Red Hat Enterprise Linux 4
System Administration Guide

For Red Hat Enterprise Linux 4
Edition 2

Red Hat Inc.

Red Hat Enterprise Linux 4 System Administration Guide

For Red Hat Enterprise Linux 4
Edition 2

Red Hat Inc.

Legal Notice

Copyright © 2008 Red Hat, Inc.

Abstract

The System Administrator's Guide documents relevant information regarding the deployment, configuration and administration of Red Hat Enterprise Linux 4. It is oriented towards system administrators with a basic understanding of the system.

Table of Contents

Introduction

Welcome to the *System Administrators Guide*.

The *System Administrators Guide* contains information on how to customize your Red Hat Enterprise 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
- File Systems Reference
- Package Management
- Network Configuration
- System Configuration
- System Monitoring

This guide assumes you have a basic understanding of your Red Hat Enterprise Linux system. If you need help installing Red Hat Enterprise Linux, refer to the *Installation Guide*. For more general information about system administration, refer to the *Introduction to System Administration*. If you need more advanced documentation such as an overview of file systems, refer to the *Reference Guide*. If you need security information, refer to the *Security Guide*.

1. Changes To This Manual

This manual has been reorganized for clarity and updated for the latest features of Red Hat Enterprise Linux 4. Some of the changes include:

`Updated Kernel Modules and Manually Updating the Kernel Chapters`

The **`Kernel Modules`** and the **`Upgrading the Kernel Manually`** chapters include updated information in regards to the 2.6 kernel. Special thanks to **`Arjan van de Ven`** for his hard work in helping to complete this chapter.

An Updated Network File System (NFS) Chapter

The *Network File System (NFS)* chapter has been revised and reorganized to include NFSv4. Special thanks to **`Steve Dickson`** for his hard work in helping to complete this chapter.

An Updated OProfile Chapter

The *OProfile* chapter has been revised and reorganized to include updated information in regards to the 2.6 kernel. Special thanks to **`Will Cohen`** for his hard work in helping to complete this chapter.

An Updated X Window System Chapter

> *The X Window System* chapter has been revised to include information on the X11R6.8 release developed by the X.Org team.

Before reading this guide, you should be familiar with the contents of the *Installation Guide* concerning installation issues, the *Introduction to System Administration* for basic administration concepts, the *System Administrators Guide* for general customization instructions, and the *Security Guide* for security related instructions. This guide contains information about topics for advanced users.

2. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the [Liberation Fonts](#) set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later include the Liberation Fonts set by default.

2.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

`Mono-spaced Bold`

Used to highlight system input, including shell commands, file names and paths. Also used to highlight keys and key combinations. For example:

> To see the contents of the file `my_next_bestselling_novel` in your current working directory, enter the `cat my_next_bestselling_novel` command at the shell prompt and press `Enter` to execute the command.

The above includes a file name, a shell command and a key, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from an individual key by the plus sign that connects each part of a key combination. For example:

> Press `Enter` to execute the command.

> Press `Ctrl`+`Alt`+`F2` to switch to a virtual terminal.

The first example highlights a particular key to press. The second example highlights a key combination: a set of three keys pressed simultaneously.

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in `mono-spaced bold`. For example:

> File-related classes include `filesystem` for file systems, `file` for files, and `dir` for directories. Each class has its own associated set of permissions.

Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialog box text; labeled buttons; check-box and radio button labels; menu titles and sub-menu titles. For example:

> Choose **System** → **Preferences** → **Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, select the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

> To insert a special character into a **gedit** file, choose **Applications** → **Accessories** → **Character Map** from the main menu bar. Next, choose **Search** → **Find...** from the **Character Map** menu bar, type the name of the character in the **Search** field and click **Next**. The character you sought will be highlighted in the **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit** → **Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in proportional bold and all distinguishable by context.

Mono-spaced Bold Italic or **Proportional Bold Italic**

Whether mono-spaced bold or proportional bold, the addition of italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

> To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh john@example.com**.

> The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount /home**.

> To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: **package-version-release**.

Note the words in bold italics above — username, domain.name, file-system, package, version and release. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new and important term. For example:

> Publican is a *DocBook* publishing system.

2.2. Pull-quote Conventions

Terminal output and source code listings are set off visually from the surrounding text.

Output sent to a terminal is set in **mono-spaced roman** and presented thus:

```
books          Desktop    documentation  drafts  mss    photos   stuff   svn
books_tests  Desktop1   downloads      images  notes  scripts  svgs
```

Source-code listings are also set in **mono-spaced roman** but add syntax highlighting as follows:

```
static int kvm_vm_ioctl_deassign_device(struct kvm *kvm,
                struct kvm_assigned_pci_dev *assigned_dev)
{
        int r = 0;
        struct kvm_assigned_dev_kernel *match;

        mutex_lock(&kvm->lock);

        match = kvm_find_assigned_dev(&kvm->arch.assigned_dev_head,
                                       assigned_dev->assigned_dev_id);
        if (!match) {
                printk(KERN_INFO "%s: device hasn't been assigned before, "
                  "so cannot be deassigned\n", __func__);
                r = -EINVAL;
                goto out;
        }

        kvm_deassign_device(kvm, match);

        kvm_free_assigned_device(kvm, match);

out:
        mutex_unlock(&kvm->lock);
        return r;
}
```

2.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.

Note

Notes are tips, shortcuts or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.

Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled 'Important' will not cause data loss but may cause irritation and frustration.

Warning

Warnings should not be ignored. Ignoring warnings will most likely cause data loss.

3. More to Come

The *System Administrators Guide* is part of Red Hat, Inc's growing commitment to provide useful and timely support to Red Hat Enterprise 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 find an error in the *System Administrators 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://bugzilla.redhat.com/bugzilla/) against the component **rh-sag**.

Be sure to mention the manual's identifier:

```
rh-sag
```

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

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

Part I. Installation-Related Information

The *Installation Guide* discusses the installation of Red Hat Enterprise Linux and some basic post-installation troubleshooting. However, advanced installation options are covered in this manual. This part provides instructions for *kickstart* (an automated installation technique) and all related tools. Use this part in conjunction with the *Installation Guide* to perform any of these advanced installation tasks.

Chapter 1. Kickstart Installations

1.1. What are Kickstart Installations?

Many system administrators would prefer to use an automated installation method to install Red Hat Enterprise Linux on their machines. To answer this need, Red Hat, Inc 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 installation.

Kickstart files can be kept on a 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 Enterprise Linux on multiple machines, making it ideal for network and system administrators.

Kickstart provides a way for users to automate a Red Hat Enterprise Linux installation.

1.2. 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 media 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.

1.3. 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 Enterprise Linux Documentation CD, using the **Kickstart Configurator** application, or writing it from scratch. The Red Hat Enterprise 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 1.4, "Kickstart Options" for a list of kickstart options. You must include the required options.
  - The **%packages** section — Refer to Section 1.5, "Package Selection" for details.
  - The **%pre** and **%post** sections — These two sections can be in any order and are not required. Refer to Section 1.6, "Pre-installation Script" and Section 1.7, "Post-installation Script" for details.
- Items that are not required can be omitted.
- Omitting any required item results 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 continues unattended (unless it finds another missing item).

 ▷ Lines starting with a pound (or hash) 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 the installation)

   ▪ Keyboard setup

   ▪ The **upgrade** keyword

   ▪ Boot loader configuration

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

1.4. 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, use the **Kickstart Configurator** application. Refer to Chapter 2, *Kickstart Configurator* 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.

autopart (optional)

> Automatically create partitions — 1 GB or more root (**/**) partition, a swap partition, and an appropriate boot partition for the architecture. One or more of the default partition sizes can be redefined with the **part** directive.

ignoredisk (optional)

> Causes the installer to ignore the specified disks. This is useful if you use autopartition and want to be sure that some disks are ignored. For example, without **ignoredisk**, attempting to deploy on a SAN-cluster the kickstart would fail, as the installer detects passive paths to the SAN that return no partition table.
>
> The **ignoredisk** option is also useful if you have multiple paths to your disks.
>
> The syntax is:

```
ignoredisk --drives=drive1,drive2,...
```

> where *driveN* is one of **sda**, **sdb**,…, **hda**,… etc.

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 is 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 (distinguished name) with **`--ldapserver=`** and **`--ldapbasedn=`**.

`--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 **`--ldapbasedn=`**.

`--ldapserver=`

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

`--ldapbasedn=`

> If you specified either **`--enableldap`** or **`--enableldapauth`**, use this option to specify the DN 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. If you enable Kerberos, you must 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. If you enable SMB, you must 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.

--smbservers=

> 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 GRUB boot loader should be installed. This option is required for both
> installations and upgrades. For upgrades, if GRUB is not the current boot loader, the boot
> loader is changed to GRUB. To preserve other boot loaders, use **bootloader --upgrade**.

--append=

> Specifies kernel parameters. To specify multiple parameters, separate them with
> spaces. For example:

```
bootloader --location=mbr --append="hdd=ide-scsi ide=nodma"
```

--driveorder

> Specify which drive is first in the BIOS boot order. For example:

```
bootloader --driveorder=sda,hda
```

--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=

> Sets the GRUB boot loader password to 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=

> Similar to `--password=` except the password should already be encrypted.

--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.

--all

> Erases all partitions from the system.

--drives=

> Specifies which drives to clear partitions from. For example, the following clears all the partitions on the first two drives on the primary IDE controller:

```
clearpart --drives=hda,hdb --all
```

--initlabel

> Initializes the disk label to the default for your architecture (for example `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.

--linux

Erases all Linux partitions.

--none (default)

Do not remove any partitions.

cmdline (optional)

Perform the installation in a completely non-interactive command line mode. Any prompts for interaction halts the install. This mode is useful on S/390 systems with the x3270 console.

device (optional)

On most PCI systems, the installation program autoprobes 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=

Mount options to use for mounting the NFS export. Any options that can be specified in **/etc/fstab** for an NFS mount are allowed. The options are listed in the **nfs(5)** man page. Multiple options are separated with a comma.

driverdisk (optional)

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

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

Alternatively, a network location can be specified for the driver diskette:

```
driverdisk --source=ftp://path/to/dd.imgdriverdisk --
source=http://path/to/dd.imgdriverdisk --source=nfs:host:/path/to/img
```

<partition>

Partition containing the driver disk.

