat Linux 8.0, the default file system assigned to the system's Linux partitions is ext3. If you upgrade from a version of Red Hat Linux that uses ext2 partitions, the installation program allows you to convert these partitions to ext3 partitions without losing data. See the appendix titled *Upgrading Your Current System* in the *Official Red Hat Linux Installation Guide* for details.

The following sections will walk you through the steps for creating and tuning ext3 partitions. If you have ext2 partitions and are running Red Hat Linux 8.0, you can skip the partitioning and formatting sections below and go directly to Section 1.3.

1.2. Creating an ext3 File System

After installation, it is sometimes necessary to create a new ext3 file system. For example, if you add a new disk drive to a Red Hat Linux system, you may want to create partition the drive and use the ext3 file system.

The steps for creating an ext3 file system are as follows:

1. Create the partition using `parted` or `fdisk`.
2. Format the partition with the ext3 file system using `mkfs`.
3. Label the partition using `e2label`.
4. Create the mount point.
5. Add the partition to `/etc/fstab`.

Refer to Chapter 5 for information on performing these steps.

1.3. Converting to an ext3 File System

The `tune2fs` program can add a journal to an existing ext2 file system without altering the data already on the partition. If the file system is already mounted while it is being transitioned, the journal will be visible as the file `.journal` in the root directory of the file system. If the file system is not mounted, the journal will be hidden and will not appear in the file system at all.

To convert an ext2 file system to ext3, log in as root and type:

```
/sbin/tune2fs -j /dev/hdbX
```

In the above command, replace `/dev/hdb` with the device name and `X` with the partition number.

After doing this, be certain to change the partition type from ext2 to ext3 in `/etc/fstab`.

If you are transitioning your root file system, you will have to use an `initrd` image (or RAM disk) to boot. To create this, run the `mkinitrd` program. For information on using the `mkinitrd` command, type `man mkinitrd`. Also make sure your GRUB or LILO configuration loads the `initrd`.

If you fail to make this change, the system will still boot, but the file system will be mounted as ext2 instead of ext3.

1.4. Reverting to an ext2 File System

Because ext3 is relatively new, some disk utilities do not yet support it. For example, you may need to shrink a partition with `resize2fs`, which does not yet support ext3. In this situation, it may be necessary to temporarily revert a file system to ext2.

To revert a partition, you must first unmount the partition by logging in as root and typing:

```
umount /dev/hdbX
```

In the above command, replace `/dev/hdb` with the device name and `X` with the partition number. For the remainder of this section, the sample commands will use `hdb1` for these values.

Next, change the file system type to ext2 by typing the following command as root:

```
/sbin/tune2fs -O ^has_journal /dev/hdb1
```

Check the partition for errors by typing the following command as root:


```
/sbin/e2fsck -y /dev/hdb1
```

Then mount the partition again as ext2 file system by typing:

```
mount -t ext2 /dev/hdb1 /mount/point
```

In the above command, replace */mount/point* with the mount point of the partition.

Next, remove the `.journal` file at the root level of the partition by changing to the directory where it is mounted and typing:

```
rm -f .journal
```

You now have an ext2 partition.

If you permanently change the partition to ext2, remember to update the `/etc/fstab` file.

2.1. What is Swap Space?

Swap space in Linux is used when the amount of physical memory (RAM) is full. If the system needs more memory resources and the physical memory is full, inactive pages in memory are moved to the swap space. While swap space can help machines with a small amount of RAM, it should not be considered a replacement for more RAM. Swap space is located on hard drives, which have a slower access time than physical memory.

Swap space can be a dedicated swap partition (recommended), a swap file, or a combination of swap partitions and swap files.

The size of your swap space should be equal to twice your computer's RAM, or 32 MB, whichever amount is larger, but no more than 2048 MB (or 2 GB).

2.2. Adding Swap Space

Sometimes it is necessary to add more swap space after installation. For example, you may upgrade the amount of RAM in your system from 64 MB to 128 MB, but there is only 128 MB of swap space. It might be advantageous to increase the amount of swap space to 256 MB if you perform memory-intensive operations or run applications that require a large amount of memory.

You have two options: add a swap partition or add a swap file. It is recommended that you add a swap partition, but sometimes that is not easy if you do not have any free space available.

To add a swap partition (assuming `/dev/hdb2` is the swap partition you want to add):

1. The hard drive can not be in use (partitions can not be mounted, and swap space can not be enabled). The easiest way to achieve this is to boot your system in rescue mode. Refer to Chapter 8 for instructions on booting into rescue mode. When prompted to mount the filesystem, select **Skip**.

Alternately, if the drive does not contain any partitions in use, you can unmount them and turn off all the swap space on the hard drive with the `swapoff` command.

2. Create the swap partition using `parted` or `fdisk`. Using `parted` is easier than `fdisk`; thus, only `parted` will be explained. To create a swap partition with `parted`:
 - At a shell prompt as root, type the command `parted /dev/hdb`, where `/dev/hdb` is the device name for the hard drive with free space.
 - At the `(parted)` prompt, type **print** to view the existing partitions and the amount of free space. The start and end values are in megabytes. Determine how much free space is on the hard drive and how much you want to allocate for a new swap partition.
 - At the `(parted)` prompt, type **mkpartfs part-type linux-swaps start end**, where `part-type` is one of primary, extended, or logical, `start` is the starting point of the partition, and `end` is the end point of the partition.



Changes take place immediately; be careful when you type.

- Exit `parted` by typing **quit**.

- Now that you have the swap partition, use the command `mkswap` to setup the swap partition. At a shell prompt as root, type the following:

```
mkswap /dev/hdb2
```

- To enable the swap partition immediately, type the following command:

```
swapon /dev/hdb2
```

- To enable it at boot time, edit `/etc/fstab` to include:

```
/dev/hdb2 swap swap defaults 0 0
```

The next time the system boots, it will enable the new swap partition.

- After adding the new swap partition and enabling it, make sure it is enabled by viewing the output of the command `cat /proc/swaps` or `free`.

To add a swap file:

- Determine the size of the new swap file and multiple by 1024 to determine the block size. For example, the block size of a 64 MB swap file is 65536.

- At a shell prompt as root, type the following command with `count` being equal to the desired block size:

```
dd if=/dev/zero of=/swapfile bs=1024 count=65536
```

- Setup the swap file with the command:

```
mkswap /swapfile
```

- To enable the swap file immediately but not automatically at boot time:

```
swapon /swapfile
```

- To enable it at boot time, edit `/etc/fstab` to include:

```
/swapfile swap swap defaults 0 0
```

The next time the system boots, it will enable the new swap file.

- After adding the new swap file and enabling it, make sure it is enabled by viewing the output of the command `cat /proc/swaps` or `free`.

2.3. Removing Swap Space

To remove a swap partition:

- The hard drive can not be in use (partitions can not be mounted, and swap space can not be enabled). The easiest way to achieve this is to boot your system in rescue mode. Refer to Chapter 8 for instructions on booting into rescue mode. When prompted to mount the filesystem, select **Skip**.

Alternately, if the drive does not contain any partitions in use, you can unmount them and turn off all the swap space on the hard drive with the `swapoff` command.

- At a shell prompt as root, execute the following command to make sure the swap partition is disabled (where `/dev/hdb2` is the swap partition):

```
swapoff /dev/hdb2
```

- Remove its entry from `/etc/fstab`.

- Remove the partition using `parted` or `fdisk`. Only `parted` will be discussed. To remove the partition with `parted`:

- At a shell prompt as root, type the command `parted /dev/hdb`, where `/dev/hdb` is the device name for the hard drive with free space.

- At the (parted) prompt, type **print** to view the existing partitions and determine the minor number of the swap partition you wish to delete.
- At the (parted) prompt, type **rm MINOR**, where *MINOR* is the minor number of the partition you want to remove.



Changes take effect immediately; you must type the correct minor number.

- Type **quit** to exit `parted`.

To remove a swap file:

1. At a shell prompt as root, execute the following command to disable the swap file (where `/swapfile` is the swap file):

```
swapoff /swapfile
```
2. Remove its entry from `/etc/fstab`.
3. Remove the actual file:

```
rm /swapfile
```

2.4. Moving Swap Space

To move swap space from one location to another, follow the steps for removing swap space, and then follow the steps for adding swap space.

Redundant Array of Independent Disks (RAID)

3.1. What is RAID?

The basic idea behind RAID is to combine multiple small, inexpensive disk drives into an array to accomplish performance or redundancy goals not attainable with one large and expensive drive. This array of drives will appear to the computer as a single logical storage unit or drive.

RAID is a method in which information is spread across several disks, using techniques such as *disk striping* (RAID Level 0), *disk mirroring* (RAID level 1), and *disk striping with parity* (RAID Level 5) to achieve redundancy, lower latency and/or increase bandwidth for reading or writing to disks, and maximize the ability to recover from hard disk crashes.

The underlying concept of RAID is that data may be distributed across each drive in the array in a consistent manner. To do this, the data must first be broken into consistently-sized *chunks* (often 32K or 64K in size, although different sizes can be used). Each chunk is then written to a hard drive in RAID according to the RAID level used. When the data is to be read, the process is reversed, giving the illusion that multiple drives are actually one large drive.

3.2. Who Should Use RAID?

Anyone who needs to keep large quantities of data on hand (such as a system administrator) would benefit by using RAID technology. Primary reasons to use RAID include:

- Enhanced speed
- Increased storage capacity using a single virtual disk
- Lessened impact of a disk failure

3.3. Hardware RAID versus Software RAID

There are two possible RAID approaches: Hardware RAID and Software RAID.

3.3.1. Hardware RAID

The hardware-based system manages the RAID subsystem independently from the host and presents to the host only a single disk per RAID array.

An example of a Hardware RAID device would be one that connects to a SCSI controller and presents the RAID arrays as a single SCSI drive. An external RAID system moves all RAID handling "intelligence" into a controller located in the external disk subsystem. The whole subsystem is connected to the host via a normal SCSI controller and appears to the host as a single disk.

RAID controllers also come in the form of cards that *act* like a SCSI controller to the operating system but handle all of the actual drive communications themselves. In these cases, you plug the drives into the RAID controller just like you would a SCSI controller, but then you add them to the RAID controller's configuration, and the operating system never knows the difference.

3.3.2. Software RAID

Software RAID implements the various RAID levels in the kernel disk (block device) code. It offers the cheapest possible solution, as expensive disk controller cards or hot-swap chassis ¹ are not required. Software RAID also works with cheaper IDE disks as well as SCSI disks. With today's fast CPUs, Software RAID performance can excel against Hardware RAID.

The MD driver in the Linux kernel is an example of a RAID solution that is completely hardware independent. The performance of a software-based array is dependent on the server CPU performance and load.

For information on configuring Software RAID in the Red Hat Linux installation program, refer to the Chapter 9.

For those interested in learning more about what Software RAID has to offer, here is a brief list of the most important features:

- Threaded rebuild process
- Kernel-based configuration
- Portability of arrays between Linux machines without reconstruction
- Backgrounded array reconstruction using idle system resources
- Hot-swappable drive support
- Automatic CPU detection to take advantage of certain CPU optimizations

3.4. RAID Levels and Linear Support

RAID supports various configurations, including levels 0, 1, 4, 5, and linear. These RAID types are defined as follows:

- *Level 0* — RAID level 0, often called "striping," is a performance-oriented striped data mapping technique. This means the data being written to the array is broken down into strips and written across the member disks of the array, allowing high I/O performance at low inherent cost but provides no redundancy. The storage capacity of a level 0 array is equal to the total capacity of the member disks in a Hardware RAID or the total capacity of member partitions in a Software RAID.
- *Level 1* — RAID level 1, or "mirroring," has been used longer than any other form of RAID. Level 1 provides redundancy by writing identical data to each member disk of the array, leaving a "mirrored" copy on each disk. Mirroring remains popular due to its simplicity and high level of data availability. Level 1 operates with two or more disks that may use parallel access for high data-transfer rates when reading but more commonly operate independently to provide high I/O transaction rates. Level 1 provides very good data reliability and improves performance for read-intensive applications but at a relatively high cost. ² The storage capacity of the level 1 array is equal to the capacity of one of the mirrored hard disks in a Hardware RAID or one of the mirrored partitions in a Software RAID.

1. A hot-swap chassis allows you to remove a hard drive without having to power-down your system.

2. RAID level 1 comes at a high cost because you write the same information to all of the disks in the array, which wastes drive space. For example, if you have RAID level 1 set up so that your root (/) partition exists on two 40G drives, you have 80G total but are only able to access 40G of that 80G. The other 40G acts like a mirror of the first 40G.

- *Level 4* — Level 4 uses parity³ concentrated on a single disk drive to protect data. It is better suited to transaction I/O rather than large file transfers. Because the dedicated parity disk represents an inherent bottleneck, level 4 is seldom used without accompanying technologies such as write-back caching. Although RAID level 4 is an option in some RAID partitioning schemes, it is not an option allowed in Red Hat Linux RAID installations.⁴ The storage capacity of Hardware RAID level 4 is equal to the capacity of member disks, minus the capacity of one member disk. The storage capacity of Software RAID level 4 is equal to the capacity of the member partitions, minus the size of one of the partitions if they are of equal size.
- *Level 5* — This is the most common type of RAID. By distributing parity across some or all of an array's member disk drives, RAID level 5 eliminates the write bottleneck inherent in level 4. The only performance bottleneck is the parity calculation process. With modern CPUs and Software RAID, that usually is not a very big problem. As with level 4, the result is asymmetrical performance, with reads substantially outperforming writes. Level 5 is often used with write-back caching to reduce the asymmetry. The storage capacity of Hardware RAID level 5 is equal to the capacity of member disks, minus the capacity of one member disk. The storage capacity of Software RAID level 5 is equal to the capacity of the member partitions, minus the size of one of the partitions if they are of equal size.
- *Linear RAID* — Linear RAID is a simple grouping of drives to create a larger virtual drive. In linear RAID, the chunks are allocated sequentially from one member drive, going to the next drive only when the first is completely filled. This grouping provides no performance benefit, as it is unlikely that any I/O operations will be split between member drives. Linear RAID also offers no redundancy and, in fact, decreases reliability — if any one member drive fails, the entire array cannot be used. The capacity is the total of all member disks.

3. Parity information is calculated based on the contents of the rest of the member disks in the array. This information can then be used to reconstruct data when one disk in the array fails. The reconstructed data can then be used to satisfy I/O requests to the failed disk before it is replaced and to repopulate the failed disk after it has been replaced.

4. RAID level 4 takes up the same amount of space as RAID level 5, but level 5 has more advantages. For this reason, level 4 is not supported.

Logical Volume Manager (LVM)

Beginning with Red Hat Linux 8.0, Logical Volume Manager (LVM) is available for hard drive allocation.

LVM is a method of allocating hard drive space into logical volumes that can be easily resized instead of partitions.

With LVM, the hard drive or set of hard drives is allocated to one or more *physical volumes*. A physical volume can not span over more than one drive.

The physical volumes are combined into *logical volume groups*, with the exception of the `/boot` partition. The `/boot` partition can not be on a logical volume group because the boot loader can not read it. If you want to have the root `/` partition on a logical volume, you will need to create a separate `/boot` partition which is not a part of a volume group.

Since a physical volume can not span over more than one drive, if you want the logical volume group to span over more than one drive, you must create one or more physical volumes per drive.

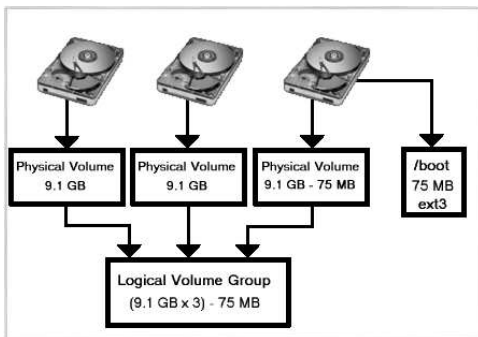


Figure 4-1. Logical Volume Group

The logical volume group is divided into *logical volumes*, which are assigned mount points such as `/home` and `/` and file system types such as `ext3`. When "partitions" reach their full capacity, free space from the logical volume group can be added to the logical volume to increase the size of the partition. When a new hard drive is added to the system, it can be added to the logical volume group, and the logical volumes that are the partitions can be expanded.




Figure 4-2. Logical Volumes

On the other hand, if a system is partitioned with the ext3 file system, the hard drive is divided into partitions of defined sizes. If a partition becomes full, it is not easy to expand the size of the partition. Even if the partition is moved to another hard drive, the original hard drive space has to be reallocated as a different partition or not used.

LVM support must be compiled into the kernel. The default kernel for Red Hat Linux 8.0 is compiled with LVM support.

To learn how to configure LVM during the Red Hat Linux installation process, refer to Chapter 10.

Managing Disk Storage

After you have installed your Red Hat Linux system, you may want to view the existing partition table, change the size of the partitions, remove partitions, or add partitions from free space or additional hard drives. The utility `parted` allows you to perform these tasks. This chapter discusses how to use `parted` to perform command file system tasks. Alternatively, you can use `fdisk` to perform most of these tasks, excluding resizing partitions. For more information on `fdisk`, refer to `man` or `info` page for `fdisk`.

If you want to view the system's disk space usage or monitor the disk space usage, refer to Section 25.3.

You must have the `parted` package installed to use the `parted` utility. To start `parted`, at a shell prompt as root, type the command `parted /dev/hdb`, where `/dev/hdb` is the device name for the drive you want to configure. You will see a `(parted)` prompt. Type `help` to view a list of available commands.

If you want to create, remove, or resize a partition, the device can not be in use (partitions can not be mounted, and swap space can not be enabled). The easiest way to achieve this is to boot your system in rescue mode. Refer to Chapter 8 for instructions on booting into rescue mode. When prompted to mount the file system, select **Skip**.

Alternately, if the drive does not contain any partitions in use, you can unmount them with the `umount` command and turn off all the swap space on the hard drive with the `swapoff` command.

Table 5-1 contains a list of commonly used `parted` commands. The sections that follow explain some of them in more detail.

Command	Description
<code>check minor-num</code>	Perform a simple check of the file system
<code>cp from to</code>	Copy file system from one partition to another; <i>from</i> and <i>to</i> are the minor numbers of the partitions
<code>help</code>	Display list of available commands
<code>mklabel label</code>	Create a disk label for the partition table
<code>mkfs minor-num file-system-type</code>	Create a file system of type <i>file-system-type</i>
<code>mkpart part-type fs-type start-mb end-mb</code>	Make a partition without creating a new file system
<code>mkpartfs part-type fs-type start-mb end-mb</code>	Make a partition and create the specified file system
<code>move minor-num start-mb end-mb</code>	Move the partition
<code>print</code>	Display the partition table
<code>quit</code>	Quit <code>parted</code>
<code>resize minor-num start-mb end-mb</code>	Resize the partition from <i>start-mb</i> to <i>end-mb</i>
<code>rm minor-num</code>	Remove the partition

Command	Description
<code>select device</code>	Select a different device to configure
<code>set minor-num flag state</code>	Set the flag on a partition; <i>state</i> is either on or off

Table 5-1. parted commands

5.1. Viewing the Partition Table

After starting `parted`, type the following command to view the partition table:

```
print
```

A table similar to the following will appear:

```
Disk geometry for /dev/hda: 0.000-9765.492 megabytes
Disk label type: msdos
Minor   Start      End        Type       Filesystem  Flags
 1       0.031      101.975    primary    ext3        boot
 2       101.975    611.850    primary    linux-swap
 3       611.851    760.891    primary    ext3
 4       760.891    9758.232   extended
 5       760.922    9758.232   logical    ext3        lba
```

The first line displays the size of the disk, the second line displays the disk label type, and the remaining output shows the partition table. In the partition table, the **Minor** number is the partition number. For example, the partition with minor number 1 corresponds to `/dev/hda1`. The **Start** and **End** values are in megabytes. The **Type** is one of primary, extended, or logical. The **Filesystem** is the file system type, which can be one of ext2, ext3, FAT, hfs, jfs, linux-swap, ntfs, reiserfs, hp-ufs, sun-ufs, or xfs. The **Flags** column lists the flags set for the partition. Available flags are boot, root, swap, hidden, raid, lvm, or lba.



Tip

To select a different device without having to restart `parted`, use the `select` command followed by the device name such as `/dev/hdb`. Then, you can view its partition table or configure it.

5.2. Creating a Partition



Warning

Do not attempt to create a partition on a device that is in use.

Before creating a partition, boot into rescue mode (or unmount any partitions on the device and turn off any swap space on the device).

Start `parted`, where `/dev/hda` is the device on which to create the partition:

```
parted /dev/hda
```

View the current partition table to determine if there is enough free space:

```
print
```

If there is not enough free space, you can resize an existing partition. Refer to Section 5.4 for details.

5.2.1. Making the Partition

From the partition table, determine the start and end points of the new partition and what partition type it should be. You can only have four primary partitions (with no extended partition) on a device. If you need more than four partitions, you can have three primary partitions, one extended partition, and multiple logical partitions within the extended. For an overview of disk partitions, refer to the appendix *An Introduction to Disk Partitions* in the *Official Red Hat Linux Installation Guide*.

For example, to create a primary partition with an ext3 file system from 1024 megabytes until 2048 megabytes on a hard drive type the following command:

```
mkpart primary ext3 1024 2048
```



Tip

If you use the `mkpartfs` command instead, the file system will be created after the partition is created. However, `parted` does not support creating an ext3 file system. Thus, if you wish to create an ext3 file system, use `mkpart` and create the file system with the `mkfs` command as described later. `mkpartfs` works for file system type `linux-swap`.

The changes start taking place as soon as you type [Enter], so review the command before executing to it.

After creating the partition, use the `print` command to confirm that it is in the partition table with the correct partition type, file system type, and size. Also remember the minor number of the new partition so that you can label it. You should also view the output of

```
cat /proc/partitions
```

to make sure the kernel recognizes the new partition.

5.2.2. Formatting the Partition

The partition still does not have a file system. Create the file system:

```
/sbin/mkfs -t ext3 /dev/hdb3
```



Warning

Formatting the partition will permanently destroy any data that currently exists on the partition.

5.2.3. Labeling the Partition

Next, give the partition a label. For example, if the new partition is `/dev/hda3` and you want to label it `/work`:

```
e2label /dev/hda3 /work
```

By default, the Red Hat Linux installation program uses the mount point of the partition as the label to make sure the label is unique. You can use any label you want.

5.2.4. Creating the Mount Point

As root, create the mount point:

```
mkdir /work
```

5.2.5. Add to `/etc/fstab`

As root, edit the `/etc/fstab` file to include the new partition. The new line should look similar to the following:

```
LABEL=/work          /work          ext3      defaults    1 2
```

The first column should contain `LABEL=` followed by the label you gave the partition. The second column should contain the mount point for the new partition, and the next column should be the file type (for example, `ext3` or `swap`). If you need more information about the format, read the man page with the command `man fstab`.

If the fourth column is the word `defaults`, the partition will be mounted at boot time. To mount the partition without rebooting, as root, type the command:

```
mount /work
```

5.3. Removing a Partition



Warning

Do not attempt to remove a partition on a device that is in use.

Before removing a partition, boot into rescue mode (or unmount any partitions on the device and turn off any swap space on the device).

Start `parted`, where `/dev/hda` is the device on which to create the partition:

```
parted /dev/hda
```

View the current partition table to determine the minor number of the partition to remove:

```
print
```


Remove the partition with the command `rm`. For example, to remove the partition with minor number 3:

```
rm 3
```

The changes start taking place as soon as you type [Enter], so review the command before committing to it.

After removing the partition, use the `print` command to confirm that it is removed from the partition table. You should also view the output of

```
cat /proc/partitions
```

to make sure the kernel knows the partition is removed.

The last step is to remove it from the `/etc/fstab` file. Find the line that declares the removed partition, and remove it from the file.

5.4. Resizing a Partition



Warning

Do not attempt to resize a partition on a device that is in use.

Before resizing a partition, boot into rescue mode (or unmount any partitions on the device and turn off any swap space on the device).

Start `parted`, where `/dev/hda` is the device on which to create the partition:

```
parted /dev/hda
```

View the current partition table to determine the minor number of the partition to remove as well as the start and end points for the partition:

```
print
```



Warning

The used space of the partition to resize must not be larger than the new size.

To resize the partition, use the `resize` command followed by the minor number for the partition, the starting place in megabytes, and the end place in megabytes. For example:

```
resize 3 1024 2048
```

After resizing the partition, use the `print` command to confirm that the partition has been resized correctly, is the correct partition type, and is the correct file system type.

After rebooting the system into normal mode, use the command `df` to make sure the partition was mounted and is recognized with the new size.

Installation-Related Information

Kickstart Installations

6.1. 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 a Red Hat Linux installation.

6.2. 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, you must:

1. Create a kickstart file.
2. Create a boot diskette with the kickstart file or make the kickstart file available on the network.
3. Make the installation tree available.
4. Start the kickstart installation.

This chapter explains these steps in detail.

6.3. Creating the Kickstart File

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, using the **Kickstart Configurator** application, or writing it from scratch. The Red Hat Linux installation program also creates a sample kickstart file based on the options that you selected during installation. It is written to the file `/root/anaconda-ks.cfg`. 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:

- Sections must be specified *in order*. Items within the sections do not have to be in a specific order unless otherwise specified. The section order is:
 - Command section — Refer to Section 6.4 for a list of kickstart options. You must include the required options.
 - The `%packages` section — Refer to Section 6.5 for details.
 - The `%pre` and `%post` sections — These two sections can be in any order and are not required. Refer to Section 6.6 and Section 6.7 for details.

- Items that are not 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
 - Language support
 - Installation method
 - Device specification (if device is needed to perform installation)
 - Keyboard setup
 - The `upgrade` keyword
 - Boot loader configuration

If any other items are specified for an upgrade, those items will be ignored (note that this includes package selection).

6.4. Kickstart Options

The following options can be placed in a kickstart file. If you prefer to use a graphical interface for creating your kickstart file, you can use the **Kickstart Configurator** application. Refer to Chapter 7 for details.



Note

If the option is followed by an equals mark (=), a value must be specified after it. In the example commands, options in brackets ([]) are optional arguments for the command.

`autostep` (optional)

Similar to `interactive` except it goes to the next screen for you. It is used mostly for debugging.

`auth` or `authconfig` (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 with the `--nisdomain=` option.

`--nisdomain=`
NIS domain name to use for NIS services.

`--nisserver=`
Server to use for NIS services (broadcasts by default).

`--useshadow` or `--enableshadow`
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 install the `nss_ldap` package. You must also specify a server and a base DN with `--ldapserver=` and `--ldappedn=`.

`--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 with `--ldapserver=` and `--ldappedn=`.

`--ldapserver=`
If you specified either `--enableldap` or `--enableldapauth`, the name of the LDAP server to use. This option is set in the `/etc/ldap.conf` file.

`--ldappedn=`
If you specified either `--enableldap` or `--enableldapauth`, 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.

`--enableldaptls`
Use TLS (Transport Layer Security) lookups. This option allows LDAP to send encrypted usernames and passwords to an LDAP server before authentication.

`--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.

--enablesmbauth

Enables authentication of users against an SMB server (typically a Samba or Windows server). SMB authentication support does not know about home directories, UIDs, or shells. So if you enable it you will need to make users' accounts known to the workstation by enabling LDAP, NIS, or Hesiod or by using the `/usr/sbin/useradd` command to make their accounts known to the workstation. To use this option, you must have the `pam_smb` package installed.

--smbserver=

The name of the server(s) to use for SMB authentication. To specify more than one server, separate the names with commas (,).

--smbworkgroup=

The name of the workgroup for the SMB servers.

--enablecache

Enables the `nscd` service. The `nscd` service caches information about users, groups, and various other types of information. Caching is especially helpful if you choose to distribute information about users and groups over your network using NIS, LDAP, or hesiod.

bootloader (required)

Specifies how the boot loader should be installed and whether the boot loader should be LILO or GRUB. This option is required for both installations and upgrades. For upgrades, if `--useLilo`

is not specified and LILO is the current bootloader, the bootloader will be changed to GRUB. To preserve LILO on upgrades, use `bootloader --upgrade`.

`--append=`

Specifies kernel parameters.

`--location=`

Specifies where the boot record is written. Valid values are the following: `mbr` (the default), `partition` (installs the boot loader on the first sector of the partition containing the kernel), or `none` (do not install the boot loader).

`--password=`

If using GRUB, sets the GRUB boot loader password the one specified with this option. This should be used to restrict access to the GRUB shell, where arbitrary kernel options can be passed.

`--md5pass=`

If using GRUB, similar to `--password=` except the password should already be encrypted.

`--useLilo`

Use LILO instead of GRUB as the boot loader.

`--linear`

If using LILO, use the `linear` LILO option; this is only for backwards compatibility (and `linear` is now used by default).

`--nolinear`

If using LILO, use the `nolinear` LILO option; `linear` is the default.

`--lba32`

If using LILO, force use of `lba32` mode instead of autodetecting.

`--upgrade`

Upgrade the existing boot loader configuration, preserving the old entries. This option is only available for upgrades.

`clearpart` (optional)

Removes partitions from the system, prior to creation of new partitions. By default, no partitions are removed.



Note

If the `clearpart` command is used, then the `--onpart` command cannot be used on a logical partition.

`--linux`

Erases all Linux partitions.

`--all`

Erases all partitions from the system.

`--drives`

Specifies which drives to clear partitions from.

`--initlabel`

Initializes the disk label to the default for your architecture (`msdos` for x86 and `gpt` for Itanium). It is useful so that the installation program does not ask if it should initialize the disk label if installing to a brand new hard drive.

`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 the installation program to install extra modules, is in this format:

```
device <type> <moduleName> --opts=<options>
```

`<type>`

Replace with either `scsi` or `eth`

`<moduleName>`

Replace with 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"
```

`deviceprobe` (optional)

Forces a probe of the PCI bus and loads modules for all the devices found if a module is available.

`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>`

Partition containing the driver disk.

`--type=`

Filesystem type (for example, `vfat`, `ext2`, or `ext3`).

`firewall` (optional)

This option corresponds to the **Firewall Configuration** screen in the installation program:

```
firewall <securitylevel> [--trust=] <incoming> [--port=]
```

<securitylevel>

Replace with one of the following levels of security:

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

--trust=

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`.

<incoming>

Replace with none or more of the following to allow the specified services through the firewall.

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

--port=

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.

`install` (optional)

Tells the system to install a fresh system rather than upgrade an existing system. This is the default mode. For installation, you must specify the type of installation from one of `cdrom`, `harddrive`, `nfs`, or `url` (for ftp or http installations).

`cdrom`

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

`harddrive`

Install from a Red Hat installation tree on a local drive, which must be either vfat or ext2.

- `--partition=`
Partition to install from (such as, sdb2).
- `--dir=`
Directory containing the `RedHat` directory of the installation tree.

For example:

```
harddrive --partition=hdb2 --dir=/tmp/install-tree
```

`nfs`

Install from the NFS server specified.

- `--server=`
Server from which to install (hostname or IP).
- `--dir=`
Directory containing the `RedHat` directory of the installation tree.

For example:

```
nfs --server=nfsserver.example.com --dir=/tmp/install-tree
```

`url`

Install from an installation tree on a remote server via FTP or HTTP.

For example:

```
url --url http://<server>/<dir>
```

or:

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

`interactive (optional)`

Uses the information provided in the kickstart file during the installation, but allow for inspection and modification of the values given. You will be presented with each screen of the installation program with the values from the kickstart file. Either accept the values by clicking **Next** or change the values and click **Next** to continue. See also `autostep`.

`keyboard (required)`

Sets system keyboard type. Here is the list of available keyboards on i386, Itanium, and Alpha machines:

```
be-latin1, be-latin2, bg, br-abnt2, cf, cz-lat2, cz-us-qwertz, de,
de-latin1, de-latin1-nodeadkeys, dk, dk-latin1, dvorak, es, et,
fi, i-latin1, fr, fr-latin0, fr-latin1, fr-pc, fr_CH, fr_CH-latin1,
gr, hu, hu101, is-latin1, it, it-ibm, it2, jp106, no, no-latin1,
pl, pt-latin1, ro, ru, ru-cpl251, ru-ms, rul, ru2, ru_win,
se-latin1, sg, sg-latin1, sk-qwerty, slovene, speakup, speakup-lt,
trq, ua, uk, us
```

`lang (required)`

Sets the language to use during installation. 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, is_IS, it_IT, ja_JP.eucJP,
ko_KR.eucKR, no_NO, pt_PT, ru_RU.koi8r, sl_SI, es_ES, sv_SE, uk_UA,
zh_CN.GB18030, zh_TW.Big5
```

langsupport (required)

Sets the language(s) to install on the system. The same language codes used with `lang` can be used with `langsupport`.

If you just want to install one language, specify it. For example, to install and use the French language `fr_FR`:

```
langsupport fr_FR
```

```
--default=
```

If you want to install language support for more than one language, you must specify a default.

For example, to install English and French and use English as the default language:

```
langsupport --default=en_US fr_FR
```

If you use `--default` with only one language, all languages will be installed with the specified language set to the default.

lilo (replaced by bootloader)



Warning

This option has been replaced by `bootloader` and is only available for backwards compatibility. Refer to `bootloader`.

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 `nonlinear` 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.

```
--lba32
```

Forces the use of `lba32` mode instead of autodetecting.

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.

logvol (optional)¹

Create a logical volume for Logical Volume Management (LVM) with the syntax:

```
logvol mountpoint --vgname=name --size=size --name=name
```

Create the partition first, create the logical volume group, and then create the logical volume. For example:

```
part pv.01 --size 3000
volgroup myvg pv.01
logvol / --vgname=myvg --size=2000 --name=rootvol
```

mouse (required)

Configures the mouse for the system, both in GUI and text modes. Options are:

`--device=`

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 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, genericusb, generic3usb,
geniusnm, geniusmps/2, geniusprops/2, geniusscrollps/2,
geniusscrollps/2+, thinking, thinkings/2, logitech,
logitechcc, logibm, logimman, logimmanps/2, logimman+,
logimman+ps/2, logimusb, microsoft, msnew, msintelli,
msintellips/2, msintelliusb, 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.

network (optional)

Configures network information for the system. If the kickstart installation does not require networking (in other words, it is not installed over NFS, HTTP, or FTP), networking is not configured for the system. If the installation does require networking and network information is not provided in the kickstart file, 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 determine its IP address dynamically. The `network` option configures networking information for kickstart installations via a network as well as for the installed system.

1. This option is new to Red Hat Linux 8.0

`--bootproto=`

One of `dhcp`, `bootp`, or `static`.

Its default to `dhcp`. `bootp` and `dhcp` are treated the same.

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. To direct a system to use DHCP:

```
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 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 and after the installation. The line for static networking is more complex, as you must include all network configuration information on one line. You must specify the IP address, netmask, gateway, and nameserver. For example: (the `\` indicates that it is all one 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 6.7) to add more name servers, if needed.

`--device=`

Used to select a specific Ethernet device for installation. Note that using `--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. For example:

```
network --bootproto=dhcp --device=eth0
```

`--ip=`

IP address for the machine to be installed.

`--gateway=`

Default gateway as an IP address.

`--nameserver=`

Primary nameserver, as an IP address.

`--nodns`

Do not configure any DNS server.

`--netmask=`

Netmask for the installed system.

`--hostname=`

Hostname for the installed system.

`part` or `partition` (required for installs, ignored for upgrades)

Creates a partition on the system.

If more than one Red Hat Linux installation exists on the system on different partitions, the installation program prompts the user and asks which installation to upgrade.



Warning

All partitions created will be formatted as part of the installation process unless `--noformat` and `--onpart` are used.

`<mntpoint>`

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

- `/<path>`

For example, `/`, `/usr`, `/home`

- `swap`

The partition will be used as swap space.

To determine the size of the swap partition automatically, use the `--recommended` option:

```
swap --recommended
```

The minimum size of the automatically-generated swap partition will be no smaller than the amount of RAM in the system and no bigger than twice the amount of RAM in the system.

- `raid.<id>`

The partition will be used for software RAID (refer to `raid`).

- `pv.<id>`

The partition will be used for LVM (refer to `logvol`).

`--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=`

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=` or `--usepart=`

Put the partition on the *already existing* device. For example:

```
partition /home --onpart hda1
```

will put `/home` on `/dev/hda1`, which must already exist.

`--ondisk=` or `--ondrive=`

Forces the partition to be created on a particular disk. For example, `--ondisk=sdb` will put the partition on the second SCSI disk on the system.

`--asprimary`

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

`--bytes-per-inode=`

Number specified 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=` (replaced by `fstype`)

This option is no longer available. Use `fstype`.

`--fstype=`

Sets the filesystem type for the partition. Valid values are `ext2`, `ext3`, `swap`, and `vfat`.

`--start=`

Specifies the starting cylinder for the partition. It requires that a drive be specified with `--ondisk=` or `ondrive=`. It also requires that the ending cylinder be specified with `--end=` or the partition size be specified with `--size=`.

`--end=`

Specifies the ending cylinder for the partition. It requires that the starting cylinder be specified with `--start=`.

`--badblocks`

Specifies that the partition should be checked for bad sectors.



Note

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

`raid` (optional)

Assembles a software RAID device. This command is of the form:

```
raid <mntpoint> --level=<level> --device=<mddevice> <partitions*>
```

`<mntpoint>`

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=`

RAID level to use (0, 1, or 5).

`--device=`

Name of the RAID device to use (such as md0 or md1). RAID devices range from md0 to md7, and each may only be used once.

`--spares=`

Specifies the number of spare drives allocated for the RAID array. Spare drives are used to rebuild the array in case of drive failure.

`--fstype=`

Sets the filesystem type for the RAID array. Valid values are ext2, ext3, swap, and vfat.

`--noformat`

Do not format the RAID array.

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
```

`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.

`rootpw` (required)

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

```
rootpw [--iscrypted] <password>
```

`--iscrypted`

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

`skipx` (optional)

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

`text` (optional)

Perform the kickstart installation in text mode. Kickstart installations are performed in graphical mode by default.

timezone (required)

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

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

```
--utc
```

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

upgrade (optional)

Tells the system to upgrade an existing system rather than install a fresh system. You must specify one of `cdrom`, `harddrive`, `nfs`, or `url` (for `ftp` and `http`) as the location of the installation tree. Refer to `install` for details.

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
```

Do not probe the monitor.

```
--card=
```

Use specified card; this card name should be from the list of cards in `/usr/share/hwdata/Cards` from the `hwdata` package. If this argument is not provided, the installation program will probe the PCI bus for the card. Since AGP is part of the PCI bus, AGP cards will be detected if supported. The probe order is determined by the PCI scan order of the motherboard.

```
--videoram=
```

Specify the amount of video RAM the video card has.

```
--monitor=
```

Use specified monitor; monitor name should be from the list of monitors in `/usr/share/hwdata/MonitorsDB` from the `hwdata` package. 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=
```

Specifies the horizontal sync frequency of the monitor.

```
--vsync=
```

Specifies the vertical sync frequency of the monitor.

```
--defaultdesktop=
```

Specify either GNOME or KDE to set the default desktop (assumes that GNOME Desktop Environment and/or KDE Desktop Environment has been installed through `%packages`).

```
--startxonboot
```

Use a graphical login on the installed system.

`--resolution=`

Specify the default resolution for the X Window System on the installed system. Valid values are 640x480, 800x600, 1024x768, 1152x864, 1280x1024, 1400x1050, 1600x1200. Be sure to specify a resolution that is compatible with the video card and monitor.

`--depth=`

Specify the default color depth for the X Window System on the installed system. Valid values are 8, 16, 24, and 32. Be sure to specify a color depth that is compatible with the video card and monitor.

`volgroup (optional)`¹

Use to create a Logical Volume Management (LVM) group with the syntax:

```
volgroup name partition
```

Create the partition first, create the logical volume group, and then create the logical volume. For example:

```
part pv.01 --size 3000
volgroup myvg pv.01
logvol / --vgname=myvg --size=2000 --name=rootvol
```

`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.

`%include`

Use the `%include /path/to/file` command to include the contents of another file in the kickstart file as though the contents were at the location of the `%include` command in the kickstart file.

6.5. Package Selection

Use the `%packages` command to begin a kickstart file section that lists the packages you would like to install (this is for installations only, as package selection during upgrades is not supported).

`--resolvedeps`

Install the listed packages and automatically resolve package dependencies.

`--ignoredeps`

Ignore the unresolved dependencies and install the listed packages without the dependencies.

Packages can be specified by group or by individual package name. The installation program defines several groups that contain related packages. See the `RedHat/base/comps.xml` file on any Red Hat Linux CD-ROM for a list of groups. Each group has an id, uservisibility value, name, description, and package list. In the package list, the packages marked as mandatory are always installed if the group is selected, the packages marked default are selected by default if the group is selected, and the packages marked optional must be specifically selected even if the group is selected to be installed.

1. This option is new to Red Hat Linux 8.0

In most cases, it is only necessary to list the desired groups and not individual packages. Note that the `Core` and `Base` groups are always selected by default, so it is not necessary to specify it in the `%packages` section.

Here is an example `%packages` selection:

```
%packages
@ X Window System
@ GNOME Desktop Environment
@ Graphical Internet
@ Sound and Video
galeon
```

As you can see, groups are specified, one to a line, starting with an `@` symbol, a space, and then the full group name as given in the `comps` file. Specify individual packages with no additional characters (the `galeon` line in the example above is an individual package).

To specify an everything installation to install all packages), add the `one` to the `%packages` section:

```
@ Everything
```

You can also specify which packages not to install from the default package list:

```
@ Games and Entertainment
-kdegames
```

6.6. Pre-installation Script

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. 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.



Note

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

```
--interpreter /usr/bin/python
```

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

6.6.1. Example

Here is an example `%pre` section:

```
;


```
%pre

#!/bin/sh

hds=""
mymedia=""

for file in /proc/ide/h*
```


```

```

do
  mymedia='cat $file/media`
  if [ $mymedia == "disk" ] ; then
    hds="$hds `basename $file`"
  fi
done

set $hds
numhd=`echo $#`

drive1=`echo $hds | cut -d' ' -f1`
drive2=`echo $hds | cut -d' ' -f2`

#Write out partition scheme based on whether there are 1 or 2 hard drives

if [ $numhd == "2" ] ; then
  #2 drives
  echo "#partitioning scheme generated in %pre for 2 drives" > /tmp/part-include
  echo "clearpart --all" >> /tmp/part-include
  echo "part /boot --fstype ext3 --size 75 --ondisk hda" >> /tmp/part-include
  echo "part / --fstype ext3 --size 1 --grow --ondisk hda" >> /tmp/part-include
  echo "part swap --recommended --ondisk $drive1" >> /tmp/part-include
  echo "part /home --fstype ext3 --size 1 --grow --ondisk hdb" >> /tmp/part-include
else
  #1 drive
  echo "#partitioning scheme generated in %pre for 1 drive" > /tmp/part-include
  echo "clearpart --all" >> /tmp/part-include
  echo "part /boot --fstype ext3 --size 75" >> /tmp/part-include
  echo "part swap --recommended" >> /tmp/part-include
  echo "part / --fstype ext3 --size 2048" >> /tmp/part-include
  echo "part /home --fstype ext3 --size 2048 --grow" >> /tmp/part-include
fi

```

This script determines the number of hard drives in the system and writes a text file with a different partitioning scheme depending on whether it has one or two drives. Instead of having a set of partitioning commands in the kickstart file, include the line:

```
%include /tmp/part-include
```

The partitioning commands selected in the script will be used.

6.7. Post-installation Script

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. This section is useful for functions such as installing additional software and configuring an additional nameserver.



Note

If you configured the network with static IP information, including a nameserver, you can access the network and resolve IP addresses in the `%post` section. If you configured the network for DHCP, the `/etc/resolv.conf` file has not been completed when the installation executes the `%post` section. You can access the network, but you can not resolve IP addresses. Thus, if you are using DHCP, you must specify IP addresses in the `%post` section.

**Note**

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/python
```

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

6.7.1. Examples

Turn services on and off:

```
/sbin/chkconfig --level 345 telnet off
/sbin/chkconfig --level 345 finger off
/sbin/chkconfig --level 345 lpd off
/sbin/chkconfig --level 345 httpd on
```

Run a script named `runme` from an NFS share:

```
mkdir /mnt/temp
mount 10.10.0.2:/usr/new-machines /mnt/temp
open -s -w -- /mnt/temp/runme
umount /mnt/temp
```

Add a user to the system:

```
/usr/sbin/useradd bob
/usr/bin/chfn -f "Bob Smith" bob
/usr/sbin/usermod -p 'kjdf$04930FTH/ ' bob
```

6.8. Making the Kickstart File Available

A kickstart file must be placed in one of two locations:

- On a boot diskette
- On a network

Normally a kickstart file is copied to the boot diskette, 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.

6.8.1. Creating a Kickstart Boot Diskette

To perform a diskette-based kickstart installation, the kickstart file must be named `ks.cfg` and must be located in the boot diskette's top-level directory. Note that the Red Hat Linux boot diskettes 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 diskette and `cp` the file over.

6.8.2. Making the Kickstart File Available on the Network

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 is an example of a line from the `dhcpcd.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 do not 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 do not 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.

6.9. Making the Installation Tree Available

The kickstart installation needs to access an *installation tree*. An installation tree is a copy of the binary Red Hat Linux CD-ROMs with the same directory structure.

If you are performing a CD-based installation, insert the Red Hat Linux CD-ROM #1 into the computer before starting the kickstart installation.

If you are performing a hard-drive installation, make sure the ISO images of the binary Red Hat Linux CD-ROMs are on a hard drive in the computer.

If you are performing a network-based (NFS, FTP, or HTTP) installation, you must make the installation tree available over the network. Refer to the *Preparing for a Network Installation* section of the *Official Red Hat Linux Installation Guide* for details.

6.10. 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 or ext2 filesystem on a floppy diskette and you boot from the Red Hat Linux CD-ROM.

An alternate boot command for booting off the Red Hat Linux CD-ROM and having the kickstart file on a vfat or ext2 filesystem on a floppy diskette is:

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

If you need to use a driver disk with kickstart, you can still have the kickstart file on a floppy disk:

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

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 in the NFS share `/mydir/ks.cfg`, the correct boot command would be `ks=nfs:server.example.com:/mydir/ks.cfg`.

```
ks=http://<server>/<path>
```

The installation program will look for the kickstart file on the HTTP server `<server>`, as file `<path>`. The installation program will use DHCP to configure the Ethernet card. For example, if your HTTP server is `server.example.com` and the kickstart file is in the HTTP directory `/mydir/ks.cfg`, the correct boot command would be `ks=http://server.example.com/mydir/ks.cfg`.

```
ks=floppy
```

The installation program looks for the file `ks.cfg` on a vfat or ext2 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.

```
ksdevice=<device>
```

The installation program will use this network device to connect to the network. For example, to start a kickstart installation with the kickstart file on an NFS server that is connected to the system through the `eth1` device, use the command `ks=nfs:<server>:/<path> ksdevice=eth1` at the `boot :` prompt.

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, verify the options you have chosen, and save the kickstart file to a desired location.

To use **Kickstart Configurator**, you must be running the X Window System. To start **Kickstart Configurator**, select the **Main Menu Button** (on the Panel) => **System Tools** => **Kickstart**, or type the command `/usr/sbin/redhat-config-kickstart`.

As you are creating a kickstart file, you can select **File** => **Preview** at any time to preview your current selections.

7.1. Basic Configuration

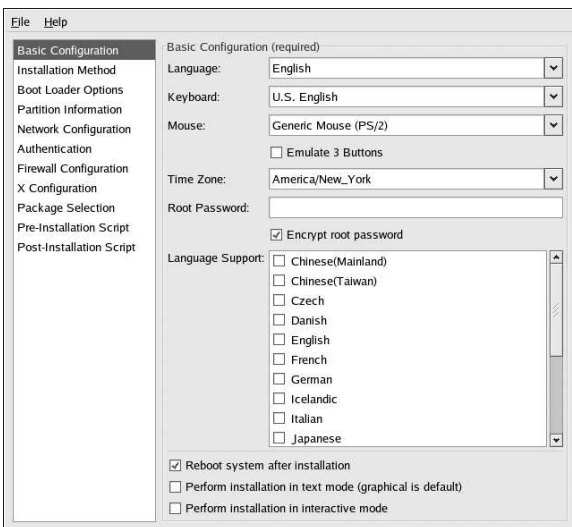


Figure 7-1. Basic Configuration

Choose the language to use during the installation and as the default language after installation from the **Language** menu.

Select the system keyboard type from the **Keyboard** menu.

Choose the mouse for the system from the **Mouse** menu. If you choose **No Mouse**, no mouse will be configured. If you choose **Probe for Mouse** the installation program will try to autodetect the mouse. Probing works for most modern mice.

If you have a two-button button mouse, you can emulate a three-button mouse by selecting **Emulate 3 Buttons**. If this option is selected, simultaneously clicking the left and right mouse buttons will be recognized as a middle mouse button click.

From the **Time Zone** menu, choose the time zone to use for the system.

Enter the desired root password for the system in the **Root Password** text entry box. If you want to save the password as an encrypted password in the file, select **Encrypt root password**. When the file is saved, the plaintext password that you typed will be encrypted and written to the kickstart file. Do not type an already encrypted password and select to encrypt it.

To install languages in addition to the one selected from the **Language** pulldown menu, check them in the **Language Support** list. The language selected from the **Language** pulldown menu is used by default after installation; however, the default can be changed with the **Language Configuration Tool** (`redhat-config-language`) after installation.

Choosing **Reboot system after installation** will reboot your system automatically after the installation is finished.

Kickstart installations are performed in graphical mode by default. To override this default and use text-mode instead, check the **Perform installation in text mode** button.

You can perform a kickstart installation in interactive mode. This means that the installation program will use all the options pre-configured in the kickstart file, but it will allow you to preview the options in each screen before you can continue to the next screen. To continue to the next screen, click the **Next** button after you have approved the settings. If you are not satisfied with the pre-configured options, you can change them before continuing the installation. If you prefer this type of installation, check the **Perform installation in interactive mode** button.

7.2. Installation Method

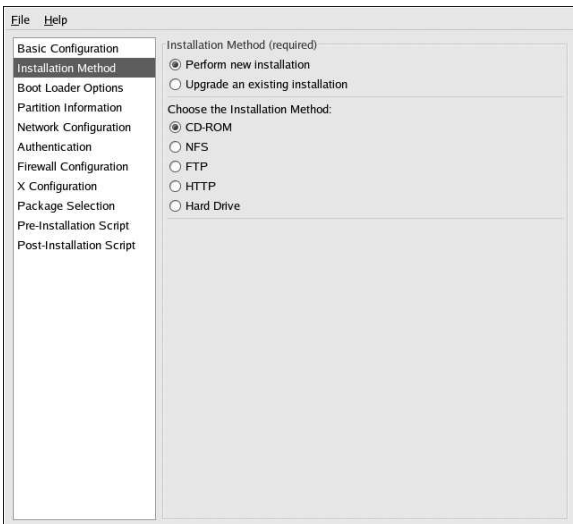


Figure 7-2. Installation Method

The **Installation Method** screen allows you to choose whether you want to perform a full installation or an upgrade. If you choose upgrade, the **Partition Information** and **Package Selection** options will be disabled. They are not supported for kickstart upgrades.

Also choose the type of kickstart installation to perform from this screen. You can choose from the following options:

- **CD-ROM** — Choose this option if you wish to install Red Hat Linux from the Red Hat Linux CD-ROMs.
- **NFS** — Choose this option if you wish to install Red Hat Linux from an NFS shared directory. Two text entry boxes for the NFS server and NFS directory will appear. Enter the fully-qualified domain name or IP address of the NFS server. For the NFS directory, enter the name of the NFS directory that contains the `RedHat` directory of the installation tree. For example, if your NFS server contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386` for the NFS directory.
- **FTP** — Choose this option if you wish to install Red Hat Linux from an FTP server. Two text entry boxes for the FTP server and FTP directory will appear. Enter the fully-qualified domain name or IP address of the FTP server. For the FTP directory, enter the name of the FTP directory that contains the `RedHat` directory. For example, if your FTP server contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386` for the FTP directory.
- **HTTP** — Choose this option if you wish to install Red Hat Linux from an HTTP server. Two text entry boxes for the HTTP server and HTTP directory will appear. Enter the fully-qualified domain name or IP address of the HTTP server. For the HTTP directory, enter the name of the HTTP directory that contains the `RedHat` directory. For example, if your HTTP server contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386` for the HTTP directory.
- **Hard Drive** — Choose this option if you wish to install Red Hat Linux from a hard drive. Two text entry boxes for hard drive partition and hard drive directory will appear. Hard drive installations require the use of ISO (or CD-ROM) images. Be sure to verify that the ISO images are intact before you start the installation. To verify them, use an `md5sum` program. Enter the hard drive partition that contains the ISO images (for example, `/dev/hda1`) in the **Hard Drive Partition** text box, and enter the directory that contains the ISO images in the **Hard Drive Directory** text box.

7.3. Boot Loader Options

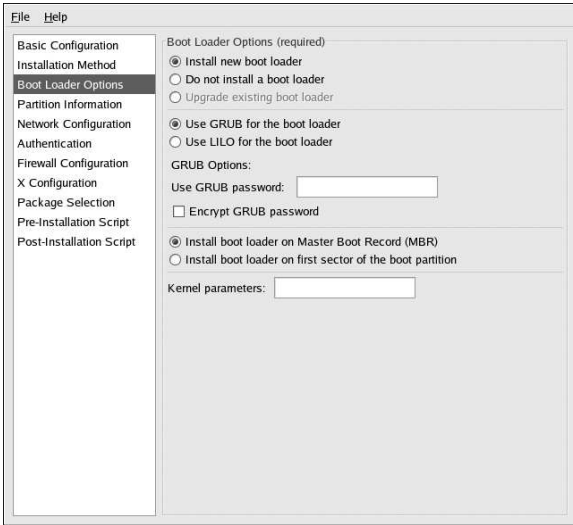


Figure 7-3. Boot Loader Options

You have the option of installing GRUB or LILO as the boot loader. If you do not want to install a boot loader, select **Do not install a boot loader**. If you choose not to install a boot loader, make sure you create a boot disk or have another way to boot (such as a third-party boot loader) your system.

If you choose to install a boot loader, you must also choose which boot loader to install (GRUB or LILO) and where to install the boot loader (the Master Boot Record or the first sector of the `/boot` partition). Install the boot loader on the MBR if you plan to use it as your boot loader. If you are using a different boot loader, install LILO or GRUB on the first sector of the `/boot` partition and configure the other boot loader to boot Red Hat Linux.

If you need to pass any special parameters to the kernel to be used when the system boots, enter them in the **Kernel parameters** text field. For example, if you have an IDE CD-ROM Writer, you can tell the kernel to use the SCSI emulation driver that must be loaded before using `cdrecord` by typing `hdd=ide-scsi` as a kernel parameter (where `hdd` is the CD-ROM device).

If you choose LILO as the boot loader, choose whether you want to use linear mode and whether you want to force the use of `lba32` mode.

If you choose GRUB as the boot loader, you can password protect it by configuring a GRUB password. Enter a password in the **Use GRUB password** text entry area. If you want to save the password as an encrypted password in the file, select **Encrypt GRUB password**. When the file is saved, the plaintext password that you typed will be encrypted and written to the kickstart file. Do not type an already encrypted password and select to encrypt it.

If you selected to **Upgrade an existing installation** on the **Installation Method** page, you can select **Upgrade existing boot loader** to upgrade the existing boot loader configuration, while preserving the old entries.

7.4. Partition Information

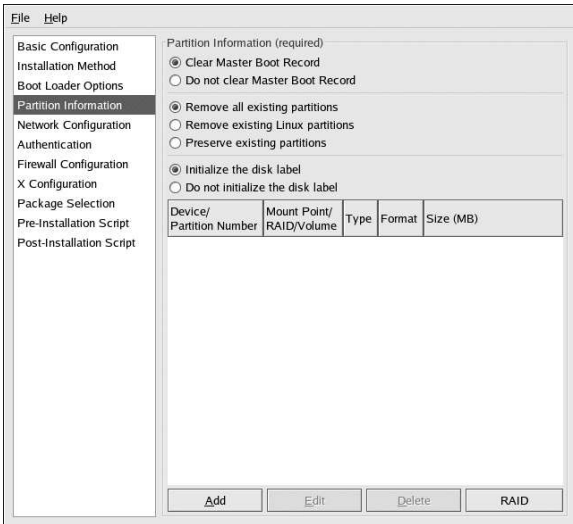


Figure 7-4. Partition Information

Select whether or not to clear the Master Boot Record (MBR). You can also choose to remove all the existing partitions, remove all the existing Linux partitions, or preserve existing partitions.

You can initialize the disk label to the default for the architecture of the system (`msdos` for x86 and `gpt` for Itanium). Select **Initialize the disk label** if you are installing on a brand new hard drive.

7.4.1. Creating Partitions

To create a partition, click the **Add** button. The **Partition Options** window shown in Figure 7-5 will appear. Choose mount point, file system type, and partition size for the new partition. Optionally, you can also choose from the following:

- In the **Additional Size Options** section, choose to make the partition a fixed size, up to a chosen size, or fill the remaining space on the hard drive. If you selected swap as the file system type, you can select to have the installation program create the swap partition with the recommended size instead of specifying a size.
- Force the partition to be created as a primary partition.
- Create the partition on a specific hard drive. For example, to make the partition on the first IDE hard disk (`/dev/hda`), specify `hda` as the drive. Do not include `/dev` in the drive name.
- Use an existing partition. For example, to make the partition on the first partition on the first IDE hard disk (`/dev/hda1`), specify `hda1` as the partition. Do not include `/dev` in the partition name.
- Format the partition as the chosen file system type.

Mount Point:

File System Type: ext3

Size (MB): 1

Additional Size Options

Fixed size

Grow to maximum of (MB):

Fill all unused space on disk

Use recommended swap size

Force to be a primary partition (asprimary)

Make partition on specific drive (ondisk)
Drive : (for example: hda or sdc)

Use existing partition (onpart)
Partition : (for example: hda1 or sdc3)

Format partition

Figure 7-5. Creating Partitions

To edit an existing partition, select the partition from the list and click the **Edit** button. The same **Partitions Options** window that appears when you add a partition appears, except it contains the values for the selected partition. Modify the partition options and click **OK**.

To delete an existing partition, select the partition from the list and click the **Delete** button.

7.4.1.1. Creating Software RAID Partitions

Read Chapter 3 to learn more about RAID and the different RAID levels. RAID 0, 1, and 5 can be configured.

To create a software RAID partition, use the following steps:

1. Click the **RAID** button.
2. Select **Create a software RAID partition**.
3. Configure the partitions as previously described, except select **Software RAID** as the file system type. Also, you must specify a hard drive on which to make the partition or an existing partition to use.

Mount Point:

File System Type: software RAID

Size (MB): 2048

Additional Size Options

Fixed size

Grow to maximum of (MB):

Fill all unused space on disk

Use recommended swap size

Force to be a primary partition (asprimary)

Make partition on specific drive (ondisk)

Drive : (for example: hda or sdc)

Use existing partition (onpart)

Partition : (for example: hda1 or sdc3)

Format partition

Figure 7-6. Creating a Software RAID Partition

Repeat these steps to create as many partitions as needed for your RAID setup. All of your partitions do not have to be RAID partitions.

After creating all the partitions needed to form a RAID device, follow these steps:

1. Click the **RAID** button.
2. Select **Create a RAID device**.
3. Select a mount point, file system type, RAID device name, RAID level, RAID members, number of spares for the software RAID device, and whether to format the partition.

Mount Point: /home

File System Type: ext3

RAID Device: md0

RAID Level: 0

Raid Members

raid.01

raid.02

Number of spares: 1

Format RAID device

Figure 7-7. Creating a Software RAID Device

4. Click **OK** to add the device to the list.

7.5. Network Configuration

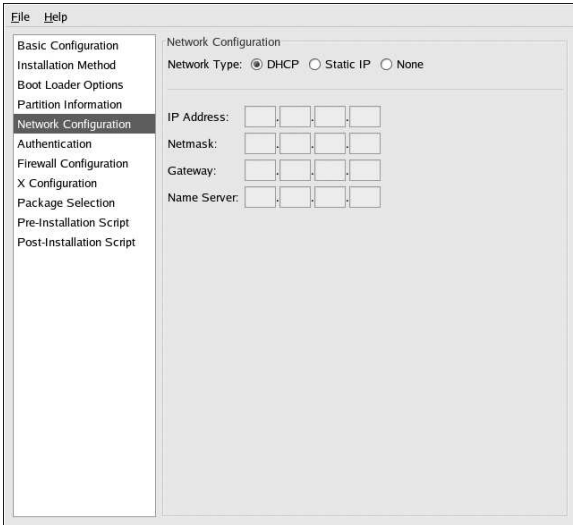


Figure 7-8. Network Configuration

There are three network configuration options: **DHCP**, **Static IP**, and **None**. If there is not an Ethernet card in the system, choose **None**.

Networking is only required if you choose a networking-type installation method (NFS, FTP, or HTTP). If you are unsure which to choose, choose **None**. Networking can always be configured after installation with the **Network Administration Tool** (`redhat-config-network`).

If you select **Static IP**, you must provide additional networking information in the table below the network types.

7.6. Authentication

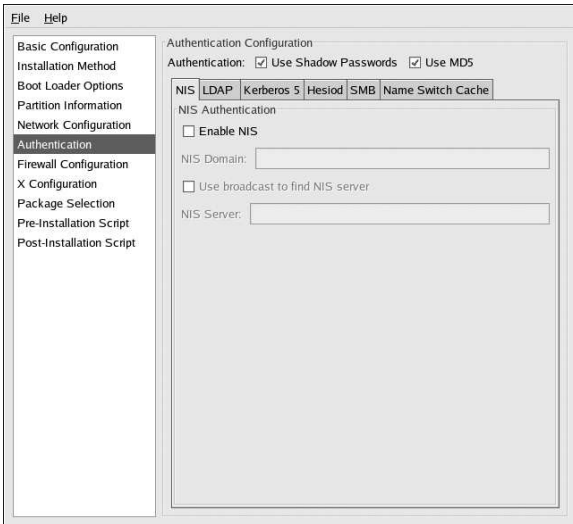


Figure 7-9. Authentication

In the **Authentication** section, select whether to use shadow passwords and MD5 encryption for user passwords. These options are highly recommended and chosen by default.

The **Authentication Configuration** options allows you to configure the following methods of authentication:

- NIS
- LDAP
- Kerberos 5
- Hesiod
- SMB
- Name Switch Cache

They are not enabled by default. To enable one or more of these methods, click the appropriate tab, click the checkbox next to **Enable**, and enter the appropriate information for the authentication method.

7.7. Firewall Configuration

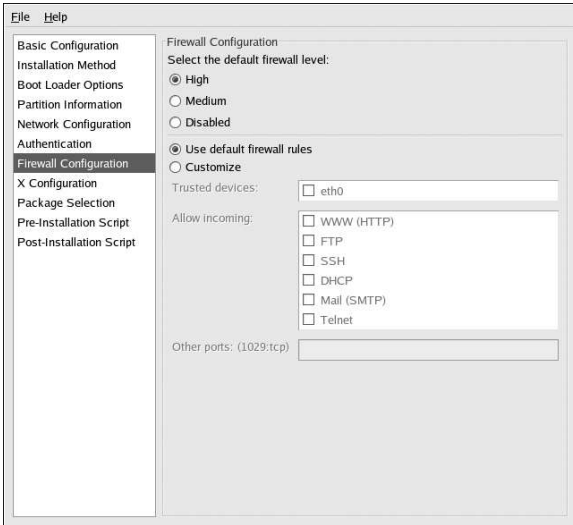


Figure 7-10. Firewall Configuration

The **Firewall Configuration** window is identical to the screen in the Red Hat Linux installation program and the **Security Level Configuration Tool**, with the same functionality. Choose between **High**, **Medium**, and **Disabled** security levels. Refer to Section 12.1 for detailed information about these security levels.

7.8. X Configuration

If you are installing the X Window System, you can configure it during the kickstart installation by checking the **Configure the X Window System** button on the **X Configuration** window as shown in Figure 7-11. If this option is not chosen, the X configuration options will be disabled and the `skipx` option will be written to the kickstart file.

7.8.1. General

The first step in configuring X is to choose the default color depth and resolution. Select them from their respective pulldown menus. Be sure to specify a color depth and resolution that is compatible with the video card and monitor for the system.

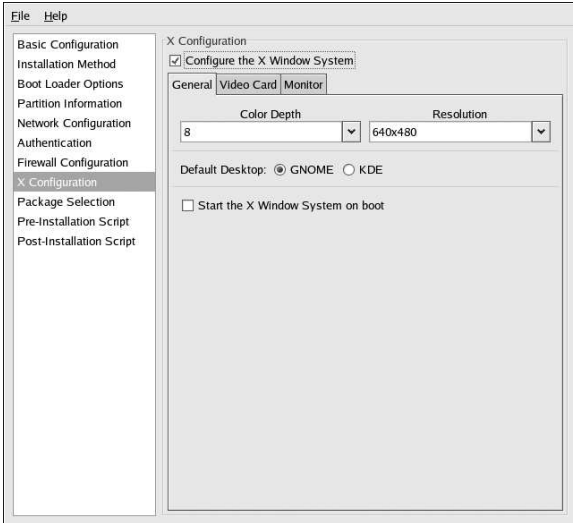


Figure 7-11. X Configuration - General

If you are installing both the GNOME and KDE desktops, you need to choose which desktop you want to be the default. If you are just installing one desktop, be sure to choose it. Once the system is installed, users can choose which desktop they want to be their default. For more information about GNOME and KDE, refer to the *Official Red Hat Linux Installation Guide* and the *Official Red Hat Linux Getting Started Guide*.

Next, choose whether to start the X Window System when the system is booted. This option will start the system in runlevel 5 with the graphical login screen. After the system is installed, this can be changed by modifying the `/etc/inittab` configuration file.

7.8.2. Video Card

Probe for video card is selected by default. Accept this default if you want the installation program to probe for the video card during installation. Probing works for most modern video cards. If you select this option and the installation program can not successfully probe the video card, the installation program will stop at the video card configuration screen. To continue the installation process, select your video card from the list and click **Next**.

Alternatively, you can select the video card from the list on the **Video Card** tab as shown in Figure 7-12. Also select the amount of video RAM the selected video card has from the **Video Card RAM** pulldown menu. These values will be used by the installation program to configure the X Window System.

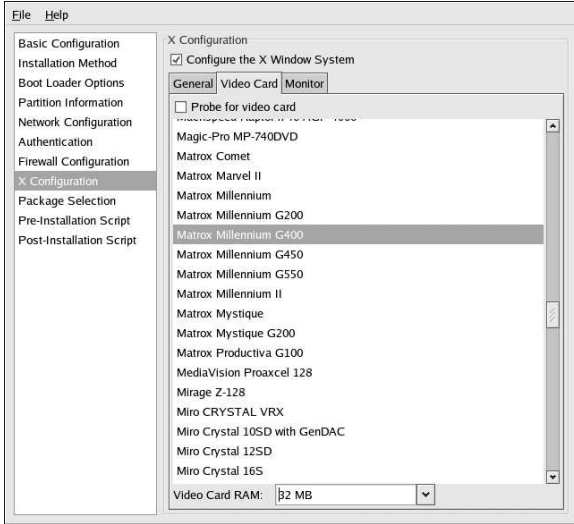


Figure 7-12. X Configuration - Video Card

7.8.3. Monitor

After configuring the video card, click on the **Monitor** tab shown in Figure 7-13.

Probe for monitor is selected by default. Accept this default if you want the installation program to probe for the monitor during installation. Probing works for most modern monitors. If you select this option and the installation program can not successfully probe the monitor, the installation program will stop at the monitor configuration screen. To continue the installation process, select your monitor from the list and click **Next**.

Alternatively, you can select your monitor from the list. You can also specify the horizontal and vertical sync rates instead of specifying a monitor by checking the **Specify hsync and vsync instead of monitor** option. This option is useful if the monitor for the system is not listed. Notice that when this option is enabled, the monitor list is disabled.

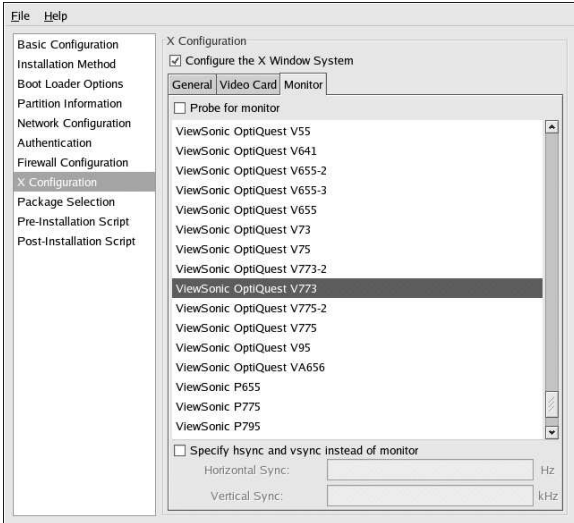


Figure 7-13. X Configuration - Monitor

7.9. Package Selection

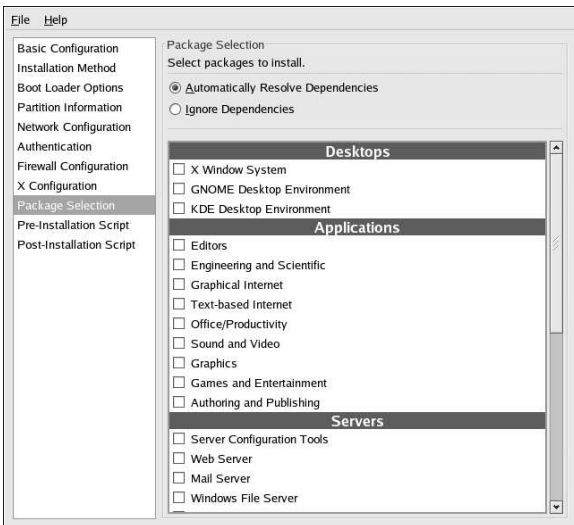


Figure 7-14. Package Selection

The **Package Selection** window allows you to choose which package groups to install.

There are also options to resolve package dependencies automatically and to ignore package dependencies.

Currently, **Kickstart Configurator** does not allow you to select individual packages. To install individual packages, modify the `%packages` section of the kickstart file after you save it.

7.10. Pre-Installation Script

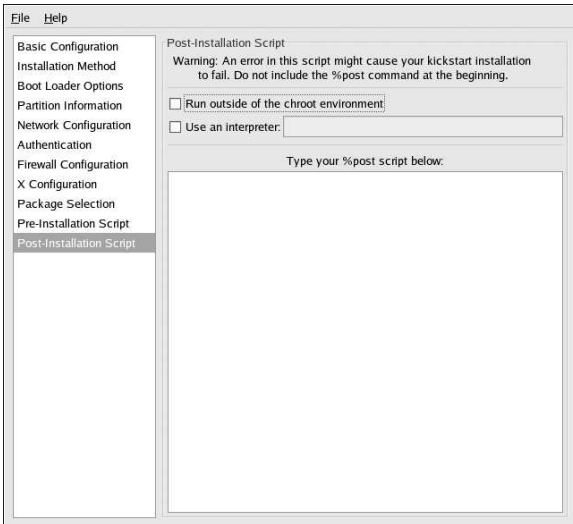


Figure 7-15. Pre-Installation Script

You can add commands to run on the system immediately after the kickstart file has been parsed and before the installation begins. If you have configured the network in the kickstart file, the network is enabled before this section is processed. If you would like to include a pre-installation script, type it in the text area.

If you want to specify a scripting language to use to execute your script, click the **Use an interpreter** button and enter the interpreter in the text box beside the button. For example, `/usr/bin/python2.2` can be specified for a Python script. This option corresponds to using `%pre --interpreter /usr/bin/python2.2` in your kickstart file.



Caution

Do not include the `%pre` command. It will be added for you.

7.11. Post-Installation Script

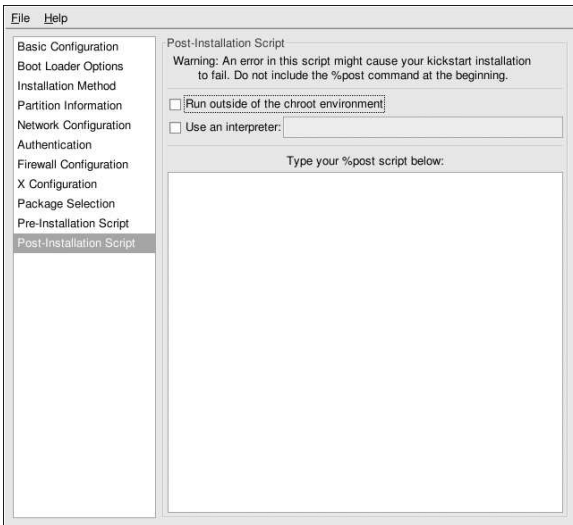


Figure 7-16. Post-Installation Script

You can also add commands to execute on the system after the installation is completed. If you have properly configured the network in the kickstart file, the network is enabled. If you would like to include a post-installation script, type it in the text area.



Caution

Do not include the `%post` command. It will be added for you.

For example, to change the message of the day for the newly installed system, add the following command to the `%post` section:

```
echo "Hackers will be punished!" > /etc/motd
```



Tip

More examples can be found at Section 6.7.1.

7.11.1. Chroot Environment

If you want your post-installation script to run outside of the chroot environment, click the checkbox next to this option on the top of the **Post-Installation** window. This is equivalent to using the `--nochroot` option in the `%post` section.

**Tip**

If you want to make any changes to the newly installed file system in the post-installation section outside of the chroot environment, you must append the directory name with `/mnt/sysimage`.

For example, if you check the **Run outside of the chroot environment** button, the previous example needs to be changed to the following:

```
echo "Hackers will be punished!" > /mnt/sysimage/etc/motd
```

7.11.2. Use an Interpreter

If you want to specify a scripting language to use to execute your script, click the **Use an interpreter** button and enter the interpreter in the text box beside the button. For example, `/usr/bin/python2.2` can be specified for a Python script. This option corresponds to using `%post --interpreter /usr/bin/python2.2` in your kickstart file.

7.12. Saving the File

After you have finished choosing your kickstart options, if you want to preview the contents of the kickstart file, select **File => Preview** from the pull-down menu.

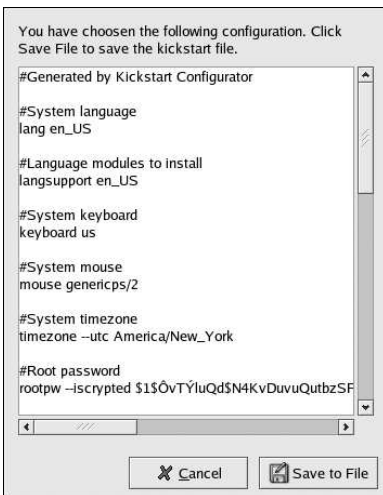


Figure 7-17. Preview

To save the kickstart file, select **File => Save File** or press `[Ctrl]-[S]`. A file dialog box will appear to allow you to select where to save the file.

After saving the file, refer to Section 6.10 for information on how to start the kickstart installation.

When things go wrong, there are ways to fix problems. However, these methods require that you understand the system well. This chapter describes how to boot into rescue mode and single user mode, where you can use your own knowledge to repair the system.

8.1. What is Rescue Mode?

Rescue mode provides the ability to boot a small Linux environment entirely from a diskette, CD-ROM, 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 take a closer look at each of these scenarios.

8.1.1. 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 they overwrite the Master Boot Record (MBR) that originally contained the GRUB or LILO boot loader. If the boot loader is overwritten in this manner, you will not be able to boot Red Hat Linux unless you can get into rescue mode.

Another common problem occurs when using a partitioning tool to resize a partition or create a new partition from free space after installation and it changes the order of your partitions. If the partition number of your / partition changes, the boot loader will not be able to find it to mount the partition. To fix this problem, boot in rescue mode and modify `/boot/grub/grub.conf` if you are using GRUB or `/etc/lilo.conf` if you are using LILO.

8.1.2. 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 (if you are using LILO as your boot loader). If you are using GRUB, you do not have to execute a command to reread the GRUB configuration file. However, if you specify an invalid root device or kernel in the GRUB configuration file, you might not know until you reboot your computer.

In 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.

8.2. Booting Rescue Mode

To boot your system in rescue mode, boot from a Red Hat Linux boot disk or the Red Hat Linux CD-ROM #1, and 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 made from the `boot.img` image. This method requires that the Red Hat Linux CD-ROM #1 be inserted as the rescue image or that the rescue image be on the hard drive as an ISO image.¹
- By booting your system from the Red Hat Linux CD-ROM #1.
- By booting from a network disk made from the `bootnet.img` or PCMCIA boot disk made from `pcmcia.img`. 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, refer to the *Official Red Hat Linux Installation Guide*.

After booting off a boot disk or Red Hat Linux CD-ROM #1 and providing a valid rescue image, you will see the following message:

```
The rescue environment will now attempt to find your Red Hat
Linux installation and mount it under the directory
/mnt/sysimage. You can then make any changes required to your
system. If you want to proceed with this step choose
'Continue'. You can also choose to mount your filesystem
read-only instead of read-write by choosing 'Read-only'.
If for some reason this process fails you can choose 'Skip'
and this step will be skipped and you will go directly to a
command shell.
```

If you select **Continue**, it will attempt to mount your filesystem under the directory `/mnt/sysimage`. If it fails to mount a partition, it will notify you. If you select **Read-Only**, it will attempt to mount your filesystem under the directory `/mnt/sysimage`, but in read-only mode. If you select **Skip**, your filesystem will not be mounted. Choose **Skip** if you think your filesystem is corrupted.

Once you have your system in rescue mode, a prompt appears on VC (virtual console) 1 and VC 2 (use the [Ctrl]-[Alt]-[F1] key combination to access VC 1 and [Ctrl]-[Alt]-[F2] to access VC 2):

```
sh-2.05a#
```

If you selected **Continue** to mount your partitions automatically and they were mounted successfully, you are in single-user mode.

To mount a Linux partition manually inside rescue mode, create a directory such as `/foo`, and type the following command:

```
mount -t ext3 /dev/hda5 /foo
```

In the above command, `/foo` is a directory that you have created and `/dev/hda5` is the partition you want to mount. If the partition is of type `ext2`, replace `ext3` with `ext2`.

If you do not know the names of your partitions, use the following command to list them:

```
fdisk -l
```

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

If your filesystem is mounted and you want to make your system the root partition, use the command `chroot /mnt/sysimage`. This is useful if you need to run commands such as `rpm` that require your root partition to be mounted as `/`. To exit the `chroot` environment, type `exit`, and you will return to the prompt.

From the `bash#` prompt, you can run many useful commands including:

anaconda	gzip	mkfs.ext2	probe
badblocks	head	mknod	ps
bash	hwclock	mkraid	python2.2
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	lsmmod	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	parted	uncpio
grep	mkdir	pico	unig
gunzip	mke2fs	ping	zcat

8.3. Booting Single-User Mode

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 single-user mode.

If you are using GRUB, use the following steps to boot into single-user mode:

1. If you have a GRUB password configured, type `p` and enter the password.
2. Select **Red Hat Linux** with the version of the kernel that you wish to boot and type `e` for edit. You will be presented with a list of items in the configuration file for the title you just selected.
3. Select the line that starts with `kernel` and type `e` to edit the line.
4. Go to the end of the line and type **single** as a separate word (press the [Spacebar] and then type **single**). Press [Enter] to exit edit mode.
5. Back at the GRUB screen, type `b` to boot into single user mode.

If you are using LILO, specify 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.

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

Software RAID Configuration

Read Chapter 3 first to learn about RAID, the differences between Hardware and Software RAID, and the differences between RAID 0, 1, and 5.

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 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.

1. On the **Disk Partitioning Setup** screen, select **Manually partition with Disk Druid**.
2. In **Disk Druid**, choose **New** to create a new partition.
3. You will not be able to enter a mount point (you will be able to do that once you have created your RAID device).
4. Choose **software RAID** from the **File System Type** pull-down menu as shown in Figure 9-1.




Figure 9-1. Creating a New RAID Partition

5. For **Allowable Drives**, select the drive(s) 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.
6. Enter the size that you want the partition to be.
7. Select **Fixed size** to make the partition the specified size, select **Fill all space up to (MB)** and enter a size in MBs to give range for the partition size, or select **Fill to maximum allowable**

size to make it grow to fill all available space on the hard disk. If you make more than one partition growable, they will share the available free space on the disk.

8. Select **Force to be a primary partition** if you want the partition to be a primary partition.
9. Select **Check for bad blocks** if you want the installation program to check for bad blocks on the hard drive before formatting it.
10. Click **OK** to return to the main screen.

Repeat 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, you can configure only the `/home` partition as a software RAID device.

Once you have all of your partitions created as **software RAID** partitions, follow these steps:

1. Select the **RAID** button on the **Disk Druid** main partitioning screen (see Figure 9-3).
2. Next, Figure 9-2 will appear, where you can make a RAID device.




Figure 9-2. Making a RAID Device

3. Enter a mount point.
4. Choose the filesystem type for the partition.
5. Choose your RAID level. You can choose from **RAID 0**, **RAID 1**, and **RAID 5**.



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).

6. Select which partitions will go into this RAID array and click **OK**.
7. A spare partition can be specified for RAID 1 and RAID 5. If a software RAID partition fails, the spare will automatically be used as a replacement. For each spare you want to specify, you must create an additional software RAID partition (in addition to the partitions for the RAID device). In the previous step, select the partitions for the RAID device and the partition(s) for the spare(s). Select the number of spares.
8. Select whether you want the partition formatted.

9. The RAID device will appear in the **Drive Summary** list as shown in Figure 9-3. At this point, you can continue with your installation process. Refer to the *Official Red Hat Linux Installation Guide* for further instructions.

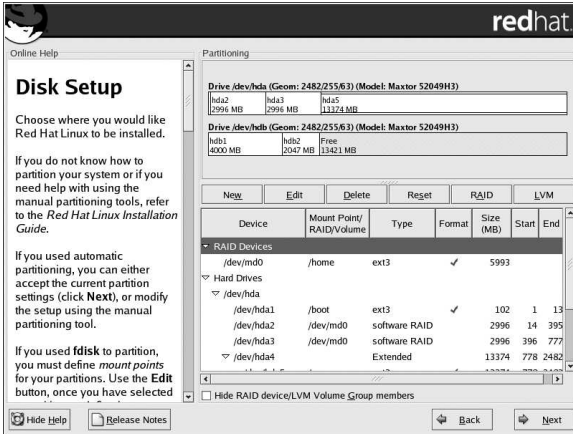


Figure 9-3. Creating a RAID Array

Chapter 10.

LVM Configuration

LVM can be configured during the graphical installation of Red Hat Linux or during a kickstart installation. You can use the utilities from the `lvm` package to create your LVM configuration, but these instructions will focus on using **Disk Druid** during the Red Hat Linux installation to complete this task.

Read Chapter 4 first to learn about LVM. An overview of the steps required to configure LVM:

- Create *physical volumes* from the hard drives.
- Create *volume groups* from the physical volumes.
- Create *logical volumes* from the volume groups and assign the logical volumes mount points.



Note

You can only edit LVM volume groups in GUI installation mode. In text installation mode, you can assign mount points to existing logical volumes.

To create a logical volume group with logical volumes during the Red Hat Linux installation:

1. On the **Disk Partitioning Setup** screen, select **Manually partition with Disk Druid**.
2. Select **New**.
3. You will not be able to enter a mount point (you will be able to do that once you have created your volume group).
4. Select **physical volume (LVM)** from the **Filesystem Type** pull-down menu as shown in Figure 10-1.




Figure 10-1. Creating a Physical Volume

5. A physical volume must be constrained to one drive. For **Allowable Drives**, select the drive on which the physical volume will be created. If you have multiple drives, all drives will be selected here, and you must deselect all but one drive.
6. Enter the size that you want the physical volume to be.

7. Select **Fixed size** to make the physical volume the specified size, select **Fill all space up to (MB)** and enter a size in MBs to give range for the physical volume size, or select **Fill to maximum allowable size** to make it grow to fill all available space on the hard disk. If you make more than one growable, they will share the available free space on the disk.
8. Select **Force to be a primary partition** if you want the partition to be a primary partition.
9. Select **Check for bad blocks** if you want the installation program to check for bad blocks on the hard drive before formatting it.
10. Click **OK** to return to the main screen.

Repeat these step to create as many physical volumes as needed for your LVM setup. For example, if you want the volume group to span over more than one drive, create a physical volume on each of the drives.

Warning

The `/boot` partition can not be on a volume group because the boot loader can not read it. If you want to have your root partition on a logical volume, you will need to create a separate `/boot` partition which is not a part of a volume group.

Once all the physical volumes are created, follow these steps:

1. Click the **LVM** button to collect the physical volumes into volume groups. A volume group is basically a collection of physical volumes. You can have multiple logical volume groups, but a physical volume can only be in one volume group.

Note

There is overhead disk space reserved in the logical volume group. The summation of the physical volumes may not equal the size of the volume group; however, the size of the logical volumes shown is correct.

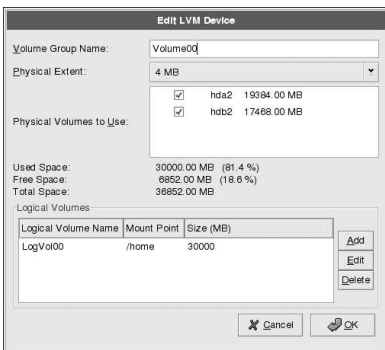


Figure 10-2. Creating an LVM Device

2. Change the **Volume Group Name** if desired.

- All logical volumes inside the volume group must be allocated in *physical extent* units. By default, the physical extent is set to 4 MB; thus, logical volume sizes must be divisible by 4 MBs. If you enter a size that is not a unit of 4 MBs, the installation program will automatically select the closest size in units of 4 MBs. It is not recommended that you change this setting.
- Select which physical volumes to use for the volume group.
- Create logical volumes with mount points such as `/home`. Remember that `/boot` can not be a logical volume. To add a logical volume, click the **Add** button in the **Logical Volumes** section. A dialog window as shown in Figure 10-3 will appear.

Figure 10-3. Creating a Logical Volume

Repeat these steps for each volume group you want to create.



Tip

You may want to leave some free space in the logical volume group so you can expand the logical volumes later.

Device	Mount Point/ RAID Volume	Type	Format	Size (MB)	Start	End
LVM Volume Groups						
LogVol00	/home	ext3	✓	30000		
Hard Disks						
dev/hda						
dev/hda1	/boot	ext3	✓	76	1	10
dev/hda2	Volume00	LVM PV	✓	19391	11	2482
dev/hdb						
dev/hdb1	swap	swap	✓	1992	1	254
dev/hdb2	Volume00	LVM PV	✓	17477	255	2482

Figure 10-4. Logical Volumes Created

Network-Related Configuration

Network Configuration

To communicate with other computers, computers need a network connection. This is accomplished by having the operating system recognize an interface card (such as Ethernet, ISDN modem, or token ring) and configuring the interface to connect to the network.

The **Network Administration Tool** can be used to configure the following types of network interfaces:

- Ethernet
- ISDN
- modem
- xDSL
- token ring
- CIPE
- wireless devices

To use the **Network Administration Tool**, you must be running the X Window System and have root privileges. To start the application, go to the **Main Menu Button** (on the Panel) => **System Settings** => **Network**, or type the command `redhat-config-network` at a shell prompt (for example, in an **XTerm** or a **GNOME terminal**).

If you prefer modifying the configuration files directly, refer to the *Official Red Hat Linux Reference Guide* for information on their location and contents.

**Tip**

Go to the Red Hat Hardware Compatibility List (<http://hardware.redhat.com/hcl/>) to determine if Red Hat Linux supports your hardware device.

11.1. Overview

To configure a network connection with the **Network Administration Tool**, perform the following steps:

1. Add the physical hardware device to the hardware list.
2. Add a network device associated with the physical hardware device.
3. Configure any hosts that can not be looked up through DNS.
4. Configure the hostname and DNS settings.

This chapter will discuss each of these steps for each type of network connection.

11.2. Establishing an Ethernet Connection

To establish an Ethernet connection, you need a network interface card (NIC), a network cable (usually a CAT5 cable), and a network to connect to. There are different speeds to networks; make sure your NIC is compatible with the network to which you want to connect.

To add an Ethernet connection, follow these steps:

1. Click the **Devices** tab.
2. Click the **Add** button.
3. Select **Ethernet connection** from the **Device Type** list, and click **Forward**.
4. If you have already added the network interface card to the hardware list, select it from the **Ethernet card** list. Otherwise, select **Other Ethernet Card** to add the hardware device.



Note

The installation program usually detects supported Ethernet devices and prompts you to configure them. If you configured any Ethernet devices during the installation, they will already appear in the hardware list on the **Hardware** tab.

5. If you selected **Other Ethernet Card**, the **Select Ethernet Adapter** window appears. Select the manufacturer and model of the Ethernet card. Select the device name. If this is the system's first Ethernet card, select **eth0** as the device name, if this is the second Ethernet card, select **eth1**, and so on. The **Network Administration Tool** also allows you to configure the resources for the NIC. Click **Forward** to continue.
6. On the **Configure Network Settings** page as shown in Figure 11-1, choose between DHCP and a static IP address. If the device receives a different IP address each time the network is started, do not specify a hostname. Click **Forward** to continue.

Figure 11-1. Ethernet Settings

7. Click **Apply** on the **Create Ethernet Device** page.

After configuring the Ethernet device, it appears in the device list as shown in Figure 11-2.



Figure 11-2. Ethernet Device

Be sure to click **Apply** to save the changes.

After adding the Ethernet device, you can edit its configuration by selecting the device from the device list and clicking **Edit**. For example, when the device is added, it is configured to start at boot time by default. You can edit its configuration to modify this setting.

When the device is added, it is not activated, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button.

11.3. Establishing an ISDN Connection

An ISDN connection is an Internet connection established with a ISDN modem card through a special phone line installed by the phone company. ISDN connections are popular in Europe.

To add an ISDN connection, follow these steps:

1. Click the **Devices** tab.
2. Click the **Add** button.
3. Select **ISDN connection** from the **Device Type** list, and click **Forward**.
4. Select the ISDN adapter from the pulldown menu. Then configure the resources and D channel protocol for the adapter. Click **Forward** to continue.

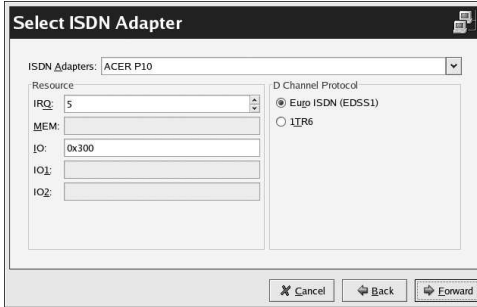


Figure 11-3. ISDN Settings

5. If your ISP is in the pre-configured list, select it. Otherwise, enter the required information about your ISP account. If you do not know the values, contact your ISP. Click **Forward**.
6. On the **Create Dialup Connection** page, click **Apply**.

After configuring the ISDN device, it appears in the device list as an `ipp0` device as shown in Figure 11-4.



Figure 11-4. ISDN Device

Be sure to click **Apply** to save the changes.

After adding the modem device, you can edit its configuration by selecting the device from the device list and clicking **Edit**. For example, when the device is added, it is configured not to start at boot time by default. Edit its configuration to modify this setting. Compression, PPP options, login name, password, and more can be changed.

When the device is added, it is not activated, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button.

11.4. Establishing a Modem Connection

A modem can be used to configure an Internet connection over an active phone line. An Internet Service Provider (ISP) account (also called a dial-up account) is required.

To add a modem connection, follow these steps:

1. Click the **Devices** tab.
2. Click the **Add** button.
3. Select **Modem connection** from the **Device Type** list, and click **Forward**.
4. If there is a modem already configured in the hardware list (on the **Hardware** tab), **Network Administration Tool** assumes you want to use it to establish a modem connection. If there is not modem already configured, it tries to detect any modems in the system. This probe might take a while.
5. After probing, the window in Figure 11-5 appears.

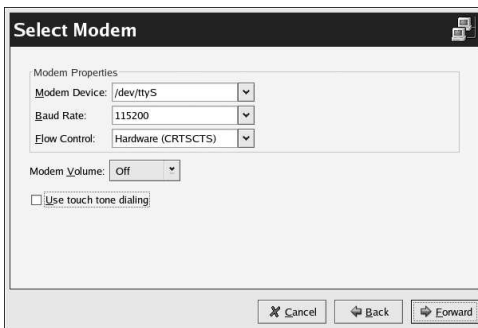


Figure 11-5. Modem Settings

6. Configure the baud rate, flow control, and modem volume. If you do not know these values, accept the defaults. If you do not have touch tone dialing, uncheck the corresponding checkbox.
7. Click **Forward**.
8. If your ISP is in the pre-configured list, select it. Otherwise, enter the required information about your ISP account. If you do not know the values, contact your ISP. Click **Forward**.
9. On the **Create Dialup Connection** page, click **Apply**.

After configuring the modem device, it appears in the device list as shown in Figure 11-6.

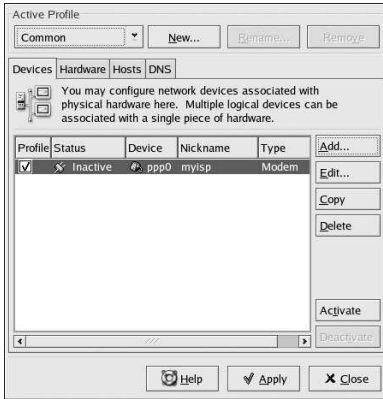


Figure 11-6. Modem Device

Be sure to click **Apply** to save the changes.

After adding the modem device, you can edit its configuration by selecting the device from the device list and clicking **Edit**. For example, when the device is added, it is configured not to start at boot time by default. Edit its configuration to modify this setting. Compression, PPP options, login name, password, and more can also be changed.

When the device is added, it is not activated, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button.

11.5. Establishing an xDSL Connection

DSL stands for Digital Subscriber Lines. There are different types of DSL such as ADSL, IDSL, and SDSL. **Network Administration Tool** uses the term xDSL to mean all types of DSL connections.

Some DSL providers require you to configure your system to obtain an IP address through DHCP with an Ethernet card. Some DSL providers require you to configure a PPPoE (Point-to-Point Protocol over Ethernet) connection with an Ethernet card. Ask your DSL provider which method to use.

If you are required to use DHCP, refer to Section 11.2 to configure your Ethernet card.

If you are required to use PPPoE, follow these steps:

1. Click the **Devices** tab.
2. Click the **Add** button.
3. Select **xDSL connection** from the **Device Type** list, and click **Forward**.
4. If your Ethernet card is already in the hardware list, select the **Ethernet Device** from the pull-down menu from the page shown in Figure 11-7. Otherwise, the **Select Ethernet Adapter** window appears.



Note

The installation program usually detects supported Ethernet devices and prompts you to configure them. If you configured any Ethernet devices during the installation, they will already appear in the hardware list on the **Hardware** tab.



Figure 11-7. xDSL Settings

5. If the **Select Ethernet Adapter** window appears, select the manufacturer and model of the Ethernet card. Select the device name. If this is the system's first Ethernet card, select **eth0** as the device name, if this is the second Ethernet card, select **eth1**, and so on. The **Network Administration Tool** also allows you to configure the resources for the NIC. Click **Forward** to continue.
6. Enter the **Provider Name**, **Login Name**, and **Password**.
7. Click **Forward**.
8. On the **Create DSL Connection** page, click **Apply**.

After configuring the DSL connect, it appears in the device list as shown in Figure 11-6.




Figure 11-8. xDSL Device

Be sure to click **Apply** to save the changes.

After adding the xDSL connection, you can edit its configuration by selecting the device from the device list and clicking **Edit**. For example, when the device is added, it is configured not to start at boot time by default. Edit its configuration to modify this setting.

When the device is added, it is not activated, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button.

11.6. Establishing a Token Ring Connection

A token ring network is a network in which all the computers are connected in a circular pattern. A *token*, or a special network packet, travels around the token ring and allows computers to send information to each other.



Tip

For more information on using token ring under Linux, refer to the *Linux Token Ring Project* website available at <http://www.linuxtr.net>.

To add a token ring connection, follow these steps:

1. Click the **Devices** tab.
2. Click the **Add** button.
3. Select **Token Ring connection** from the **Device Type** list, and click **Forward**.
4. If you have already added the token ring card to the hardware list, select it from the **Ethernet card** list. Otherwise, select **Other Tokenring Card** to add the hardware device.
5. If you selected **Other Tokenring Card**, the **Select Token Ring Adapter** window as shown in Figure 11-9 appears. Select the manufacturer and model of the adapter. Select the device name. If this is the system's first token ring card, select **tr1**, if this is the second token ring card, select **tr1**, and so on. The **Network Administration Tool** also allows the user to configure the resources for the adapter. Click **Forward** to continue.

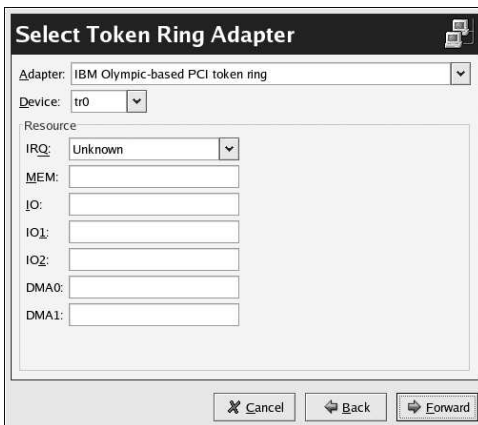


Figure 11-9. Token Ring Settings

6. On the **Configure Network Settings** page, choose between DHCP and static IP address. You may specify a hostname for the device. If the device receives a dynamic IP address each time the network is started, do not specify a hostname. Click **Forward** to continue.
7. Click **Apply** on the **Create Ethernet Device** page.

After configuring the token ring device, it appears in the device list as shown in Figure 11-10.

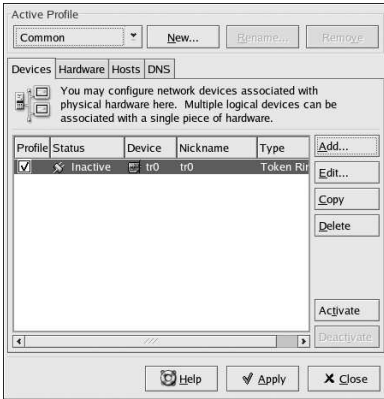


Figure 11-10. Token Ring Device

Be sure to click **Apply** to save the changes.

After adding the device, you can edit its configuration by selecting the device from the device list and clicking **Edit**. For example, you can configure whether the device is started at boot time.

When the device is added, it is not activated, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button.

11.7. Establishing a CIPE Connection

CIPE stands for Crypto IP Encapsulation. It is used to configure an IP tunneling device. For example, CIPE can be used to grant access from the outside world into a Virtual Private Network (VPN). If you need to setup a CIPE device, contact your system administrator for the correct values.

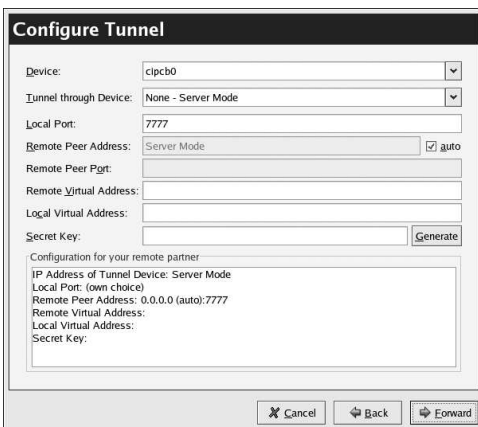


Figure 11-11. CIPE Settings

11.8. Establishing a Wireless Connection

Wireless Ethernet devices are becoming increasingly popular. The configuration is similar to the Ethernet configuration except that it allows you to configure ESSID, mode, frequency, channel, transmit rate, and key for your wireless device.

To add a wireless Ethernet connection, follow these steps:

1. Click the **Devices** tab.
2. Click the **Add** button.
3. Select **Wireless connection** from the **Device Type** list, and click **Forward**.
4. If you have already added the wireless network interface card to the hardware list, select it from the **Ethernet card** list. Otherwise, select **Other Ethernet Card** to add the hardware device.



Note

The installation program usually detects supported wireless Ethernet devices and prompts you to configure them. If you configured them during the installation program, they will already appear in the hardware list on the **Hardware** tab.

5. If you selected **Other Ethernet Card**, the **Select Ethernet Adapter** window appears. Select the manufacturer and model of the Ethernet card and the device. If this is the first Ethernet card for the system, select **eth0**, if this is the second Ethernet card for the system, select **eth1**, and so on. The **Network Administration Tool** also allows the user to configure the resources for the wireless network interface card. Click **Forward** to continue.
6. On the **Configure Wireless Connection** page as shown in Figure 11-12, configure the ESSID, mode, frequency, channel, transmit rate, and key for your wireless device.

Figure 11-12. Wireless Settings

7. On the **Configure Network Settings** page, choose between DHCP and static IP address. You may specify a hostname for the device. If the device receives a dynamic IP address each time the network is started, do not specify a hostname. Click **Forward** to continue.
8. Click **Apply** on the **Create Wireless Device** page.

After configuring the wireless device, it appears in the device list as shown in Figure 11-13.



Figure 11-13. Wireless Device

Be sure to click **Apply** to save the changes.

After adding the wireless device, you can edit its configuration by selecting the device from the device list and clicking **Edit**. For example, you can configure the device to activate at boot time.

When the device is added, it is not activated, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button.

11.9. Managing Hosts

The **Hosts** tab allows you to add, edit, or remove hosts from the `/etc/hosts` file. This file contains IP addresses and their corresponding hostnames.

When your system tries to resolve a hostname to an IP address or determine the hostname for an IP address, it refers to the `/etc/hosts` file before using the name servers (if you are using the default Red Hat Linux configuration). If the IP address is listed in the `/etc/hosts` file, the name servers are not used. If your network contains computers whose IP addresses are not listed in DNS, it is recommended that you add them to the `/etc/hosts` file.

To add an entry to the `/etc/hosts` file, click **Add** in the **Hosts** tab, provide the requested information, and click **OK**. Click **Apply** to write the entry to the file.



Warning

Do not remove the `localhost` entry.



Figure 11-14. Hosts Configuration



Tip

To change lookup order, edit the `/etc/host.conf` file. The line `order hosts, bind` specifies that the `/etc/hosts` takes precedence over the name servers. Changing the line to `order bind, hosts` configures your system to resolve hostnames and IP addresses using the name servers first. If the IP address can not be resolved through the name servers, your system looks for the IP address in the `/etc/hosts` file.

11.10. Managing DNS Settings

The **DNS** tab allows you to configure the system's hostname, domain, name servers, and search domain. Name servers are used to look up other hosts on the network.




Figure 11-15. DNS Configuration

**Note**

The name servers section does not configure the system to be a name server.

If the DNS server names are retrieved from DHCP (or retrieved from the ISP of a modem connection) do not add primary, secondary, or tertiary DNS servers.

11.11. Activating Devices

Network devices can be configured to activate at boot time or not to start at boot time. For example, a network device for a modem connection is usually not configured to start at boot time; whereas, an Ethernet connection is usually configured to activate at boot time. If your network device is configured not to start at boot time, you can use **Red Hat Control Network** to activate it after boot time. To start it, select **Main Menu Button** (on the Panel) => **System Tools** => **Network Device Control** or type the command `redhat-control-network`.

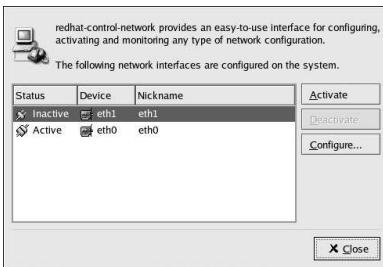


Figure 11-16. Activating Devices

To activate a device, select it from the list and click the **Activate** button. To stop the device, select it from the list and click **Deactivate**.

11.12. Working with Profiles

Multiple logical network devices can be created for each physical hardware device. For example, if you have one Ethernet card in your system (eth0), you can create logical network devices with different nicknames and different configuration options, all associated with eth0.

Logical network devices are different from device aliases. Logical network devices associated with the same physical device must exist in different profiles and can not be activated simultaneously. Device aliases are also associated with the same physical hardware device, but device aliases associated with the same physical hardware can be activated at the same time. Refer to Section 11.13 for details about creating device aliases.

Profiles can be used to create multiple configuration sets for different networks. A configuration set can include logical devices as well as hosts and DNS settings. After configuring the profiles, you can use the **Network Administration Tool** to switch back and forth between them.

By default, there is one profile called **Common**. To create a new profile, click the **New** button in the **Active Profile** frame. Enter a unique name for the profile.

After creating a new profile, if all the devices are not listed for all the profiles, add them by click the **Add** button. If a device already exists for the physical device, use the **Copy** button to copy the existing

device. If you use the **Add** button, a network alias will be created, which is not correct. The device name should not end with a colon followed by a number.

In the list of devices, there is a column of checkboxes labeled **Profile**. For each profile, you can check or uncheck devices. Only the checked devices are included for the currently selected profile.

For example, Figure 11-17 shows a profile called **Office** with the logical device **eth0_office**. It is configured to activate the first Ethernet card using DHCP.



Figure 11-17. Office Profile

Notice that the **Home** profile as shown in Figure 11-18 activates the **eth0_home** logical device, which is associated with eth0 and is configured to use a static IP address.



Figure 11-18. Home Profile

You can also configure eth0 to activate in the **Office** profile only and only activate a ppp (modem) device in the **Home** profile. Another example is to have the **Common** profile activate eth0 and an **Away** profile activate a ppp device for use while traveling.

A profile can not be activated at boot time. Only the devices in the **Common** profile, which are set to activate at boot time are activated at boot time. After the system as booted, execute the following command to enable a profile (replace `<profilename>` with the name of the profile):

```
redhat-config-network-cmd --profile <profilename>
```

11.13. Device Aliases

Device aliases are virtual devices associated with the same physical hardware, but they can be activated at the same time to have different IP addresses. They are commonly represented as the device name followed by a colon and a number (for example, `eth0:1`). They are useful if you want to have more than one IP address for a system, but the system only has one network card.

If you have configured a device such as `eth0`, click the **Add** button in **Network Administration Tool** to create an alias for the device. Select the network device and configure the network settings. The alias will appear in the device list with a device name followed by a colon and the alias number.



Warning

If you are configuring an Ethernet device to have an alias, neither the device nor the alias can be configured to use DHCP. You must configure the IP addresses manually.

Figure 11-19 shows an example of one alias for the `eth0` device. Notice the `eth0:1` device — the first alias for `eth0`. The second alias for `eth0` would have the device name `eth0:2`, and so on. To modify the settings for the device alias such as whether to activate it at boot time and the alias number, select it from the list and click the **Edit** button.



Figure 11-19. Network Device Alias Example

Select the alias and click the **Activate** button to activate the alias. If you have configured multiple profiles, select which profiles in which to include it.

To verify that the alias has been activated, use the command `/sbin/ifconfig`. The output should show the device and the device alias with different IP address:

```
eth0      Link encap:Ethernet  HWaddr 00:A0:CC:60:B7:G4
          inet addr:192.168.100.5  Bcast:192.168.100.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:161930 errors:1 dropped:0 overruns:0 frame:0
          TX packets:244570 errors:0 dropped:0 overruns:0 carrier:0
          collisions:475 txqueuelen:100
          RX bytes:55075551 (52.5 Mb)  TX bytes:178108895 (169.8 Mb)
          Interrupt:10 Base address:0x9000

eth0:1    Link encap:Ethernet  HWaddr 00:A0:CC:60:B7:G4
          inet addr:192.168.100.42  Bcast:192.168.100.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          Interrupt:10 Base address:0x9000

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:5998 errors:0 dropped:0 overruns:0 frame:0
          TX packets:5998 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:1627579 (1.5 Mb)  TX bytes:1627579 (1.5 Mb)
```


Basic Firewall Configuration

Just as a firewall in a building attempts to prevent a fire from spreading, a computer firewall attempts to prevent computer viruses from spreading to your computer and to prevent unauthorized users from accessing your computer. A firewall exists between your computer and the network. It determines which services on your computer remote users on the network can access. A properly configured firewall can greatly increase the security of your system. It is recommended that you configure a firewall for any Red Hat Linux system with an Internet connection.

12.1. Security Level Configuration Tool

During the **Firewall Configuration** screen of the Red Hat Linux installation, you were given the option to choose a high, medium, or no security level as well as allow specific devices, incoming services, and ports.

After installation, you can change the security level of your system by using **Security Level Configuration Tool**. If you prefer a wizard-based application, refer to Section 12.2.

To start the application, select **Main Menu Button** (on the Panel) => **System Settings** => **Security Level** or type the command `redhat-config-securitylevel` from a shell prompt (for example, in an XTerm or a GNOME terminal).

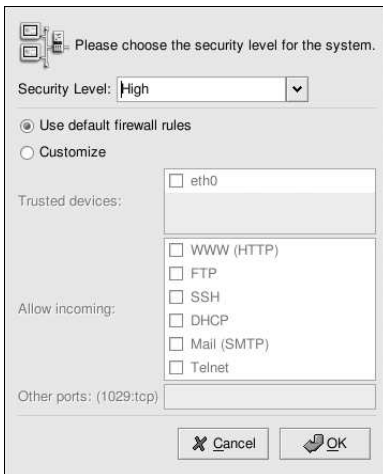


Figure 12-1. Security Level Configuration Tool

Select the desired security level from pulldown menu.

High

If you choose **High**, your system will not accept connections (other than the default settings) that are not explicitly defined by you. By default, only the following connections are allowed:

- DNS replies

- DHCP — so any network interfaces that use DHCP can be properly configured

If you choose **High**, your firewall will not allow the following:

- Active mode FTP (passive mode FTP, used by default in most clients, should still work)
- IRC DCC file transfers
- RealAudio™
- Remote X Window System clients

If you are connecting your system to the Internet, but do not plan to run a server, this is the safest choice. If additional services are needed, you can choose **Customize** to allow specific services through the firewall.



Note

If you select a medium or high firewall, network authentication methods (NIS and LDAP) will not work.

Medium

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

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

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



Note

If you select a medium or high firewall, network authentication methods (NIS and LDAP) will not work.

No Firewall

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

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

Trusted Devices

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

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

Allow Incoming

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

DHCP

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

SSH

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

Telnet

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

WWW (HTTP)

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

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

Mail (SMTP)

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

FTP

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

Other ports

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

You must have the `iptables` service enabled and running to activate the security level. Refer to Section 12.3 for details.

12.2. GNOME Lokkit

GNOME Lokkit allows you to configure firewall settings for an average user by constructing basic `iptables` networking rules. Instead of having to write the rules, this program asks you a series of questions about how you use your system and then writes it for you in the file `/etc/sysconfig/iptables`.

You should not try to use **GNOME Lokkit** to generate complex firewall rules. It is intended for average users who want to protect themselves while using a modem, cable, or DSL Internet connection. To configure specific firewall rules, refer to the *Firewalling with `iptables`* chapter in the *Official Red Hat Linux Reference Guide*.

To disable specific services and deny specific hosts and users, refer to Chapter 13.

To start **GNOME Lokkit**, type the command `gnome-lokkit` at a shell prompt as root. If you do not have the X Window System installed or if you prefer a text-based program, use the command `lokkit` to start the text-mode version of **GNOME Lokkit**.

12.2.1. Basic



Figure 12-2. Basic

After starting the program, choose the appropriate security level for your system:

- **High Security** — This option disables almost all network connects except DNS replies and DHCP so that network interfaces can be activated. IRC, ICQ, and other instant messaging services as well as RealAudio™ will not work without a proxy.
- **Low Security** — This option will not allow remote connections to the system, including NFS connections and remote X Window System sessions. Services that run below port 1023 will not accept connections, including FTP, SSH, Telnet, and HTTP.
- **Disable Firewall** — This option does not create any security rules. It is recommended that this option only be chosen if the system is on a trusted network (not on the Internet), if the system is

behind a larger firewall, or if you write your own custom firewall rules. If you choose this option and click **Next**, proceed to Section 12.3. The security of your system will not be changed.

12.2.2. Local Hosts

If there are Ethernet devices on the system, the **Local Hosts** page allows you to configure whether the firewall rules apply to connection requests sent to each device. If the device connects the system to a local area network behind a firewall and does not connect directly to the Internet, select **Yes**. If the Ethernet card connects the system to a cable or DSL modem, it is recommended that you select **No**.



Figure 12-3. Local Hosts

12.2.3. DHCP

If you are using DHCP to activate any Ethernet interfaces on the system, you must say **Yes** to the DHCP question. If you say no, you will not be able to establish a connect using the Ethernet interface. Many cable and DSL Internet providers require you to use DHCP to establish an Internet connection.

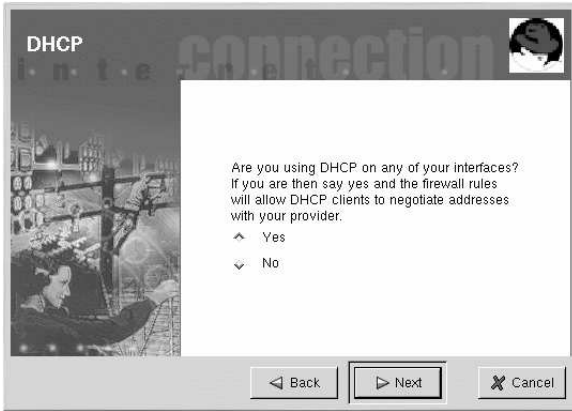


Figure 12-4. DHCP

12.2.4. Configuring Services

GNOME Lokkit also allows you to turn common services on and off. If you answer **Yes** to configuring services, you are prompted about the following services:

- **Web Server** — Choose this option if you want people to connect to a Web server such as Apache running on your system. You do not need to choose this option if you want to view pages on your own system or on other servers on the network.
- **Incoming Mail** — Choose this option if your system needs to accept incoming mail. You do not need this option if you retrieve email using IMAP, POP3, or fetchmail.
- **Secure Shell** — Secure Shell, or SSH, is a suite of tools for logging into and executing commands on a remote machine over an encrypted connection. If you need to access your machine remotely through ssh, select this option.
- **Telnet** — Telnet allows you to log into your machine remotely; however, it is not secure. It sends plain text (including passwords) over the network. It is recommended that you use SSH to log into your machine remotely. If you are required to have telnet access to your system, select this option.

To disable other services that you do not need, use **Serviceconf** (see Section 13.3) or **ntsysv** (see Section 13.4), or **chkconfig** (see Section 13.5).

12.2.5. Activating the Firewall

Clicking **Finish** will write the firewall rules to `/etc/sysconfig/iptables` and start the firewall by starting the `iptables` service.

It is highly recommended that you run **GNOME Lokkit** from the machine, not from a remote X session. If you disable remote access to your system, you will no longer be able to access it or disable the firewall rules.

Click **Cancel** if you do not want to write the firewall rules.

12.2.5.1. Mail Relay

A mail relay is a system that allows other systems to send email through it. If your system is a mail relay, someone can possibly use it to spam others from your machine.

If you chose to enable mail services, after you click **Finish** on the **Activate the Firewall** page, you will be prompted to check for mail relay. If you choose **Yes** to check for mail relay, **GNOME Lokkit** will attempt to connect to the *Mail Abuse Prevention System* website at <http://www.mail-abuse.org/> and run a mail relay test program. The results of the test will be displayed when it is finished. If your system is open to mail relay, it is highly recommended that you configure Sendmail to prevent it.

You must have the `iptables` service enabled and running to activate the firewall. Refer to Section 12.3 for details.

12.3. Activating the `iptables` Service

The firewall rules will only be active if the `iptables` service is running. To manually start the service, use the command:

```
/sbin/service iptables restart
```

To ensure that it is started when the system is booted, issue the command:

```
/sbin/chkconfig --level 345 iptables on
```

You can also use **Serviceconf** to activate `iptables`. See Section 13.3.

You can not run the `ipchains` service along with the `iptables` service. To make sure the `ipchains` service is disabled, execute the command:

```
/sbin/chkconfig --level 345 ipchains off
```


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 are running a Web server). However, if you do not need to provide a service, you should turn it off — this will minimize your exposure to possible bug exploits.

There are several different methods for managing access to system services. You must decide which method you would 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 will talk about more later in this section) and the services in the `/etc/rc.d` hierarchy can be configured to start or stop using three different applications:

- **Services Configuration Tool** — a graphical application that displays a description of each service, displays whether each service is started at boot time (for runlevels 3, 4, and 5), and allows you to start, stop, and restart each service.
- **ntsysv** — a text-based application that allows you to configure which services are started at boot time for each runlevel. Changes do not take effect immediately. Services can not be started, stopped, or restarted using this program.
- **chkconfig** — a command-line utility that allows you to turn services on and off for the different runlevels. Changes do not take effect immediately for non-`xinetd` services. Non-`xinetd` services can not be started, stopped, or restarted using this utility.

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`.

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

On the other hand, the benefit of using `iptables` is flexibility. For example, if you need a customized solution which provides certain hosts access to certain services, `iptables` can provide it for you. See the *Official Red Hat Linux Reference Guide* for more information about `iptables`.

Alternatively, if you are 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, it will then configure a simple firewall for you. Refer to Chapter 12 for more information. You can also use the **Services Configuration Tool** (`redhat-config-servicelevel`). It allows you to select the security level for your system, similar to the **Service Level** screen in the Red Hat Linux installation program.

13.1. Runlevels

Before you can configure access to services, you must understand Linux runlevels. A runlevel is a state, or *mode*, that is defined by the services listed in the directory `/etc/rc.d/rc<x>.d`, where `<x>` is the number of the runlevel.

Red Hat Linux uses the following runlevels:

- 0 — Halt
- 1 — Single-user mode
- 2 — Not used (user-definable)
- 3 — Full multi-user mode
- 4 — Not used (user-definable)
- 5 — Full multi-user mode (with an X-based login screen)
- 6 — Reboot

If you configured the X Window System during the Red Hat Linux installation program, you had the option of choosing a graphical or text login screen. If you chose a text login screen, you are operating in runlevel 3. If you chose a graphical login screen, you are operating in runlevel 5.

The default runlevel can be changed by modifying the `/etc/inittab` file, which contains a line near the top of the file similar to the following:

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

Change the number in this line to the desired runlevel. The change will not take effect until you reboot the system.

To change the runlevel immediately, use the command `telinit` followed by the runlevel number. You must be root to use this command.

13.2. TCP Wrappers

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. As the names imply, `hosts.allow` contains a list of rules that allow clients to access the network services controlled by `xinetd`, and `hosts.deny` contains rules to deny access. The `hosts.allow` file takes precedence over the `hosts.deny` file. Permissions to grant or deny access can be based on individual IP address (or hostnames) or on a pattern of clients. See the *Official Red Hat Linux Reference Guide* and the `hosts_access` man page for details.

13.2.1. xinetd

To control access to Internet services, use `xinetd`, which is 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 access only to particular hosts, to deny access to particular hosts, to 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 will notice upon inspection of the file that it only 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. You can edit any of the `xinetd` configuration files or change its enabled status using **Services Configuration Tool**, `ntsysv`, or `chkconfig`. For a list of network services controlled by `xinetd` list of the contents of the `/etc/xinetd.d` directory with the command `ls /etc/xinetd.d`.

13.3. Services Configuration Tool

Services Configuration Tool is a graphical application developed by Red Hat to configure which SysV services in `/etc/rc.d/init.d` are started at boot time (for runlevels 3, 4, and 5) and which `xinetd` services are enabled. It also allows you to start, stop, and restart SysV services as well as restart `xinetd`.

To start **Services Configuration Tool** from the desktop, go to the **Main Menu Button** (on the Panel) => **Server Settings => Services** or type the command `redhat-config-services` at a shell prompt (for example, in an **XTerm** or a **GNOME terminal**).



Figure 13-1. Services Configuration Tool

Services Configuration Tool displays the current runlevel as well as which runlevel you are currently editing. To edit a different runlevel, select **Edit Runlevel** from the pulldown menu and select runlevel 3, 4, or 5. Refer to Section 13.1 for a description of runlevels.

Services Configuration Tool lists the services from `/etc/rc.d/init.d` as well as the services controlled by `xinetd`. Click on a service to display a brief description of that service at the bottom of the window.

To start, stop, or restart a service immediately, select the service and choose the action from the **Actions** pulldown menu. You can also select the service and click the start, stop, or restart button on the toolbar.

If you select an `xinetd` service such as `telnet`, the **Start**, **Stop**, and **Restart** buttons will not be active. If you change the **Start at Boot** value of an `xinetd` service, you must click the **Save Changes** button to restart `xinetd` and disable/enable the `xinetd` services that you changed.

To enable a service at boot time for the currently selected runlevel, check the checkbox beside the name of the service under the **Start at Boot** column. After configuring the runlevel, you must apply the changes. Select **File => Save Changes** from the pulldown menu or click the **Save Changes** button.

Warning

When you save changes to `xinetd` services, `xinetd` is restarted. When you save changes to other services, the runlevel is reconfigured, but the changes do not take effect immediately.

If you check or uncheck the **Start at Boot** value for a service in `/etc/rc.d/init.d`, the **Save Changes** button will become active. Click it to reconfigure the currently selected runlevel. The changes do not affect the system immediately. For example, assume you are configuring runlevel 3. If you change the **Start at Boot** value for the `anacron` service from checked to unchecked and then click the **Save Changes** button, the runlevel 3 configuration changes so that `anacron` is not started at boot time. However, runlevel 3 is not reinitialized, so `anacron` is still running. Select one of following options at this point:

1. Stop the `anacron` service — Stop the service by selecting it from the list and clicking the **Stop the selected service** button. A message will be displayed stating that the service was stopped successfully.
2. Re-initialize the runlevel — Reinitialize the runlevel by going to a shell prompt (such as an **XTerm** or **GNOME terminal**) and typing the command `telinit 3` (where 3 is the runlevel number). This option is recommended if you change the **Start at Boot** value of more than one service and want to activate the changes immediately.
3. Do nothing else — You do not have to stop the `anacron` service. You can wait until the system is rebooted for the service to stop. The next time the system is booted, the runlevel will be initialized without the `anacron` service running.

13.4. ntsysv

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) is used to configure current runlevel. If you want to configure a different runlevel, use something like `ntsysv --levels 016`. (In this example, you would 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.

Warning

Changes do not take effect immediately after using `ntsysv`. 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`.

13.5. chkconfig

The `chkconfig` command can also be used to activate and deactivate services. If you use the `chkconfig --list` command, you will 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 will see a section for the services managed by `xinetd`).

If you use `chkconfig --list` to query a service managed by `xinetd`, you will 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 will see the service's settings for each runlevel, like the following:

```
$ 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, to turn `nscd` off in runlevels 3, 4, and 5, use the command:

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

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



Warning

Services managed by `xinetd` are immediately affected by `chkconfig`. For example, if `xinetd` is running, `finger` is disabled, and the command `chkconfig finger on` is executed, `finger` is immediately enabled without having to restart `xinetd` manually. Changes for other services do not take effect immediately after using `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.