`--type=`

File system type (for example, vfat or ext2).

`firewall (optional)`

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

```
firewall --enabled|--disabled [--trust=] <device> [--port=]
```

`--enabled`

Reject incoming connections that are not in response to outbound requests, such as DNS replies or DHCP requests. If access to services running on this machine is needed, you can choose to allow specific services through the firewall.

`--disabled`

Do not configure any iptables rules.

`--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 one or more of the following to allow the specified services through the firewall.

- **`--ssh`**
- **`--telnet`**
- **`--smtp`**
- **`--http`**
- **`--ftp`**

`--port=`

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

`firstboot (optional)`

Determine whether the **Setup Agent** starts the first time the system is booted. If enabled, the

firstboot package must be installed. If not specified, this option is disabled by default.

--enable

The **Setup Agent** is started the first time the system boots.

--disable

The **Setup Agent** is not started the first time the system boots.

--reconfig

Enable the **Setup Agent** to start at boot time in reconfiguration mode. This mode enables the language, mouse, keyboard, root password, security level, time zone, and networking configuration options in addition to the default ones.

halt (optional)

Halt the system after the installation has successfully completed. This is similar to a manual installation, where anaconda displays a message and waits for the user to press a key before rebooting. During a kickstart installation, if no completion method is specified, the **reboot** option is used as default.

The **halt** option is roughly equivalent to the **shutdown -h** command.

For other completion methods, refer to the **poweroff**, **reboot**, and **shutdown** kickstart options.

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 **cdrom**, **harddrive**, **nfs**, or **url** (for FTP or HTTP installations). The **install** command and the installation method command must be on separate lines.

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 are 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. Refer to the **autostep** command.

keyboard (required)

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

```
be-latin1, bg, br-abnt2, cf, cz-lat2, cz-us-qwertz, de,
de-latin1, de-latin1-nodeadkeys, dk, dk-latin1, dvorak, es, et,
fi, fi-latin1, fr, fr-latin0, fr-latin1, fr-pc, fr_CH, fr_CH-latin1,
gr, hu, hu101, is-latin1, it, it-ibm, it2, jp106, la-latin1, mk-utf,
no, no-latin1, pl, pt-latin1, ro_win, ru, ru-cp1251, ru-ms, ru1, ru2,
ru_win, sg, sg-latin1, sk-qwerty, slovene, speakup,
speakup-lt, sv-latin1, sg, sg-latin1, sk-querty, slovene, trq, ua,
uk, us, us-acentos
```

The file **/usr/lib/python2.2/site-packages/rhpl/keyboard_models.py** also contains this list and is part of the **rhpl** package.

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
```

The file **/usr/share/system-config-language/locale-list** provides a list of the valid language codes in the first column of each line and is part of the **system-config-language** package.

langsupport (required)

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

To install one language, specify it. For example, to install and use the French language **fr_FR**:

```
langsupport fr_FR
```

--default=

If language support for more than one language is specified, a default must be identified.

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 are installed with the specified language set to the default.

logvol (optional)

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

```
logvol <mntpoint> --vgname=<name> --size=<size> --name=<name><options>
```

The options are as follows:

--noformat

Use an existing logical volume and do not format it.

--useexisting

Use an existing logical volume and reformat it.

--pesize

Set the size of the physical extents.

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
```

For a detailed example of **logvol** in action, refer to Section 1.4.1, "Advanced Partitioning Example".

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 are 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, genericwheelps/2, genericusb, generic3usb, genericwheelusb,
geniusnm, geniusnmps/2, geniusprops/2, geniusscrollps/2, geniusscrollps/2+,
thinking, thinkingps/2, logitech, logitechcc, logibm, logimman,
logimmanps/2, logimman+, logimman+ps/2, logimmusb, microsoft, msnew,
msintelli, msintellips/2, msintelliusb, msbm, mousesystems, mmseries,
mmhittab, sun, none
```

This list can also be found in the **/usr/lib/python2.2/site-packages/rhpl/mouse.py** file, which is part of the **rhpl** package.

If the mouse command is given without any arguments, or it is omitted, the installation program attempts to automatically detect 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 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.

--bootproto=

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

It defaults 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 are 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 this should be read as one continuous 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 also configure multiple nameservers here. To do so, specify them as a comma-delimited list in the command line. For example:

```
network --bootproto=static --ip=10.0.2.15 --netmask=255.255.255.0
\
            --gateway=10.0.2.254 --nameserver
192.168.2.1,192.168.3.1
```

`--device=`

Used to select a specific Ethernet device for installation. Note that using **`--device=`** is not effective unless the kickstart file is a local file (such as **`ks=floppy`**), since the installation program configures 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.

`--nostorage`

> Do not auto-probe storage devices such as ISCI, IDE or RAID.

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

> Creates a partition on the system.
>
> If more than one Red Hat Enterprise 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 are formatted as part of the installation process unless **`--noformat`** and **`--onpart`** are used.

> For a detailed example of **`part`** in action, refer to [Section 1.4.1, "Advanced Partitioning Example"](#).
>
> ***`<mntpoint>`***
>
> > The ***`<mntpoint>`*** is where the partition is mounted and must be of one of the following forms:
> >
> > - ***`/<path>`***
> >
> >   For example, **`/`**, **`/usr`**, **`/home`**
> >
> > - **`swap`**
> >
> >   The partition is 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 is no smaller than the amount of RAM in the system and no larger than twice the amount of RAM in the system.
> >
> >   Using the **`recommended`** option yields a limitation of 8GB for the swap partition. If you want to create a larger swap partition, specify the correct size in the kickstart file or create the partitions manually.
> >
> > - **`raid.<id>`**
> >
> >   The partition is used for software RAID (refer to **`raid`**).
> >
> > - **`pv.<id>`**

The partition is 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
```

puts **/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** puts the partition on the second SCSI disk on the system.

--asprimary

Forces automatic allocation of the partition as a primary partition, or the partitioning fails.

--type= (replaced by fstype)

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

--fstype=

Sets the file system 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=**.

> **Note**
>
> If partitioning fails for any reason, diagnostic messages appear on virtual console 3.

poweroff (optional)

Shut down and power off the system after the installation has successfully completed. Normally during a manual installation, anaconda displays a message and waits for the user to press a key before rebooting. During a kickstart installation, if no completion method is specified, the **reboot** option is used as default.

The **poweroff** option is roughly equivalent to the **shutdown -p** command.

> **Note**
>
> The **poweroff** option is highly dependent on the system hardware in use. Specifically, certain hardware components such as the BIOS, APM (advanced power management), and ACPI (advanced configuration and power interface) must be able to interact with the system kernel. Contact your manufacturer for more information on you system's APM/ACPI abilities.

For other completion methods, refer to the **halt**, **reboot**, and **shutdown** kickstart options.

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 file system 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 file system type for the RAID array. Valid values are ext2, ext3, swap, and vfat.

`--noformat`

Use an existing RAID device and do not format the RAID array.

`--useexisting`

Use an existing RAID device and reformat it.

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
```

For a detailed example of **`raid`** in action, refer to Section 1.4.1, "Advanced Partitioning Example".

`reboot` (optional)

Reboot after the installation is successfully completed (no arguments). Normally during a manual installation, anaconda displays a message and waits for the user to press a key before rebooting.

The **`reboot`** option is roughly equivalent to the **`shutdown -r`** command.

> **Note**
>
> Use of the **reboot** option *may* result in an endless installation loop, depending on the installation media and method.
> The **reboot** option is the default completion method if no other methods are explicitly specified in the kickstart file.

For other completion methods, refer to the **halt**, **poweroff**, and **shutdown** kickstart options.

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.

selinux (optional)

Sets the system's SELinux mode to one of the following arguments:

--enforcing

Enables SELinux with the default targeted policy being enforced.

> **Note**
>
> If the **selinux** option is not present in the kickstart file, SELinux is enabled and set to **--enforcing** by default.

--permissive

Outputs warnings only based on the SELinux policy, but does not actually enforce the policy.

--disabled

Disables SELinux completely on the system.

For complete information regarding SELinux for Red Hat Enterprise Linux, refer to the *Red Hat, Inc.*

shutdown (optional)

Shut down the system after the installation has successfully completed. During a kickstart installation, if no completion method is specified, the **reboot** option is used as default.

The **shutdown** option is roughly equivalent to the **shutdown** command.

For other completion methods, refer to the **halt**, **poweroff**, and **reboot** kickstart options.

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 must 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. The list of cards can also be found on the **X Configuration** screen of the **Kickstart Configurator**. If this argument is not provided, the installation program probes the PCI bus for the card. Since AGP is part of the PCI bus, AGP cards are detected if supported. The probe order is determined by the PCI scan order of the motherboard.

--videoram=

Specifies 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. The list of monitors can also be found on the **X Configuration** screen of the **Kickstart Configurator**. 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)

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

```
volgroup <name><partition><options>
```

The options are as follows:

--noformat

Use an existing volume group and do not format it.

--useexisting

Use an existing volume group and reformat it.

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
```

For a detailed example of **volgroup** in action, refer to Section 1.4.1, "Advanced Partitioning Example".

zerombr (optional)

If **zerombr** is specified, and **yes** is its sole argument, any invalid partition tables found on disks are initialized. This destroys 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.

1.4.1. Advanced Partitioning Example

The following is a single, integrated example showing the **clearpart**, **raid**, **part**, **volgroup**, and **logvol** kickstart options in action:

```
clearpart --drives=hda,hdc --initlabel

# Raid 1 IDE config
part raid.11    --size 1000     --asprimary     --ondrive=hda
part raid.12    --size 1000     --asprimary     --ondrive=hda
part raid.13    --size 2000     --asprimary     --ondrive=hda
part raid.14    --size 8000                     --ondrive=hda
part raid.15    --size 1 --grow                 --ondrive=hda

part raid.21    --size 1000     --asprimary     --ondrive=hdc
part raid.22    --size 1000     --asprimary     --ondrive=hdc
part raid.23    --size 2000     --asprimary     --ondrive=hdc
part raid.24    --size 8000                     --ondrive=hdc
part raid.25    --size 1 --grow                 --ondrive=hdc

# You can add --spares=x
raid /          --fstype ext3 --device md0 --level=RAID1 raid.11 raid.21
raid /safe      --fstype ext3 --device md1 --level=RAID1 raid.12 raid.22
raid swap       --fstype swap --device md2 --level=RAID1 raid.13 raid.23
raid /usr       --fstype ext3 --device md3 --level=RAID1 raid.14 raid.24
raid pv.01      --fstype ext3 --device md4 --level=RAID1 raid.15 raid.25

# LVM configuration so that we can resize /var and /usr/local later
volgroup sysvg pv.01
logvol /var             --vgname=sysvg  --size=8000     --name=var
logvol /var/freespace   --vgname=sysvg  --size=8000     --name=freespacetouse
logvol /usr/local       --vgname=sysvg  --size=1 --grow --name=usrlocal
```

This advanced example implements LVM over RAID, as well as the ability to resize various directories for future growth.