13.6. Additional Resources

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

13.6.1. Installed Documentation

- `man ntsysv` — The `ntsysv` man page.
- `man chkconfig` — The `chkconfig` man page.
- `man xinetd` — The `xinetd` man page.
- `man xinetd.conf` — The man page for the `xinetd.conf` configuration file.
- `man 5 hosts_access` — The man page for the format of host access control files (in section 5 of the man pages).

13.6.2. Useful Websites

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

14.3. 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.

14.3.1. 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 a remote machine as well as 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' (RSA) to the list of known hosts.
```

Next, you will 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 do not specify a username 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 username@penguin.example.net
```

You can also use the syntax `ssh -l 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.

14.3.2. Using the `scp` Command

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

The general syntax to transfer a local file to a remote system is as follows:

```
scp localfile username@tohostname:/newfilename
```

The `localfile` specifies the source, and `username@tohostname:/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 as follows:

```
scp username@tohostname:/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/
```

14.3.3. 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` man page for a list of these commands. To read the man page, execute the command `man sftp` at a shell prompt. The `sftp` utility is only available in OpenSSH version 2.5.0p1 and higher.

14.3.4. Generating Key Pairs

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

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

Starting with OpenSSH version 3.0, `~/.ssh/authorized_keys2`, `~/.ssh/known_hosts2`, and `/etc/ssh_known_hosts2` are obsolete. SSH Protocol 1 and 2 share the `~/.ssh/authorized_keys`, `~/.ssh/known_hosts`, and `/etc/ssh/ssh_known_hosts` files.

Red Hat Linux 8.0 uses SSH Protocol 2 and RSA keys by default.



Tip

If you reinstall Red Hat Linux but want to save your generated key pair, backup the `.ssh` directory in your home directory. After reinstalling, copy this directory back to your home directory. This process can be done for all users on your system, including root.

14.3.4.1. Generating an RSA Key Pair for Version 2

Use the following steps to generate a RSA key pair for version 2 of the SSH protocol. This is the default starting with OpenSSH 2.9.

1. To generate a RSA key pair to work with version 2 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.

The public key is written to `~/.ssh/id_rsa.pub`. The private key is written to `~/.ssh/id_rsa`. Never distribute your private key to anyone.

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_keys` on the machine to which you want to connect. If the file `~/.ssh/authorized_keys` does not exist, you can copy the file `~/.ssh/id_rsa.pub` to the file `~/.ssh/authorized_keys` on the other machine.
4. If you are running GNOME, skip to Section 14.3.4.4. If you are not running the X Window System, skip to Section 14.3.4.5.

14.3.4.2. Generating a DSA Key Pair for Version 2

Use the following steps to generate a DSA key pair for version 2 of the SSH Protocol.

1. To generate a DSA key pair to work with version 2 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.



Tip

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.

The public key is written to `~/.ssh/id_dsa.pub`. The private key is written to `~/.ssh/id_dsa`. It is important never to give anyone the private key.

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_keys` on the machine to which you want to connect. If the file `~/.ssh/authorized_keys` does not exist, you can copy the file `~/.ssh/id_dsa.pub` to the file `~/.ssh/authorized_keys` on the other machine.
4. If you are running GNOME, skip to Section 14.3.4.4. If you are not running the X Window System, skip to Section 14.3.4.5.

14.3.4.3. 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 8.0 systems, you do not need an RSA version 1.3 or RSA version 1.5 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 -t rsa1
```

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

The public key is written to `~/.ssh/identity.pub`. The private key is written to `~/.ssh/identity`. Do not give anyone the private key.

2. Change the permissions of your `.ssh` directory and your key 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` does not exist, you can copy the file `~/.ssh/identity.pub` to the file `~/.ssh/authorized_keys` on the remote machine.
4. If you are running GNOME, skip to Section 14.3.4.4. If you are not running GNOME, skip to Section 14.3.4.5.

14.3.4.4. 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 Section 14.3.4.5.

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

1. You will 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 Linux CD-ROM set, from a Red Hat FTP mirror site, or using Red Hat Network.
2. If you do not have an `~/.Xclients` file, run `switchdesk` to create it. In your `~/.Xclients` file, find the following line:


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

 Change the line so that it instead reads:


```
exec /usr/bin/ssh-agent $HOME/.Xclients-default
```
3. Select **Main Menu Button** (on the Panel) => **Extras** => **Preferences** => **Sessions**, and click on the **Startup Programs** tab. 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 **Close** to exit the program.
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`.

14.3.4.5. 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 Section 14.3.4.4), 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. 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 more than one key pair configured, you will be prompted for each one.

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

14.4. 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.

14.4.1. Installed Documentation

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

14.4.2. 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.

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 the client's local hard drive. Red Hat Linux can be both an NFS server and an NFS client, which means that it can export file systems to other systems and mount file systems exported from other machines.

15.1. 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 directory of the NFS file system (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. Users work as if the directory is on their local machines.

15.2. Mounting NFS File Systems

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

```
mount shadowman.example.com:/misc/export /misc/local
```



Warning

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

In this command, `shadowman.example.com` is the hostname of the NFS fileserver, `/misc/export` is the directory that `shadowman` is exporting, and `/misc/local` is the location to mount the file system on the local machine. After the `mount` command runs (and if the client has proper permissions from the `shadowman.example.com` NFS server) the client user can execute the command `ls /misc/local` to display a listing of the files in `/misc/export` on `shadowman.example.com`.

15.2.1. Mounting NFS File Systems using `/etc/fstab`

An alternate way to mount an NFS share from another machine is to add a line to the `/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 the NFS share is to be mounted. 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,wsz=8192,timeo=14,intr
```

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

15.2.2. Mounting NFS File Systems 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 file systems to be mounted under this mount point. For example, the `/etc/auto.misc` file might define mount points in the `/misc` 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 `/misc/myproject` on your machine, add the following line to `auto.master`:

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

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

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

The first field in `/etc/auto.misc` is the name of the `/misc` 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.



Note

The directory `/misc` must exist on the local file system. There should be no subdirectories in `/misc` on the local file system.

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

```
/sbin/service autofs restart
```

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

```
/sbin/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:

```
/sbin/service autofs reload
```

To learn how to configure autofs to start at boot time, refer to Chapter 13 for information on managing services.

15.3. Exporting NFS File Systems

Sharing files from an NFS server is known as exporting the directories. The **NFS Server Configuration Tool** can be used to configure a system as an NFS server.

To use the **NFS Server Configuration Tool**, you must be running the X Window System. To start the application, select **Main Menu Button** (on the Panel) => **Server Settings** => **NFS Server**, or type the command `redhat-config-nfs`.




Figure 15-1. NFS Server Configuration Tool

To add an NFS share, click the **Add** button. The dialog box shown in Figure 15-2 will appear.

The **Basic** tab requires the following information:

- **Directory** — Specify the directory to share, such as `/tmp`.
- **Host(s)** — Specify the host(s) to which to share the directory. Refer to Section 15.3.2 for an explanation of possible formats.
- **Basic permissions** — Specify whether the directory should have read-only or read/write permissions.

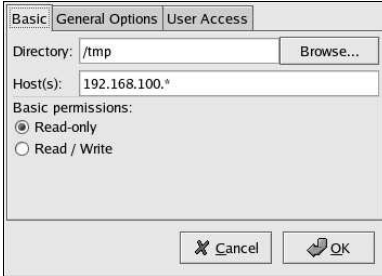


Figure 15-2. Add Share

The **General Options** tab allows the following options to be configured:

- **Allow connections from port 1024 and higher** — Services started on port numbers less than 1024 must be started as root. Select this option to allow the NFS service to be started by a user other than root. This option corresponds to `insecure`.
- **Allow insecure file locking** — Do not require a lock request. This option corresponds to `insecure_locks`.
- **Disable subtree checking** — If a subdirectory of a file system is exported, but the entire file system is not exported, the server checks to see if the requested file is in the subdirectory exported. This check is called *subtree checking*. Select this option to disable subtree checking. If the entire file system is exported, selecting to disable subtree checking can increase the transfer rate. This option corresponds to `no_subtree_check`.

- **Sync write operations on request** — Enabled by default, this option does not allow the server to reply to requests before the changes made by the request are written to the disk. This option corresponds to `sync`. If this is not selected, the `async` option is used.
- **Force sync of write operations immediately** — Do not delay writing to disk. This option corresponds to `no_wdelay`.

The **User Access** tab allows the following options to be configured:

- **Treat remote root user as local root** — By default, the user and group IDs of the root user are both 0. Root squashing maps the user ID 0 and the group ID 0 to the user and group IDs of anonymous so that root on the client does not have root privileges on the NFS server. If this option is selected, root is not mapped to anonymous, and root on a client has root privileges to exported directories. Selecting this option can greatly decrease the security of the system. Do not select it unless it is absolutely necessary. This option corresponds to `no_root_squash`.
- **Treat all client users as anonymous users** — If this option is selected, all user and group IDs are mapped to the anonymous user. This option corresponds to `all_squash`.
- **Specify local user ID for anonymous users** — If **Treat all client users as anonymous users** is selected, this option lets you specify a user ID for the anonymous user. This option corresponds to `anonuid`.
- **Specify local group ID for anonymous users** — If **Treat all client users as anonymous users** is selected, this option lets you specify a group ID for the anonymous user. This option corresponds to `anongid`.

To edit an existing NFS share, select the share from the list, and click the **Properties** button. To delete an existing NFS share, select the share from the list, and click the **Delete** button.

After adding, modifying, or deleting shares, the `nfs` service must be restarted for the changes to take effect. To apply the changes to the configuration file and restart the service, click the **Apply** button. If the `nfs` service is not already running, you will be prompted to start the service. Select **Yes** to start the daemon. The daemon must be running to export the configured directories. The old configuration file will be saved as `/etc/exports.bak`, and the new configuration will be written to `/etc/exports`.

NFS Server Configuration Tool reads and writes directly to the `/etc/exports` configuration file. Thus, the file can be modified manually after using the tool, and the tool can be used after modifying the file manually (provided the file was modified with correct syntax).

15.3.1. Command Line Configuration

If you prefer editing configuration files using a text editor or if you do have the X Window System installed, you can modify the configuration file directly.

The `/etc/exports` file controls what directories the NFS server exports. Its format is as follows:

```
directory hostname(options)
```

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

```
/misc/export    speedy.redhat.com
```

would allow users from `speedy.redhat.com` to mount `/misc/export` with the default read-only permissions, but:

```
/misc/export    speedy.redhat.com(rw)
```


would allow users from `speedy.redhat.com` to mount `/misc/export` with read/write privileges.

Refer to Section 15.3.2 for an explanation of possible hostname formats.

Refer to the *Official Red Hat Linux Reference Guide* for a list of options that can be specified.



Caution

Be careful with spaces in the `/etc/exports` file. If there are no spaces between the hostname and the options in parentheses, the options apply only to the hostname. If there is a space between the hostname and the options, the options apply to the rest of the world. For example, examine the following lines:

```
/misc/export speedy.redhat.com(rw)
/misc/export speedy.redhat.com (rw)
```

The first line grants users from `speedy.redhat.com` read-write access and denies all other users. The second line grants users from `speedy.redhat.com` read-only access (the default) and allows the rest of the world read-write access.

Each time you change `/etc/exports`, you must inform the NFS daemon of the change, or reload the configuration file with the following command:

```
/sbin/service nfs reload
```

15.3.2. Hostname Formats

The host(s) can be in the following forms:

- Single machine — A fully qualified domain name (that can be resolved by the server), hostname (that can be resolved by the server), or an IP address
- Series of machines specified with wild cards — Use the `*` or `?` character to specify a string match. For example, `192.168.100.*` specifies any IP address that begins with `192.168.100`. When specifying wild cards in fully qualified domain names, dots (`.`) are not included in the wild card. For example, `*.example.com` includes `one.example.com` but does not include `one.two.example.com`.
- IP networks — Use `a.b.c.d/z`, where `a.b.c.d` is the network and `z` is the number of bits in the netmask (for example `192.168.0.0/24`). Another acceptable format is `a.b.c.d/netmask`, where `a.b.c.d` is the network and `netmask` is the netmask (for example, `192.168.100.8/255.255.255.0`).
- Netgroups — In the format `@group-name`, where `group-name` is the NIS netgroup name.

15.3.3. Starting and Stopping the Server

On the server that is exporting NFS file systems, the `nfs` service must be running.

View the status of the NFS daemon with the following command:

```
/sbin/service nfs status
```

Start the NFS daemon with the following command:

```
/sbin/service nfs start
```

Stop the NFS daemon with the following command:

```
/sbin/service nfs stop
```

To start the `nfs` service at boot time, use the command:

```
/sbin/chkconfig --level 345 nfs on
```

You can also use `chkconfig`, `ntsysv` or the **Services Configuration Tool** to configure which services start at boot time. Refer to Chapter 13 for details.

15.4. Additional Resources

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

15.4.1. Installed Documentation

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

15.4.2. Useful Websites

- <http://www.tldp.org/HOWTO/NFS-HOWTO/index.html> — The *Linux NFS-HOWTO* from the Linux Documentation Project.

15.4.3. Related Books

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

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.

16.1. 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 15. If you want to share printers between Red Hat Linux machines only refer to Chapter 26.

16.2. 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 8.0 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.

16.2.1. Samba Passwords

In Red Hat Linux 8.0 encrypted passwords are enabled by default because it is more secure. If encrypted passwords are not used, plain text passwords are used, which can be intercepted by someone using a network packet sniffer. It is recommended that encrypted passwords be used.

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
```

If the system uses NIS, type the following command:

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

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

2. Change the permissions of the Samba password file so that only root has read and write permissions:

```
chmod 600 /etc/samba/smbpasswd
```

3. The script does not copy user passwords to the new file. To set each Samba user's password, use the command (replace `username` with each user's username):

```
smbpasswd username
```

A Samba user account will not be active until a Samba password is set for it.

4. Encrypted passwords must be enabled in the Samba configuration file. In the file `smb.conf`, verify that the following lines are not commented out:

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

5. Make sure the `smb` service is started by typing the command `service smb restart` at a shell prompt.
6. If you want the `smb` service to start automatically, use `ntsysv`, `chkconfig`, or **Services Configuration Tool** to enable it at runtime. Refer to Chapter 13 for details.



Tip

Read `/usr/share/doc/samba-<version>/docs/htmldocs/ENCRYPTION.html` to learn more about encrypted passwords. (replace `<version>` with the version number of Samba that you have installed).

The `pam_smbpass` PAM module can be used to sync users' Samba passwords with their system passwords when the `passwd` command is used. If a user invokes the `passwd` command, the password he uses to log in to the Red Hat Linux system as well as the password he must provide to connect to a Samba share are changed.

To enable this feature, add the following line to `/etc/pam.d/system-auth` below the `pam_cracklib.so` invocation:

```
password required /lib/security/pam_smbpass.so nullok use_authtok try_first_pass
```

16.3. 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.

You can also use **Nautilus** to view available Samba shares on your network. On the GNOME desktop, go to the **Main Menu Button** (on the Panel) => **Programs** => **Applications** => **Nautilus** to open a **Nautilus** window. Type **smb:** in the **Location:** bar.

As shown in Figure 16-1, you will see an icon for each available SMB workgroups on your network. To access one, double-click the icon for it.




Figure 16-1. SMB Browser in Nautilus

If the SMB share you are connecting to requires a user name and password combination, you must specify them in the **Location:** bar using the following syntax (replace *user*, *password*, *servername*, and *sharename* with the appropriate values):

```
smb://user:password@servername/sharename/
```

16.4. Additional Resources

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

16.4.1. 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

16.4.2. Useful Websites

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

Dynamic Host Configuration Protocol (DHCP)

Dynamic Host Configuration Protocol (DHCP) is network protocol for automatically assigning TCP/IP information to client machines. Each DHCP client connects to the centrally-located DHCP server which returns that client's network configuration including IP address, gateway, and DNS servers.

17.1. Why Use DHCP?

DHCP is useful for fast delivery of client network configuration. When configuring the client system, the administrator can choose DHCP and not have to enter an IP address, netmask, gateway, or DNS servers. The client retrieves this information from the DHCP server. DHCP is also useful if an administrator wants to change the IP addresses of a large number of systems. Instead of reconfiguring all the systems, he can just edit one DHCP configuration file on the server for the new set of IP address. If the DNS servers for an organization changes, the changes are made on the DHCP server, not on the DHCP clients. Once the network is restarted on the clients (or the clients are rebooted), the changes will take effect.

Furthermore, if a laptop or any type of mobile computer is configured for DHCP, it can be moved from office to office without being reconfigured as long as each office has a DHCP server that allows it to connect to the network.

17.2. Configuring a DHCP Server

You can configure a DHCP server using the configuration file `/etc/dhcpd.conf`.

DHCP also uses the file `/var/lib/dhcp/dhcpd.leases` to store the client lease database. Refer to Section 17.2.2 for more information.

17.2.1. Configuration File

The first step in configuring a DHCP server is to create the configuration file that stores the network information for the clients. Global options can be declared for all clients, or options can be declared for each client system.

The configuration file can contain any extra tabs or blank lines for easier formatting. The keywords are case-insensitive, and lines beginning with a hash mark (#) are considered comments.

Two DNS update schemes are currently implemented — the ad-hoc DNS update mode and the interim DHCP-DNS interaction draft update mode. If and when these two are accepted as part of the IETF standards process, there will be a third mode — the standard DNS update method. The DHCP server must be configured to use one of the two current schemes. Version 3.0b2p111 and previous version used the ad-hoc mode; however, it has been depreciated. If you want to keep the same behavior, add the following line to the top of the configuration file:

```
ddns-update-style ad-hoc;
```

To use the recommended mode, add the following line to the top of the configuration file:

```
ddns-update-style interim;
```

Read the `dhcpd.conf` man page for details about the different modes.

There are two types of statements in the configuration file:

- Parameters — state how to perform a task, whether to perform a task, or what network configuration options to send to the client.
- Declarations — describe the topology of the network, describe the clients, provide addresses for the clients, or apply a group of parameters to a group of declarations.

Some parameters must start with the `option` keyword and are referred to as options. Options configure DHCP options; whereas, parameters configure values that are not optional or control how the DHCP server behaves.

Parameters (including options) declared before a section enclosed in curly brackets (`{ }`) are considered global parameters. Global parameters apply to all the sections below it.



Important

If you change the configuration file, the changes will not take effect until you restart the DHCP daemon with the command `service dhcpd restart`.

In Example 17-1, the `routers`, `subnet-mask`, `domain-name`, `domain-name-servers`, and `time-offset` options are used for any `host` statements declared below it.

As shown in Example 17-1, you can declare a `subnet`. You must include a `subnet` declaration for every subnet in your network. If you do not, the DHCP server will fail to start.

In this example, there are global options for every DHCP client in the subnet and a `range` declared. Clients are assigned an IP address within the `range`.

```
subnet 192.168.1.0 netmask 255.255.255.0 {
    option routers                192.168.1.254;
    option subnet-mask            255.255.255.0;

    option domain-name           "example.com";
    option domain-name-servers   192.168.1.1;

    option time-offset            -18000;      # Eastern Standard Time

    range 192.168.1.10 192.168.1.100;
}
```

Example 17-1. Subnet Declaration

All subnets that share the same physical network should be declared within a `shared-network` declaration as shown in Example 17-2. Parameters within the `shared-network` but outside the enclosed `subnet` declarations are considered global parameters. The name of the `shared-network` should be a descriptive title for the network such as `test-lab` to describe all the subnets in a test lab environment.

```
shared-network name {
    option domain-name            "test.redhat.com";
    option domain-name-servers   ns1.redhat.com, ns2.redhat.com;
    option routers                192.168.1.254;
    more parameters for EXAMPLE shared-network
    subnet 192.168.1.0 netmask 255.255.255.0 {
        parameters for subnet
        range 192.168.1.1 192.168.1.31;
    }
}
```



```

    }
    subnet 192.168.1.32 netmask 255.255.255.0 {
        parameters for subnet
        range 192.168.1.33 192.168.1.63;
    }
}

```

Example 17-2. Shared-network Declaration

As demonstrated in Example 17-3, the `group` declaration can be used to apply global parameters to a group of declarations. You can group shared networks, subnets, hosts, or other groups.

```

group {
    option routers                192.168.1.254;
    option subnet-mask            255.255.255.0;

    option domain-name            "example.com";
    option domain-name-servers    192.168.1.1;

    option time-offset             -18000;      # Eastern Standard Time

    host apex {
        option host-name          "apex.example.com";
        hardware ethernet 00:A0:78:8E:9E:AA;
        fixed-address 192.168.1.4;
    }

    host raleigh {
        option host-name          "raleigh.example.com";
        hardware ethernet 00:A1:DD:74:C3:F2;
        fixed-address 192.168.1.6;
    }
}

```

Example 17-3. Group Declaration

To configure a DHCP server that leases a dynamic IP address to a system within a subnet, modify Example 17-4 with your values. It declares a default lease time, maximum lease time, and network configuration values for the clients. This example assigns IP addresses in the `range 192.168.1.10 and 192.168.1.100` to client systems.

```

default-lease-time 600;
max-lease-time 7200;
option subnet-mask 255.255.255.0;
option broadcast-address 192.168.1.255;
option routers 192.168.1.254;
option domain-name-servers 192.168.1.1, 192.168.1.2;
option domain-name "example.com";

subnet 192.168.1.0 netmask 255.255.255.0 {
    range 192.168.1.10 192.168.1.100;
}

```

Example 17-4. Range Parameter

To assign an IP address to a client based on the MAC address of the network interface card, use the `hardware ethernet` parameter within a `host` declaration. As demonstrated in Example

17-5, the `host apex` declaration specifies that the network interface card with the MAC address 00:A0:78:8E:9E:AA always receives the IP address 192.168.1.4.

Notice that you can also use the optional parameter `host-name` to assign a host name to the client.

```
host apex {
    option host-name "apex.example.com";
    hardware ethernet 00:A0:78:8E:9E:AA;
    fixed-address 192.168.1.4;
}
```

Example 17-5. Static IP Address using DHCP



Tip

You can use the sample configuration file in Red Hat Linux 8.0 as a starting point and then add your own custom configuration options to it. Copy it to its proper location with the command

```
cp /usr/share/doc/dhcp-<version-number>/dhcpd.conf.sample /etc/dhcpd.conf
```

(where `<version-number>` is the DHCP version you are using).

For a complete list of option statements and what they do, refer to the `dhcp-options` man page.

17.2.2. Lease Database

On the DHCP server, the file `/var/lib/dhcp/dhcpd.leases` stores the DHCP client lease database. This file should not be modified by hand. DHCP lease information for each recently assigned IP address is automatically stored in the lease database. The information includes the length of the lease, to whom the IP address has been assigned, the start and end dates for the lease, and the MAC address of the network interface card that was used to retrieve the lease.

All times in the lease database are in Greenwich Mean Time (GMT), not local time.

The lease database is recreated from time to time so that it is not too large. First, all known leases are saved in a temporary lease database. The `dhcpd.leases` file is renamed `dhcpd.leases~`, and the temporary lease database is written to `dhcpd.leases`.

The DHCP daemon could be killed or the system could crash after the lease database has been renamed to the backup file but before the new file has been written. If this happens, there is no `dhcpd.leases` file that is required to start the service. Do not create a new lease file if this occurs. If you do, all the old leases will be lost and cause many problems. The correct solution is to rename the `dhcpd.leases~` backup file to `dhcpd.leases` and then start the daemon.

17.2.3. Starting and Stopping the Server



Important

Before you start the DHCP server for the first time, it will fail unless there is an existing `dhcpd.leases` file. Use the command `touch /var/lib/dhcp/dhcpd.leases` to create the file if it does not exist.

To start the DHCP service, use the command `/sbin/service dhcpd start`. To stop the DHCP server, use the command `/sbin/service dhcpd stop`. If you want the daemon to start automatically at boot time, see Chapter 13 for information on how to manage services.

If you have more than one network interface attached to the system, but you only want the DHCP server to start on one of the interface, you can configure the DHCP server to start only on that device. In `/etc/sysconfig/dhcpd`, add the name of the interface to the list of `DHCPDARGS`:

```
# Command line options here
DHCPDARGS=eth0
```

This is useful if you have a firewall machine with two network cards. One network card can be configured as a DHCP client to retrieve an IP address to the Internet. The other network card can be used as a DHCP server for the internal network behind the firewall. Specifying only the network card connected to the internal network makes the system more secure because users can not connect to the daemon via the Internet.

Other command line options that can be specified in `/etc/sysconfig/dhcpd` include:

- `-p <portnum>` — Specify the udp port number on which `dhcpd` should listen. The default is port 67. The DHCP server transmits responses to the DHCP clients at a port number one greater than the udp port specified. For example, if you accept the default of port 67, the server listens on port 67 for requests and responses to the client on port 68. If you specify a port here and use the DHCP relay agent, you must specify the same port on which the DHCP relay agent should listen. See Section 17.2.4 for details.
- `-f` — Run the daemon as a foreground process. This is mostly used for debugging.
- `-d` — Log the DHCP server daemon to the standard error descriptor. This is mostly used for debugging. If this is not specified, the log is written to `/var/log/messages`.
- `-cf filename` — Specify the location of the configuration file. The default location is `/etc/dhcpd.conf`.
- `-lf filename` Specify the location of the lease database file. If a lease database file already exists, it is very important that the same file be used every time the DHCP server is started. It is strongly recommended that this option only be used for debugging purposes on non-production machines. The default location is `/var/lib/dhcp/dhcpd.leases`.
- `-q` — Do not print the entire copyright message when starting the daemon.

17.2.4. DHCP Relay Agent

The DHCP Relay Agent (`dhcrelay`) allows you to relay DHCP and BOOTP requests from a subnet with no DHCP server on it to one or more DHCP servers on other subnets.

When a DHCP client requests information, the DHCP Relay Agent forwards the request to the list of DHCP servers specified when the DHCP Relay Agent is started. When a DHCP server returns a reply, the reply is broadcast or unicast on the network that sent the original request.

The DHCP Relay Agent listens for DHCP requests on all interfaces unless the interfaces are specified in `/etc/sysconfig/dhcrelay` with the `INTERFACES` directive.

To start the DHCP Relay Agent, use the command `service dhcrelay start`.

17.3. Configuring a DHCP Client

The first step for configuring a DHCP client is to make sure the kernel recognizes the network interface card. Most cards are recognized during the installation process, and the system is configured to use the correct kernel module for the card. If you install a card after installation, **Kudzu**¹ should recognize it and prompt you to configure the corresponding kernel module for it. Be sure to check the Red Hat Linux Hardware Compatibility List available at <http://hardware.redhat.com/hcl/>. If the network card is not configured by the installation program or **Kudzu** and you know which kernel module to load for it, refer to Chapter 30 for details on loading kernel modules.

To configure a DHCP client manually, you need to modify the `/etc/sysconfig/network` file to enable networking and the configuration file for each network device in the `/etc/sysconfig/network-scripts` directory. In this directory, each device should have a configuration file named `ifcfg-eth0` where `eth0` is the network device name.

The `/etc/sysconfig/network` file should contain the following line:

```
NETWORKING=yes
```

You might have more information in this file, but the `NETWORKING` variable must be set to `yes` if you want networking to start at boot time.

The `/etc/sysconfig/network-scripts/ifcfg-eth0` file should contain the following lines:

```
DEVICE=eth0
BOOTPROTO=dhcp
ONBOOT=yes
```

You need a configuration file for each device that you want to configure to use DHCP.

If you prefer a graphical interface for configuring a DHCP client, refer to Chapter 11 for details on using **Network Configurator** to configure a network interface to use DHCP.

17.4. Additional Resources

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

17.4.1. Installed Documentation

- `dhcpd` man page — describes how the DHCP daemon works
- `dhcpd.conf` man page — explains how to configure the DHCP configuration file; includes some examples
- `dhcpd.leases` man page — explains how to configure the DHCP leases file; includes some examples
- `dhcp-options` man page — explains the syntax for declaring DHCP options in `dhcpd.conf`; includes some examples
- `dhcrelay` man page — explains the DHCP Relay Agent and its configuration options.

1. **Kudzu** is a hardware probing tool run at system boot time to determine what hardware has been added or removed from the system.

17.4.2. Useful Websites

- <http://www.linuxdoc.org/HOWTO/mini/DHCP/index.html> — *DHCP mini-HOWTO* from the Linux Documentation Project

Apache HTTP Server Configuration

The Apache HTTP Server has been updated to version 2.0, and the package has been renamed `httpd`. The configuration options have changed. If you want to migrate an existing configuration file by hand, refer to the migration guide at `/usr/share/doc/httpd-<ver>/migration.html` or the *Official Red Hat Linux Reference Guide* for details.

If you configured the Apache HTTP Server with the **HTTP Configuration Tool** in previous versions of Red Hat Linux and then performed an upgrade, you can use it the application to migrate the configuration file to the new format for version 2.0. Start the **HTTP Configuration Tool**, make any changes to the configuration, and save it. The configuration file saved will be compatible with version 2.0.

The **HTTP Configuration Tool** allows you to configure the `/etc/httpd/conf/httpd.conf` configuration file for the Apache HTTP Server. It does not use the old `srm.conf` or `access.conf` configuration files; leave them empty. Through the graphical interface, you can configure 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 **HTTP Configuration Tool**. If additional modules are installed, they can not be configured using this tool.

The **HTTP Configuration Tool** requires the X Window System and root access. To start the application, go to the **Main Menu Button** => **Server Settings** => **HTTP Server** or type the command `redhat-config-httpd` at a shell prompt (for example, in an XTerm or GNOME Terminal).



Caution

Do not edit the `/etc/httpd/conf/httpd.conf` configuration file by hand if you wish to use this tool. **HTTP 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 **HTTP Configuration Tool**, you cannot use this tool.

The general steps for configuring the Apache HTTP Server using the **HTTP 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 **HTTP Configuration Tool**.

18.1. Basic Settings

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




Figure 18-1. 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, the Web server 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 the server will accept incoming requests. This option corresponds to the `Listen` directive in `httpd.conf`. By default, Red Hat configures the Apache HTTP Server to listen to port 80 for non-secure Web communications. Click the **Add** button to define additional ports on which to accept requests. A window as shown in Figure 18-2 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-2.0/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 18-2. Available Addresses

**Tip**

If you set the server 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.

18.2. 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 18-3 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.

18.2.1. 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.

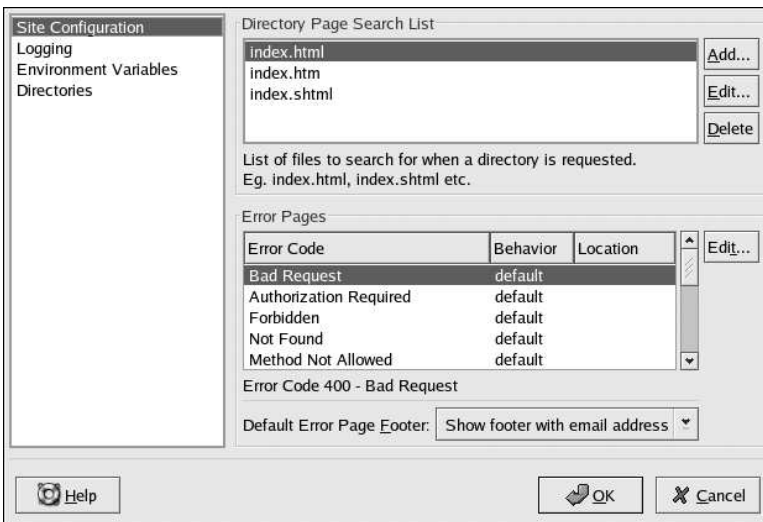


Figure 18-3. Site Configuration

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 does not 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 HTTP Server 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 HTTP 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 location 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 footer at the bottom of all error pages along with the email address of the website maintainer specified by the `ServerAdmin` directive. Refer to Section 18.3.1.1 for information about configuring the `ServerAdmin` directive.
- **Show footer** — Display just the default footer at the bottom of error pages.
- **No footer** — Do not display a footer at the bottom of error pages.

18.2.2. Logging

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

The screenshot shows the 'Logging' configuration window. On the left, a sidebar lists 'Site Configuration', 'Logging' (selected), 'Environment Variables', and 'Directories'. The main window is titled 'Transfer Log' and contains the following settings:

- Transfer Log:**
 - Log to File: logs/access_log (with a 'Browse...' button)
 - Log to Program: (empty field)
 - Use System Log: (empty field)
 - Use custom logging facilities
 - Custom Log String: (empty field)
- Error Log:**
 - Log to File: logs/error_log (with a 'Browse...' button)
 - Log to Program: (empty field)
 - Use System Log: (empty field)
- Log Level:** Error (dropdown menu)
- Reverse DNS Lookup:** Reverse Lookup (dropdown menu)

At the bottom of the window, there are three buttons: 'Help', 'OK', and 'Cancel'.

Figure 18-4. 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-2.0/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**.

18.2.3. Environment Variables

The Apache HTTP Server 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 module.

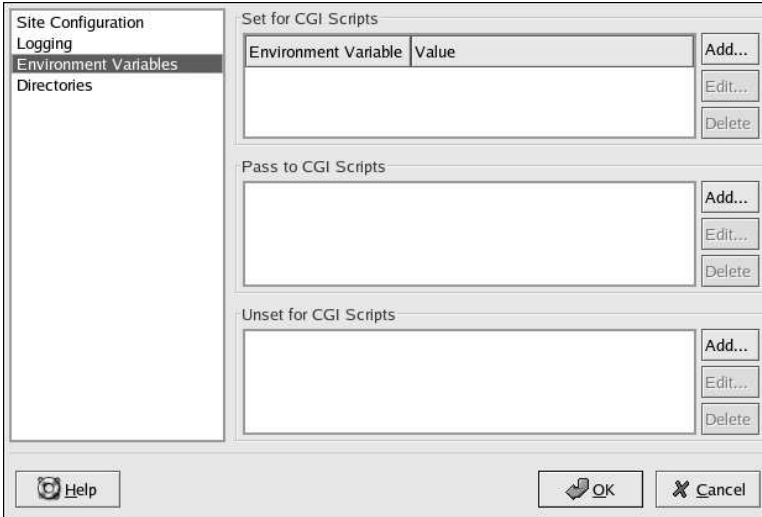


Figure 18-5. Environment Variables

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 18.2.3 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 the server 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.

18.2.4. Directories

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

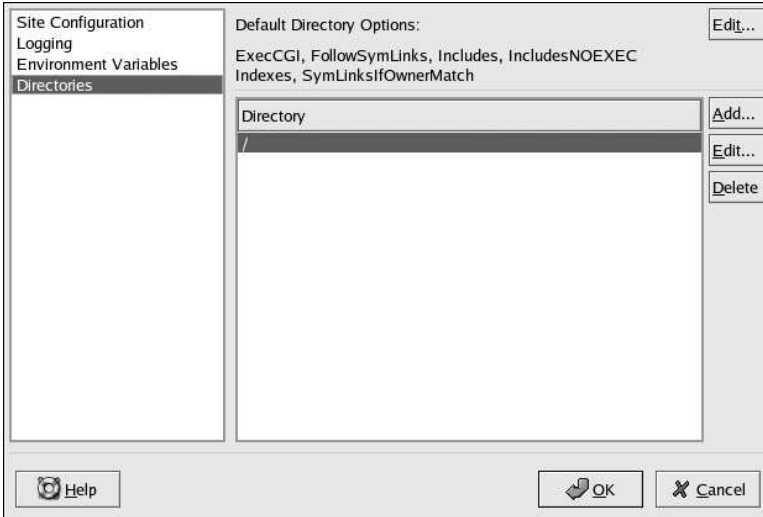


Figure 18-6. Directories

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 18-7 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 18-7. Directory Settings

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

18.3. Virtual Hosts Settings

You can use the **HTTP 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 Web 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 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.

HTTP Configuration Tool includes a default virtual host as shown in Figure 18-8.




Figure 18-8. Virtual Hosts

<http://httpd.apache.org/docs-2.0/vhosts/> and the Apache HTTP Server documentation on your machine provides more information about virtual hosts.

18.3.1. Adding and Editing a Virtual Host

To add a virtual host, click the **Virtual Hosts** tab and then click the **Add** button. You can also edit a virtual host by selecting it in the list and clicking the **Edit** button.

18.3.1.1. 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 **HTTP 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. Before Red Hat Linux 7, the Apache HTTP Server provided with Red Hat Linux used `/home/httpd/html` as the `DocumentRoot`. In Red Hat Linux 8.0, however, the default `DocumentRoot` is `/var/www/html`.

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

You should only configure one default virtual host (remember that there is one setup by default). 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.

IP based Virtual Host

If you choose **IP based Virtual Host**, a window 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.

Name based Virtual Host

If you choose **Name based Virtual Host**, a window 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.

18.3.1.2. SSL



Note

You can not 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 HTTP Server is not configured with SSL support, communications between an Apache HTTP 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 **HTTP Configuration Tool** you must allow access through port 443 under the **Main** tab => **Available Addresses**. Refer to Section 18.1 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 18-9. 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 Chapter 19.

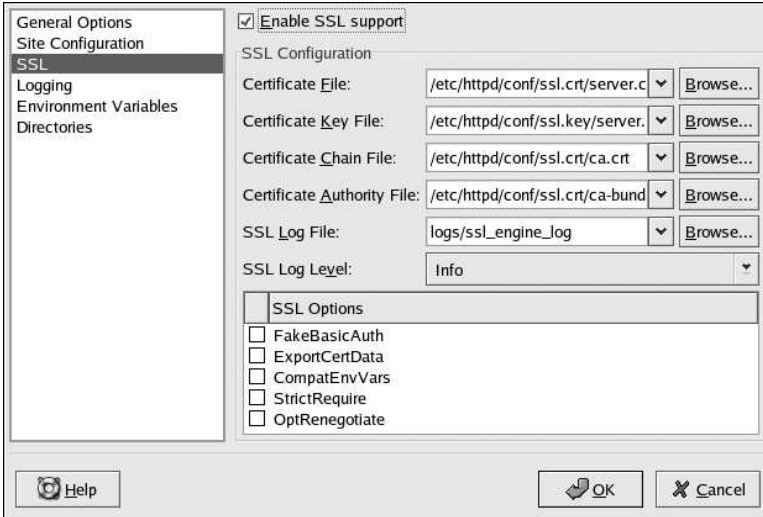


Figure 18-9. SSL Support

18.3.1.3. 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 18.2 for details on these options.