1.5. 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).

Packages can be specified by group or by individual package name. The installation program defines several groups that contain related packages. Refer to the **RedHat/base/comps.xml** file on the first Red Hat Enterprise Linux CD-ROM for a list of groups. Each group has an id, user visibility 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.

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 them in the **%packages** section.

Here is an example **%packages** selection:

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

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.xml** file. Groups can also be specified using the id for the group, such as **gnome-desktop**. Specify individual packages with no additional characters (the **dhcp** line in

the example above is an individual package).

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

```
-autofs
```

The following options are available for the **%packages** option:

--resolvedeps

> Install the listed packages and automatically resolve package dependencies. If this option is not specified and there are package dependencies, the automated installation pauses and prompts the user. For example:

```
%packages --resolvedeps
```

--ignoredeps

> Ignore the unresolved dependencies and install the listed packages without the dependencies. For example:

```
%packages --ignoredeps
```

--ignoremissing

> Ignore the missing packages and groups instead of halting the installation to ask if the installation should be aborted or continued. For example:

```
%packages --ignoremissing
```

1.6. 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 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.

1.6.1. 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-includ
   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 are used.

> **Note**
>
> The pre-installation script section of kickstart *cannot* manage multiple install trees or source media. This information must be included for each created ks.cfg file, as the pre-installation script occurs during the second stage of the installation process.

1.7. 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 do 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 file system 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.

1.7.1. 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 -o nolock 10.10.0.2:/usr/new-machines /mnt/temp open -s
-w -- /mnt/temp/runme umount /mnt/temp
```

> **Note**
>
> NFS file locking is *not* supported while in kickstart mode, therefore **-o nolock** is required when mounting an NFS mount.

Add a user to the system:

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

1.8. Making the Kickstart File Available

A kickstart file must be placed in one of the following locations:

- On a boot diskette
- On a boot CD-ROM
- 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.

1.8.1. Creating Kickstart Boot Media

Diskette-based booting is no longer supported in Red Hat Enterprise Linux. Installations must use CD-ROM or flash memory products for booting. However, the kickstart file may still reside on a diskette's top-level directory, and must be named **ks.cfg**.

To perform a CD-ROM-based kickstart installation, the kickstart file must be named **ks.cfg** and must be located in the boot CD-ROM's top-level directory. Since a CD-ROM is read-only, the file must be added to the directory used to create the image that is written to the CD-ROM. Refer to the *Installation Guide* for instructions on creating boot media; however, before making the **file.iso** image file, copy the **ks.cfg** kickstart file to the **isolinux/** directory.

To perform a pen-based flash memory kickstart installation, the kickstart file must be named **ks.cfg** and must be located in the flash memory's top-level directory. Create the boot image first, and then copy the **ks.cfg** file.

For example, the following transfers a boot image to the pen drive (**/dev/sda**) using the **dd** command:

```
dd if=diskboot.img of=/dev/sda bs=1M
```

> **Note**
>
> Creation of USB flash memory pen drives for booting is possible, but is heavily dependent on system hardware BIOS settings. Refer to your hardware manufacturer to see if your system supports booting to alternate devices.

1.8.2. 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 Enterprise Linux. The BOOTP/DHCP server provides 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 attempts an NFS mount of the file's path, and copies 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 **`dhcpd.conf`** file for the DHCP server:

```
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 file name 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 file name the client searches for is:

```
<ip-addr>-kickstart
```

The **`<ip-addr>`** section of the file name should be replaced with the client's IP address in dotted decimal notation. For example, the file name 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 attempts to use the server that answered the BOOTP/DHCP request as its NFS server. If you do not specify a path or file name, the client system tries to mount **`/kickstart`** from the BOOTP/DHCP server and tries to find the kickstart file using the same **`<ip-addr>-kickstart`** file name as described above.

1.9. Making the Installation Tree Available

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

If you are performing a CD-based installation, insert the Red Hat Enterprise 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 Enterprise 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 *Installation Guide* for details.

1.10. Starting a Kickstart Installation

To begin a kickstart installation, you must boot the system from boot media you have made or the Red Hat Enterprise Linux CD-ROM #1, and enter a special boot command at the boot prompt. The installation program looks for a kickstart file if the **ks** command line argument is passed to the kernel.

CD-ROM #1 and Diskette

The **linux ks=floppy** command also works if the **ks.cfg** file is located on a vfat or ext2 file system on a diskette and you boot from the Red Hat Enterprise Linux CD-ROM #1.

An alternate boot command is to boot off the Red Hat Enterprise Linux CD-ROM #1 and have the kickstart file on a vfat or ext2 file system on a diskette. To do so, enter the following command at the **boot:** prompt:

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

With Driver Disk

If you need to use a driver disk with kickstart, specify the **dd** option as well. For example, to boot off a boot diskette and use a driver disk, enter the following command at the **boot:** prompt:

```
linux ks=floppy dd
```

Boot CD-ROM

If the kickstart file is on a boot CD-ROM as described in <u>Section 1.8.1, "Creating Kickstart Boot Media"</u>, insert the CD-ROM into the system, boot the system, and enter the following command at the **boot:** prompt (where **ks.cfg** is the name of the kickstart file):

```
linux ks=cdrom:/ks.cfg
```

Other options to start a kickstart installation are as follows:

ks=nfs:<server>:/<path>

The installation program looks for the kickstart file on the NFS server *<server>*, as file *<path>*. The installation program uses 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 looks for the kickstart file on the HTTP server *<server>*, as file *<path>*. The installation program uses 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 file system on the diskette in **/dev/fd0**.

ks=floppy:/<*path*>

The installation program looks for the kickstart file on the diskette in **/dev/fd0**, as file **<*path*>**.

ks=hd:<*device*>:/<*file*>

The installation program mounts the file system on **<*device*>** (which must be vfat or ext2), and look for the kickstart configuration file as **<*file*>** in that file system (for example, **ks=hd:sda3:/mydir/ks.cfg**).

ks=file:/<*file*>

The installation program tries to read the file **<*file*>** from the file system; no mounts are done. This is normally used if the kickstart file is already on the **initrd** image.

ks=cdrom:/<*path*>

The installation program looks for the kickstart file on CD-ROM, as file **<*path*>**.

ks

If **ks** is used alone, the installation program configures the Ethernet card to use DHCP. The kickstart file is read from the "bootServer" from the DHCP response as if it is an NFS server sharing the kickstart file. By default, the bootServer is the same as the DHCP server. The name of the kickstart file is one of the following:

- If DHCP is specified and the boot file begins with a **/**, the boot file provided by DHCP is looked for on the NFS server.
- If DHCP is specified and the boot file begins with something other then a **/**, the boot file provided by DHCP is looked for in the **/kickstart** directory on the NFS server.
- If DHCP did not specify a boot file, 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 uses 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.

Chapter 2. Kickstart Configurator

Kickstart Configurator allows you to create or modify a kickstart file using a graphical user interface, so that you do not have to remember the correct syntax of the file.

To use **Kickstart Configurator**, you must be running the X Window System. To start **Kickstart Configurator**, select **Applications** (the main menu on the panel) => **System Tools** => **Kickstart**, or type the command **/usr/sbin/system-config-kickstart**.

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

To start with an existing kickstart file, select **File** => **Open** and select the existing file.

2.1. Basic Configuration
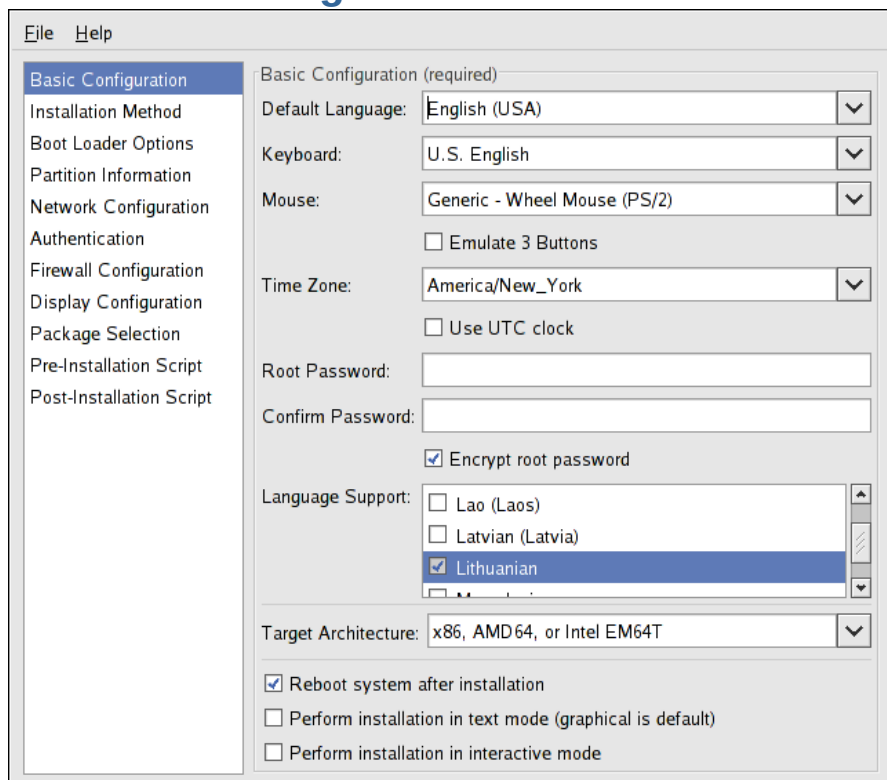
Figure 2.1. Basic Configuration

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

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

Choose the mouse for the system from the **Mouse** menu. If **No Mouse** is selected, no mouse is configured. If **Probe for Mouse** is selected, the installation program tries to automatically detect the mouse. Probing works for most modern mice.

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

From the **Time Zone** menu, choose the time zone to use for the system. To configure the system to use UTC, select **Use UTC clock**.

Enter the desired root password for the system in the **Root Password** text entry box. Type the same password in the **Confirm Password** text box. The second field is to make sure you do not mistype the password and then realize you do not know what it is after you have completed the installation. To save the password as an encrypted password in the file, select **Encrypt root password**. If the encryption option is selected, when the file is saved, the plain text password that you typed are encrypted and written to the kickstart file. Do not type an already encrypted password and select to encrypt it. Because a kickstart file is a plain text file that can be easily read, it is recommended that an encrypted password be used.

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

Choosing **Target Architecture** specifies which specific hardware architecture distribution is used during installation.

Choosing **Reboot system after installation** reboots 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, select the **Perform installation in text mode** option.

You can perform a kickstart installation in interactive mode. This means that the installation program uses all the options pre-configured in the kickstart file, but it allows you to preview the options in each screen before continuing to the next screen. To continue to the next screen, click the **Next** button after you have approved the settings or change them before continuing the installation. To select this type of installation, select the **Perform installation in interactive mode** option.

2.2. Installation Method
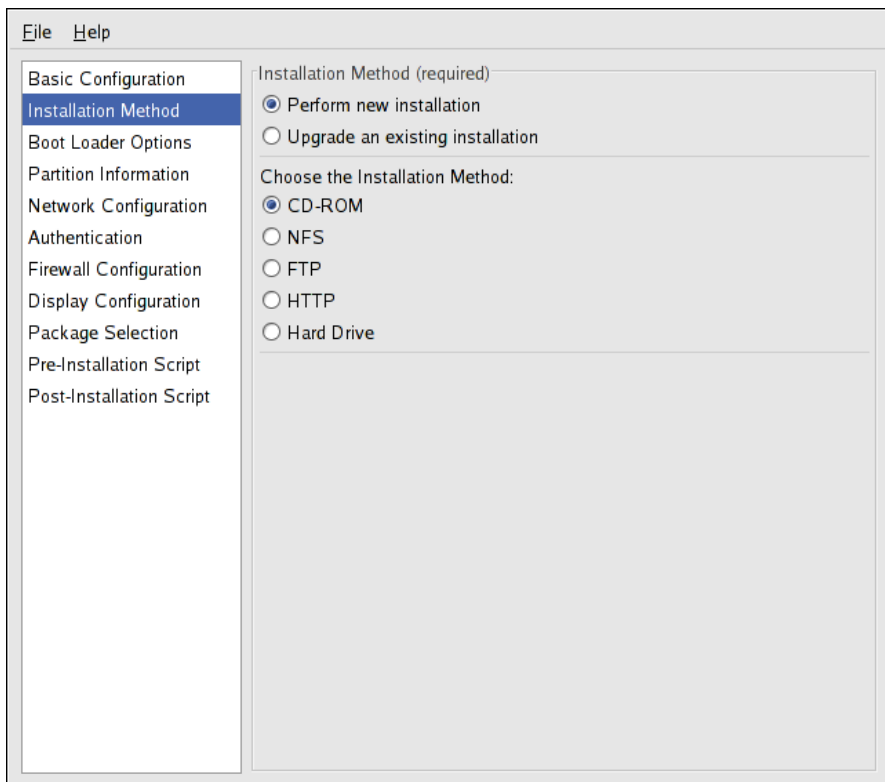
Figure 2.2. Installation Method
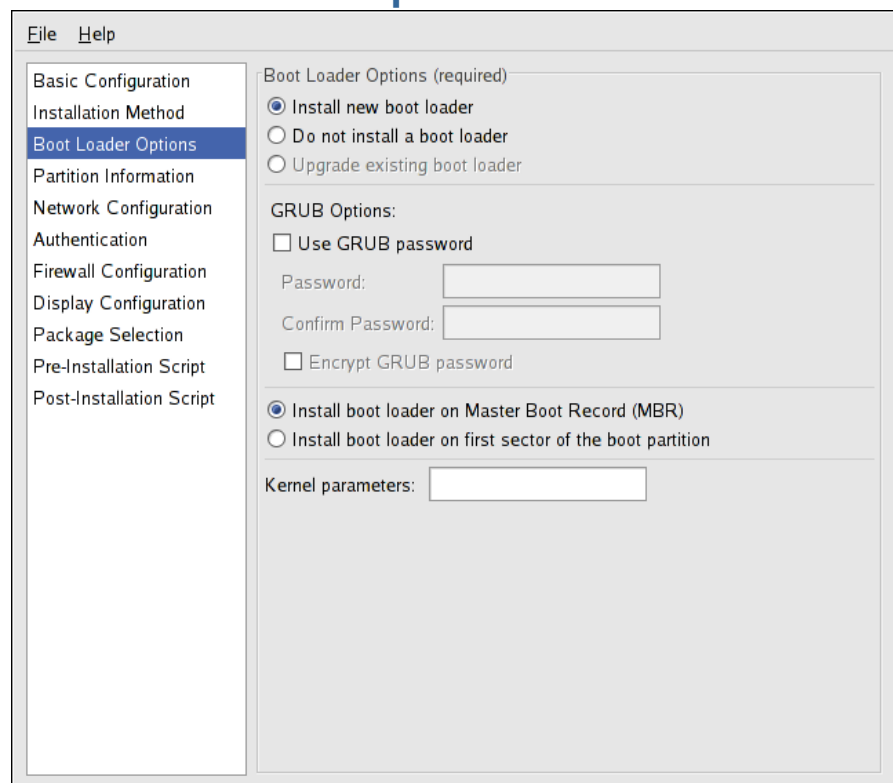
The **Installation Method** screen allows you to choose whether to perform a new installation or an upgrade. If you choose upgrade, the **Partition Information** and **Package Selection** options are disabled. They are not supported for kickstart upgrades.

Choose the type of kickstart installation or upgrade screen from the following options:

- **CD-ROM** — Choose this option to install or upgrade from the Red Hat Enterprise Linux CD-ROMs.
- **NFS** — Choose this option to install or upgrade from an NFS shared directory. In the text field for the the NFS server, enter a fully-qualified domain name or IP address. For the NFS directory, enter the name of the NFS directory that contains the **RedHat** directory of the installation tree. For example, if the NFS server contains the directory **/mirrors/redhat/i386/RedHat/**, enter **/mirrors/redhat/i386/** for the NFS directory.
- **FTP** — Choose this option to install or upgrade from an FTP server. In the FTP server text field, enter a fully-qualified domain name or IP address. For the FTP directory, enter the name of the FTP directory that contains the **RedHat** directory. For example, if the FTP server contains the directory **/mirrors/redhat/i386/RedHat/**, enter **/mirrors/redhat/i386/** for the FTP directory. If the FTP server requires a username and password, specify them as well.
- **HTTP** — Choose this option to install or upgrade from an HTTP server. In the text field for the HTTP server, enter the fully-qualified domain name or IP address. For the HTTP directory, enter the name of the HTTP directory that contains the **RedHat** directory. For example, if the HTTP server contains the directory **/mirrors/redhat/i386/RedHat/**, enter **/mirrors/redhat/i386/** for the HTTP directory.
- **Hard Drive** — Choose this option to install or upgrade from a hard drive. 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 as well as the **linux mediacheck** boot option as discussed in the *Installation Guide*. Enter the hard drive partition that contains the ISO images (for example, **/dev/hda1**) in the **Hard Drive Partition** text box. Enter

the directory that contains the ISO images in the **Hard Drive Directory** text box.

2.3. Boot Loader Options

Figure 2.3. Boot Loader Options

GRUB is the default boot loader for Red Hat Enterprise Linux. 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 diskette or have another way to boot your system, such as a third-party boot loader.

You must choose 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.

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 configuring **hdd=ide-scsi** as a kernel parameter (where **hdd** is the CD-ROM device).

You can password protect the GRUB boot loader by configuring a GRUB password. Select **Use GRUB password**, and enter a password in the **Password** field. Type the same password in the **Confirm Password** text field. To save the password as an encrypted password in the file, select **Encrypt GRUB password**. If the encryption option is selected, when the file is saved, the plain text password that you typed are encrypted and written to the kickstart file. If type an already encrypted password, unselect to encrypt it.

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

2.4. Partition Information
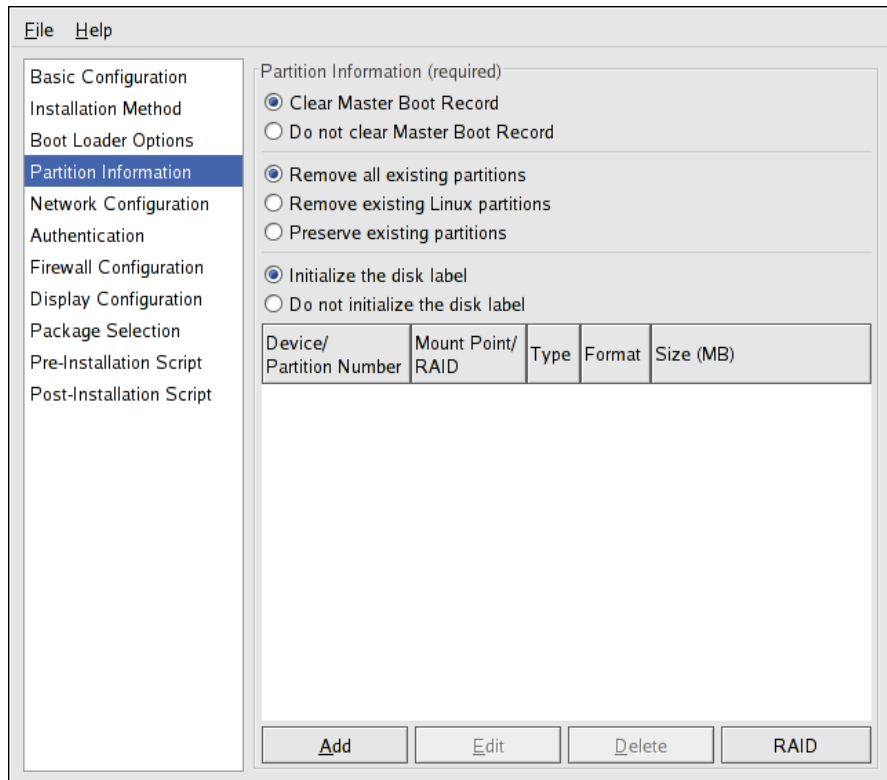
Figure 2.4. Partition Information

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

To initialize the disk label to the default for the architecture of the system (for example, **msdos** for x86 and **gpt** for Itanium), select **Initialize the disk label** if you are installing on a brand new hard drive.