18.4. Server Settings

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




Figure 18-10. Server Configuration

The **Lock File** value corresponds to the `LockFile` directive. This directive sets the path to the lockfile used when the server 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. The Apache HTTP Server 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 are doing, do not set the `User` directive 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 not 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` directive. `Group` sets the group under which the server will answer requests. The default group is also `apache`.

18.5. Performance Tuning

Click on the **Performance Tuning** tab to configure the maximum number of child server processes you want and to configure the Apache HTTP Server 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 18-11. 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 processes 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. 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.

18.6. Saving Your Settings

If you do not want to save your Apache HTTP Server configuration settings, click the **Cancel** button in the bottom right corner of the **HTTP 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 HTTP Server configuration settings, click the **OK** button in the bottom right corner of the **HTTP Configuration Tool** window. A dialog window 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.

If this is the first time that you have used the **HTTP Configuration Tool**, you will see a dialog window warning you that the configuration file has been manually modified. If the **HTTP 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`.



Important

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

18.7. Additional Resources

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

18.7.1. Installed Documentation

- Apache HTTP Server documentation — If you have the `httpd-manual` package installed and the Apache HTTP Server daemon (`httpd`) running, you can view the Apache HTTP Server documentation. Open a Web browser, and go to the URL `http://localhost` on the server that is running the Apache HTTP Server. Then, click the **Documentation** link.
- `/usr/share/docs/httpd-<version>` — The *Apache Migration HOWTO* document contains a list of changes from version 1.3 to version 2.0 as well as information about how to migrate the configuration file manually.

18.7.2. Useful Websites

- <http://www.apache.org> — *The Apache Software Foundation*.
- <http://httpd.apache.org/docs-2.0/> — *Apache HTTP Server Version 2.0 User's Guide*.
- <http://localhost/manual/index.html> — After starting the Apache HTTP Server on your local system, you can view the *Apache HTTP Server Version 2.0 User's Guide* using this URL.
- <http://www.redhat.com/support/docs/apache.html> — Red Hat Support maintains a list of useful Apache HTTP Server links.
- <http://www.redhat.com/support/docs/faqs/RH-apache-FAQ/book1.html> — The Red Hat Linux Apache Centralized Knowledgebase compiled by Red Hat.

18.7.3. Related Books

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

Apache HTTP Secure Server Configuration

19.1. Introduction

This chapter provides basic information on the Apache HTTP Server with the `mod_ssl` security module enabled to use the OpenSSL library and toolkit. The combination of these three components, provided with Red Hat Linux, will be referred to in this chapter as the secure Web server or just as the secure server.

The `mod_ssl` module is a security module for the Apache HTTP Server. The `mod_ssl` module uses the tools provided by the OpenSSL Project to add a very important feature to the Apache HTTP Server — the ability to encrypt communications. In contrast, using regular HTTP, communications between a browser and a Web server are sent in plaintext, which could be intercepted and read by someone along the route between the browser and the server.

This chapter is not meant to be complete and exclusive documentation for any of these programs. When possible, this guide will point you to appropriate places where you can find more in-depth documentation on particular subjects.

This chapter will show you how to install these programs. You will also learn the steps necessary to generate a private key and a certificate request, how to generate your own self-signed certificate, and how to install a certificate to use with your secure Web server.

The configuration for `mod_ssl` has moved from `/etc/httpd/conf/httpd.conf` to `/etc/httpd/conf.d/ssl.conf`. For this file to be loaded, and hence for `mod_ssl` to work, you must have the statement `Include conf.d/*.conf` in `/etc/httpd/conf/httpd.conf`.

19.2. An Overview of Security-Related Packages

To enable the secure server, you need to have the following packages installed at a minimum:

`httpd`

The `httpd` package contains the `httpd` daemon and related utilities, configuration files, icons, Apache HTTP Server modules, man pages and other files used by the Apache HTTP Server.

`mod_ssl`

The `mod_ssl` package includes the `mod_ssl` module, which provides strong cryptography for the Apache HTTP Server via the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols.

`openssl`

The `openssl` package contains the OpenSSL toolkit. The OpenSSL toolkit implements the SSL and TLS protocols and also includes a general purpose cryptography library.

Additionally, other software packages included with Red Hat Linux can provide certain security functionalities (but are not required by the secure server to function):

httpd-devel

The `httpd-devel` package contains the Apache HTTP Server include files, header files and the APXS utility. You will need all of these if you intend to load any extra modules, other than the modules provided with this product. Please see the *Official Red Hat Linux Reference Guide* for more information on loading modules onto your secure Web server using Apache's DSO functionality.

If you do not intend to load other modules onto your Apache server, you do not need to install this package.

httpd-manual

The `httpd-manual` package contains the Apache Project's *Apache User's Guide* in HTML format. This manual is also available on the Web at <http://httpd.apache.org/docs-2.0/>.

OpenSSH packages

The OpenSSH packages provide the OpenSSH set of network connectivity tools for logging into and executing commands on a remote machine. OpenSSH tools encrypt all traffic (including passwords), so you can avoid eavesdropping, connection hijacking, and other attacks on the communications between your machine and the remote machine.

The `openssh` package includes core files needed by both the OpenSSH client programs and the OpenSSH server. The `openssh` package also contains `scp`, a secure replacement for `rcp` (for copying files between machines) and `ftp` (for transferring files between machines).

The `openssh-askpass` package supports the display of a dialog window which prompts for a password during use of the OpenSSH agent with RSA authentication.

The `openssh-askpass-gnome` package contains a GNOME GUI desktop environment dialog window which is displayed when OpenSSH programs prompt for a password. If you are running GNOME and using OpenSSH utilities, you should install this package.

The `openssh-server` package contains the `sshd` secure shell daemon and related files. The secure shell daemon is the server side of the OpenSSH suite, and must be installed on your host if you want to allow SSH clients to connect to your host.

The `openssh-clients` package contains the client programs needed to make encrypted connections to SSH servers, including the following: `ssh`, a secure replacement for `rsh`; and `slogin`, a secure replacement for `rlogin` (for remote login) and `telnet` (for communicating with another host via the TELNET protocol).

For more information about OpenSSH, see Chapter 14 and the OpenSSH website at <http://www.openssh.com>.

openssl-devel

The `openssl-devel` package contains the static libraries and the include file needed to compile applications with support for various cryptographic algorithms and protocols. You need to install this package only if you are developing applications which include SSL support — you do not need this package to use SSL.

stunnel

The `stunnel` package provides the Stunnel SSL wrapper. Stunnel supports the SSL encryption of TCP connections, so it can provide encryption for non-SSL aware daemons and protocols (such as POP, IMAP and LDAP) without requiring any changes to the daemon's code.

Table 19-1 displays the location of the secure server packages and additional security-related packages within the package groups provided by Red Hat Linux. This table also tells you whether each package is optional or not for the installation of a secure Web server.

Package Name	Located in Group	Optional?
httpd	System Environment/Daemons	no
mod_ssl	System Environment/Daemons	no
openssl	System Environment/Libraries	no
httpd-devel	Development/Libraries	yes
httpd-manual	Documentation	yes
openssh	Applications/Internet	yes
openssh-askpass	Applications/Internet	yes
openssh-askpass-gnome	Applications/Internet	yes
openssh-clients	Applications/Internet	yes
openssh-server	System Environment/Daemons	yes
openssl-devel	Development/Libraries	yes
stunnel	Applications/Internet	yes

Table 19-1. Security Packages

19.3. An Overview of Certificates and Security

Your secure Web server provides security using a combination of the Secure Sockets Layer (SSL) protocol and (in most cases) a digital certificate from a Certificate Authority (CA). SSL handles the encrypted communications and the mutual authentication between browsers and your secure Web server. The CA-approved digital certificate provides authentication for your secure Web server (the CA puts its reputation behind its certification of your organization's identity). When your browser is communicating using SSL encryption, you will see the `https://` prefix at the beginning of the Uniform Resource Locator (URL) in the navigation bar.

Encryption depends upon the use of keys (think of them as secret encoder/decoder rings in data format). In conventional or symmetric cryptography, both ends of the transaction have the same key, which they use to decode each other's transmissions. In public or asymmetric cryptography, two keys co-exist: a public key and a private key. A person or an organization keeps their private key a secret, and publishes their public key. Data encoded with the public key can only be decoded with the private key; data encoded with the private key can only be decoded with the public key.

To set up your secure server, you will use public cryptography to create a public and private key pair. In most cases, you will send your certificate request (including your public key), proof of your company's identity, and payment to a CA. The CA will verify the certificate request and your identity, and then send back a certificate for your secure Web server.

A secure server uses a certificate to identify itself to Web browsers. You can generate your own certificate (called a "self-signed" certificate) or you can get a certificate from a Certificate Authority or CA. A certificate from a reputable CA guarantees that a website is associated with a particular company or organization.

Alternatively, you can create your own self-signed certificate. Note, however, that self-signed certificates should not be used in most production environments. Self-signed certificates will not be automatically accepted by a user's browser — the user will be asked by the browser if they want to accept the certificate and create the secure connection. See Section 19.5 for more information on the differences between self-signed and CA-signed certificates.

Once you have a self-signed certificate or a signed certificate from the CA of your choice, you will need to install it on your secure Web server.

19.4. Using Pre-Existing Keys and Certificates

If you already have an existing key and certificate (for example, if you are installing the secure Web server to replace another company's secure Web server product), you will probably be able to use your existing key and certificate with the secure Web server. In the following two situations, you will not be able to use your existing key and certificate:

- *If you are changing your IP address or domain name* — You can not use your old key and certificate if you are changing your IP address or domain name. Certificates are issued for a particular IP address and domain name pair. You will need to get a new certificate if you are changing your IP address or domain name.
- *If you have a certificate from VeriSign and you are changing your server software* — VeriSign is a widely used CA. If you already have a VeriSign certificate for another purpose, you may have been considering using your existing VeriSign certificate with your new secure Web server. However, you will not be allowed to, because VeriSign issues certificates for one particular server software and IP address/domain name combination.

If you change either of those parameters (for example, if you previously used another secure Web server product and now you want to use the secure Web server), the VeriSign certificate you obtained to use with the previous configuration will not work with the new configuration. You will need to obtain a new certificate.

If you have an existing key and certificate that you can use, you will not have to generate a new key and obtain a new certificate. However, you may need to move and rename the files which contain your key and certificate.

Move your existing key file to:

```
/etc/httpd/conf/ssl.key/server.key
```

Move your existing certificate file to:

```
/etc/httpd/conf/ssl.crt/server.crt
```

After you have moved your key and certificate, skip to Section 19.9.

If you are upgrading from the Red Hat Secure Web Server versions 1.0 and 2.0, your old key (`httpsd.key`) and certificate (`httpsd.crt`) will be located in `/etc/httpd/conf/`. You will need to move and rename your key and certificate, so that the secure Web server can use them. Use the following two commands to move and rename your key and certificate files:

```
mv /etc/httpd/conf/httpsd.key /etc/httpd/conf/ssl.key/server.key
mv /etc/httpd/conf/httpsd.crt /etc/httpd/conf/ssl.crt/server.crt
```

Then start your secure Web server with the command:

```
/sbin/service httpd start
```

For a secure server, you will be prompted to enter your password. After you type it in and press [Enter], the server will start.

You should not need to get a new certificate, if you are upgrading from a previous version of the secure Web server.

19.5. Types of Certificates

If you installed your secure Web server using the Red Hat Linux installation program, a random key and a test certificate are generated and put into the appropriate directories. Before you begin using your secure server, however, you will need to generate your own key and obtain a certificate which correctly identifies your server.

You need a key and a certificate to operate your secure Web server — which means that you can either generate a self-signed certificate or purchase a CA-signed certificate from a CA. What are the differences between the two?

A CA-signed certificate provides two important capabilities for your server:

- Browsers will (usually) automatically recognize the certificate and allow a secure connection to be made, without prompting the user.
- When a CA issues a signed certificate, they are guaranteeing the identity of the organization that is providing the Web pages to the browser.

If your secure server is being accessed by the public at large, your secure Web server needs a certificate signed by a CA, so that people who visit your website know that the website is owned by the organization who claims to own it. Before signing a certificate, a CA verifies that the organization requesting the certificate was actually who they claimed to be.

Most Web browsers that support SSL have a list of CAs whose certificates they will automatically accept. If a browser encounters a certificate whose authorizing CA is not in the list, the browser will ask the user to choose whether to accept or decline the connection.

You can generate a self-signed certificate for your secure Web server, but be aware that a self-signed certificate will not provide the same functionality as a CA-signed certificate. A self-signed certificate will not be automatically recognized by users' browsers, and a self-signed certificate does not provide any guarantee concerning the identity of the organization that is providing the website. A CA-signed certificate provides both of these important capabilities for a secure server. If your secure server will be used in a production environment, you will probably need a CA-signed certificate.

The process of getting a certificate from a CA is fairly easy. A quick overview is as follows:

1. Create an encryption private and public key pair.
2. Create a certificate request based on the public key. The certificate request contains information about your server and the company hosting it.
3. Send the certificate request, along with documents proving your identity, to a CA. We cannot tell you which certificate authority to choose. Your decision may be based on your past experiences, or on the experiences of your friends or colleagues, or purely on monetary factors.

To see a list of CAs, click on the **Security** button on your Web browser toolbar or on the padlock icon at the bottom left of the screen, then click on **Signers** to see a list of certificate signers from whom your browser will accept certificates. You can also search the Web for CAs. Once you have decided upon a CA, you will need to follow the instructions they provide on how to obtain a certificate from them.

4. When the CA is satisfied that you are indeed who you claim to be, they will send you a digital certificate.
5. Install this certificate on your Web server, and begin handling secure transactions.

Whether you are getting a certificate from a CA or generating your own self-signed certificate, the first step is to generate a key. See Section 19.6 for instructions on how to generate a key.

19.6. Generating a Key

You must be root to generate a key.

First, `cd` to the `/etc/httpd/conf` directory. Remove the fake key and certificate that were generated during the installation with the following commands:

```
rm ssl.key/server.key
rm ssl.crt/server.crt
```

Next, you need to create your own random key. Change to the `/usr/share/ssl/certs` directory, and type in the following command:

```
make genkey
```

Your system will display a message similar to the following:

```
umask 77 ; \
/usr/bin/openssl genrsa -des3 1024 > /etc/httpd/conf/ssl.key/server.key
Generating RSA private key, 1024 bit long modulus
.....+++++
.....+++++
e is 65537 (0x10001)
Enter PEM pass phrase:
```

You now need to type in a password. For best security, your password should contain at least eight characters, include numbers and/or punctuation, and not be a word in a dictionary. Also, remember that your password is case sensitive.



Note

You will need to remember and enter this password every time you start your secure Web server, so do not forget it.

You will be asked to re-type the password, to verify that it is correct. Once you have typed it in correctly, `/etc/httpd/conf/ssl.key/server.key`, containing your key, will be created.

Note that if you do not want to type in a password every time you start your secure Web server, you will need to use the following two commands instead of `make genkey` to create the key.

Use the following command:

```
/usr/bin/openssl genrsa 1024 > /etc/httpd/conf/ssl.key/server.key
```

to create your key. Then use this command:

```
chmod go-rwx /etc/httpd/conf/ssl.key/server.key
```

to make sure that the permissions are set correctly on your key.

After you use the above commands to create your key, you will not need to use a password to start your secure Web server.

**Caution**

Disabling the password feature for your secure Web server is a security risk. It is NOT recommend that you disable the password feature for your secure Web server.

The problems associated with not using a password are directly related to the security maintained on the host machine. For example, if an unscrupulous individual compromises the regular UNIX security on the host machine, that person could obtain your private key (the contents of your `server.key` file). The key could be used to serve Web pages that will appear to be from your Web server.

If UNIX security practices are rigorously maintained on the host computer (all operating system patches and updates are installed as soon as they are available, no unnecessary or risky services are operating, and so on), the secure Web server's password may seem unnecessary. However, since your secure Web server should not need to be re-booted very often, the extra security provided by entering a password is a worthwhile effort in most cases.

The `server.key` file should be owned by the root user on your system and should not be accessible to any other user. Make a backup copy of this file and keep the backup copy in a safe, secure place. You need the backup copy because if you ever lose the `server.key` file after using it to create your certificate request, your certificate will no longer work and the CA will not be able to help you. Your only option would be to request (and pay for) a new certificate.

If you are going to purchase a certificate from a CA, continue to Section 19.7. If you are generating your own self-signed certificate, continue to Section 19.8.

19.7. Generating a Certificate Request to Send to a CA

Once you have created a key, the next step is to generate a certificate request which you will need to send to the CA of your choice. Make sure you are in the `/usr/share/ssl/certs` directory, and type in the following command:

```
make certreq
```

Your system will display the following output and will ask you for your password (unless you disabled the password option):

```
umask 77 ; \  
/usr/bin/openssl req -new -key /etc/httpd/conf/ssl.key/server.key  
-out /etc/httpd/conf/ssl.csr/server.csr  
Using configuration from /usr/share/ssl/openssl.cnf  
Enter PEM pass phrase:
```

Type in the password that you chose when you were generating your key. Your system will display some instructions and then ask for a series of responses from you. Your inputs will be incorporated into the certificate request. The display, with example responses, will look like this:

```
You are about to be asked to enter information that will be incorporated  
into your certificate request.  
What you are about to enter is what is called a Distinguished Name or a  
DN.  
There are quite a few fields but you can leave some blank  
For some fields there will be a default value,  
If you enter '.', the field will be left blank.  
-----  
Country Name (2 letter code) [GB]:US
```

```

State or Province Name (full name) [Berkshire]:North Carolina
Locality Name (eg, city) [Newbury]:Raleigh
Organization Name (eg, company) [My Company Ltd]:Test Company
Organizational Unit Name (eg, section) []:Testing
Common Name (your name or server's hostname) []:test.example.com
Email Address []:admin@example.com
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:

```

The default answers appear in brackets [] immediately after each request for input. For example, the first information required is the name of the country where the certificate will be used, shown like the following:

```
Country Name (2 letter code) [GB]:
```

The default input, in brackets, is GB. To accept the default, press [Enter], or fill in your country's two letter code.

You will have to type in the rest of the inputs (State or Province Name, Locality Name, Organization Name, Organizational Unit Name, Common Name, and Email address). All of these should be self-explanatory, but you need to follow these guidelines:

- Do not abbreviate the locality or state. Write them out (for example, St. Louis should be written out as Saint Louis).
- If you are sending this CSR to a CA, be very careful to provide correct information for all of the fields, but especially for the Organization Name and the Common Name. CAs check the information provided in the CSR to determine whether your organization is responsible for what you provided as the Common Name. CAs will reject CSRs which include information they perceive as invalid.
- For Common Name, make sure you type in the *real* name of your secure Web server (a valid DNS name) and not any aliases which the server may have.
- The Email Address should be the email address for the webmaster or system administrator.
- Avoid any special characters like @, #, &, !, etc. Some CAs will reject a certificate request which contains a special character. So, if your company name includes an ampersand (&), spell it out as "and" instead of "&."
- Do not use either of the extra attributes (A challenge password and An optional company name). To continue without entering these fields, just press [Enter] to accept the blank default for both inputs.

When you have finished entering your information, the file `/etc/httpd/conf/ssl.csr/server.csr` will be created. This file is your certificate request, ready to send to your CA.

After you have decided on a CA, follow the instructions they provide on their website. Their instructions will tell you how to send your certificate request, any other documentation that they require, and your payment to them.

After you have fulfilled the CA's requirements, they will send a certificate to you (usually by email). Save (or cut and paste) the certificate that they send you as `/etc/httpd/conf/ssl.crt/server.crt`.

19.8. Creating a Self-Signed Certificate

You can create your own self-signed certificate. Please note that a self-signed certificate will not provide the security guarantees provided by a CA-signed certificate. See Section 19.5 for more details about certificates.

If you would like to make your own self-signed certificate, you will first need to create a random key using the instructions provided in Section 19.6. Once you have a key, make sure you are in the `/usr/share/ssl/certs` directory, and type the following command:

```
make testcert
```

You will see the following output and you will be prompted for your password (unless you generated a key without a password):

```
umask 77 ; \  
/usr/bin/openssl req -new -key /etc/httpd/conf/ssl.key/server.key  
-x509 -days 365 -out /etc/httpd/conf/ssl.crt/server.crt  
Using configuration from /usr/share/ssl/openssl.cnf  
Enter PEM pass phrase:
```

After you enter your password (or without a prompt if you created a key without a password), you will be asked for more information. The computer's output and a set of inputs looks like the following (you will need to provide the correct information for your organization and host):

```
You are about to be asked to enter information that will be incorporated  
into your certificate request.  
What you are about to enter is what is called a Distinguished Name or a  
DN.  
There are quite a few fields but you can leave some blank  
For some fields there will be a default value,  
If you enter '.', the field will be left blank.  
-----  
Country Name (2 letter code) [GB]:US  
State or Province Name (full name) [Berkshire]:North Carolina  
Locality Name (eg, city) [Newbury]:Raleigh  
Organization Name (eg, company) [My Company Ltd]:My Company, Inc.  
Organizational Unit Name (eg, section) []:Documentation  
Common Name (your name or server's hostname) []:myhost.example.com  
Email Address []:myemail@example.com
```

After you provide the correct information, a self-signed certificate will be created in `/etc/httpd/conf/ssl.crt/server.crt`. You will need to restart your secure server after generating the certificate with following the command:

```
/sbin/service httpd restart
```

19.9. Testing Your Certificate

When the secure server is installed by the Red Hat Linux installation program, a random key and a generic certificate are installed, for testing purposes. You can connect to your secure server using this certificate. For any purposes other than testing, however, you need to get a certificate from a CA or generate a self-signed certificate. See Section 19.5 if you need more information on the different types of certificates available.

If you have purchased a certificate from a CA or generated a self-signed certificate, you should have a file named `/etc/httpd/conf/ssl.key/server.key`, containing your key, and a file named

/etc/httpd/conf/ssl.crt/server.crt, containing your certificate. If your key and certificate are somewhere else, move them to these directories. If you changed any of the default locations or filenames for the secure Web server in your Apache HTTP Server configuration files, you should put these two files in the appropriate directory, based on your modifications.

Once these CA files have been moved, restart your server with the command:

```
/sbin/service httpd restart
```

If your key file is encrypted, you will be asked for the password. Type in your password to start your server.

Point your Web browser to your server's home page. The URL to access your secure Web server will look like this:

```
https://your_domain
```

**Note**

Note the "s" after "http." The https: prefix is used for secure HTTP transactions.

If you are using a CA-signed certificate from a well-known CA, your browser will probably automatically accept the certificate (without prompting you for input) and create the secure connection. Your browser will not automatically recognize a test or a self-signed certificate, because the certificate is not signed by a CA. If you are not using a certificate from a CA, follow the instructions provided by your browser to accept the certificate. You can just accept the defaults by clicking **Next** until the dialogs are finished.

Once your browser accepts the certificate, your secure Web server will show you a default home page.

19.10. Accessing Your Secure Server

To access your secure server, use a URL like this:

```
https://your_domain
```

Note that URLs which are intended to connect to your secure Web server should begin with the https: protocol designator instead of the more common http: protocol designator.

Your non-secure server can be accessed using an URL like this:

```
http://your_domain
```

The standard port for secure Web communications is port 443. The standard port for non-secure Web communications is port 80. The secure Web server default configuration listens on both of the two standard ports. Therefore, you will not need to specify the port number in a URL (the port number is assumed).

However, if you configure your server to listen on a non-standard port (i.e., anything besides 80 or 443), you will need to specify the port number in every URL which is intended to connect to the server on the non-standard port.

For example, you may have configured your server so that you have a virtual host running non-secured on port 12331. Any URLs intended to connect to that virtual host must specify the port number in the URL. The following URL example will attempt to connect to a non-secure Web server listening on port 12331:

`http://your_domain:12331`

**Note**

Some of the example URLs used in this manual may need to be changed, depending upon whether you are accessing your secure Web server or your non-secure Web server. Please view all URLs in this manual as general examples and not as explicit instructions that will work under all circumstances.

19.11. Additional Resources

Refer to Section 18.7 for additional references about the Apache HTTP Server.

19.11.1. Installed Documentation

- `mod_ssl` documentation — Open a Web browser, and go to the URL `http://localhost/manual/mod/mod_ssl/` on the server that is running the Apache HTTP Server.

19.11.2. Useful Websites

- Mailing list — You can subscribe to the `redhat-secure-server` mailing list at `http://www.redhat.com/mailling-lists`.
You can also subscribe to the `redhat-secure-server` mailing list by emailing `<redhat-secure-server-request@redhat.com>` and include the word *subscribe* in the subject line.
- `http://www.modssl.org` — The `mod_ssl` website is the definitive source for information about `mod_ssl`. The website includes a wealth of documentation, including a *User Manual* at `http://www.modssl.org/docs`.

19.11.3. Related Books

- *Apache: The Definitive Guide*, 2nd edition, by Ben Laurie and Peter Laurie, O'Reilly & Associates, Inc.

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 the **Bind Configuration Tool** (`redhat-config-bind`) to configure basic BIND server zones. The **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.



Important

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**, add them to `/etc/named.custom`.

The **Bind Configuration Tool** requires the X Window System and root access. To start the **Bind Configuration Tool**, go to the **Main Menu Button** (on the Panel) => **Server Settings** => **Domain Name Service** or type the command `redhat-config-bind` at a shell prompt (for example, in an XTerm or GNOME-terminal).




Figure 20-1. Bind Configuration Tool

The **Bind Configuration Tool** configures the default zone directory to be `/var/named`. All zone files specified are relative to this directory. The **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.

The **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 20-1.

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 cause the `named` service reload the configuration files. You can also choose **File** => **Quit** and click **Save and quit**.

20.1. 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 20-2 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`. It is preset to the domain name with `.zone` appended to it.
- **Contact** — Email address of the main contact for the master zone.
- **Primary Nameserver (SOA)** — State of authority (SOA) record. This specifies the nameserver that is the best resource of information for this domain.
- **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 nameservers for the zone will retrieve the latest data. The **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 20-2. Adding a Forward Master Zone

The configuration shown in Figure 20-2 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      ns.example.com.  root.localhost (
                        2 ; serial
                        28800 ; refresh
                        7200 ; retry
                        604800 ; expire
                        86400 ; ttl
                        )

ns     IN      NS       1.2.3.4

```

After configuring the Forward Master Zone, click **OK** to return to the main window as shown in Figure 20-1. 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.

20.2. 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 20-3, 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. **Contact** — Email address of the main contact for the master zone.
4. **File Name** — File name of DNS database file in the `/var/named` directory.
5. **Primary Nameserver (SOA)** — State of authority (SOA) record. This specifies the nameserver that is the best resource of information for this domain.
6. **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 nameservers for the zone will retrieve the latest data. The **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.
7. **Time Settings** — The **Refresh**, **Retry**, **Expire**, and **Minimum TTL** (Time to Live) values that are stored in the DNS database file.
8. **Nameservers** — Add, edit, and delete name servers for the reverse master zone. At least one nameserver is required.
9. **Reverse Address Table** — List of IP addresses within the reverse master zone and their hostnames. For example, for the reverse master zone `192.168.10`, you can add `192.168.10.1` in the **Reverse Address Table** with the hostname `one.example.com`. The hostname must end with a period (`.`) to specify that it is a full hostname.

Reverse Master Zone

IP Address: 192.168.10

Reverse IP Address: 10.168.192.in-addr.arpa

Contact: root@localhost

File Name: 10.168.192.in-addr.arpa.zone

Primary Nameserver (SOA):

Serial Number: 1

Nameservers

Reverse Address Table

Address	Host or Domain

Figure 20-3. Adding a Reverse Master Zone

The configuration shown in Figure 20-3 creates the following entry in `/etc/named.conf`:

```
zone "10.168.192.in-addr.arpa" {
    type master;
    file "10.168.192.in-addr.arpa.zone";
};
```

It also creates the file `/var/named/10.168.192.in-addr.arpa.zone` with the following information:

```
$TTL 86400
@      IN      SOA      ns.example.com. root.localhost (
        2 ; serial
        28800 ; refresh
        7200 ; retry
        604800 ; expire
        86400 ; ttk
        )

@      IN      NS       ns2.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 20-1. 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.

20.3. 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 20-4, with the following options:

- **Name** — The domain name that was entered in the previous window.
- **Masters List** — The nameserver 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`.

The screenshot shows a dialog box with a light gray background and a thin border. It contains three text input fields stacked vertically. The first field is labeled 'Name:' and contains the text 'slave.example.com'. The second field is labeled 'Masters List:' and contains '1.2.3.4'. The third field is labeled 'File Name:' and contains 'slave.example.com.zone'. Below the fields are two buttons: 'Cancel' with a red 'X' icon and 'OK' with a hand icon pointing to the button.

Figure 20-4. Adding a Slave Zone

The configuration shown in Figure 20-4 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 20-1. From the pulldown menu, choose **File** => **Apply** to write the `/etc/named.conf` configuration file and have the daemon reload the configuration files.

Mail Transport Agent (MTA) Configuration

A *Mail Transport Agent* (MTA) is essential for sending email from a Red Hat Linux system. The *Mail User Agent* (MUA) such as **Mozilla Mail**, **Mutt**, **Pine**, and **Evolution** is used to read and compose email. When a user sends an email from an MUA, the messages are handed off to the MTA, which sends the message to a series of MTAs until it reaches its destination.

Even if a user does not plan to send email from the system, some automated tasks or system programs might use the `/bin/mail` command to send email containing log messages to the root user of the local system.

Red Hat Linux 8.0 provides two MTAs: Sendmail and Postfix. If both are installed, `sendmail` is the default MTA. **Mail Transport Agent Switcher** allows a user to select either `sendmail` or `postfix` as the default MTA for the system.

To start the **Mail Transport Agent Switcher**, select **Main Menu Button** (on the Panel) => **Extras** => **System Settings** => **Mail Transport Agent Switcher**, or type the command `redhat-switchmail` at a shell prompt (for example, in an XTerm or GNOME terminal).

The program automatically detect if the X Window System is running. If it is running, the program starts in graphical mode as shown in Figure 21-1. If X is not detected, it starts in text-mode. To force **Mail Transport Agent Switcher** to run in text-mode, use the command `redhat-switchmail-nox`.



Figure 21-1. Mail Transport Agent Switcher

If you selected **Postfix**, you must make sure the `sendmail` service is stopped and the `postfix` service is started:

```
/sbin/service sendmail stop
/sbin/service postfix start
```

If you selected **Sendmail**, you must make sure the `postfix` service is stopped and the `sendmail` service is started:

```
/sbin/service postfix stop
/sbin/service sendmail start
```

To enable or disable the services at boot time, you must configure the runlevel with **Services Configuration Tool**, `ntsysv`, or `chkconfig`. Refer to Chapter 13 for details.

For more information about email protocols and MTAs, refer to the *Official Red Hat Linux Reference Guide*.

System Configuration

Console Access

When normal (non-root) users log into a computer locally, they are given two types of special permissions:

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

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

In contrast, *every* user who logs in at the console will be allowed to run programs that accomplish tasks normally restricted to the root user. If X is running, these actions can be included as menu items in a graphical user interface. As shipped, the console-accessible programs include `halt`, `poweroff`, and `reboot`.

22.1. Disabling Shutdown Via Ctrl-Alt-Del

By default, `/etc/inittab` specifies that your system is set to shutdown and reboot the system in response to a `[Ctrl]-[Alt]-[Del]` key combination used at the console. If you would like to completely disable this ability, you will need to comment out the following line in `/etc/inittab` by putting a hash mark (`#`) in front of it:

```
ca::ctrlaltdel:/sbin/shutdown -t3 -r now
```

Alternatively, you may just want to allow certain non-root users the right to shutdown the system from the console using `[Ctrl]-[Alt]-[Del]`. You can restrict this privilege to certain users, by taking the following steps:

1. Add a `-a` option to the `/etc/inittab` line shown above, so that it reads:

```
ca::ctrlaltdel:/sbin/shutdown -a -t3 -r now
```

The `-a` flag tells `shutdown` to look for the `/etc/shutdown.allow` file, which you will create in the next step.

2. Create a file named `shutdown.allow` in `/etc`. The `shutdown.allow` file should list the usernames of any users who are allowed to shutdown the system using `[Ctrl]-[Alt]-[Del]`. The format of the `/etc/shutdown.allow` file is a list of usernames, one per line, like the following:

```
stephen
jack
sophie
```

According to this example `shutdown.allow` file, `stephen`, `jack`, and `sophie` are allowed to shutdown the system from the console using `[Ctrl]-[Alt]-[Del]`. When that key combination is used, the `shutdown -a` in `/etc/inittab` checks to see if any of the users in `/etc/shutdown.allow` (or root) are logged in on a virtual console. If one of them is, the shutdown of the system will continue; if not, an error message will be written to the system console instead.

For more information on `shutdown.allow` see the `shutdown` man page.

22.2. Disabling Console Program Access

In order to disable access by users to console programs, you should run this command as root:

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

In environments where the console is otherwise secured (BIOS and boot loader passwords are set, [Ctrl]-[Alt]-[Delete] is disabled, the power and reset switches are disabled, and so forth), you may not want to allow any user at the console to run `poweroff`, `halt`, and `reboot`, which are accessible from the console by default.

To remove these abilities, run the following commands as root:

```
rm -f /etc/security/console.apps/poweroff
rm -f /etc/security/console.apps/halt
rm -f /etc/security/console.apps/reboot
```

22.3. Disabling All Console Access

The PAM `pam_console.so` module manages console file permissions and authentication. (See the *Official Red Hat Linux Reference Guide* for more information on configuring PAM.) If you want to disable all console access, including program and file access, comment out all lines that refer to `pam_console.so` in the `/etc/pam.d` directory. As root, the following script will do the trick:

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

22.4. Defining the Console

The `pam_console.so` module uses the `/etc/security/console.perms` file to determine the permissions for users at the system console. The syntax of the file is very flexible; you can edit the file so that these instructions no longer apply. However, the default file has a line that looks like this:

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

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

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

22.5. Making Files Accessible From the Console

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

```
<floppy>=/dev/fd[0-1]* \
/dev/floppy/*
<sound>=/dev/dsp* /dev/audio* /dev/midi* \
/dev/mixer* /dev/sequencer \
```

```

/dev/sound/*
<cdrom>=/dev/cdrom* /dev/cdwriter*

```

You can add your own lines to this section, if necessary. Make sure that any lines you add refer to the appropriate device. For example, you could add the following line:

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

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

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

```

<console> 0660 <floppy> 0660 root.floppy
<console> 0600 <sound> 0640 root
<console> 0600 <cdrom> 0600 root.disk

```

and add a line like:

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

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

22.6. Enabling Console Access for Other Applications

If you wish to make other applications accessible to console users, you will have to do a bit more work.