2.4.1. Creating Partitions

To create a partition, click the **Add** button. The **Partition Options** window shown in Figure 2.5, "Creating Partitions" appears. Choose the 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.

Figure 2.5. Creating Partitions
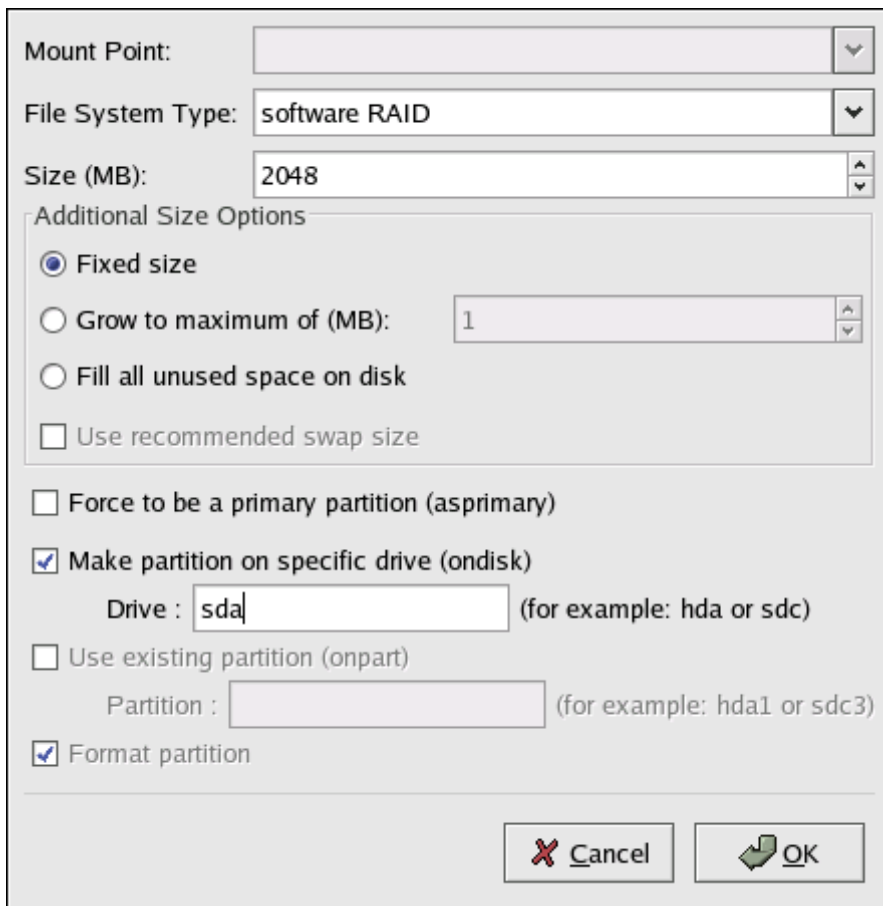
To edit an existing partition, select the partition from the list and click the **Edit** button. The same `Partition Options` window appears as when you chose to add a partition as shown in Figure 2.5, "Creating Partitions", except it reflects the values for the selected partition. Modify the partition options and click **OK**.

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

2.4.1.1. Creating Software RAID Partitions

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 specify an existing partition to use.

Figure 2.6. 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 RAID device.
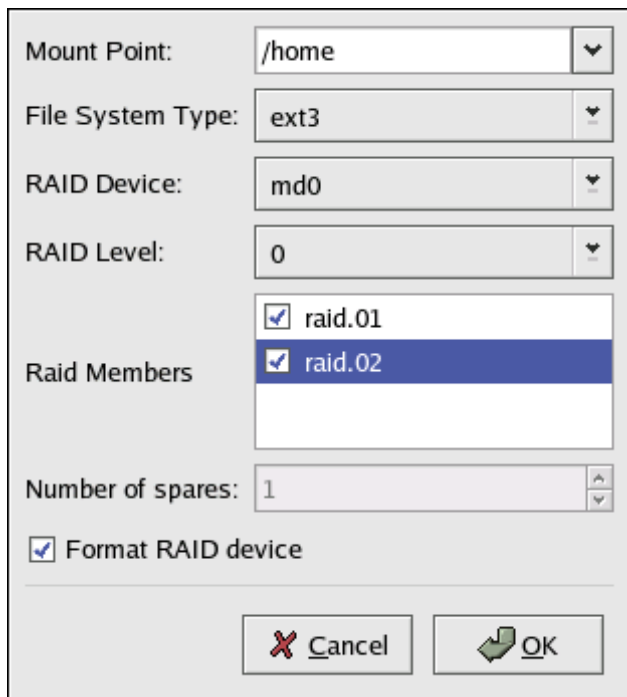
Figure 2.7. Creating a Software RAID Device

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

2.5. Network Configuration
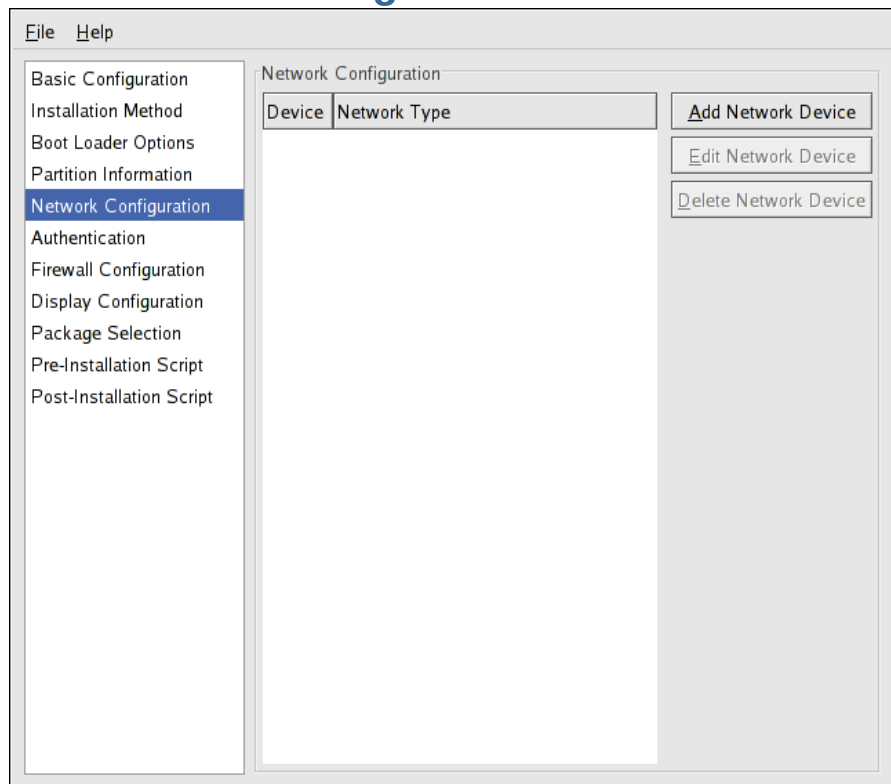
Figure 2.8. Network Configuration

If the system to be installed via kickstart does not have an Ethernet card, do not configure one on the

Network Configuration page.

Networking is only required if you choose a networking-based installation method (NFS, FTP, or HTTP). Networking can always be configured after installation with the **Network Administration Tool** (**system-config-network**). Refer to Chapter 17, *Network Configuration* for details.

For each Ethernet card on the system, click **Add Network Device** and select the network device and network type for the device. Select **eth0** to configure the first Ethernet card, **eth1** for the second Ethernet card, and so on.

2.6. Authentication
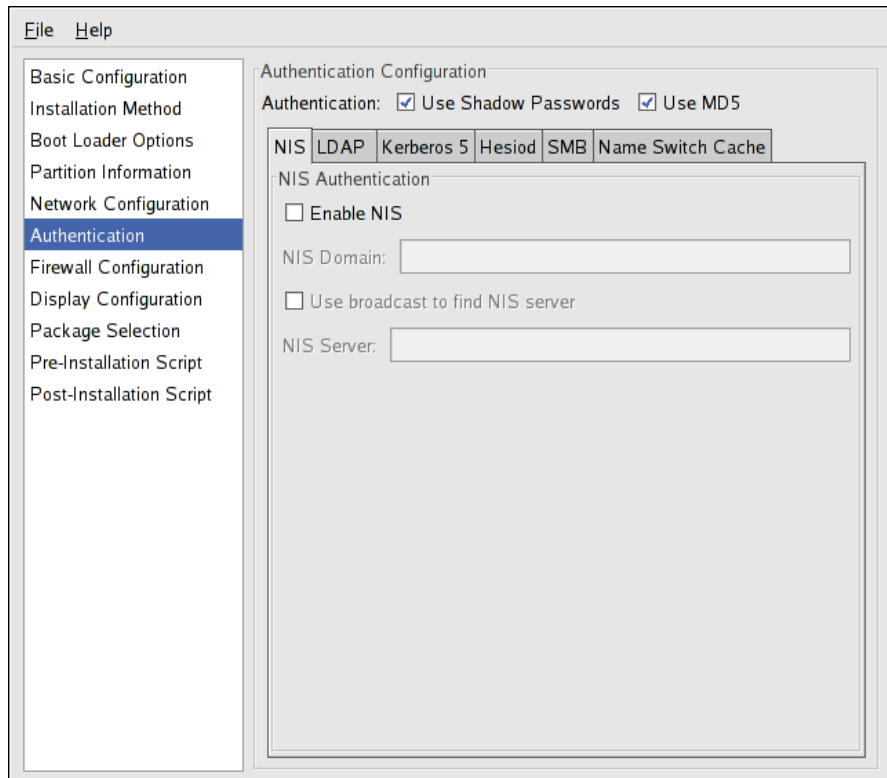
Figure 2.9. Authentication

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 allow you to configure the following methods of authentication:

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

These methods 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. Refer to Chapter 26, *Authentication Configuration* for more information about the options.

2.7. Firewall Configuration

The **Firewall Configuration** window is similar to the screen in the installation program and the **Security Level Configuration Tool**.
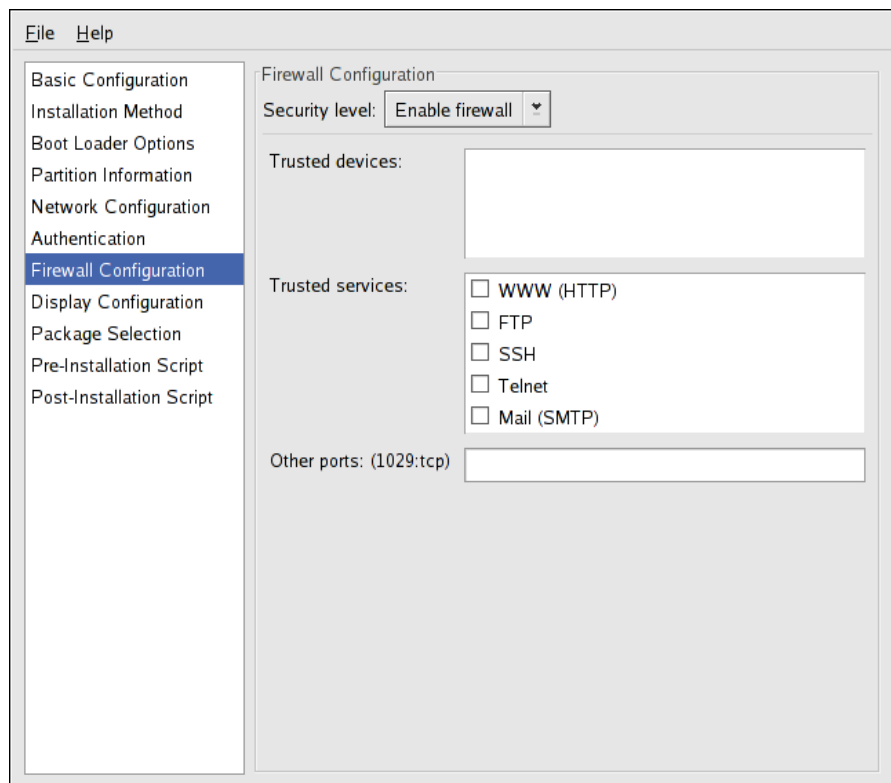
Figure 2.10. Firewall Configuration

If **Disable firewall** is selected, the system allows complete access to any active services and ports. No connections to the system are refused or denied.

Selecting **Enable firewall** configures the system to reject incoming connections that are not in response to outbound requests, such as DNS replies or DHCP requests. If access to services running on this machine is required, you can choose to allow specific services through the firewall.

Only devices configured in the **Network Configuration** section are listed as available **Trusted devices**. Connections from any devices selected in the list are accepted by the system. For example, if **eth1** only receives connections from internal system, you might want to allow connections from it.

If a service is selected in the **Trusted services** list, connections for the service are accepted and processed by the system.

In the **Other ports** text field, list any additional ports that should be opened for remote access. Use the following format: **port:protocol**. For example, to allow IMAP access through the firewall, specify **imap:tcp**. Specify numeric ports can also be specified; to allow UDP packets on port 1234 through the firewall, enter **1234:udp**. To specify multiple ports, separate them with commas.

2.7.1. SELinux Configuration

Although configuration for SELinux is not specified in the **Kickstart Configurator**, kickstart enables SELinux in **enforcing** mode by default if the **selinux** parameter is omitted from the kickstart file.

2.8. Display 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** option on the **Display Configuration** window as shown in Figure 2.11, "X Configuration - General". If this option is not chosen, the X configuration options are disabled and the **skipx** option is written to the kickstart file.

2.8.1. 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.
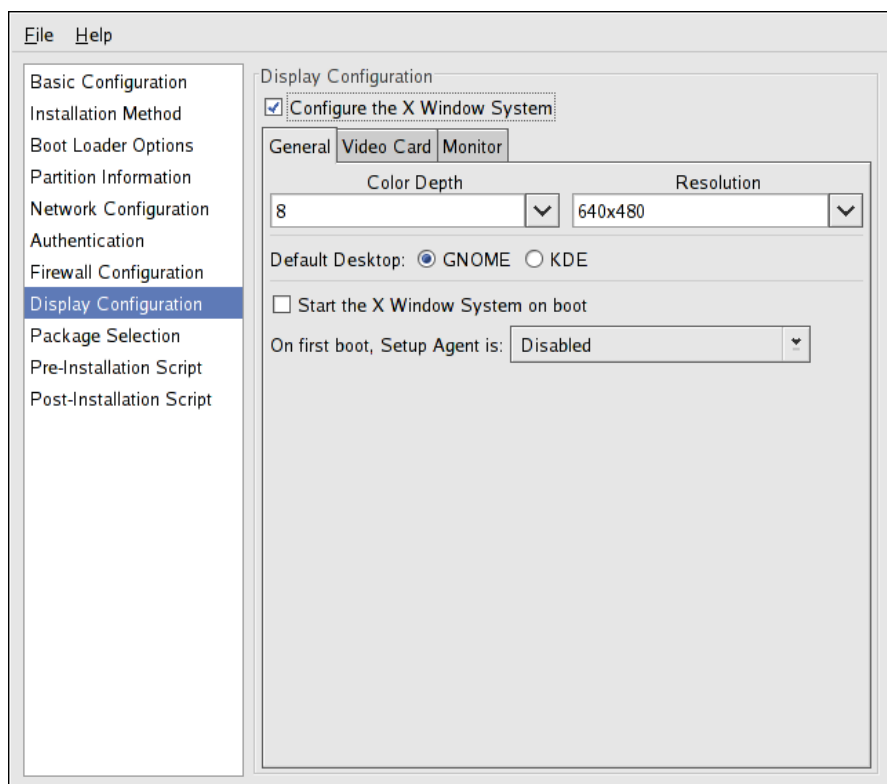
Figure 2.11. X Configuration - General

If you are installing both the GNOME and KDE desktops, you must choose which desktop should be the default. If only one desktop is to be installed, be sure to choose it. Once the system is installed, users can choose which desktop they want to be their default.

Next, choose whether to start the X Window System when the system is booted. This option starts 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.

Also select whether to start the **Setup Agent** the first time the system is rebooted. It is disabled by default, but the setting can be changed to enabled or enabled in reconfiguration mode. Reconfiguration mode enables the language, mouse, keyboard, root password, security level, time zone, and networking configuration options in addition to the default ones.

2.8.2. Video Card

`Probe for video card` is selected by default. Accept this default to have the installation program probe for the video card during installation. Probing works for most modern video cards. If this option is selected and the installation program cannot successfully probe the video card, the installation program stops 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 2.12, "X Configuration - Video Card". Specify the amount of video RAM the selected video card has from the **Video Card RAM** pulldown menu. These values are used by the installation program to configure the X Window System.
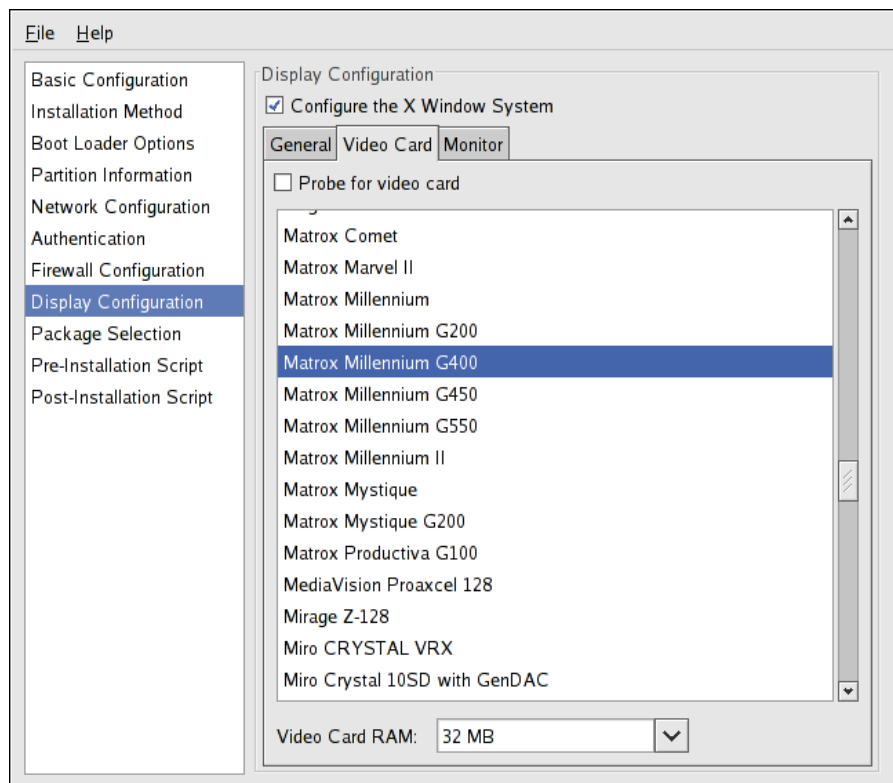
Figure 2.12. X Configuration - Video Card