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

1. Create a link from the name of your application, such as our sample `foo` program, to the `/usr/bin/consolehelper` application:

```

cd /usr/bin
ln -s consolehelper foo

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

```

touch /etc/security/console.apps/foo

```
3. Create a PAM configuration file for the `foo` service in `/etc/pam.d/`. An easy way to do this is to start with a copy of the `halt` service's PAM configuration file, and then modify the file if you want to change the behavior:

```

cp /etc/pam.d/halt /etc/pam.d/foo

```


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

In the PAM configuration file, an application can be configured to use the `pam_timestamp` module to remember (cache) a successful authentication attempt. When an application is started and proper authentication is provided (the root password), a timestamp file is created. By default, a successful authentication is cached for five minutes. During this time, any other application that is configured to use `pam_timestamp` and run from the same session is automatically authenticated for the user — the user does not have to enter the root password again.

This module is included in the `pam` package. To enable this feature, the PAM configuration file in `etc/pam.d/` must include the following lines:

```
auth sufficient /lib/security/pam_timestamp.so
session optional /lib/security/pam_timestamp.so
```

The first line that begins with `auth` should be after any other `auth sufficient` lines, and the line that begins with `session` should be after any other `session optional` lines.

If an application configured to use `pam_timestamp` is successfully authenticated from the **Main Menu Button** (on the Panel), the  icon is displayed in the notification area of the panel if you are running the GNOME desktop environment. After the authentication expires (the default is five minutes), the icon disappears.

The user can select to forget the cached authentication by clicking on the icon and selecting the option to forget authentication.

22.7. The `floppy` Group

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

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

Now, user `fred` will now be able to access the system's diskette drive from the console.

Time and Date Configuration

Time and Date Properties Tool allows the user to change the system date and time, to configure the time zone used by the system, and to setup the Network Time Protocol (NTP) daemon to synchronize the system clock with a time server.

You must be running the X Window System and have root privileges. To start the application from the desktop go to the **Main Menu Button** (on the Panel) => **System Settings** => **Date & Time** or type the command `redhat-config-date` at a shell prompt (for example, in an XTerm or a GNOME terminal).

23.1. Time and Date Properties

As shown in Figure 23-1, the first tabbed window that appears is for configuring the system date and time and the NTP daemon (`ntpd`).

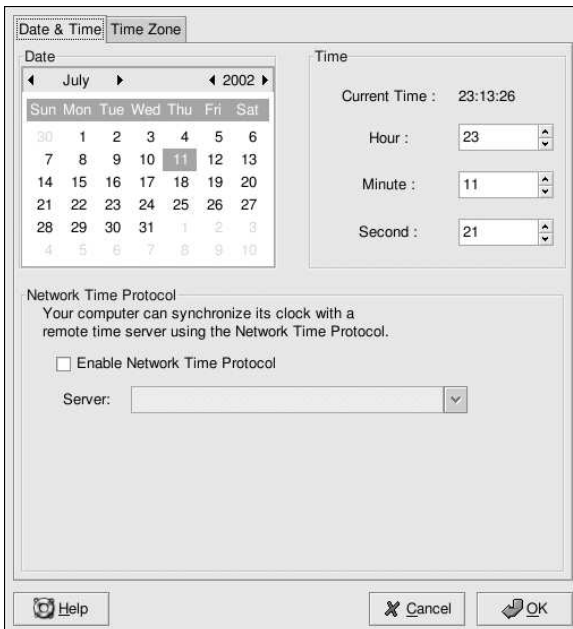


Figure 23-1. Time and Date Properties

To change the date, use the arrows to the left and right of the month to change the month. Use the arrows to the left and right of the year to change the year, and click on the day of the week to change the day of the week. Changes will not take place until you click the **Apply** button.

To change the time, use the up and down arrow buttons beside the **Hour**, **Minute**, and **Second** in the **Time** section. Changes will not take place until you click the **Apply** button.

**Note**

Changing the date and time will change the system clock as well as the hardware clock. Clicking **Apply** or **OK** is equivalent to executing the `date` and `hwclock` commands with the selected date and time.

The Network Time Protocol (NTP) daemon synchronizes the system clock with a remote time server or time source (such as a satellite). The application allows you to configure a NTP daemon to synchronize your system clock with a remote server. To enable this feature, click the **Enable Network Time Protocol** button. This will enable the **Server** pulldown menu. You can choose one of the predefined servers or type a server name in the pulldown menu. Your system will not start synchronizing with the NTP server until you click **Apply**. After you click **Apply**, the configuration will be saved and the NTP daemon will be started (or restarted if it is already running). If you want this daemon to start automatically at boot time, you need to execute the command `/sbin/chkconfig --level 345 ntpd on` to enable `ntpd` for runlevels 3, 4, and 5.

For more information on NTP, read the NTP documentation available in the `/usr/share/doc/ntp-<version>` directory.

The NTP server is written to the `/etc/ntp.conf` and `/etc/ntp/step-tickers` files.

Clicking the **Apply** button will apply any changes that you have made to the date and time, the NTP daemon settings, and the time zone settings. Clicking the **OK** button will apply the changes and then exit the program.

**Warning**

If you configured a medium or high security level during installation or with **Security Level Configuration Tool**, the firewall rules will block the connection to the NTP port. To allow NTP to work, run the **Security Level Configuration Tool**, select **Customize**, and add `udp:ntp` to the other ports.

23.2. Time Zone Configuration

To configure the system time zone, click the **Time Zone** tab. The time zone can be changed by either using the interactive map or by choosing the desired time zone from the list below the map. To use the map, click on the city that represents the desired time zone. A red **X** will appear and the time zone selection will change in the list below the map. Click **Apply** to save the changes. Click **OK** to apply the changes and exit the program.

If your system clock is set to use UTC, select the **System clock uses UTC** option. UTC stands for the universal time zone, also known as Greenwich mean time (GMT). Other time zones are determined by adding or subtracting from the UTC time.

User and Group Configuration

User Manager allows you to view, modify, add, and delete local users and groups.

To use **User Manager**, you must be running the X Window System, have root privileges, and have the `redhat-config-users` RPM package installed. To start **User Manager** from the desktop, go to the **Main Menu Button** (on the Panel) => **System Settings** => **Users & Groups** or type the command `redhat-config-users` at a shell prompt (for example, in an XTerm or a GNOME terminal).

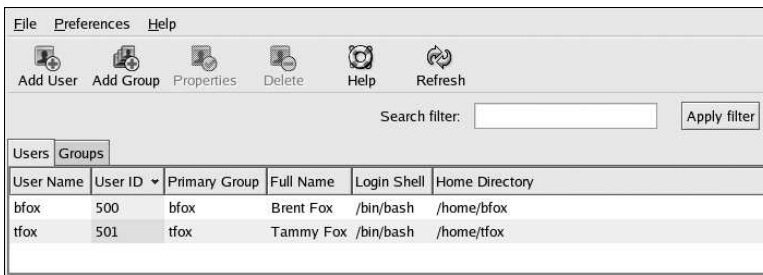


Figure 24-1. User Manager

To view a list of all local users on the system, click the **Users** tab. To view a list of all local groups on the system, click the **Groups** tab.

If you need to find a specific user or group, type the first few letters of the name in the **Filter by** field. Press [Enter] or click the **Apply filter** button. The filtered list will be displayed.

To sort the users or groups, click on the column name. The users or groups will be sorted by the value of that column.

Red Hat Linux reserves user IDs above 500 for system users. By default, **User Manager** does not display system users. To view all users, including the system users, uncheck **Preferences** => **Filter system users and groups** from the pull-down menu.

For more information on users and groups, refer to the *Official Red Hat Linux Reference Guide*.

24.1. Adding a New User

To add a new user, click the **New User** button. A window as shown in Figure 24-2 will appear. Type the username and full name for the new user in the appropriate fields. Type the user's password in the **Password** and **Confirm Password** fields. The password must be at least six characters.



The longer the user's password, the more difficult it is for someone else to guess it and log in to the user's account without permission. It is also recommended that the password not be based on a word and that the password be a combination of letters, numbers, and special characters.

Select a login shell. If you are not sure which shell to select, accept the default value of `/bin/bash`. The default home directory is `/home/username`. You can change the home directory that is created for the user, or you can choose not to create the home directory by unselecting **Create home directory**.

Red Hat Linux uses a *user private group* (UPG) scheme. The UPG scheme does not add or change anything in the standard UNIX way of handling groups; it simply offers a new convention. Whenever you create a new user, by default, a unique group with the same name as the user is created. If you do not want to create this group, unselect **Create a private group for the user**.

To specify a user ID for the user, select **Specify user ID manually**. If the option is not selected, the next available user ID starting with number 500 will be assigned to the new user. Red Hat Linux reserves user IDs above 500 for system users.

Click **OK** to create the user.

The screenshot shows a 'New User' dialog box with the following fields and options:

- User Name:
- Full Name:
- Password:
- Confirm Password:
- Login Shell: (dropdown menu)
- Create home directory
 - Home Directory:
- Create a private group for the user
- Specify user ID manually
- UID: (spinners)
- Buttons: and

Figure 24-2. New User

To configure more advanced user properties such as password expiration, modify the user's properties after adding the user. Refer to Section 24.2 for more information.

To add the user to more user groups, click on the **User** tab, select the user, and click **Properties**. In the **User Properties** window, select the **Groups** tab. Select the groups that you want the user to be a member of, select the primary group for the user, and click **OK**.

24.2. Modifying User Properties

To view the properties of an existing user, click on the **Users** tab, select the user from the user list, and click **Properties** from the button menu (or choose **File => Properties** from the pull-down menu). A window similar to Figure 24-3 will appear.

User Data	Account Info	Password Info	Groups
User Name:	tfox		
Full Name:	Tammy Fox		
Password:	[masked]		
Confirm Password:	[masked]		
Home Directory:	/home/tfox		
Login Shell:	/bin/bash		

Figure 24-3. User Properties

The **User Properties** window is divided into tabbed pages:

- **User Data** — Basic user information configured when you added the user. Use this tab to change the user's full name, password, home directory, or login shell.
- **Account Info** — Select **Enable account expiration** if you want the account to expire on a certain date. Enter the date in the provided fields. Select **User account is locked** to lock the user account so that the user can not log in to the system.
- **Password Info** — This tab shows the date that the user last changed his password. To force the user to change his password after a certain number of days, select **Enable password expiration**. You can also set the number of days before the user is allowed to change his password, the number of days before the user is warned to change his password, and days before the account become inactive.
- **Groups** — Select the groups that you want the user to be a member of and the user's primary group.

24.3. Adding a New Group

To add a new user group, click the **New Group** button. A window similar to Figure 24-4 will appear. Type the name of the new group to create. To specify a group ID for the new group, select **Specify group ID manually** and select the GID. Red Hat Linux reserves group IDs lower than 500 for system groups.

Click **OK** to create the group. The new group will appear in the group list.

Group Name:	mygroup
<input checked="" type="checkbox"/> Specify group ID manually	
GID:	500

Figure 24-4. New Group

To add users to the group, refer to Section 24.4.

24.4. Modifying Group Properties

To view the properties of an existing group, select the group from the group list and click **Properties** from the button menu (or choose **File => Properties** from the pull-down menu). A window similar to Figure 24-3 will appear.



Figure 24-5. Group Properties

The **Group Users** tab displays which users are members of the group. Select additional users to add them to the group, and unselect users to remove from the group. Click **OK** or **Apply** to modify the users in the group.

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, the amount of available hard drive space, 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.

25.1. System Processes

The `ps ax` 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.

The `ps` output can be long. To prevent it from scrolling off the screen, you can pipe it through `less`:

```
ps aux | less
```

You can use the `ps` command in combination with the `grep` command to see if a process is running. For example, to determine if `emacs` is running, use the following command:

```
ps ax | grep emacs
```

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:

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.

Table 25-1. Interactive `top` commands



Tip

Application such as **Mozilla** and **Nautilus** are *thread-aware* — multiple threads are created to handle multiple users or multiple requests, and each thread is given a process ID. By default, `ps` and `top` only display the main (initial) thread. To view all threads, use the command `ps -m` or type `[Shift]-[H]` in `top`.

If you prefer a graphical interface for `top`, you can use the **GNOME System Monitor**. To start it from the desktop, select **Main Menu Button** (on the Panel) => **System Tools** => **System Monitor** or type `gnome-system-monitor` at a shell prompt from within the X Window System. Then select the **Process Listing** tab.

The **GNOME System Monitor** allows you to search for process in the list of running process as well as view all processes, your processes, or active processes.

To learn more about a process, select it and click the **More Info** button. Details about the process will be displayed at the bottom of the window.

To stop a process, select it and click **End Process**. This function is useful for processes that have stopped responding to user input.

To sort by the information in a specific column, click on the name of the column. The column that the information is sorted by appears in a darker gray color.

By default, the **GNOME System Monitor** does not display threads. To change this preferences, select **Edit** => **Preferences**, click the **Process Listing** tab, and select **Show Threads**. The preferences also allows you to configure the update interval, what type of information to display about each process by default, and the colors of the system monitor graphs.




Figure 25-1. GNOME System Monitor

25.2. 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 is 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 prefer a graphical interface for `free`, you can use the **GNOME System Monitor**. To start it from the desktop, go to the **Main Menu Button** (on the Panel) => **System Tools** => **System Monitor** or type `gnome-system-monitor` at a shell prompt from within X Window System. Then choose the **System Monitor** tab.

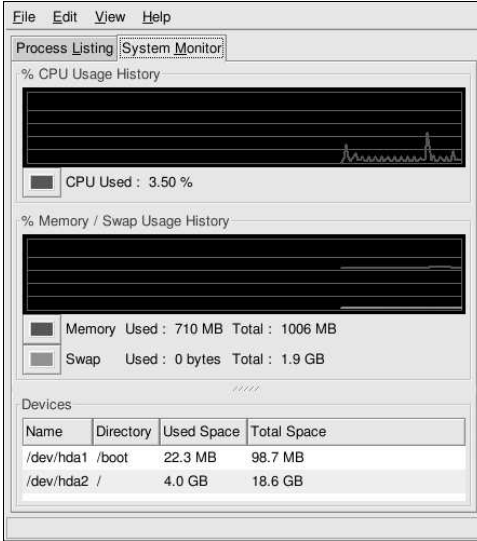


Figure 25-2. GNOME System Monitor

25.3. File Systems

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
none           256796         0         256796    0% /dev/shm
```

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
none           251M    0   250M  0% /dev/shm
```

In the list of partitions, there is an entry for `/dev/shm`. This entry represents the system's virtual memory file system.

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 the totals for 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.

To view the system's partitions and disk space usage in a graphical format, use the **System Monitor** tab as shown at the bottom of Figure 25-2.

25.3.1. Monitoring File Systems

Red Hat Linux provides a utility called `diskcheck` that monitors the amount of free disk space on the system. Based on the configuration file, it will send email to the system administrator when one or more disk drives reach a specified capacity. To use this utility, you must have the `diskcheck` RPM package installed.

This utility is run as an hourly cron ¹ task.

The following variables can be defined in `/etc/diskcheck.conf`:

- `defaultCutoff` — When disk drives reach this percent capacity, it will be reported. For example, if `defaultCutoff = 90`, an email will be sent when the monitored disk drives reach 90% capacity.
- `cutoff[/dev/partition]` — Override the `defaultCutoff` for the partition. For example, if `cutoff['/dev/hda3'] = 50` is specified, `diskcheck` will alert the system administrator when the partition `/dev/hda3` reaches 50% capacity.
- `cutoff[/mountpoint]` — Override the `defaultCutoff` for the mount point. For example, if `cutoff['/home'] = 50` is specified, `diskcheck` will alert the system administrator when the mount point `/home` reaches 50% capacity.
- `exclude` — Specify one or more partitions for `diskcheck` to ignore. For example, if `exclude = "/dev/sda2 /dev/sda4"` is specified, `diskcheck` will not alert the system administrator if `/dev/sda2` or `/dev/sda4` reaches the specified cutoff percentage.
- `ignore` — Specify one or more file system types to ignore in the format `-x filesystem-type`. For example, if `ignore = "-x nfs -x iso9660"` is specified, the system administrator will not be alerted about `nfs` or `iso9660` file systems reaching capacity.
- `mailTo` — Email address of the system administrator to alert when partitions and mount points reach the specified capacity. For example, if `mailTo = "webmaster@example.com"` is specified, `webmaster@example.com` will be emailed alerts.
- `mailFrom` — Specify the identity of the email sender. This is useful if the system administrator wants to filter the mail from `diskcheck`. For example, if `mailFrom = "Disk Usage Monitor"` is specified, email will be sent to the system administrator with the sender `Disk Usage Monitor`.
- `mailProg` — Specify the mail program to use to send email alerts. For example, if `mailProg = "/usr/sbin/sendmail"` is specified, `Sendmail` will be used as the mail program.

You do not have to restart a service if you change the configuration file because it is read each time the cron task is run. You must have the `crond` service running for cron tasks to be executed. To determine if the daemon is running, use the command `/sbin/service crond status`. It is recommended that you start the service at boot time. Refer to Chapter 13 for details on starting the cron service automatically at boot time.

1. Refer to Chapter 27 for more information on cron.

25.4. Hardware

If you are having trouble configuring your hardware or just want to know what hardware is in your system, you can use the **Hardware Browser** application to display the hardware that can be probed. To start the program from the desktop, select **Main Menu Button => System Tools => Hardware Browser** or type `hwbrowser` at a shell prompt. As shown in Figure 25-3, it displays your CD-ROM drives, floppy disks, hard drives and their partitions, network devices, pointing devices, system devices, and video cards. Click on the category name in the left menu, and the information will be displayed.

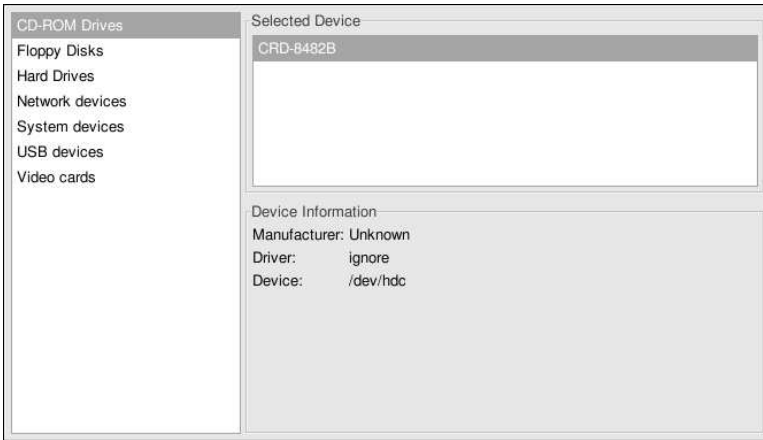


Figure 25-3. Hardware Browser

You can also use the `lspci` command to list all PCI devices. Use the command `lspci -v` for more verbose information or `lspci -vv` for very verbose output.

For example, `lspci` can be used to determine the manufacturer, model, and memory size of a system's video card:

```
01:00.0 VGA compatible controller: Matrox Graphics, Inc. MGA G400 AGP (rev 04) (prog-if 00 [VGA])
Subsystem: Matrox Graphics, Inc. Millennium G400 Dual Head Max
Flags: medium devsel, IRQ 16
Memory at f4000000 (32-bit, prefetchable) [size=32M]
Memory at fcffc000 (32-bit, non-prefetchable) [size=16K]
Memory at fc000000 (32-bit, non-prefetchable) [size=8M]
Expansion ROM at 80000000 [disabled] [size=64K]
Capabilities: [dc] Power Management version 2
Capabilities: [f0] AGP version 2.0
```

The `lspci` is also useful to determine the network card in your system if you do not know the manufacturer or model number.

25.5. Additional Resources

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

25.5.1. Installed Documentation

- `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.

Printer Configuration

Printer Configuration Tool allows users to configure a printer in Red Hat Linux. It helps maintain the `/etc/printcap` configuration file, print spool directories, and print filters.

Starting with version 8.0, Red Hat Linux ships with two printer systems. **Printer Configuration Tool** configures the printing system called LPRng. LPRng is also the default printing system. This chapter focuses on using **Printer Configuration Tool** to configure LPRng. For more information on the alternate printing system called CUPS, refer to Section 26.12.

To use **Printer Configuration Tool**, you must have root privileges. To start the application, select **Main Menu Button** (on the Panel) => **System Settings** => **Printing**, or type the command `redhat-config-printer`. This command automatically determines whether to run the graphical or text-based version.

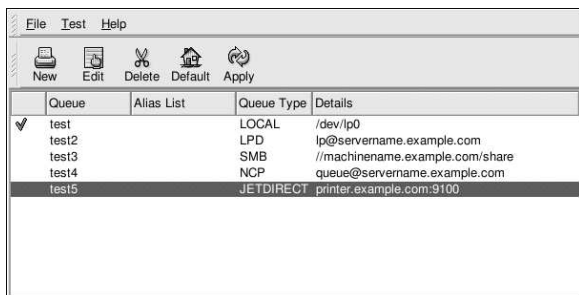
You can also force **Printer Configuration Tool** to run as a text-based application by using the command `redhat-config-printer-tui` from a shell prompt.

To add a printer using the command line version, refer to Section 26.10.

Important

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 **Printer Configuration Tool**, edit the `/etc/printcap.local` file. The entries in `/etc/printcap.local` are not displayed in the **Printer Configuration Tool** 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 this application. Each time a new configuration file is generated, the old file is saved as `/etc/printcap.old`.



Queue	Alias List	Queue Type	Details
<input checked="" type="checkbox"/> test		LOCAL	/dev/lp0
test2		LPD	lp@servername.example.com
test3		SMB	//machinename.example.com/share
test4		NCP	queue@servername.example.com
test5		JETDIRECT	printer.example.com:9100

Figure 26-1. Printer Configuration Tool

Five types of print queues can be configured:

- **Local Printer** — a printer attached directly to your computer through a parallel or USB port. In the main printer list as shown in Figure 26-1, 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 (for example, a printer attached to another Red Hat Linux system on your network). In the main printer list as shown in Figure 26-1, the **Queue Type** for a remote UNIX printer is set to **LPD**.
- **Windows Printer (SMB)** — 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 26-1, 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 26-1, the **Queue Type** for a remote Novell printer is set to **NCP**.
- **JetDirect Printer** — a printer connected directly to the network through HP JetDirect instead of to a computer. In the main printer list as shown in Figure 26-1, 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. The changes are not written to the `/etc/printcap` configuration file until the printer daemon (`lpd`) is restarted. 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.

26.1. 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 **New** button in the main **Printer Configuration Tool** window. The window shown in Figure 26-2 will appear. Click **Next** to proceed.

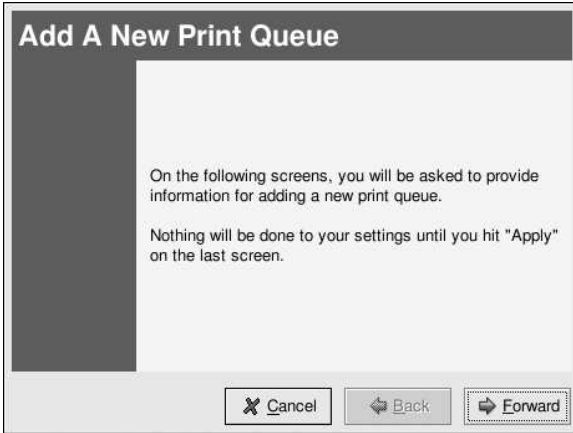


Figure 26-2. Adding a Printer

You will then see the screen shown in Figure 26-3. Enter a unique name for the printer in the **Queue Name** text field. This can be any descriptive name for your printer. The printer name 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 **_**.

Select **Local Printer** from the **Queue Type** menu, and click **Next**.




Figure 26-3. Adding a Local Printer

Printer Configuration Tool attempts to detect your printer device and model as shown in Figure 26-4. If your printer device is not shown, click **Custom Device**. Type the name of your printer device and click **OK** to add it to the printer device list. A printer device attached to the parallel port is usually referred to as `/dev/lp0`. A printer device attached to the USB port is usually referred to as `/dev/usb/lp0`. If your printer model does not appear, you will be given the opportunity to select it in the next step. Click **Next** to continue.

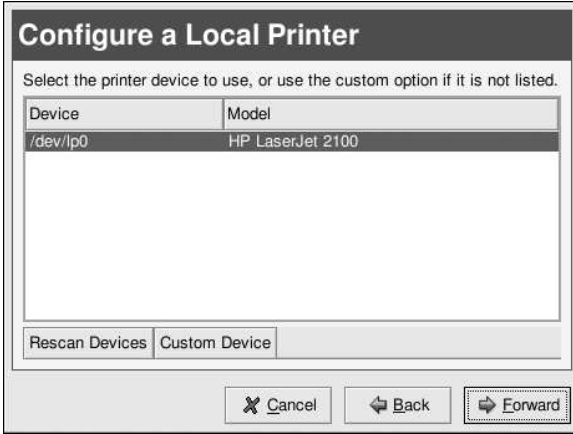


Figure 26-4. Local Printer Device

If **Printer Configuration Tool** detected your printer model, it will display the recommended print driver. Skip to Section 26.6 to continue.

26.2. 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 **New** button in the main **Printer Configuration Tool** window. The window shown in Figure 26-2 will appear. Click **Next** to proceed.

You will then see the screen shown in Figure 26-5. Enter a unique name for the printer in the **Queue Name** text field. The printer name 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 **_**.

Select **Unix Printer** from the **Queue Type** menu, and click **Next**.

Set the Print Queue Name and Type

Enter the Queue's name, and select the Queue's Type.
Valid names can contain the characters "a-z", "A-Z", "0-9", "-", and "_".
They must begin with letters.

Queue Name:
test2

Queue Type

<input type="radio"/> Local Printer	LOCAL
<input checked="" type="radio"/> Unix Printer	LPD
<input type="radio"/> Windows Printer	SMB
<input type="radio"/> Novell Printer	NCP
<input type="radio"/> JetDirect Printer	JETDIRECT

Cancel Back Forward

Figure 26-5. Adding a Remote Printer

Text fields for the following options appears as shown in Figure 26-6:

- **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.

Click **Next** to continue.

Configure a Unix Print Queue

Configure the server and queue of your remote lpd server. If no queue is specified, you will use the default queue on the remote server. Select strict RFC1179 compliance if you are having problems connecting to an older lpd server.

Server: servername.example.com

Queue: lp

Strict RFC1179 Compliance

Cancel Back Forward

Figure 26-6. Choosing the Printer Server

The next step is to select the type of printer that is connected to the remote system. Skip to Section 26.6 to continue.



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.

26.3. Adding a Samba (SMB) Printer

To add a printer which is accessed using the SMB protocol (such as a printer attached to a Microsoft Windows system) click the **New** button in the main **Printer Configuration Tool** window. The window shown in Figure 26-2 will appear. Click **Next** to proceed.

You will see the screen shown in Figure 26-7. Enter a unique name for the printer in the **Queue Name** text field. The printer name 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 **_**.

Select **Windows Printer** from the **Queue Type** menu, and click **Next**. If the printer is attached to a Microsoft Windows system, choose this queue type.

Figure 26-7. Adding a SMB Printer

Text fields for the following options appear as shown in Figure 26-8:

- **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 log in as to access the printer. This user must exist on the Windows system, and the user must have permission to access the printer. The default 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.

Click **Next** to continue.

Configure a Windows Print Queue

Configure the SMB share of your remote printer.

Share: //machinename/sharename User: guest

Host IP: 192.168.1.9 Password: *****

Workgroup: devel Translate \n => \r\n

Figure 26-8. Choosing the Print Server

The next step is to select the type of printer that is connected to the remote SMB system. Skip to Section 26.6 to continue.



Note

If you require a username and password for an SMB (LAN Manager) print queue, they are stored unencrypted in the spool directory, which only root or lp may read. Thus, it is possible for others to learn the username and password if they have root access. To avoid this, the username and password to access 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.

26.4. Adding a Novell NetWare (NCP) Printer

To add a Novell NetWare (NCP) printer, click the **New** button in the main **Printer Configuration Tool** window. The window shown in Figure 26-1 will appear. Click **Next** to proceed.

You will see the screen shown in Figure 26-9. Enter a unique name for the printer in the **Queue Name** text field. The printer name 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 **_**.

Select **Novell Printer** from the **Queue Type** menu, and click **Next**.

Set the Print Queue Name and Type

Enter the Queue's name, and select the Queue's Type.
Valid names can contain the characters "a-z", "A-Z", "0-9", "-", and "_".
They must begin with letters.

Queue Name:
test4

Queue Type:

Local Printer LOCAL

Unix Printer LPD

Windows Printer SMB

Novell Printer NCP

JetDirect Printer JETDIRECT

Cancel Back Forward

Figure 26-9. Adding an NCP Printer

Text fields for the following options appear below the **Queue Type** menu as shown in Figure 26-10:

- **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 log in as to access the printer.
- **Password** — The password for the user specified in the **User** field above.

Configure a Novell Print Queue

Configure the remote server and queue of your NCP printer.

Server: servername.example.com User: username

Queue: queue Password: *****

Cancel Back Forward

Figure 26-10. Choosing the Print Server

The next step is to select the type of printer that is connected to the remote NCP system. Skip to Section 26.6 to continue.

**Note**

If you require a username and password for a NCP (NetWare) print queue, they are stored 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.

26.5. Adding a JetDirect Printer

To add a JetDirect printer, click the **New** button in the main **Printer Configuration Tool** window. The window shown in Figure 26-1 will appear. Click **Next** to proceed.

You will see the screen shown in Figure 26-11. Enter a unique name for the printer in the **Queue Name** text field. The printer name 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 **_**.

Select **JetDirect Printer** from the **Queue Type** menu, and click **Next**.

Set the Print Queue Name and Type

Enter the Queue's name, and select the Queue's Type.
Valid names can contain the characters "a-z", "A-Z", "0-9", "-", and "_".
They must begin with letters.

Queue Name:
test5

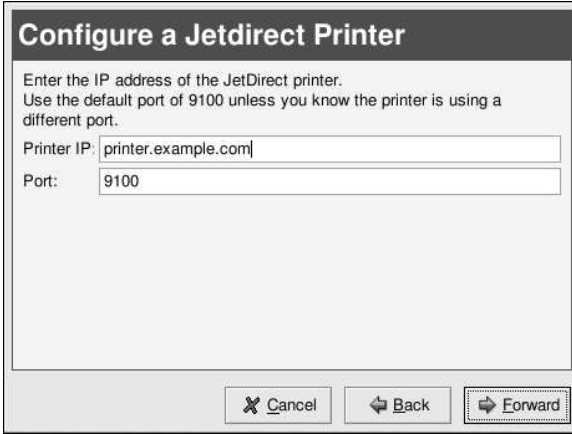
Queue Type

<input type="radio"/> Local Printer	LOCAL
<input type="radio"/> Unix Printer	LPD
<input type="radio"/> Windows Printer	SMB
<input type="radio"/> Novell Printer	NCP
<input checked="" type="radio"/> JetDirect Printer	JETDIRECT

Figure 26-11. Adding a JetDirect Printer

Text fields for the following options appear below the **Queue Type** menu as shown in Figure 26-12:

- **Printer IP** — The hostname or IP address of the JetDirect printer.
- **Port** — The port on the JetDirect printer that is listening for print jobs. The default port is 9100.



Configure a Jetdirect Printer

Enter the IP address of the JetDirect printer.
Use the default port of 9100 unless you know the printer is using a different port.

Printer IP:

Port:

Figure 26-12. Choosing a Print Server

The next step is to select the type of printer that is connected to the JetDirect system. Skip to Section 26.6 to continue.

26.6. Selecting the Print Driver and Finishing

After selecting the queue type of the printer, the next step in adding a printer is to select the print driver.

You will see a window similar to Figure 26-13. If you are configuring a local printer and the model was autodetected, the recommended driver is autoselected and marked with an asterisk (*). If it was not autodetected, select the driver from the list. The printers are divided by manufacturers. Click the arrow beside the manufacturer for your printer. Find your printer from the expanded list, and click the arrow beside the printer name. A list of drivers for your printer will appear. Select one.



Tip

To learn more about the print drivers, go to http://www.linuxprinting.org/printer_list.cgi.



Tip

You can select a different print driver after adding a printer by starting **Printer Configuration Tool**, selecting the printer from the list, clicking **Edit**, clicking the **Driver** tab, selecting a different print driver, and applying the changes.

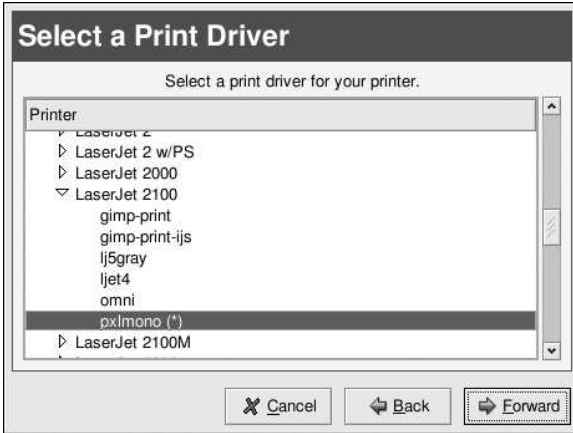


Figure 26-13. Selecting a Print Driver

The print driver processes the data that you want to print into a format the printer can understand. Since a local printer is attached directly to your computer, you need to select a print driver to process the data that is sent to the printer.

If you are configuring a remote printer (LPD, SMB, or NCP), the remote print server usually has its own print driver. If you select an additional print driver on your local computer, the data will be filtered more than once, and the data will be converted to a format that the printer can not understand.

To make sure the data is not filtered more than once, first try selecting **Raw Print Queue** or **Postscript Printer** if you are configuring a remote printer. After applying the changes, print a test page to test this configuration. If the test fails, the remote print server might not have a print driver configured. Try selecting a print driver according to the manufacturer and model of the remote printer, applying the changes, and printing a test page.

26.6.1. Confirming Printer Configuration

The last step is to confirm your printer configuration. Click **Apply** to add the print queue if the settings are correct. Click **Back** to modify the printer configuration.

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 26.7 for details.

If you need to print characters beyond the basic ASCII set (including those used for languages such as Japanese), you need to go to your driver options and select **Prerender Postscript**. Refer to Section 26.8 for details. You can also configure options such as paper size if you edit the print queue after adding it.


26.7. 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 select the appropriate test page from the **Test** pulldown menu.


If you change the print driver or modify the driver options, you should print a test page to test the different configuration.

26.8. Modifying Existing Printers

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.

After adding your printer(s), you can edit settings by selecting the printer from the printer list and clicking the **Edit** button. The tabbed window shown in Figure 26-14 will appear. The window contains the current values for the printer that you selected to edit. Make any changes, and click **OK**. Click **Apply** in the main **Printer Configuration Tool** window to save the changes and restart the printer daemon.



Figure 26-14. Editing a Printer

26.8.1. Names and Aliases

If you want to rename a printer, change the value of **Queue Name** in the **Names and Aliases** tab. Click **OK** to return to the main window. The name of the printer should change in the printer list. Click **Apply** to save the change and restart the printer daemon.

A printer alias is an alternate name for a printer. To add an alias for an existing printer, click the **Add** button in the **Name and Aliases** tab, enter the name of the alias, and click **OK**. Click **OK** again to return to the main window. Click **Apply** to save the aliases and restart the printer daemon. A printer can have more than one alias.

26.8.2. Queue Type

The **Queue Type** tab shows queue type that you selected when adding the printer and its settings. You can change the queue type of the printer or just change the settings. After making modifications, click **OK** to return to the main window. Click **Apply** to save the change and restart the printer daemon.

Depending on which queue type you choose, you will see different different options. Refer to the appropriate section on adding a printer for a description of the options.

26.8.3. Driver

The **Driver** tab shows which print driver is currently being used. This is the same list that you used when adding the printer. If you change the print driver, click **OK** to return to the main window. Click **Apply** to save the change and restart the printer daemon.

26.8.4. Driver Options

The **Driver Options** tab displays advanced printer options. Options vary for each print driver. Common options include:

- **Send Form-Feed (FF)** should be selected if the last page of your print job is not ejected from the printer (for example, the form feed light flashes). If this does not work, try selecting **Send End-of-Transmission (EOT)** instead. Some printers require both **Send Form-Feed (FF)** and **Send End-of-Transmission (EOT)** to eject the last page.
- **Send End-of-Transmission (EOT)** if sending a form-feed does not work. Refer to **Send FF** above.
- **Assume Unknown Data is Text** should be selected if your print driver does not recognize some of the data sent to it. Only select it if you are having problems printing. If this option is selected, the print driver will assume that any data that it can not recognize is text and try to print it as text. If you select this option and **Convert Text to Postscript**, the print driver will assume the unknown data is text and then convert it to PostScript.
- **Prerender Postscript** should be selected if you are printing characters beyond the basic ASCII set but they are not printing correctly (such as Japanese characters). This option will prerender non-standard PostScript fonts so that they are printed correctly.

If your printer does not support the fonts you are trying to print, try selecting this option. For example, you should select this option if you are printing Japanese fonts to a non-Japanese printer.

Extra time is required to perform this action. Do not choose it unless you are having problems printing the correct fonts.

You should also select this option if your printer can not handle PostScript level 3. This option converts it to PostScript level 1.

- **Convert Text to Postscript** is selected by default. If your printer can print plain text, try unselecting this when printing plain text documents to decrease the time it takes to print.
- **Page Size** allows you to select the paper size for your printer such as US Letter, US Legal, A3, and A4.
- **Effective Filter Locale** defaults to **C**. If you are printing Japanese characters, select **ja_JP**. Otherwise, accept the default of **C**.
- **Media Source** defaults to **Printer default**. Change this option to use paper from a different tray.

If you modify the driver options, click **OK** to return to the main window. Click **Apply** to save the change and restart the printer daemon.

26.9. Saving the Configuration File

When you save your printer configuration using **Printer Configuration Tool**, it creates its own configuration file that is used to create the `/etc/printcap` file that the printer daemon (`lpd`) reads. You can use the command line options to save or restore this file. If you save your `/etc/printcap` file and overwrite your existing `/etc/printcap` file with the saved file, your printer configuration will not be restored. Each time the printer daemon is restarted, it creates a new `/etc/printcap` file from the special **Printer Configuration Tool** configuration file. If you have configured a backup system for your configuration files, you should use the following method to save your printer configuration. If you added any custom settings in the `/etc/printcap.local` file, you should save it as part of your backup system also.

To save your printer configuration, type this command as root:

```
/usr/sbin/redhat-config-printer-tui --Xexport > settings.xml
```

Your configuration is saved to the file `settings.xml`.

If you save this file, you can restore your printer settings. This is useful if your printer configuration is deleted, you reinstall Red Hat Linux and do not have your printer configuration file anymore, or you want to use the same printer configuration on multiple systems. To restore the configuration, type this command as root:

```
/usr/sbin/redhat-config-printer-tui --Ximport < settings.xml
```

If you already have a configuration file (you have configured one or more printers on the system already) and you try to import another configuration file, the existing configuration file will be overwritten. If you want to keep your existing configuration and add the configuration in the saved file, you can merge the files with the following command (as root):

```
/usr/sbin/redhat-config-printer-tui --Ximport --merge < settings.xml
```

Your printer list will then consist of the printers you configured on the system as well as the printers you imported from the saved configuration file. If the imported configuration file has a print queue with the same name as an existing print queue on the system, the print queue from the imported file will override the existing printer.

After importing the configuration file (with or without the `merge` command), you must restart the printer daemon with the command `/sbin/service lpd restart` or by starting **Printer Configuration Tool** and clicking **Apply**.

26.10. Command Line Configuration

If you do not have X installed and you do not want to use the text-based version, you can add a printer using the command line. This method is useful if you want to add a printer from a script or in the post of a kickstart installation.

26.10.1. Adding a Printer

To add a printer:

```
redhat-config-printer-tui --Xadd-local options
```

Options:

```
--device=node
```

The device node to use. For example, `/dev/lp0`. (required)

`--make=make`

The IEEE 1284 MANUFACTURER string or the printer manufacturer's name as in the foomatic database if the manufacturer string is not available. (required)

`--model=model`

The IEEE 1284 MODEL string or the printer model in the foomatic database if the model string is not available. (required)

`--name=name`

The name to be given to the new queue. (optional) If one is not given, a name based on the device node (such as "lp0") will be used.

`--as-default`

Set this as the default queue. (optional)

After adding the printer, use the following command to start/restart the printer daemon, `lpd`:

```
service lpd restart
```

26.10.2. Removing a Printer

You can also remove a printer queue through the command line.

To remove a printer queue:

```
redhat-config-printer-tui --Xremove-local options
```

Options:

`--device=node`

The device node used (for example, `/dev/lp0`). Required.

`--make=make`

The IEEE 1284 MANUFACTURER string, or (if none is available) the printer manufacturer's name as in the foomatic database. Required.

`--model=model`

The IEEE 1284 MODEL string, or (if none is available) the printer model as in the foomatic database. Required.

After removing the printer from the **Printer Configuration Tool** configuration, restart the printer daemon for the changes to take effect:

```
service lpd restart
```

If you removed all printers and do not want to run the printer daemon anymore, execute the following command:

```
service lpd stop
```

26.11. Managing Your Print Jobs

When you send a print job to the printer daemon such as printing text file from **Emacs** or printing an image from **The GIMP**, the print job is added to the print spool queue. The print spool queue is a list of print jobs that have been sent to the printer and information about each print request such as the status of the request, the username of the person who sent the request, the hostname of the system that sent the request, the job number, and more. To view the list of print jobs in the print spool, open a shell prompt and type the command `lpq`. The last few lines will look similar to the following:

```
Rank   Owner/ID           Class Job Files      Size Time
active user@localhost+902 A    902 sample.txt    2050 01:20:46
```

Example 26-1. Example of `lpq` output

If you want to cancel a print job, find the job number of the request with the command `lpq` and then use the command `lprm job number`. For example, `lprm 902` would cancel the print job in Example 26-1. You must have proper permissions to cancel a print job. You can not cancel print jobs that were started by other users unless you are logged in as root on the machine to which the printer is attached.

You can also print a file directly from a shell prompt. For example, the command `lpr sample.txt` will print the text file `sample.txt`. The print filter determines what type of file it is and converts it a format the printer can understand.

26.12. Configuring the CUPS Printing System

CUPS (Common *UNIX* Printing System) can be used instead of the default LPRng printing system. Some of the advantages of CUPS include:

- Support for IPP (next generation network printing protocol)
- Autodetection of network printers
- Web interface configuration tool
- Support for PPD printer description files
- Support for a wide-range of printers

26.12.1. Switching Print Systems

To use the CUPS printing system instead of LPRng, run the **Printer System Switcher** application. Start it by selecting to the **Main Menu Button** (on the Panel) => **Extras** => **System Settings** => **Printer System Switcher**, or type the command `redhat-switch-printer` at a shell prompt (for example, in an XTerm or GNOME terminal).

The program automatically detects if the X Window System is running. If it is running, the program starts in graphical mode as shown in Figure 26-15. If X is not detected, it starts in text-mode. To force it to run in text-mode, use the command `redhat-switch-printer-nox`.

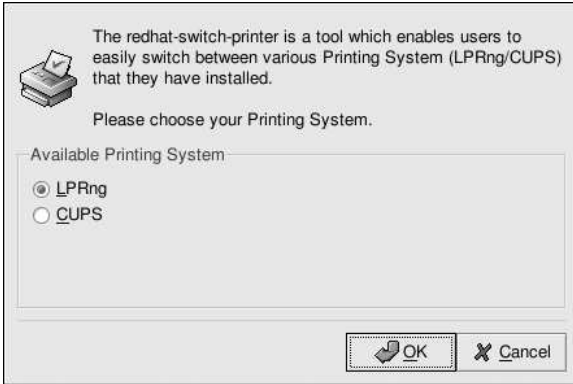


Figure 26-15. Printer System Switcher

If you selected **CUPS**, you must make sure the `lpd` service is stopped and the `cups` service is started:

```
/sbin/service lpd stop
/sbin/service cups start
```

If you selected **LPRng**, you must make sure the `cups` service is stopped and the `lpd` service is started:

```
/sbin/service cups stop
/sbin/service lpd start
```

Also use `chkconfig`, `ntsysv`, or **Services Configuration Tool** to configure your system to start the `cups` service automatically and disable the `lpd` service. Refer to Chapter 13 for details.

26.12.2. CUPS Configuration Interface

After starting the `cups` daemon, open a Web browser and connect to the URL `http://localhost:631` as shown in Figure 26-16.

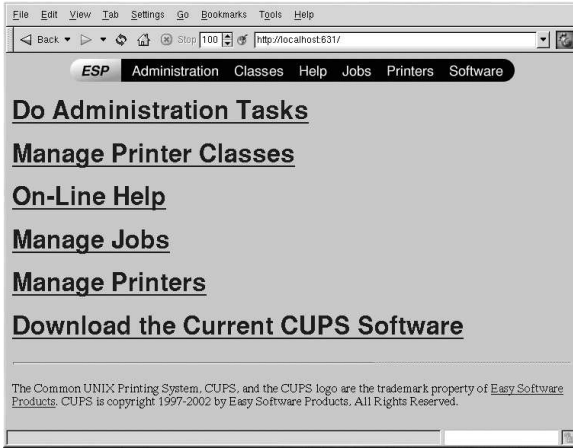


Figure 26-16. CUPS Configuration Tool

To add a printer, click **Manage Printers**, and then click the **Add Printer** button. For more information, click the **Help** button.

26.13. Additional Resources

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

26.13.1. Installed Documentation

- `man printcap` — The manual page for the `/etc/printcap` printer configuration file.
- `man lpr` — The manual page for the `lpr` command that allows you to print files from the command line.
- `man lpd` — The manual page for the LPRng printer daemon.
- `man lprm` — The manual page on the command line utility to remove print jobs from the LPRng spool queue.
- `man mpage` — The manual page on the command line utility to print multiple pages on one sheet of paper.
- `man cupsd` — The manual page for the CUPS printer daemon.
- `man cupsd.conf` — The manual page for the CUPS printer daemon configuration file.
- `man classes.conf` — The manual page for the class configuration file for CUPS.

26.13.2. Useful Websites

- <http://www.linuxprinting.org> — *GNU/Linux Printing* contains a large amount information about printing in Linux.
- <http://www.cups.org/> — Documentation, FAQs, and newsgroups about CUPS.

Automated Tasks

In Linux, tasks can be configured to run automatically within a specified period of time, on a specified date, or when the system load average is below a specified number. Red Hat Linux comes preconfigured to run important system tasks to keep the system updated. For example, the `slocate` database used by the `locate` command is updated daily. A system administrator can use automated tasks to perform periodic backups, monitor the system, run custom scripts, and more.

Red Hat Linux comes with four automated tasks utilities: `cron`, `anacron`, `at`, and `batch`.

27.1. Cron

Cron is a daemon that can be used to schedule the execution of recurring tasks according to a combination of the time, day of the month, month, day of the week, and week.

Cron assumes that the system is on continuously. If the system is not on when a task is scheduled, it is not executed. To configure tasks based on time periods instead of exact times, refer to Section 27.2. To schedule one-time tasks, refer to Section 27.3.

To use the cron service, you must have the `vixie-cron` RPM package installed, and the `crond` service must be running. To determine if the package is installed, use the `rpm -q vixie-cron` command. To determine if the service is running, use the command `/sbin/service crond status`.

27.1.1. Configuring Cron Tasks

The main configuration file for cron, `/etc/crontab`, contains the following lines:

```
SHELL=/bin/bash
PATH=/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root
HOME=/

# run-parts
01 * * * * root run-parts /etc/cron.hourly
02 4 * * * root run-parts /etc/cron.daily
22 4 * * 0 root run-parts /etc/cron.weekly
42 4 1 * * root run-parts /etc/cron.monthly
```

The first four lines are variables used to configure the environment in which the cron tasks are run. The value of the `SHELL` variable tells the system which shell environment to use (in this example the bash shell), and the `PATH` variable defines the path used to execute commands. The output of the cron tasks are emailed to the username defined with the `MAILTO` variable. If the `MAILTO` variable is defined as an empty string (`MAILTO=""`), email will not be sent. The `HOME` variable can be used to set the home directory to use when executing commands or scripts.

Each line in the `/etc/crontab` file represents a task and has the format:

```
minute hour day month dayofweek command
```

- `minute` — any integer from 0 to 59
- `hour` — any integer from 0 to 23
- `day` — any integer from 1 to 31 (must be a valid day if a month is specified)

- `month` — any integer from 1 to 12 (or the short name of the month such as `jan`, `feb`, and so on)
- `dayofweek` — any integer from 0 to 7, where 0 or 7 represents Sunday (or the short name of the week such as `sun`, `mon`, and so on)
- `command` — the command to execute. The command can either be a command such as `ls /proc >> /tmp/proc` or the command to execute a custom script that you wrote.

For any of the above values, an asterisk (*) can be used to specify all valid values. For example, an asterisk for the month value means execute the command every month within the constraints of the other values.

A hyphen (-) between integers specifies a range of integers. For example, `1-4` means the integers 1, 2, 3, and 4.

A list of values separated by commas (,) specifies a list. For example, `3, 4, 6, 8` indicates those four specific integers.

The forward slash (/) can be used to specify step values. The value of an integer can be skipped within a range by following the range with `/<integer>`. For example, `0-59/2` can be used to define every other minute in the minute field. Step values can also be used with an asterisk. For instance, the value `*/3` can be used in the month field to run the task every third month.

Any lines that begin with a hash mark (#) are comments and are not processed.

As you can see from the `/etc/crontab` file, it uses the `run-parts` script to execute the scripts in the `/etc/cron.hourly`, `/etc/cron.daily`, `/etc/cron.weekly`, and `/etc/cron.monthly` directories on an hourly, daily, weekly, or monthly basis respectively. The files in these directory should be shell scripts.

If a cron task needs to be executed on a schedule other than hourly, daily, weekly, or monthly, it can be added to the `/etc/cron.d` directory. All files in this directory use the same syntax as `/etc/crontab`. Refer to Example 27-1 for examples.