2.8.3. Monitor

After configuring the video card, click on the **Monitor** tab as shown in Figure 2.13, "X Configuration - Monitor".
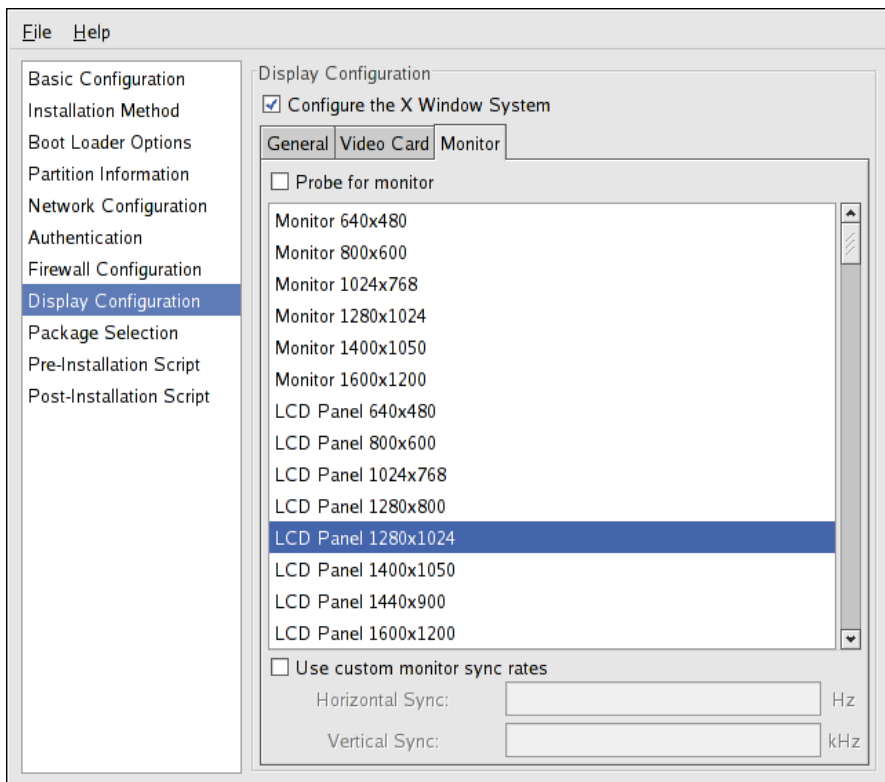
Figure 2.13. X Configuration - Monitor

Probe for monitor is selected by default. Accept this default to have the installation program probe for the monitor during installation. Probing works for most modern monitors. If this option is selected and the installation program cannot successfully probe the monitor, the installation program stops 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 selecting a specific monitor by checking the **Specify hysnc 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.

2.9. Package Selection
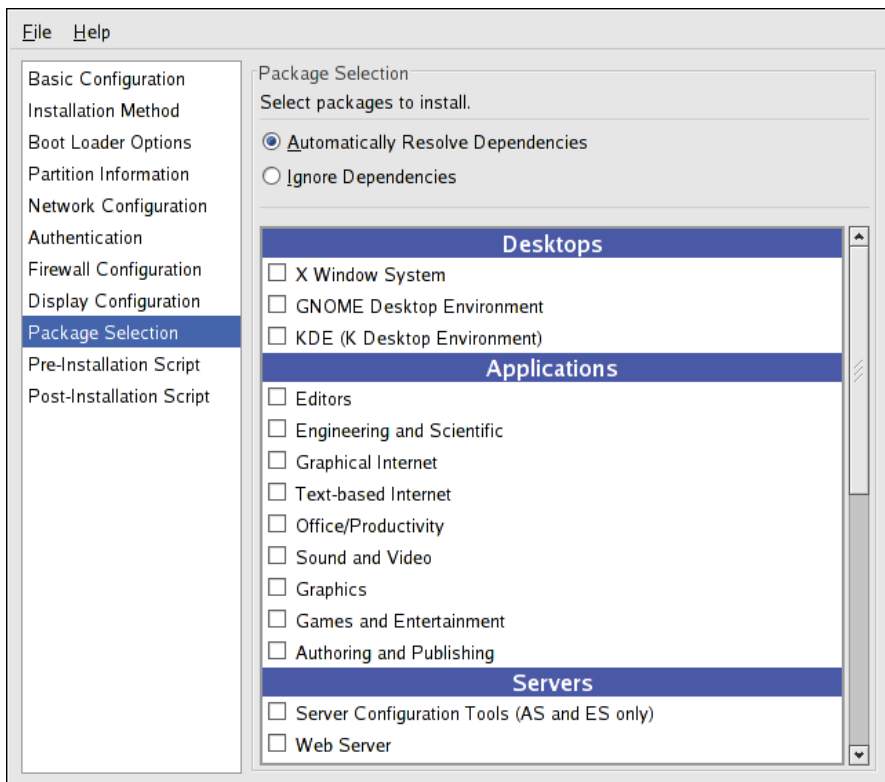
Figure 2.14. `Package Selection`

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

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

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. Refer to Section 1.5, "Package Selection" for details.

2.10. Pre-Installation Script
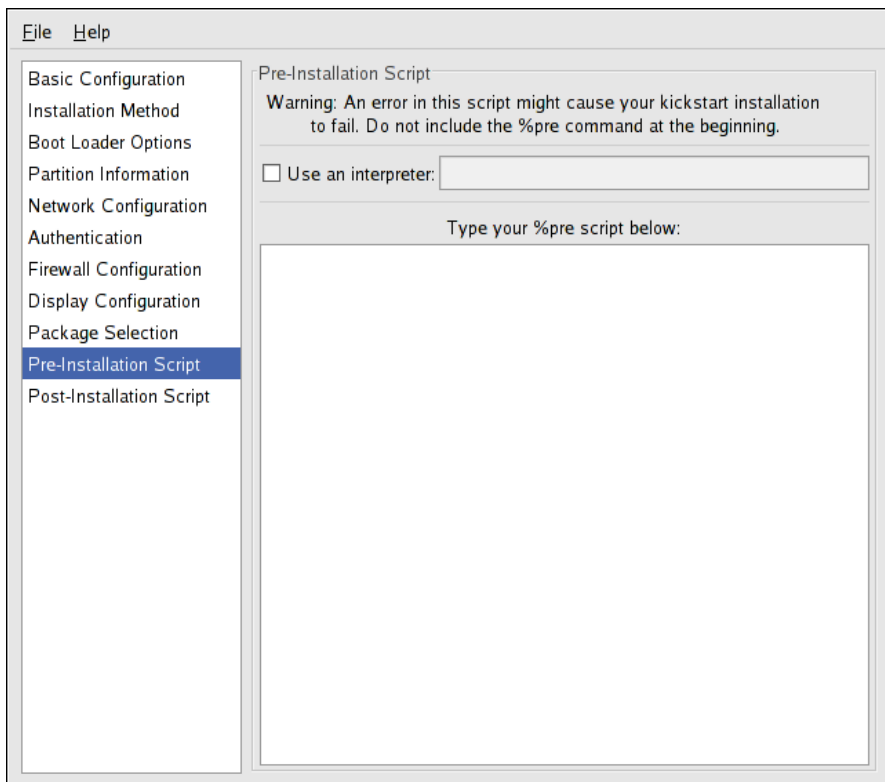
Figure 2.15. Pre-Installation Script

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. To include a pre-installation script, type it in the text area.

To specify a scripting language to use to execute the script, select the **Use an interpreter** option and enter the interpreter in the text box beside it. 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.

> ⚠️ **Warning**
>
> Do not include the **%pre** command. It is added for you.

2.11. Post-Installation Script
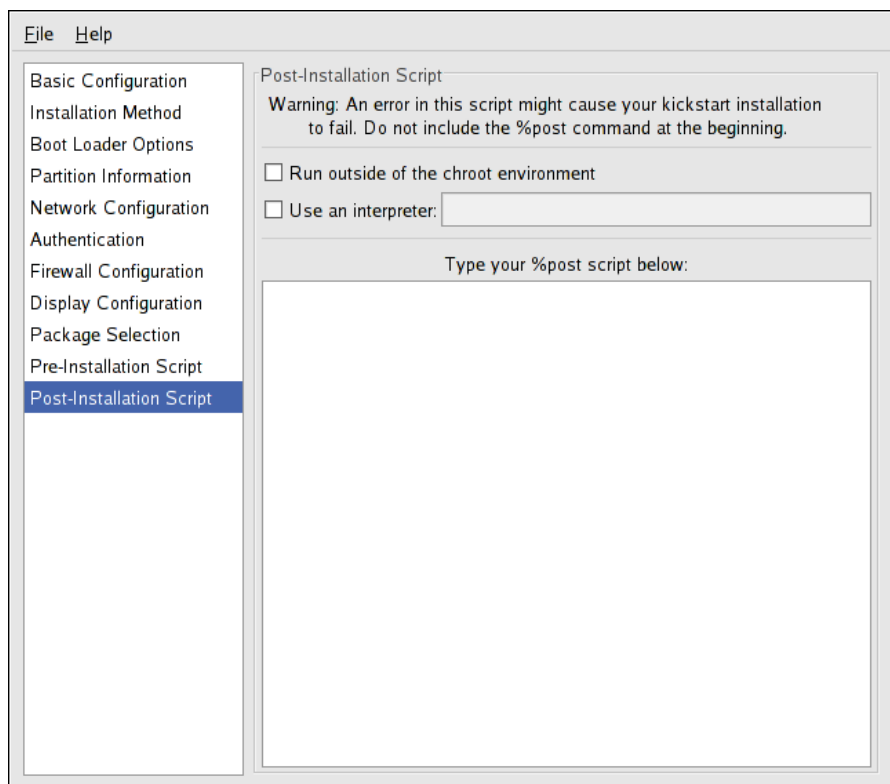
Figure 2.16. Post-Installation Script

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

> ⚠️ **Warning**
>
> Do not include the **%post** command. It is 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 in .

2.11.1. Chroot Environment

To run the post-installation script 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.

To make changes to the newly installed file system, within the post-installation section, but outside of the chroot environment, you must prepend the directory name with **/mnt/sysimage/**.

For example, if you select **Run outside of the chroot environment**, the previous example must be changed to the following:

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

2.11.2. Use an Interpreter

To specify a scripting language to use to execute the script, select the **Use an interpreter** option and enter the interpreter in the text box beside it. 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.