```
# record the memory usage of the system every monday
# at 3:30AM in the file /tmp/meminfo
30 3 * * mon cat /proc/meminfo >> /tmp/meminfo
# run custom script the first day of every month at 4:10AM
10 4 1 * * /root/scripts/backup.sh
```

Example 27-1. Crontab Examples

Users other than root can configure cron tasks by using the `crontab` utility. All user-defined crontabs are stored in the `/var/spool/cron` directory and are executed using the usernames of the users that created them. To create a crontab as a user, login as that user and type the command `crontab -e` to edit the user's crontab using the editor specified by the `VISUAL` or `EDITOR` environment variable. The file uses the same format as `/etc/crontab`. When the changes to the crontab are saved, the crontab is stored according to username and written to the file `/var/spool/cron/username`.

The cron daemon checks the `/etc/crontab` file, the `/etc/cron.d/` directory, and the `/var/spool/cron` directory every minute for any changes. If any changes are found, they are loaded into memory. Thus, the daemon does not need to be restarted if a crontab file is changed.

27.1.2. Controlling Access to Cron

The `/etc/cron.allow` and `/etc/cron.deny` files are used to restrict access to cron. The format of both access control files is one username on each line. Whitespace is not permitted in either file. The cron daemon (`crond`) does not have to be restarted if the access control files are modified. The access control files are read each time a user tries to add or delete a cron task.

The root user can always use cron, regardless of the user names listed in the access control files.

If the file `cron.allow` exists, only users listed in it are allowed to use cron, and the `cron.deny` file is ignored.

If `cron.allow` does not exist, all users listed in `cron.deny` are not allowed to use cron.

27.1.3. Starting and Stopping the Service

To start the cron service, use the command `/sbin/service crond start`. To stop the service, use the command `/sbin/service crond stop`. It is recommended that you start the service at boot time. Refer to Chapter 13 for details on starting the cron service automatically at boot time.

27.2. Anacron

Anacron is a task scheduler similar to cron except that it does not require the system to run continuously. It can be used to run the daily, weekly, and monthly jobs usually run by cron.

To use the Anacron service, you must have the `anacron` RPM package installed and the `anacron` service must be running. To determine if the package is installed, use the `rpm -q anacron` command. To determine if the service is running, use the command `/sbin/service anacron status`.

27.2.1. Configuring Anacron Tasks

Anacron tasks are listed in the configuration file `/etc/anacrontab`. Each line in the configuration file corresponds to a task and has the format:

```
period delay job-identifier command
```

- `period` — frequency (in days) to execute the command
- `delay` — delay time in minutes
- `job-identifier` — description of the task, used in Anacron messages and as the name of the job's timestamp file, can contain any non-blank characters (except slashes).
- `command` — command to execute

For each task, Anacron determines if the task has been executed within the period specified in the `period` field of the configuration file. If it has not been executed within the given period, Anacron executes the command specified in the `command` field after waiting the number of minutes specified in the `delay` field.

After the task is completed, Anacron records the date in a timestamp file in the `/var/spool/anacron` directory. Only the date is used (not the time), and the value of the `job-identifier` is used as the filename for the timestamp file.

Environment variables such as `SHELL` and `PATH` can be defined at the top of `/etc/anacrontab` as with the cron configuration file.

The default configuration file looks similar to the following:

```
# /etc/anacrontab: configuration file for anacron
# See anacron(8) and anacrontab(5) for details.

SHELL=/bin/sh
PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin
```

```
# These entries are useful for a Red Hat Linux system.
1      5      cron.daily      run-parts /etc/cron.daily
7      10     cron.weekly     run-parts /etc/cron.weekly
30     15     cron.monthly    run-parts /etc/cron.monthly
```

Figure 27-1. Default anacrontab

As you can see in Figure 27-1, anacron for Red Hat Linux is configured to make sure the daily, weekly, and monthly cron tasks are run.

27.2.2. Starting and Stopping the Service

To start the anacron service, use the command `/sbin/service anacron start`. To stop the service, use the command `/sbin/service anacron stop`. It is recommended that you start the service at boot time. Refer to Chapter 13 for details on starting the anacron service automatically at boot time.

27.3. At and Batch

While cron and anacron are used to schedule recurring tasks, the `at` command is used to schedule a one-time task at a specific time. The `batch` command is used to schedule a one-time task to be executed when the systems load average drops below 0.8.

To use `at` or `batch` you must have the `at` RPM package installed, and the `atd` service must be running. To determine if the package is installed, use the `rpm -q at` command. To determine if the service is running, use the command `/sbin/service atd status`.

27.3.1. Configuring At Jobs

To schedule a one-time job at a specific time, type the command `at time`, where *time* is the time to execute the command.

The argument *time* can be one of the following:

- `HH:MM` format — For example, `04:00` specifies 4:00AM. If the time is already past, it is executed at the specified time the next day.
- `midnight` — Specifies 12:00AM.
- `noon` — Specifies 12:00PM.
- `teatime` — Specifies 4:00PM.
- `month-name day year` format — For example, `January 15 2002` specifies the 15th day of January in the year 2002. The year is optional.
- `MMDDYY`, `MM/DD/YY`, or `MM.DD.YY` formats — For example, `011502` for the 15th day of January in the year 2002.
- `now + time` — *time* is in minutes, hours, days, or weeks. For example, `now + 5 days` specifies that the command should be executed at the same time in five days.

The time must be specified first, followed by the optional date. For more information about the time format, read the `/usr/share/doc/at-<version>/timespec` text file.

After typing the `at` command with the time argument, the `at>` prompt is displayed. Type the command to execute, press [Enter], and type Ctrl-D. More than one command can be specified by typing each command followed by the [Enter] key. After typing all the commands, press [Enter] to go to a blank

line and type Ctrl-D. Alternatively, a shell script can be entered at the prompt, pressing [Enter] after each line in the script, and typing Ctrl-D on a blank line to exit. If a script is entered, the shell used is the shell set in the user's SHELL environment, the user's login shell, or `/bin/sh` (whichever is found first).

If the set of commands or script tries to display information to standard out, the output is emailed to the user.

Use the command `atq` to view pending jobs. Refer to Section 27.3.3 for more information.

Usage of the `at` command can be restricted. Refer to Section 27.3.5 for details.

27.3.2. Configuring Batch Jobs

To execute a one-time task when the load average is below 0.8, use the `batch` command.

After typing the `batch` command, the `at>` prompt is displayed. Type the command to execute, press [Enter], and type Ctrl-D. More than one command can be specified by typing each command followed by the [Enter] key. After typing all the commands, press [Enter] to go to a blank line and type Ctrl-D. Alternatively, a shell script can be entered at the prompt, pressing [Enter] after each line in the script, and typing Ctrl-D on a blank line to exit. If a script is entered, the shell used is the shell set in the user's SHELL environment, the user's login shell, or `/bin/sh` (whichever is found first). As soon as the load average is below 0.8, the set of commands or script is executed.

If the set of commands or script tries to display information to standard out, the output is emailed to the user.

Use the command `atq` to view pending jobs. Refer to Section 27.3.3 for more information.

Usage of the `batch` command can be restricted. Refer to Section 27.3.5 for details.

27.3.3. Viewing Pending Jobs

To view pending `at` and `batch` jobs, use the `atq` command. It displays a list of pending jobs, with each job on a line. Each line is in the format job number, date, hour, job class, and username. Users can only view their own jobs. If the root user executes the `atq` command, all jobs for all users are displayed.

27.3.4. Additional Command Line Options

Additional command line options for `at` and `batch` include:

Option	Description
<code>-f</code>	Read the commands or shell script from a file instead of specifying them at the prompt.
<code>-m</code>	Send email to the user when the job has been completed.
<code>-v</code>	Display the time that the job will be executed.

Table 27-1. `at` and `batch` Command Line Options

27.3.5. Controlling Access to At and Batch

The `/etc/at.allow` and `/etc/at.deny` files can be used to restrict access to the `at` and `batch`

commands. The format of both access control files is one username on each line. Whitespace is not permitted in either file. The `at` daemon (`atd`) does not have to be restarted if the access control files are modified. The access control files are read each time a user tries to execute the `at` or `batch` commands.

The root user can always execute `at` and `batch` commands, regardless of the access control files.

If the file `at.allow` exists, only users listed in it are allowed to use `at` or `batch`, and the `at.deny` file is ignored.

If `at.allow` does not exist, all users listed in `at.deny` are not allowed to use `at` or `batch`.

27.3.6. Starting and Stopping the Service

To start the `at` service, use the command `/sbin/service atd start`. To stop the service, use the command `/sbin/service atd stop`. It is recommended that you start the service at boot time. Refer to Chapter 13 for details on starting the cron service automatically at boot time.

27.4. Additional Resources

To learn more about configuring automated tasks, refer to the following resources.

27.4.1. Installed Documentation

- `cron` man page — overview of `cron`.
- `crontab` man pages in sections 1 and 5 — The man page in section 1 contains an overview of the `crontab` file. The man page in section 5 contains the format for the file and some example entries.
- `/usr/share/doc/at-<version>/timespec` contains more detailed information about the times that can be specified for cron jobs.
- `anacron` man page — description of `anacron` and its command line options.
- `anacrontab` man page — brief overview of the `anacron` configuration file.
- `/usr/share/doc/anacron-<version>/README` describes `Anacron` and why it is useful.
- `at` man page — description of `at` and `batch` and their command line options.

Log Files

Log files are files that contain messages about the system, including the kernel, services, and applications running on it. There are different log files for different information. For example, there is a default system log file, a log file just for security messages, and a log file for cron tasks.

Log files can be very useful if you are trying to troubleshoot a problem with the system such as trying to load a kernel driver or if you are looking for unauthorized log in attempts to the system. This chapter discusses where to find log files, how to view log files, and what to look for in log files.

Some log files are controlled by a daemon called `syslogd`. A list of log messages maintained by `syslogd` can be found in the `/etc/syslog.conf` configuration file.

28.1. Locating Log Files

Most log files are located in the `/var/log` directory. Some applications such as `httpd` and `samba` have a directory within `/var/log` for their log files.

Notice the multiple files in the log file directory with numbers after them. These are created when the log files are rotated. Log files are rotated so their file sizes do not become too large. The `logrotate` package contains a cron task that automatically rotates log files according to the `/etc/logrotate.conf` configuration file and the configuration files in the `/etc/logrotate.d` directory. By default, it is configured to rotate every week and keep four weeks worth of previous log files.

28.2. Viewing Log Files

Most log files are in plain text format. You can view them with any text editor such as **Vi** or **Emacs**. Some log files are readable by all users on the system; however, you must be logged in as root to read most log files.

To view system log files in an interactive, real-time application, use the **Log Viewer**. To start the application, go to the **Main Menu Button** (on the Panel) => **System Tools** => **System Logs**, or type the command `redhat-logviewer` at a shell prompt.

The application only displays log files that exist; thus, your list might differ from the one shown in Figure 28-1. To view the complete list of log files that it can view, refer to the configuration file, `/etc/sysconfig/redhat-logviewer`.

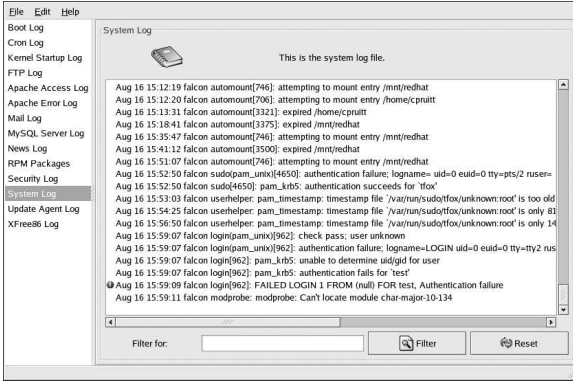


Figure 28-1. Log Viewer

By default, the currently viewable log file is refreshed every 30 seconds. To change the refresh rate, select **Edit => Preferences** from the pull-down menu. The window shown in Figure 28-2 will appear. In the **Log Files** tab, click the up and down arrows beside the refresh rate to change it. Click **Close** to return to the main window. The refresh rate is changed immediately. To refresh the currently viewable file manually, select **File => Refresh Now** or press [Ctrl]-[R].

To filter the contents of the log file for keywords, type the word or words you are looking for the **Filter for** text field, and click **Filter**. Click **Reset** to reset the contents.

You can also change where the application looks for the log files from the **Log Files** tab. Select the log file from the list, and click the **Change Location** button. Type the new location of the log file or click the **Browse** button to locate the file location using a file selection dialog. Click **OK** to return to the preferences, and click **Close** to return to the main window.

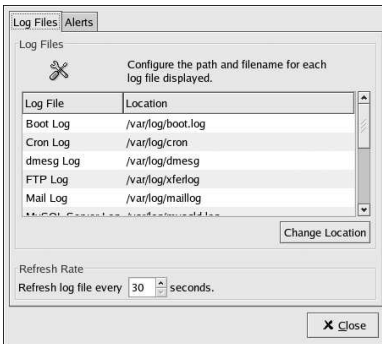


Figure 28-2. Log File Locations

28.3. Examining Log Files

Log Viewer can be configured to display an alert icon beside lines that contain key alert words. To add alerts words, select **Edit => Preferences** from the pull-down menu, and click on the **Alerts** tab.

Click the **Add** button to add an alert word. To delete an alert word, select the word from the list, and click **Delete**.

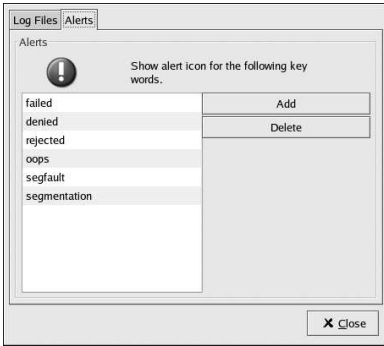


Figure 28-3. Alerts

Upgrading the Kernel

The Red Hat Linux kernel is custom built by the Red Hat kernel team to ensure its integrity and compatibility with supported hardware. Before Red Hat releases a kernel, it must pass a rigorous set of quality assurance tests.

Official Red Hat Linux kernels are packaged in RPM format so that they are easy to upgrade and verify. For example, the `kernel` RPM package creates the `initrd` image; it is not necessary to use the `mkinitrd` command after installing a different kernel if you install the kernel from the Red Hat RPM package. It also modifies the boot loader configuration file to include the new kernel if either GRUB or LILO is installed. (You do need to set the new kernel as the default kernel to boot.)

This chapter discusses the steps necessary to upgrade the kernel on an x86 system only.



Warning

Building your own custom kernel is not supported by the Red Hat Linux Installation Support Team. For more information on building a custom kernel from the source code, refer to Appendix A.

29.1. The 2.4 Kernel

Red Hat Linux ships with a custom 2.4 kernel, which offers the following features:

- The directory for the kernel source is `/usr/src/linux-2.4` instead of `/usr/src/linux`.
- Support for the ext3 filesystem.
- Multi-processor (SMP) support.
- USB support.
- Preliminary support for IEEE 1394, also referred to as FireWire™, devices.

29.2. Preparing to Upgrade

Before you upgrade your kernel, you must take a few precautionary steps. The first step is to make sure you have a working boot diskette for your system in case a problem occurs. If the boot loader is not configured properly to boot the new kernel, you will not be able to boot your system unless you have a boot diskette.

To create a boot diskette for your system, you need to determine which version of the kernel you are currently running. Execute the following command:

```
uname -r
```

You must be root to create a boot diskette for your system. Login as root at a shell prompt, and type the following command (where `kernelversion` is the output of the `uname -r` command):

```
/sbin/mkbootdisk kernelversion
```

**Tip**

Refer to the man page for `mkbootdisk` for more options.

Reboot your machine with the boot diskette and verify that it works before continuing.

Hopefully, you will not have to use the diskette, but you should store it in a safe place just in case.

To determine which kernel packages you have installed, execute the following command at a shell prompt:

```
rpm -qa | grep kernel
```

The output will contain some or all of the following packages, depending on what type of installation you performed (your version numbers and packages may differ):

```
kernel-2.4.18-7.95
kernel-debug-2.4.18-7.95
kernel-source-2.4.18-7.95
kernel-doc-2.4.18-7.95
kernel-pcmcia-cs-3.1.27-12
kernel-smp-2.4.18-7.95
```

From the output, you can determine which packages you need to download for the kernel upgrade. For a single processor system, the only required package is the `kernel` package.

If you have a computer with more than one processor, you need the `kernel-smp` package that contains support for multiple processors. It is recommended that you also install the `kernel` package in case the multi-processor kernel does not work properly for your system.

If you have a computer with more than four gigabytes of memory, you need the `kernel-bigmem` package. Again, it is recommended that you also install the `kernel` package for debugging purposes. The `kernel-bigmem` package is only built for the i686 architecture.

If you are upgrading the kernel on a laptop or are using PCMCIA, the `kernel-pcmcia-cs` package is also required.

You do not need the `kernel-source` package unless you plan to recompile the kernel yourself or plan to perform kernel development. The `kernel-doc` package contains kernel development documentation and is not required. The `kernel-util` package includes utilities that can be used to control the kernel or the system's hardware and is not required.

Red Hat builds kernels that are optimized for different x86 versions. The options are `athlon` for AMD Athlon™ and AMD Duron™ systems, `i686` for Intel® Pentium® II, Intel® Pentium® III, and Intel® Pentium® 4 systems, and `i586` for Intel® Pentium® and AMD K6™ systems. If you do not know the version of your x86 system, use the kernel built for the i386 version; it is built for all x86-based systems.

The x86 version of the RPM package is included in the file name. For example, `kernel-2.4.18-7.95.athlon.rpm` is optimized for AMD Athlon™ and AMD Duron™ systems and `kernel-2.4.18-7.95.i686.rpm` is optimized for Intel® Pentium® II, Intel® Pentium® III, and Intel® Pentium® 4 systems. When you have determined which packages you need to upgrade your kernel, select the proper architecture for the `kernel`, `kernel-smp`, and `kernel-bigmem` packages. Use the i386 versions of the other packages.

29.3. Downloading the Upgraded Kernel

There are several ways to determine if there is an updated kernel available for your system.

- Go to <http://www.redhat.com/apps/support/errata/>, choose the version of Red Hat Linux you are using, and view the errata for it. Kernel errata are usually under the **Security Advisories** section. From the list of errata, click the kernel errata to view the detailed errata report for it. In the errata report, there is a list of required RPM packages and a link to download them from the Red Hat FTP site. You can also download them from a Red Hat FTP mirror site. A list of mirror sites is available at <http://www.redhat.com/download/mirror.html>.
- Use Red Hat Network. You can use Red Hat Network to download the kernel RPM packages and then manually upgrade to the latest kernel. Or, if you have elected to let the **Red Hat Update Agent** upgrade packages for you, Red Hat Network can download the latest kernel, upgrade the kernel on your system, create an initial RAM disk if needed, and configure the boot loader to boot the new kernel. All you have to do is reboot into the new kernel. For more information, refer to the *Red Hat Network User Reference Guide* available at <http://www.redhat.com/docs/manuals/RHNetwork/>.

If there is an updated kernel for the version of Red Hat Linux you are running, download the appropriate packages using one of these methods. If you used Red Hat Network to upgrade your kernel automatically, you are finished — just reboot your system to use the new kernel. If you just downloaded the RPM packages from the Red Hat Linux errata page or from Red Hat Network, proceed to Section 29.4.

29.4. Performing the Upgrade

Now that you have the necessary kernel RPM packages, you can upgrade your existing kernel. At a shell prompt as root, change to the directory that contains the kernel RPM packages and follow these steps.



Important

It is strongly recommended that you keep the old kernel in case you have problems with the new kernel.

Use the `-i` argument with the `rpm` command if you want to keep the old kernel. If you use the `-U` option to upgrade the `kernel` package, it will overwrite the currently installed kernel (the kernel version and x86 version might vary):

```
rpm -ivh kernel-2.4.18-7.95.i386.rpm
```

If the system is a multi-processor system, install the `kernel-smp` packages as well (the kernel version and x86 version might vary):

```
rpm -ivh kernel-smp-2.4.18-7.95.i386.rpm
```

If the system is i686-based and contains more than 4 gigabytes of RAM, install the `kernel-bigmem` package built for the i686 architecture as well (the kernel version might vary):

```
rpm -ivh kernel-bigmem-2.4.18-7.95.i686.rpm
```

If you plan to upgrade the `kernel-source`, `kernel-docs`, or `kernel-utils` packages, you probably do not need to keep the older versions. Use the following commands to upgrade these packages (the versions might vary):

```
rpm -Uvh kernel-source-2.4.18-7.95.i386.rpm
rpm -Uvh kernel-docs-2.4.18-7.95.i386.rpm
rpm -Uvh kernel-utils-2.4.18-7.95.i386.rpm
```

If you are using PCMCIA (for example, a laptop), you also need to install the `kernel-pcmcia-cs` and keep the old version. If you use the `-i` switch, it will probably return a conflict because the older kernel needs this package to boot with PCMCIA support. To work around this, use the `--force` switch as follows (the version might vary):

```
rpm -ivh --force kernel-pcmcia-cs-3.1.24-2.i386.rpm
```

If you are using the ext3 file system or a SCSI controller, you need an initial RAM disk. The purpose of the initial RAM disk is to allow a modular kernel to have access to modules that it might need to boot from before the kernel has access to the device where the modules normally reside.

The initial RAM disk is created by using the `mkinitrd` command. However, the Red Hat `kernel RPM` package performs this step for you. To verify that it was created, use the command `ls -l /boot`. You should see the file `initrd-2.4.18-7.95.img` (the version should match the version of the kernel you just installed).

Now that you have installed the new kernel, you need to configure the boot loader to boot the new kernel. Refer to Section 29.5 for details.

29.5. Configuring the Boot Loader

The `kernel RPM` package configures the GRUB or LILO boot loader to boot the newly installed kernel if either boot loader is installed. However, it does not configure the boot loader to boot the new kernel by default.

It is always a good idea to confirm that the boot loader has been configured correctly. This is a crucial step. If the boot loader is configured incorrectly, you will not be able to boot your system. If this happens, boot your system with the boot diskette you created earlier and try configuring the boot loader again.

29.5.1. GRUB

If you selected GRUB as your boot loader, confirm that the file `/boot/grub/grub.conf` contains a `title` section with the same version as the `kernel` package you just installed (if you installed the `kernel-smp` and/or `kernel-bigmem`, you will have a section for it as well):

```
# NOTICE: You have a /boot partition. This means that
#           all kernel paths are relative to /boot/
default=0
timeout=30
splashimage=(hd0,0)/grub/splash.xpm.gz
title Red Hat Linux (2.4.18-7.95)
    root (hd0,0)
    kernel /vmlinuz-2.4.18-7.95 ro root=/dev/hda3
    initrd /initrd-2.4.18-7.95.img
```

If you created a separate `/boot` partition, the paths to the kernel and `initrd` image are relative to the `/boot` partition.

To configure GRUB to boot the new kernel by default, change the value of the `default` variable to the title section number for the title section that contains the new kernel. The count starts with 0. For example, if the new kernel is the second title section, set `default` to `1`.

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

29.5.2. LILO

If you selected LILO as your boot loader, confirm that the file `/etc/lilo.conf` contains an `image` section with the same version as the `kernel` package you just installed:

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

image=/boot/vmlinuz-2.4.18-7.95
  label=linux
    initrd=initrd-2.4.18-7.95.img
  read-only
  root=/dev/hda5
```

The `kernel RPM` package activates the changes by running the command `/sbin/lilo`.

To configure LILO to boot the new kernel by default, set the `default` variable to the value of `label` in the `image` section for the new kernel. You must run the `/sbin/lilo` command as `root` to enable the changes. After running it, you will see output similar to the following:

```
Added linux *
```

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

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

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.

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 modules need to be loaded at boot time. 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** runs when the system boots and 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 the 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.

30.1. Kernel Module Utilities

You can also use a group of commands to list, load, or unload kernel modules. These commands are useful if you want to try different modules or see if a module has been loaded successfully.

The command `/sbin/lsmmod` displays a list of currently loaded modules.

Module	Size	Used by
sr_mod	15264	0 (autoclean)
mga	95984	1
agpgart	23392	3
nfs	79008	1 (autoclean)
lockd	52464	1 (autoclean) [nfs]
sunrpc	61328	1 (autoclean) [nfs lockd]
autofs	11264	4 (autoclean)
3c59x	25344	1 (autoclean)
ipchains	38976	0 (unused)
ide-scsi	8352	0
scsi_mod	95104	2 [sr_mod ide-scsi]
ide-cd	26848	0
cdrom	27232	0 [sr_mod ide-cd]
usb-uhci	20720	0 (unused)
usbcore	49664	1 [usb-uhci]

Example 30-1. Example `lsmmod` output

As you can see in Example 30-1, `lsmod` displays the size, use count, and referring modules for each module currently loaded.

To load a kernel module, you can use the command `/sbin/insmod` followed by the kernel module name. By default, `insmod` tries to load the module from the `/lib/modules/<kernel-version>/kernel/drivers` subdirectories. There is a subdirectory for each type of module, such as the `net` subdirectory for network interface drivers. Some kernel modules have module dependencies — other modules must be loaded first for it to load. To resolve these dependencies, you can either load the module dependencies and then load the module you want, or you can use the command `/sbin/modprobe` followed by the module name to load the module along with its dependencies.

For example, the command

```
/sbin/modprobe tulip
```

loads the `tulip` network interface module.

To unload kernel modules, use the command `/sbin/rmmod` followed by the module name. The `rmmod` utility will only unload modules that are not in use and that are not a dependency of other modules in use.

For example, the command

```
/sbin/rmmod tulip
```

unloads the `tulip` network interface module.

Another useful kernel module utility is `modinfo`. You can use the command `/sbin/modinfo` to display information about a kernel module. The general syntax is:

```
/sbin/modinfo [options] <module>
```

Options include `-d` that displays a brief description of the module and `-p` that lists the parameters the module supports. For a complete list of options, refer to the `modinfo` man page (`man modinfo`).

30.2. Additional Resources

For more information on kernel modules and their utilities, refer to the following resources.

30.2.1. Installed Documentation

- `lsmod` man page — description and explanation of its output.
- `insmod` man page — description and list of command line options.
- `modprobe` man page — description and list of command line options.
- `rmmod` man page — description and list of command line options.
- `modinfo` man page — description and list of command line options.
- `/usr/src/linux-2.4/Documentation/modules.txt` — how to compile and use kernel modules.

Package Management

Package Management with RPM

The RPM 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 **Package Management Tool** to perform many RPM commands. Refer to Chapter 32 for details.

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 between "pristine" sources and your patches along with build instructions eases the maintenance of the package as new versions of the software are released.



Note

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

31.1. 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 needed 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.

31.2. 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 31.5 for more information on RPM.

31.2.1. Finding RPM Packages

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/apps/support/errata/>
- A Red Hat FTP Mirror Site available at <http://www.redhat.com/download/mirror.html>
- Red Hat Network — See Chapter 33 for more details on Red Hat Network

31.2.2. 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 logging in as root and typing the following command at a shell prompt:

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

If installation is successful, you will see the following:

```
Preparing...                               ##### [100%]
 1:foo                                       ##### [100%]
```

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.

Starting with version 4.1 of RPM, the signature of a package is checked when installing or upgrading a package. If verifying the signature fails, you will see an error message such as:

```
error: V3 DSA signature: BAD, key ID 0352860f
```

If it is a new, header-only, signature, you will see an error message such as:

```
error: Header V3 DSA signature: BAD, key ID 0352860f
```

If you do not have the appropriate key installed to verify the signature, the message will contain `NOKEY` such as:

```
warning: V3 DSA signature: NOKEY, key ID 0352860f
```

Refer to Section 31.3 for more information on checking a package's signature.



Note

If you are installing a kernel package, you should use `rpm -ivh` instead. Refer to Chapter 29 for details.

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

31.2.2.1. Package Already Installed

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

```
Preparing...                               ##### [100%]
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 `--replacepks` option, which tells RPM to ignore the error:

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

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.

31.2.2.2. 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 will see:

```
Preparing...                               ##### [100%]
file /usr/bin/foo from install of foo-1.0-1 conflicts with file from package bar-2.0.20
```

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

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

31.2.2.3. 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 will see:

```
Preparing...                               ##### [100%]
error: Failed dependencies:
    bar.so.2 is needed by foo-1.0-1
Suggested resolutions:
    bar-2.0.20-3.i386.rpm
```

If you are installing an official Red Hat, it will usually suggest the package(s) need to resolve the dependency. Find this package on the Red Hat Linux CD-ROMs or from the Red Hat FTP site (or mirror), and add it to the command:

```
rpm -ivh foo-1.0-1.i386.rpm bar-2.0.20-3.i386.rpm
```

If installation of both packages is successful, you will see:

```
Preparing... ##### [100%]
 1:foo ##### [ 50%]
 2:bar ##### [100%]
```

If it does not suggest a package to resolve the dependency, you can try the `--redhatprovides` option to determine which package contains the required file. You need the `rpmdb-redhat` package installed to use this options.

```
rpm -q --redhatprovides bar.so.2
```

If the package that contains `bar.so.2` is in the installed database from the `rpmdb-redhat` package, the name of the package will be displayed:

```
bar-2.0.20-3.i386.rpm
```

If you want to force the installation anyway (a bad idea since the package probably will not run correctly), use the `--nodeps` option.

31.2.3. 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:

```
Preparing... ##### [100%]
error: removing these packages would break dependencies:
       foo is needed by bar-2.0.20-3.i386.rpm
```

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.

31.2.4. 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
```

What you do not see above is that RPM automatically uninstalls 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:

```
package foo-2.0-1 (which is newer than foo-1.0-1) is already installed
```

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

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

31.2.5. 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
```

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 delete 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.

31.2.6. 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`:

```
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/ls`).
- `-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.

31.2.7. 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:

- `S` — 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.

31.3. 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 (`<rpm-file>` with filename of the RPM package):

```
rpm -K --nogpg <rpm-file>
```

You will see the message `<rpm-file>: md5 OK`. This brief message means that the file was not corrupted by the download. To see a more verbose message, replace `-K` with `-Kvv` in the command.

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 will 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.

31.3.1. Importing Keys

To verify official Red Hat packages, you must import the Red Hat GPG key. To do so, execute the following command at a shell prompt:

```
rpm --import /usr/share/rhn/RPM-GPG-KEY
```

To display a list of all keys installed for RPM verification, execute the command:

```
rpm -qa gpg-pubkey*
```

For the Red Hat key, the output will include:

```
gpg-pubkey-db42a60e-37ea5438
```

To display details about a specific key, use `rpm -qi` followed by the output from the previous command:

```
rpm -qi gpg-pubkey-db42a60e-37ea5438
```

31.3.2. Verifying Signature of Packages

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

```
rpm -K <rpm-file>
```

If all goes well, you will see the message: `md5 gpg OK`. That means that the signature of the package has been verified and that it is not corrupt.



Tip

For more information about GnuPG, see Appendix B.

31.4. 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-3.5.8-18
```

- 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 do not know what it does. To find information about it, use the following command:

```
rpm -qip sndconfig-0.68-3.i386.rpm
```

The output would look like the following:

```
Name       : sndconfig                Relocations: (not relocateable)
Version    : 0.68                      Vendor: Red Hat
Release    : 3                        Build Date: Sun 23 Jun 2002 08:22:52 PM EDT
Install date: Mon 01 Jul 2002 08:40:06 AM EDT  Build Host: perf90.perf.redhat.com
Group      : Applications/Multimedia    Source RPM: sndconfig-0.68-3.src.rpm
Size       : 619097                    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 that sets up the configuration files
you will need to use a sound card with a Red Hat Linux system.
Sndconfig can be used to set the proper sound type for programs that
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-0.68-3.i386.rpm
```

The output will look like the following:

```
/sbin/sndconfig
/usr/sbin/sndconfig
/usr/share/locale/bs/LC_MESSAGES/sndconfig.mo
/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/eu_ES/LC_MESSAGES/sndconfig.mo
/usr/share/locale/fi/LC_MESSAGES/sndconfig.mo
/usr/share/locale/fr/LC_MESSAGES/sndconfig.mo
/usr/share/locale/gl/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/ja/LC_MESSAGES/sndconfig.mo
/usr/share/locale/ko/LC_MESSAGES/sndconfig.mo
/usr/share/locale/no/LC_MESSAGES/sndconfig.mo
/usr/share/locale/pl/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/locale/wa/LC_MESSAGES/sndconfig.mo
/usr/share/locale/zh/LC_MESSAGES/sndconfig.mo
/usr/share/locale/zh_CN.GB2312/LC_MESSAGES/sndconfig.mo
/usr/share/locale/zh_TW.Big5/LC_MESSAGES/sndconfig.mo
/usr/share/locale/zh_TW/LC_MESSAGES/sndconfig.mo
/usr/share/man/man8/sndconfig.8.gz
/usr/share/sndconfig
/usr/share/sndconfig/sample.au
/usr/share/sndconfig/sample.midi
/usr/share/sndconfig/sample2.au
```

These are just a few examples. As you use it, you will find many more uses for RPM.

31.5. 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.

31.5.1. 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.

31.5.2. Useful Websites

- <http://www.rpm.org/> — The RPM website.
- <http://www.redhat.com/mailling-lists/rpm-list/> — 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.

31.5.3. 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/docs/books/>.

Package Management Tool

During installation, users select an installation type such as **Workstation** or **Server**. Software packages are installed based on this selection. Because people use their computers differently, users might want to install or remove packages after installation. The **Package Management Tool** allows users to perform these actions.

To start the application, go to the **Main Menu Button** (on the Panel) => **System Settings** => **Packages**, or type the command `redhat-config-packages` at shell prompt.

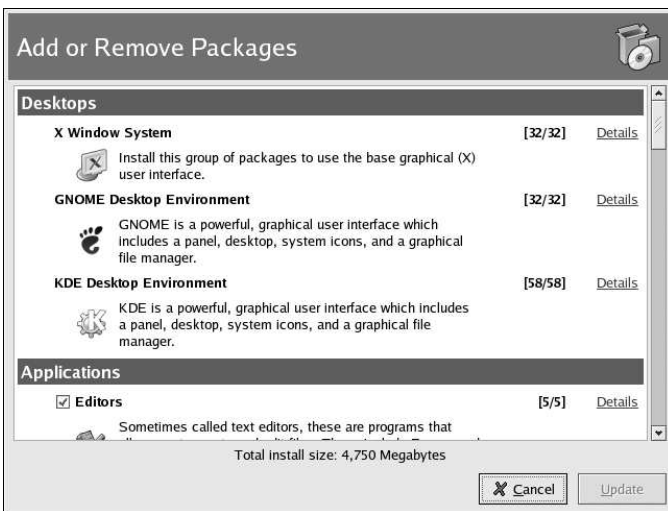


Figure 32-1. Package Management Tool

The interface for this application is similar to the one used during installation. Packages are divided into package groups, which contain a list of *standard packages* and *extra packages* that share common functionality. For example, the **Graphical Internet** group contains a Web browser, email client, and other graphical programs used to connected to the Internet. The standard packages can not be selected for removal unless the entire package group is removed. The extra packages are optional packages that can be selected for installation or removal, as long as the package group is selected.

The main window shows a list of package groups. If the package group has a checkmark in the checkbox beside it, packages from that group are currently installed. To view the individual packages list for a group, click the **Details** button beside it. The individual packages with a checkmark beside them are currently installed.

32.1. Installing Packages

To install the standard packages in a package group that is not currently installed, check the checkbox beside it. To customize the packages to be installed within the group, click the **Details** button beside it. The list of standard and extra packages is displayed, as shown in Figure 32-2. Clicking on the package

name displays the disk space required to install the package at the bottom of the window. Checking the checkbox beside the package name marks it for installation.

You can also select individual packages from already installed package groups by click the **Details** button and checking any of the extra packages not already installed.

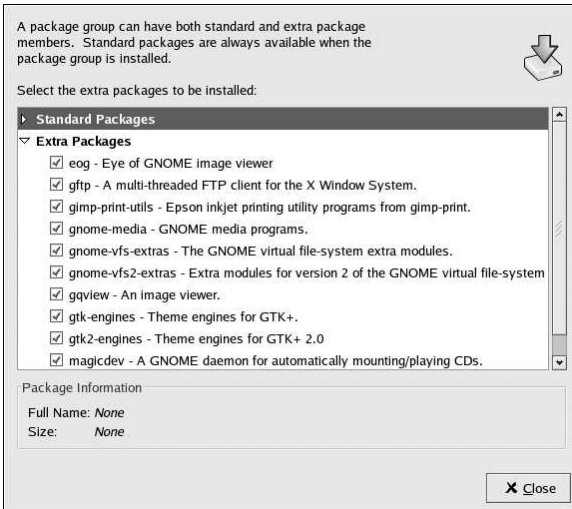


Figure 32-2. Individual Package Selection

After selecting package groups and individual packages to install, click the **Update** button on the main window. The application will then compute the amount of disk space required to install the packages as well as any package dependencies and display a summary window. If there are package dependencies, they will be automatically added to the list of packages to install. Click the **Show Details** button to view the complete list of packages to be installed.




Figure 32-3. Package Installation Summary

Click **Continue** to start the installation process. When it is finished, an **Update Complete** message will appear.

**Tip**

If you use **Nautilus** to browse the files and directories on your computer, you can also use it to install packages. In **Nautilus**, go to the directory that contains an RPM package (they usually end in `.rpm`), and double-click on the RPM icon.

32.2. Removing Packages

To remove all the package installed within a package group, uncheck the checkbox beside it. To remove individual packages, click the **Details** button beside the package group and uncheck the individual packages.

When you are finished selecting packages to remove, click the **Update** button in the main window. The application computes the amount of disk space that will be freed as well as the software package dependencies. If other packages depend on the packages you selected to remove, they will be automatically added to the list of packages to be removed. Click the **Show Details** button to view the list of packages to be removed.

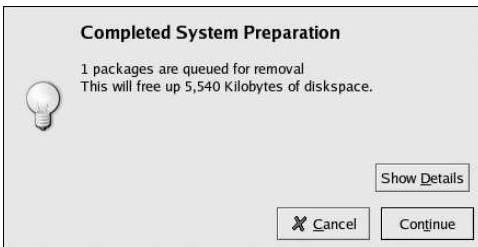


Figure 32-4. Package Removal Summary

Click **Continue** to start the removal process. When it is finished, an **Update Complete** message will appear.

**Tip**

You can combine the installation and removal of packages by selecting package groups/packages to be installed/removed and then clicking **Update**. The **Completed System Preparation** window will display the number of packages to be installed and removed.

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 the RHN website available at <http://rhn.redhat.com/>.

Red Hat Network saves Red Hat Linux users time because they receive email when updated packages are released. Users do not have to search the Web for updated packages or security alerts. By default, Red Hat Network installs the packages as well. Users do not have to learn how to use RPM or worry about resolving software package dependencies. RHN does it all.

Each Red Hat Network account comes with:

- Errata Alerts — Learn when Security Alerts, Bug Fix Alerts, and Enhancement Alerts are issued for all the systems in your network through the RHN website
- Automatic email notifications — Receive an email notification when an Errata Alert is issued for your system
- Scheduled Errata Updates — Schedule delivery of Errata Updates with optional automatic installation
- Package installation — Schedule package installation on one or more systems with the click of a button
- **Red Hat Update Agent** — Use the **Red Hat Update Agent** to download the latest software packages for your system with optional package installation
- Red Hat Network website — Manage multiple systems, downloaded individual packages, and schedule actions such as Errata Updates.

To start using Red Hat Network, follow these three basic steps:

1. Create a System Profile using one of the following methods:
 - Registering the system with RHN during the **Setup Agent** the first time your system boots after installation.
 - Select **Main Menu Button => System Tools => Red Hat Network** on your desktop.
 - Execute the command `up2date` from a shell prompt.
2. Log in to RHN at <http://rhn.redhat.com/> and entitle the system to a service offering. Everyone receives a free Red Hat Network account for one system. Additional accounts can be purchased.
3. Start scheduling updates through the RHN website or download and install Errata Updates with the **Red Hat Update Agent**.

For more detailed instructions, read the *Red Hat Network User Reference Guide* available at <http://www.redhat.com/docs/manuals/RHNetwork/>.

Appendixes

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 Chapter 30. Unless you are recompiling a customized kernel for your system, you will not see many changes in how kernel modules are handled.

A.1. 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 A.4 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 2.4.18-7.95 as the kernel version. Your kernel version might differ. To determine your kernel version, type the command `uname -r`. Replace 2.4.18-7.95 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 did not 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 2.4.18-7.95). Once done, test the boot disk to make sure that it will boot the system.
2. You must have the `kernel-source` package installed. Issue the command `rpm -q kernel-source` to determine the package version, if it is installed. If it is not installed, install them from one of the Red Hat Linux CD-ROMs 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 31 for information on installing RPM packages.
3. Open a shell prompt and change to the directory `/usr/src/linux-2.4`. All commands from this point forward must be executed 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 an existing configuration file that works (`/usr/src/linux-2.4/.config`) and you want to use, back it up to a different directory before running this command and copy it back afterward.

- Now you need 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.

If you want to use the settings of a default Red Hat Linux kernel, copy the the configuration file from the `/usr/src/linux-2.4/configs` directory to `/usr/src/linux-2.4/.config`. Then, run the `make xconfig` command and only make the desired changes. Be sure to save your changes to the configuration file.

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 are using the default Red Hat Linux 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 do not want.



Note

To use `kmod` (see Chapter 30 for details) and kernel modules you must answer **Yes** to `kmod` support and `module version` (`CONFIG_MODVERSIONS`) support during the configuration.

- After creating a `/usr/src/linux-2.4/.config` file, use the command `make dep` to set up all the dependencies correctly.
- Use the command `make clean` to prepare the source tree for the build.
- It is recommended that you give the custom kernel you are building a modified version number 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`.

By default, `/usr/src/linux-2.4/Makefile` includes the word `custom` at the end of the line beginning with `EXTRAVERSION`. Appending the string will allow you to have the old working kernel and the new kernel, version 2.4.18-7.95custom, on your system at the same time.

To give the kernel an "unique" name, you can also append the date to the end of the string.
- Build the kernel with `make bzImage`.
- Build any modules you configured with `make modules`.
- Use the command `make modules_install` to install the kernel modules (even if you did not build any). Make sure that you type the underscore (`_`). This will install the kernel modules into the directory path `/lib/modules/KERNELVERSION/kernel/drivers` (where

`KERNELVERSION` is the version specified in the `Makefile`). In the example it would be `/lib/modules/2.4.18-7.95custom/kernel/drivers/`.

- Use `make install` to copy your new kernel and its associated files to the proper directories.

In addition to installing the kernel files in the `/boot` directory, this command also executes the `/sbin/new-kernel-pkg` script that builds a new `initrd` image and adds new entries to the boot loader configuration file.

If you have a SCSI adapter and you compiled the SCSI driver as a module or if you built your kernel with `ext3` support as a module (the default in Red Hat Linux), the `initrd` image is required.

- Even though the `initrd` image and boot loader modifications are made for you, you should verify that they were done correctly. Refer to Section A.2 and Section A.3 for details.

A.2. Making an `initrd` Image

An `initrd` image is needed to load a SCSI module at boot time or if you compiled the kernel with `ext3` support as a module. To verify that a new `initrd` file was created, view the contents of the `/boot` directory. You should see the file `initrd-2.4.18-7.95custom.img`, where `2.4.18-7.95custom` is the name of the kernel you just built.

If it does not exist, use the `/sbin/mkinitrd` shell script to create it:

```
/sbin/mkinitrd /boot/initrd-2.4.18-7.95custom.img 2.4.18-7.95custom
```

In the above example, `/boot/initrd-2.4.18-7.95custom.img` is the file name of the new `initrd` image. `2.4.18-7.95custom` 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.

A.3. Configuring the Boot Loader

Making sure the boot loader configuration file has been correctly modified is a crucial step. If the file is modified incorrectly, you may not be able to boot your system. If this happens, boot your system with the boot diskette you created earlier and try configuring the boot loader again. If your boot diskette does not work, refer to Chapter 8 for more information about rescue mode.

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. During the installation of Red Hat Linux 8.0, you had the option to choose either GRUB or LILO as your boot loader. Refer to the appropriate section that follows.

A.3.1. GRUB

If you selected GRUB as your boot loader, the `new-kernel-pkg` script should have modified `/boot/grub/grub.conf` to include a section for the new kernel.

The default GRUB configuration file looks similar to the following:

```
# NOTICE: You have a /boot partition. This means that
#           all kernel paths are relative to /boot/
default=0
timeout=30
splashimage=(hd0,0)/grub/splash.xpm.gz

title Red Hat Linux (2.4.18-7.95)
    root (hd0,0)
    kernel /vmlinuz-2.4.18-7.95 ro root=/dev/hda3
```

```
initrd /initrd-2.4.18-7.95.img
```

If you created a separate `/boot` partition, the paths to the kernel and `initrd` image are relative to the `/boot` partition.

By default, Red Hat Linux uses Red Hat Linux and the kernel version in parentheses to differentiate between different kernels for GRUB to boot. In our example, the new `/boot/grub/grub.conf` file created by the `new-kernel-pkg` script would look like the following:

```
# NOTICE: You have a /boot partition. This means that
#           all kernel paths are relative to /boot/
default=1
timeout=30
splashimage=(hd0,0)/grub/splash.xpm.gz

title Red Hat Linux (2.4.18-7.95custom)
    root (hd0,0)
    kernel /vmlinuz-2.4.18-7.95custom ro root=/dev/hda3
    initrd /initrd-2.4.18-7.95custom.img

title Red Hat Linux (2.4.18-7.95)
    root (hd0,0)
    kernel /vmlinuz-2.4.18-7.95 ro root=/dev/hda3
    initrd /initrd-2.4.18-7.95.img
```

Notice that the default boot entry is set to 1. The script does not change the default kernel to boot; it only adds a new section for the new kernel.

If the file is not modified correctly and you did not receive any error messages from `make install`, add the new section manually.

After rebooting, selecting the new kernel from the list, and seeing that the new kernel works, you can make your new kernel the default. Either place its section first or change the default entry number to the appropriate number (remember that it starts counting with 0). For GRUB, you do not need to run any commands after modifying the configuration file.

A.3.2. LILO

If you selected LILO as your boot loader, the `new-kernel-pkg` script should have modified `/boot/lilo.conf` to include a section for the new kernel and run `/sbin/lilo`.

The default LILO configuration file looks similar to the following:

```
prompt
timeout=50
default=linux
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
message=/boot/message
lba32

image=/boot/vmlinuz-2.4.18-7.95
    label=linux
        initrd=initrd-2.4.18-7.95.img
    read-only
    append="root=LABEL=/"
```

The modified `/etc/lilo.conf` should look similar to the following:

```

prompt
timeout=50
default=linux
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
message=/boot/message
lba32

image=/boot/vmlinuz-2.4.18-7.95custom
    label=2.4.18-7.95custom
    initrd=initrd-2.4.18-7.95custom.img
    read-only
    append="root=LABEL=/"

image=/boot/vmlinuz-2.4.18-7.95
    label=linux
    initrd=initrd-2.4.18-7.95.img
    read-only
    append="root=LABEL=/"

```

If the file is not modified correctly and you did not receive any error messages from `make install`, add the new section manually.

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

```

LILO version 21.4-4, Copyright (C) 1992-1998 Werner Almesberger
'lba32' extensions Copyright (C) 1999,2000 John Coffman

Reading boot sector from /dev/hda
Merging with /boot/boot.b
Mapping message file /boot/message
Boot image: /boot/vmlinuz-2.4.18-7.95custom
Added 2.4.18-7.95custom *
Boot image: /boot/vmlinuz-2.4.18-7.95
Added linux
Writing boot sector.

```

Be sure the messages contains `Writing boot sector`. The `*` after `linux` means that the section labeled `linux` is the default kernel that LILO will boot.

From now on, when the system boots you will see the old and new entries.

To boot the new kernel, select it and press [Enter].

You can begin testing your new kernel by rebooting your computer and watching the messages to ensure your hardware is detected properly. If it is successful, you can change LILO to boot the new kernel by default by changing the label after `default=` in the configuration file and running the command `lilo -v`.

A.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` to include the line `append=nomodules` or append the kernel line in `grub.conf` with `nomodules`.

A.5. Additional Resources

For more information on the Linux kernel refer to these resources.

A.5.1. Installed Documentation

- `/usr/src/linux-2.4/Documentation` — Advanced documentation on the Linux kernel and its modules. These documents are written for people interested in contributing to the kernel source code and understanding how the kernel works.

A.5.2. Useful Websites

- <http://www.redhat.com/mirrors/LDP/HOWTO/Kernel-HOWTO.html> — *The Linux Kernel HOWTO* from the Linux Documentation Project
- <http://www.kernel.org/pub/linux/docs/lkml/> — The linux-kernel mailing list.

Getting Started with Gnu Privacy Guard

B.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 8.0 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.



Important

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 will 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 B.8.

B.2. Warning Messages

When executing GnuPG commands, you will probably see the message:

```
gpg: Warning: using insecure memory!
```

This warning is because non-root users can not lock memory pages. If users could lock memory pages, they could perform out-of-memory denial of service attacks; thus, it is a possible security problem. For details, refer to <http://www.gnupg.org/faq.html#q6.1>.

If you upgraded from a previous version of GnuPG, you might see the message:

```
gpg: WARNING: --honor-http-proxy is a deprecated option.
gpg: please use "--keyserver-options honor-http-proxy" instead
```

This warning is because your `~/.gnupg/options` file contains the line:

```
honor-http-proxy
```

Version 1.0.7 prefers a different syntax. Change the line to the following:

```
keyserver-options honor-http-proxy
```

B.3. 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 (not as root).

You will see an introductory screen, with key options, including one recommended option (the default), similar to the following:

```
gpg (GnuPG) 1.0.7; Copyright (C) 2002 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)
- (5) RSA (sign only)

```
Your selection?
```

In fact, most of the screens which require you to choose an option will list the default option, within parentheses. You can accept the default options by pressing [Enter].

In the first screen, you should accept the default option: (1) DSA and ElGamal. This option will allow you to create a digital signature and encrypt (and decrypt) with two types of technologies. Type **1** and then press [Enter].

Next, choose the key size, or how long the key should be. Generally, the longer the key, the more resistant against attacks your messages will be. The default size, 1024 bits, should be sufficiently strong for most users, so press [Enter].

The next option asks you to specify how long you want your key to be valid. Usually, the default (0 = key does not expire) is fine. If you do choose an expiration date, remember that anyone with whom you exchanged your public key will also have to be informed of its expiration and supplied with a new public key. If you do not choose an expiration date, you will be asked to confirm your decision. Press [y] to confirm your decision.

Your next task is to provide a user ID that consists of your name, your email address, and an optional comment. When you are finished, you will be presented with a summary of the information you entered.

Once you accept your choices, you will have to enter a passphrase.



Tip

Like your account passwords, a good passphrase is essential for optimal security in GnuPG. For example, mix your passphrase with uppercase and lowercase letters, use numbers, or punctuation marks.

Once you enter and verify your passphrase, your keys will be generated. You will see a message similar to the following:

```
We need to generate a lot of random bytes. It is a good idea to perform
some other action (type on the keyboard, move the mouse, utilize the
disks) during the prime generation; this gives the random number
generator a better chance to gain enough entropy.
++++.++++.+++++.....+++++.++++.++++.++++.+++++
+++.....+++++
```

When the activity on the screen ceases, your new keys will be made and placed in the directory `.gnupg` in your home directory. To list your keys, use the command:

```
gpg --list-keys
```

You will see something similar to the following:

```
/home/newuser/.gnupg/pubring.gpg
-----
pub 1024D/B7085C8A 2000-06-18 Your Name <you@yourisp.net>
sub 1024g/E12AF9C4 2000-06-18
```

If you created a GnuPG key with version 1.0.6 or lower, exported your private key, and imported it into a new one, you have to explicitly trust your own key to sign items with version 1.0.7. To trust your key, type the following command (replace `<user-id>`):

```
gpg --edit-key <user-id>
```

At the `Command>` prompt type **trust** and select `5 = I trust ultimately` to trust your own key.

B.4. 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.

**Note**

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:

```
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, 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 diskette and store it someplace safe.

The output will look similar to the following:

```
sec 1024D/823D25A9 2000-04-26 Your Name <you@yourisp.net>
Create a revocation certificate for this key?
```

Press [Y] to create a revocation certificate for the listed key. Next, you will be asked to select the reason for revocation and provide an optional description. After confirming the reason, enter the passphrase you used to generate the key.

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 do not know how to copy a file to a diskette in Red Hat Linux, see the *Official Red Hat Linux Getting Started Guide*.)

B.5. 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 command:

```
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/PotnYcZRgLT
z6syBBWs8JB4xt5V09iJSGAMPUE8Jpdn2aRXPapdoDw179LM8Rq6r+gwCg5Zza
pGNlkqFu24Wm5wC1zg4QTbMD/3MJCSxfL99Ek5HXcB3yhj+o0LmIrGAVBgoWdrRd
BIGjQQFhV1NSwC8YhN/4nGHwPaTxgEtnb4CI1wI/G3DK9oLYMyRJinkGJ6XYfP3b
cCQmqATDF5ugIamddi tnw7deXqn/eavaMxRXJM/RQsGjJyVpbAO2OqKe6L6Inb5H
kjcZA/9obTm499dDMRQ/CNR92fA5pr0zriy/ziLUow+cqI59nt+bEb9nYlMfmUN6
SW0jCH+pIQH51erV+EookyOyq3ocUdjeRYF/d2j19xmeSyL2H3tDvnuE6vggqFU/N
sdvby4B2Iku7S/h06W6GPQAE+pzdyX9vs+Pnf8osu7W3j60WprQkUGF1bCBHYWxs
YWdoZXIGPHhdWxnYWxsQHJLZGhhdC5jb20+iFYEEeECABYFAjkHP3UECwoEAwMV
AwIDFgIBAheAAAJEJECmvGCPsWpMjQaonF2zvRgdR/8or9pBhu95zeSnbk7AKCm
/uxVS0a5KoN7J61/1vEwx11poLkBDQ5Bz+MEQA8ztcWRJjW8cHCgLaE402jyqQ
37gDT/n4V566nU+YItzDFScVmgMuFRzhibLb1fO9TpZzxEbSF3T6p9hLLnHCQ1bD
HRsKfh0eJYMMqB3+HyUpNeqCMEEd9AnWD9F4rQtO7Pes38sV01X00SvsTyMG9wEB
vSNZk+Rl+phA55r1s8cAAwUEAJjqazvk0bgFrw1OPG9m7fEEd1vPSV6HSA0fvz4w
c7ckfpuxg/URQNF3TJA00Acprk8Gg8J2CtebAyR/sP5IsrK511luGdk+10M85FpT
/cen2OdJtToAF/6fGnIkeCeP1O5aWTbDgdAUHBRykpDWU3GJ7NS6923fVg5khQWg
uwrAiEYEGBECAAYFAjkHP4wACgkQkQkA8YI9JamliwCfXox/HjlorMKnQRJkeBcZ
iLyPH1QAoI33Ft/0HBqLtqdtP4vWYQRb1bjW
=BMEc
-----END PGP PUBLIC KEY BLOCK-----

```

B.5.1. 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.

A keyserver is a repository on the Internet which can store and distribute your public key to anyone who requests it. Many keysevers 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.



Tip

Because most keysevers are synchronized, sending your public key to one keyserver is usually as good as sending it to them all. You can, however, locate different keysevers. One place to begin your search for keysevers and more information is *Keyserver.Net* available at <http://www.keyserver.net>.

You can send your public key from either the shell prompt or from a browser; 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 \(http://www.keyserver.net\)](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 B.5) 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 B-1).

- 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.)



```
File Edit View Terminal Go Help
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: GnuPG v1.0.7 (GNU/Linux)
Comment: For info see http://www.gnupg.org

mQGIBDKHP3URBACKWGsYh43pkXU9wj/X1G67K8/DSr185r7dNHNfLL/ewil10k2
q8saWJn26QZP5DVgdtUJMDHf16kQTAt9Nz0qbcVrxLYNfgeBsvkHF/P0tnYcZRGL
tz6syBBW8jB4xt5U09iJSGAMPUQE8Jpdn2ARXPAdoDw179LM8Rq6e+gnCg5Zza
pgNlkgFu24WMSClzq4TbMD/3MJCSxFL99EK5HXcB3hj+o0LtrGAVGowdrRd
B1GjQJQFhV1NSwC8Yhn/4nG#paTvgEtnb4CI1wI/G3DK9o1YMyRJinkGj6XYfP3b
cCQmqATDF5ugIAnddi1tnw7deXgn/eavaMxRXJM/RQsgJjYvpaA020gke6L6Inb5H
kjcZA/9obTm499dMRQ/CNR92fA5pr0zriy/z1LUow+cgI59nt+bEb9nY1mfuNG
SWOjCH+plQHS1erV+Eookyoyq3ocUdJeRYF/dzj19xmeSYLZH3TDvnuE6vggFU/N
sdvby4B21ku7S/h06w6GPQAE+pzdyX9vS+PnF8osu7W3j60wprQkUGF1bCBHYWxs
YwdoZXiGPHBhdwmxYwxsQHJLZghdC5jB20+1FYEExECABYFAjkHF3UECwoEAwMV
AwIDFGlBAheAAoJEJECmvGCP5WpMjQoANF2zvRgdr/8or9pBhu95zeSnbk7AKCm
/uXVS0a5k0n7J61/1vEwx11poLkBDQq58z+MEAQA8zctwRj jw8cHcGLaE402jyqQ
37gDT/n4V566nt+Y1tzDFScVngMuFRzhibLbf09TpZzxEbSF3T6p9hLLnHCQ1bd
HRsKfhoEjYMMgB3+HyUpNeqCMEEd9AnWD9P4rQt07Pes38v01X00Sv5TyMG9wEB
vSNzk+Rl+phA5r1s8cAAwUEAJjgazvk0BgFwr10PG9m7fEed1VSV6HSA0Fvz4w
c7ckFpuxg/URQNf3TJ400Acprk8Gg8J2ctebAyR/sP51srK511LuGdk+10M85FpT
/cen20dJTtoAF/6fOnIkeCp105AwTbDgdAUHHRykpDWU3GJ7NS6923fvG5khQWg
uwrAI5EY6GEBEAAyFAjkHP4wAcgkQkQkA8yI9JJanliwCFXox/HjlorMKnQRJkeBcZ
iLlyPHlQAoI33Ft/OHBqLtqdtP4vWYQRb1bjw

mykey.asc
```

Figure B-1. Copying Your Public Key

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.

B.6. Importing a Public Key

The other end of key exchange is 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 email 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.

After downloading a key and saving it to the file `key.asc`, use the following command to add it to your keyring.

```
gpg --import key.asc
```

Another way to save a key is to use a browser's **Save As** feature. If you are using a browser such as **Mozilla**, and you locate a key at a keyserver, you can save the page as a text file (go to **File => Save Page As**). In the drop-down box next to **Files of Type**, choose **Text Files (*.txt)**. 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 the following command:

```
gpg --import newkey.txt
```

The output will look similar to the following:

```
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.

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.

B.7. 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.

B.8. 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.

B.8.1. Installed Documentation

- `man gpg` and `info gpg` — Quick Reference of GnuPG commands and options.

B.8.2. 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.

B.8.3. 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

Symbols

/dev/shm, 192
/etc/auto.master, 126
/etc/exports, 128
/etc/fstab, 16, 125
/etc/hosts.lpd, 202
/etc/httpd/conf/httpd.conf, 143
/etc/named.custom, 169
/etc/printcap, 197
/etc/printcap.local, 197
/etc/sysconfig/dhcpd, 139
/etc/sysconfig/iptables, 108, 110
/proc directory, 195
/var/spool/cron, 216

A

anacron
 additional resources, 220
Apache HTTP Secure Server
 accessing, 166
 books, 167
 certificate
 authorities, 161
 choosing a CA, 161
 creation of request, 163
 installing, 165
 moving it after an upgrade, 160
 pre-existing, 160
 self-signed, 165
 test vs. signed vs. self-signed, 161
 testing, 165
 connecting to, 166
 explanation of security, 159
 installed docs, 167
 installing, 157
 key
 generating, 162
 packages, 157
 port numbers, 166
 providing a certificate for, 159
 security
 explanation of, 159
 upgrading from 1.0 or 2.0, 160
 URLs, 166
 URLs for, 166
 websites, 167
Apache HTTP Server
 (See HTTP Configuration Tool)
 additional resources, 156
 related books, 156
 securing, 159

APXS, 158
at, 218
 additional resources, 220
autofs, 126
 /etc/auto.master, 126
Automated Tasks, 215

B

batch, 218
 additional resources, 220
BIND configuration, 169
 adding a forward master zone, 170
 adding a reverse master zone, 171
 adding a slave zone, 173
 applying changes, 169
 default directory, 169
boot diskette, 225
booting
 emergency mode, 77
 rescue mode, 75
 single-user mode, 77

C

CA
 (See Apache HTTP Secure Server)
chkconfig, 116
CIPE connection
 (See network connection)
command line options
 printing from, 212
configuration
 console access, 179
 NFS, 125
console
 making files accessible from, 180
console access
 configuring, 179
 defining, 180
 disabling, 180
 disabling all, 180
 enabling, 181
conventions
 document, viii
Cron, 215
 additional resources, 220
 configuration file, 215
 example crontabs, 216
 user-defined tasks, 216
crontab, 215
CtrlAltDel
 shutdown, disabling, 179

D

- date configuration, 183
- dateconfig
 - (See Time and Date Properties Tool)
- decryption
 - with GnuPG, 261
- devel package, 158
- df, 192
- DHCP, 135
 - additional resources, 140
 - client configuration, 140
 - command line options, 139
 - connecting to, 140
 - dhcpd.conf, 135
 - dhcpd.leases, 138
 - dhcrelay, 139
 - global parameters, 136
 - group, 137
 - options, 136
 - reasons for using, 135
 - Relay Agent, 139
 - server configuration, 135
 - shared-network, 136
 - starting the server, 138
 - stopping the server, 138
 - subnet, 136
- dhcpd.conf, 135
- dhcpd.leases, 138
- dhcrelay, 139
- disk storage
 - parted
 - (See parted)
- diskcheck, 193
- documentation
 - finding installed, 243
- DSA keys
 - generating, 122
- DSOs
 - loading, 158
- du, 192
- Dynamic Host Configuration Protocol
 - (See DHCP)

E

- e2fsck, 16
- e2label, 32
- encryption
 - with GnuPG, 261
- Ethernet connection
 - (See network configuration)
- exporting NFS file Systems, 126
- exports, 128
- ext2
 - reverting from ext3, 16

ext3

- converting from ext2, 16
- creating, 16
- features, 15

F

- feedback, xi
- file systems, 192
 - ext2
 - (See ext2)
 - ext3
 - (See ext3)
 - LVM
 - (See LVM)
 - monitoring, 193
 - NFS
 - (See NFS)
- firewall configuration
 - (See GNOME Lokkit)
- floppy group, use of, 182
- free, 191
- ftp, 119

G

- GNOME Lokkit
 - activating the firewall, 110
 - basic firewall configuration, 108
 - configuring common services, 110
 - DHCP, 109
 - iptables service, 111
 - local hosts, 109
 - mail relay, 111
- GNOME System Monitor, 190
- gnome-lokkit
 - (See GNOME Lokkit)
- gnome-system-monitor, 190
- Gnu Privacy Guard
 - (See GnuPG)
- GnuPG
 - additional resources, 267
 - checking RPM package signatures, 241
 - digital signatures, 267
 - exporting public key, 264
 - to keyserver, 265
 - generating a keypair, 262
 - generating a revocation certificate, 263
 - importing a public key, 266
 - insecure memory warning, 261
 - introduction, 261
- GPG
 - (See GnuPG)
- group configuration
 - adding groups, 187

- filtering list of groups, 185
- modify groups for a user, 186
- modify users in groups, 188
- modifying group properties, 187
- viewing list of groups, 185

groups

- (See group configuration)

- floppy, use of, 182

H

hardware

- viewing, 194

Hardware Browser, 194

Hardware RAID

- (See RAID)

HTTP Configuration Tool

- directives

- (See HTTP directives)

- error log, 146

- modules, 143

- transfer log, 146

HTTP directives

- DirectoryIndex, 145

- ErrorDocument, 145

- ErrorLog, 147

- Group, 154

- HostnameLookups, 147

- KeepAlive, 155

- KeepAliveTimeout, 155

- Listen, 144

- LogFormat, 147

- LogLevel, 147

- MaxClients, 155

- MaxKeepAliveRequests, 155

- Options, 145

- ServerAdmin, 144

- ServerName, 144

- Timeout, 155

- TransferLog, 146

- User, 154

httpd, 143

hwbrowser, 194

I

information

- about your system, 189

initrd, 257

insmod, 232

installation

- kickstart

- (See kickstart installations)

- LVM, 83

- software RAID, 79

Internet connection

- (See network configuration)

introduction, vii

ISDN connection

- (See network configuration)

K

kernel

- building, 255, 259

- custom, 255, 259

- downloading, 226

- initrd image for, 257

- large memory support, 226

- modular, 255

- modules, 231

- monolithic, 259

- multiple processor support, 226

- upgrading, 225

kernel modules

- listing, 231

- loading, 232

- unload, 232

kickstart

- how the file is found, 57

Kickstart Configurator, 59

- %post script, 73

- %pre script, 72

- authentication options, 67

- basic options, 59

- boot loader, 62

- boot loader options, 62

- firewall configuration, 68

- installation method selection, 60

- interactive, 60

- keyboard, 59

- language, 59

- language support, 60

- mouse, 59

- network configuration, 66

- package selection, 71

- partitioning, 63

- software RAID, 64

- preview, 59

- reboot, 60

- root password, 60

- encrypt, 60

- saving, 74

- text-mode installation, 60

- time zone, 59

- X configuration, 68

kickstart file

- %include, 52

- %post, 54

- %pre, 53

- auth, 38
- authconfig, 38
- autostep, 38
- bootloader, 40
- clearpart, 41
- creating, 38
- device, 42
- deviceprobe, 42
- diskette-based, 56
- driverdisk, 42
- firewall, 43
- format of, 37
- include contents of another file, 52
- install, 43
- installation methods, 43
- interactive, 44
- keyboard, 44
- lang, 44
- langsupport, 45
- lilo, 45
- lilocheck, 46
- logvol, 46
- mouse, 46
- network, 46
- network-based, 56, 57
- options, 38
- package selection specification, 52
- part, 48
- partition, 48
- post-installation configuration, 54
- pre-installation configuration, 53
- raid, 49
- reboot, 50
- rootpw, 50
- skipx, 50
- text, 50
- timezone, 51
- upgrade, 51
- volgroup, 52
- what it looks like, 37
- xconfig, 51
- zerombr, 52
- kickstart installations, 37
 - diskette-based, 56
 - file format, 37
 - file locations, 55
 - installation tree, 57
 - LVM, 46
 - network-based, 56, 57
 - starting, 57

L

- loading kernel modules, 231
- log files, 221
 - (See Also Log Viewer)
 - description, 221
 - examining, 222
 - locating, 221
 - rotating, 221
 - syslogd, 221
 - viewing, 221
- Log Viewer
 - alerts, 222
 - filtering, 222
 - log file locations, 222
 - refresh rate, 222
 - searching, 222
- logical volume, 27, 85
- logical volume group, 27, 83
- Logical Volume Manager
 - (See LVM)
- logrotate, 221
- lpd, 198
- LPRng, 197
- lsmod, 231
- lspci, 194
- LVM, 27
 - configuring LVM during installation, 83
 - explanation of, 27
 - logical volume, 27, 85
 - logical volume group, 27, 83
 - physical extent, 84
 - physical volume, 27, 83
 - with kickstart, 46

M

- Mail Transport Agent
 - (See MTA)
- Mail Transport Agent Switcher, 175
- Maximum RPM, 245
- memory usage, 191
- mkfs, 31
- mkpart, 31
- modem connection
 - (See network configuration)
- modprobe, 232
- modules.conf, 231
- mounting
 - NFS file systems, 125
- MTA
 - Mail Transport Agent Switcher, 175
 - setting default, 175

N

named.conf, 169

neat
(See network configuration)

netcfg
(See network configuration)

Network Administration Tool
(See network configuration)

network configuration

- activating devices, 101
- CIPE connection, 97
 - activating, 99
- device aliases, 103
- Ethernet connection, 90
 - activating, 91
- ISDN connection, 91
 - activating, 92
- logical network devices, 101
- managing /etc/hosts, 99
- managing DNS Settings, 100
- managing hosts, 99
- modem connection, 93
 - activating, 94
- overview, 89
- profiles, 101
- token ring connection, 96
 - activating, 97
- wireless connection, 98
- xDSL connection, 94
 - activating, 95

Network File System
(See NFS)

Network Time Protocol
(See NTP)

NFS

- /etc/fstab, 125
- additional resources, 130
- autofs
(See autofs)
- command line configuration, 128
- configuration, 125
- exporting, 126
- hostname formats, 129
- mounting, 125
- starting the server, 129
- status of the server, 129
- stopping the server, 129

NFS Server Configuration Tool, 126

NTP

- configuring, 184
 - ntp.conf, 184
 - ntpd, 184
 - step-tickers, 184
- ntpd, 184
- ntsysv, 116

O

O'Reilly & Associates, Inc., 130, 156, 268

OpenSSH, 119

- additional resources, 124
- client, 120
 - scp, 120
 - sftp, 121
 - ssh, 120
- DSA keys
 - generating, 122
- generating key pairs, 121
- RSA keys
 - generating, 121
- RSA Version 1 keys
 - generating, 122
- server, 119
 - /etc/ssh/ssh_config, 119
 - starting and stopping, 119
- ssh-add, 124
- ssh-agent, 123
 - with GNOME, 123
- ssh-keygen
 - DSA, 122
 - RSA, 121
 - RSA Version 1, 122

OpenSSL

- additional resources, 124

P

Package Management Tool, 247

- installing packages, 247
- removing packages, 249

packages

- dependencies, 237
- determining file ownership with, 242
- finding deleted files from, 242
- freshening with RPM, 239
- installing, 236
 - with Package Management Tool, 247
- locating documentation for, 243
- obtaining list of files, 244
- preserving configuration files, 239
- querying, 240
- querying uninstalled, 243
- removing, 238
 - with Package Management Tool, 249
- tips, 242
- upgrading, 239
- verifying, 240

pam_smbpass, 132

pam_timestamp, 181

parted, 29

- creating partitions, 30

- overview, 29
- removing partitions, 32
- resizing partitions, 33
- selecting device, 30
- table of commands, 29
- viewing partition table, 30
- partition table
 - viewing, 30
- partitions
 - creating, 30
 - formatting
 - mkfs, 31
 - labeling
 - e2label, 32
 - making
 - mkpart, 31
 - removing, 32
 - resizing, 33
 - viewing list, 30
- PCI devices
 - listing, 194
- physical extent, 84
- physical volume, 27, 83
- postfix, 175
- printconf
 - (See printer configuration)
- printer configuration, 197
 - cancel print job, 212
 - command line options, 210
 - add a printer, 210
 - remove a printer, 211
 - restore configuration, 210
 - save configuration, 210
 - creating an alias, 208
 - CUPS printing system, 212
 - configuration interface, 213
 - default printer, 208
 - delete existing printer, 208
 - driver options, 209
 - Assume Unknown Data is Text, 209
 - Convert Text to Postscript, 209
 - Effective Filter Locale, 209
 - Media Source, 209
 - Page Size, 209
 - Prerender Postscript, 209
 - Send End-of-Transmission (EOT), 209
 - Send Form-Feed (FF), 209
 - edit driver, 209
 - edit existing printer, 208
 - exporting settings, 210
 - importing settings, 210
 - local printer, 198
 - LPRng, 197
 - managing print jobs, 212
 - modifying existing printers, 208
 - Novell NetWare (NCP) printer, 203

- overriding a printer, 208
- printing from the command line, 212
- remote UNIX printer, 200
- rename existing device, 208
- Samba (SMB) printer, 202
- save configuration to file, 210
- strict RFC1179 compliance, 201
- test page, 207
- text-based application, 197
- viewing print spool, 212
- Printer Configuration Tool
 - (See printer configuration)
- printtool
 - (See printer configuration)
- processes, 189
- ps, 189

R

- RAID, 23
 - configuring software RAID, 79
 - explanation of, 23
 - Hardware RAID, 23
 - level 0, 24
 - level 1, 24
 - level 4, 24
 - level 5, 24
 - levels, 24
 - reasons to use, 23
 - Software RAID, 23
- RAM, 191
- rcp, 120
- Red Hat Network, 251
- redhat-config-apache
 - (See HTTP Configuration Tool)
- redhat-config-date
 - (See Time and Date Properties Tool)
- redhat-config-kickstart
 - (See Kickstart Configurator)
- redhat-config-network
 - (See network configuration)
- redhat-config-packages
 - (See Package Management Tool)
- redhat-config-printer
 - (See printer configuration)
- redhat-config-securitylevel
 - (See Security Level Configuration Tool)
- redhat-config-time
 - (See Time and Date Properties Tool)
- redhat-config-users
 - (See user configuration and group configuration)
- redhat-logviewer
 - (See Log Viewer)
- redhat-switch-printer
 - (See Printer System Switcher)

redhat-switchmail
 (See Mail Transport Agent Switcher)

redhat-switchmail-nox
 (See Mail Transport Agent Switcher)

rescue mode, 75
 definition of, 75
 from CD, diskette, 75
 using, 75
 utilities available, 77

resize2fs, 16

rmmmod, 232

RPM, 235
 additional resources, 244
 book about, 245
 checking package signatures, 241
 dependencies, 237
 design goals, 235
 determining file ownership with, 242
 documentation with, 243
 file conflicts
 resolving, 237
 finding deleted files with, 242
 freshen, 239
 freshening packages, 239
 GnuPG, 241
 graphical interface, 247
 installing, 236
 with Package Management Tool, 247
 md5sum, 241
 preserving configuration files, 239
 querying, 240
 querying for file list, 244
 querying uninstalled packages, 243
 tips, 242
 uninstalling, 238
 with Package Management Tool, 249
 upgrading, 239
 using, 236
 verifying, 240
 website, 245

RPM Package Manager
 (See RPM)

RSA keys
 generating, 121

RSA Version 1 keys
 generating, 122

runlevels, 113

S

Samba, 131
 additional resources, 133
 configuration, 131
 smb.conf, 131
 encrypted passwords, 132
 pam_smbpass, 132
 reasons for using, 131
 share
 connecting to, 133
 connecting to with Nautilus, 133
 syncing passwords with passwd, 132
 with Windows 2000, 132
 with Windows NT 4.0, 132

scp
 (See OpenSSH)

security, 113
 security level
 (See Security Level Configuration Tool)

Security Level Configuration Tool
 customize incoming services, 106
 customize trusted devices, 106
 iptables service, 111
 security levels
 high, 105
 medium, 106
 no firewall, 106

sendmail, 175

services
 controlling access to, 113

Services Configuration Tool, 115

sftp
 (See OpenSSH)

shutdown
 disablingCtrlAltDel , 179

SMB protocol, 131

smb.conf, 131

Software RAID
 (See RAID)

ssh
 (See OpenSSH)

ssh-add, 124

ssh-agent, 123
 with GNOME, 123

striping
 RAID fundamentals, 23

swap space, 19
 adding, 19
 explanation of, 19
 moving, 21
 recommended size, 19
 removing, 20

syslogd, 221

system information
 file systems, 192

- /dev/shm, 192
 - monitoring, 193
- gathering, 189
- hardware, 194
- memory usage, 191
- processes, 189
 - currently running, 189

T

- TCP wrappers, 114
- telinit, 114
- telnet, 119
- time configuration, 183
 - synchronize with NTP server, 184
- time zone configuration, 184
- timetool
 - (See Time and Date Properties Tool)
- token ring connection
 - (See network configuration)
- top, 189
- tune2fs
 - converting to ext3 with, 16
 - reverting to ext2 with, 16

U

- user configuration
 - adding users, 185
 - adding users to groups, 187
 - changing full name, 187
 - changing home directory, 187
 - changing login shell, 187

- changing password, 187
- filtering list of users, 185
- locking user accounts, 187
- modify groups for a user, 186
- modifying users, 186
- password expiration, 187
- setting user account expiration, 187
- viewing list of users, 185

User Manager

- (See user configuration)

users

- (See user configuration)

V

VeriSign

- using existing certificate, 160

volume group, 27, 83

W

Windows

- file and print sharing, 131

Windows 2000

- connecting to shares using Samba, 132

Windows NT 4.0

- connecting to shares using Samba, 132

X

xDSL connection

- (See network configuration)

xinetd, 114



Colophon

The Official Red Hat Linux manuals are written in DocBook SGML v4.1 format. The HTML and PDF formats are produced using custom DSSSL stylesheets and custom jade wrapper scripts.

Marianne Pecci <goddess@ipass.net> created the admonition graphics (note, tip, important, caution, and warning). They may be redistributed with written permission from Marianne Pecci and Red Hat, Inc..

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

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

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

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

Johnray Fuller — Primary Writer/Maintainer of the *Official Red Hat Linux Reference Guide*; Co-writer/Co-maintainer of the *Official Red Hat Linux Security Guide*

John Ha — Primary Writer/Maintainer to the *Official Red Hat Linux Getting Started Guide*; Co-writer/Co-maintainer of the *Official Red Hat Linux Security Guide*