2.12. Saving the File

To review the contents of the kickstart file after you have finished choosing your kickstart options, select **File** => **Preview** from the pull-down menu.
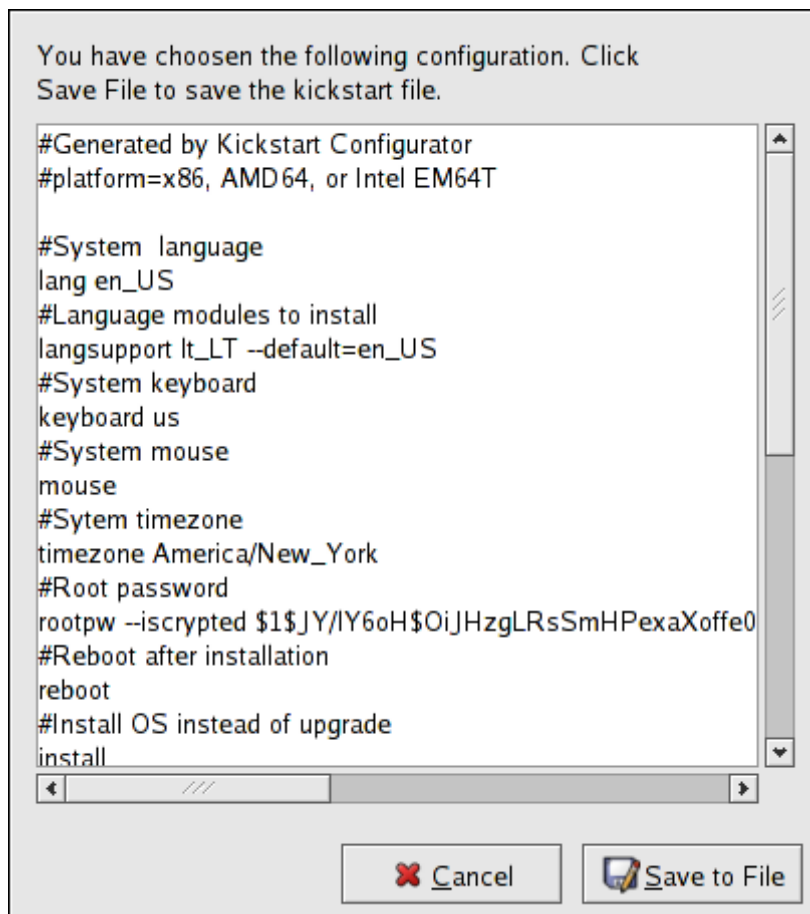
Figure 2.17. Preview

To save the kickstart file, click the **Save to File** button in the preview window. To save the file without previewing it, select **File** => **Save File** or press **Ctrl**+**S**. A dialog box appears. Select where to save the file.

After saving the file, refer to Section 1.10, "Starting a Kickstart Installation" for information on how to start the kickstart installation.

Chapter 3. PXE Network Installations

Red Hat Enterprise Linux allows for installation over a network using the NFS, FTP, or HTTP protocols. A network installation can be started from a boot CD-ROM, a bootable flash memory drive, or by using the **askmethod** boot option with the Red Hat Enterprise Linux CD #1. Alternatively, if the system to be installed contains a network interface card (NIC) with Pre-Execution Environment (PXE) support, it can be configured to boot from files on another networked system rather than local media such as a CD-ROM.

For a PXE network installation, the client's NIC with PXE support sends out a broadcast request for DHCP information. The DHCP server provides the client with an IP address, other network information such as name server, the IP address or hostname of the **tftp** server (which provides the files necessary to start the installation program), and the location of the files on the **tftp** server. This is possible because of PXELINUX, which is part of the **syslinux** package.

The following steps must be performed to prepare for a PXE installation:

1. Configure the network (NFS, FTP, HTTP) server to export the installation tree.
2. Configure the files on the **tftp** server necessary for PXE booting.
3. Configure which hosts are allowed to boot from the PXE configuration.
4. Start the **tftp** service.
5. Configure DHCP.
6. Boot the client, and start the installation.

3.1. Setting up the Network Server

First, configure an NFS, FTP, or HTTP server to export the entire installation tree for the version and variant of Red Hat Enterprise Linux to be installed. Refer to the section *Preparing for a Network Installation* in the *Installation Guide* for detailed instructions.

3.2. PXE Boot Configuration

The next step is to copy the files necessary to start the installation to the **tftp** server so they can be found when the client requests them. The **tftp** server is usually the same server as the network server exporting the installation tree.

To copy these files, run the **Network Booting Tool** on the NFS, FTP, or HTTP server. A separate PXE server is not necessary.

For the command line version of these instructions, refer to Section 3.2.1, "Command Line Configuration".

To use the graphical version of the **Network Booting Tool**, you must be running the X Window System, have root privileges, and have the **system-config-netboot** RPM package installed. To start the **Network Booting Tool** from the desktop, go to **Applications** (the main menu on the panel) => **System Settings** => **Server Settings** => **Network Booting Service**. Or, type the command **system-config-netboot** at a shell prompt (for example, in an **XTerm** or a **GNOME terminal**).

If starting the **Network Booting Tool** for the first time, select **Network Install** from the **First Time Druid**. Otherwise, select **Configure** => **Network Installation** from the pulldown menu, and then click **Add**. The dialog in Figure 3.1, "Network Installation Setup" is displayed.

Figure 3.1. Network Installation Setup

- **Operating system identifier** — Provide a unique name using one word to identify the Red Hat Enterprise Linux version and variant. It is used as the directory name in the **/tftpboot/linux-install/** directory.

- **Description** — Provide a brief description of the Red Hat Enterprise Linux version and variant.

- **Selects protocol for installation** — Selects NFS, FTP, or HTTP as the network installation type depending on which one was configured previously. If FTP is selected and anonymous FTP is not being used, uncheck **Anonymous FTP** and provide a valid username and password combination.

- **Kickstart** — Specify the location of the kickstart file. The file can be a URL or a file stored locally (diskette). The kickstart file can be created with the **Kickstart Configurator**. Refer to Chapter 2, *Kickstart Configurator* for details.

- **Server** — Provide the IP address or domain name of the NFS, FTP, or HTTP server.

- **Location** — Provide the directory shared by the network server. If FTP or HTTP was selected, the directory must be relative to the default directory for the FTP server or the document root for the HTTP server. For all network installations, the directory provided must contain the **RedHat/** directory of the installation tree.

After clicking **OK**, the **initrd.img** and **vmlinuz** files necessary to boot the installation program are transfered from **images/pxeboot/** in the provided installation tree to **/tftpboot/linux-install/<os-identifier>/** on the **tftp** server (the one you are running the **Network Booting Tool** on).

3.2.1. Command Line Configuration

If the network server is not running X, the **pxeos** command line utility, which is part of the **system-config-netboot** package, can be used to configure the **tftp** server files :

```
pxeos -a -i "<description>" -p <NFS|HTTP|FTP> -D 0 -s client.example.com \ -L
<net-location> -k <kernel> -K <kickstart><os-identifer>
```

The following list explains the options:

- **-a** — Specifies that an OS instance is being added to the PXE configuration.

- **-i** "**<description>**" — Replace "**<description>**" with a description of the OS instance. This corresponds to the **Description** field in Figure 3.1, "Network Installation Setup".

- **-p** <NFS|HTTP|FTP> — Specify which of the NFS, FTP, or HTTP protocols to use for installation. Only one may be specified. This corresponds to the **Select protocol for installation** menu in Figure 3.1, "Network Installation Setup".

- **-D** <0|1> — Specify "**0**" which indicates that it is *not* a diskless configuration since **pxeos** can be used to configure a diskless environment as well.

- **-s***client.example.com* — Provide the name of the NFS, FTP, or HTTP server after the **-s** option. This corresponds to the **Server** field in Figure 3.1, "Network Installation Setup".

- **-L***<net-location>* — Provide the location of the installation tree on that server after the **-L** option. This corresponds to the **Location** field in Figure 3.1, "Network Installation Setup".

- **-k***<kernel>* — Provide the specific kernel version of the server installation tree for booting.

- **-K***<kickstart>* — Provide the location of the kickstart file, if available.

- ***<os-identifer>*** — Specify the OS identifier, which is used as the directory name in the **/tftpboot/linux-install/** directory. This corresponds to the **Operating system identifier** field in Figure 3.1, "Network Installation Setup".

If FTP is selected as the installation protocol and anonymous login is not available, specify a username and password for login, with the following options before ***<os-identifer>*** in the previous command:

```
-A 0 -u <username> -p <password>
```

For more information on command line options available for the **pxeos** command, refer to the **pxeos** man page.

3.3. Adding PXE Hosts

After configuring the network server, the interface as shown in Figure 3.2, "Add Hosts" is displayed.
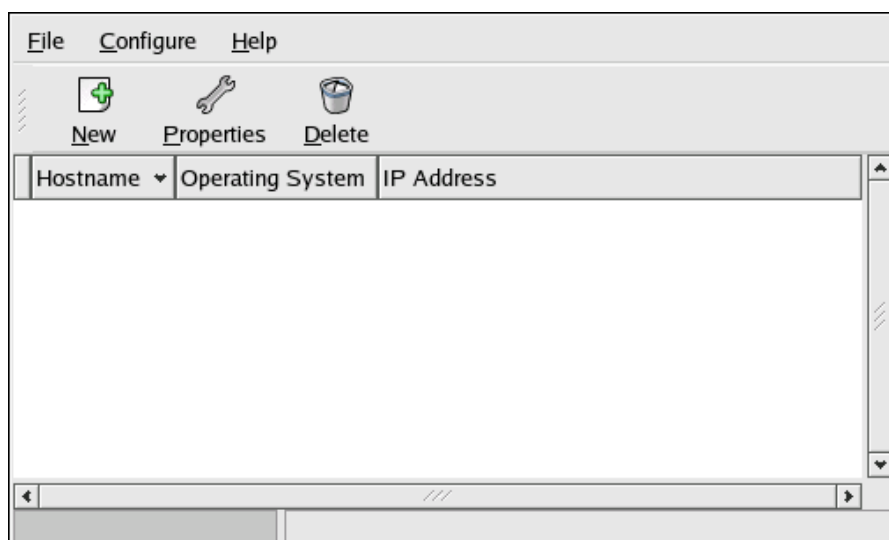
Figure 3.2. Add Hosts

The next step is to configure which hosts are allowed to connect to the PXE boot server. For the command line version of this step, refer to Section 3.3.1, "Command Line Configuration".

To add hosts, click the **New** button.

Figure 3.3. Add a Host

Enter the following information:

- **Hostname or IP Address/Subnet** — The IP address, fully qualified hostname, or a subnet of systems that should be allowed to connect to the PXE server for installations.
- **Operating System** — The operating system identifier to install on this client. The list is populated from the network install instances created from the **Network Installation Dialog**.
- **Serial Console** — This option allows use of a serial console.
- **Kickstart File** — The location of a kickstart file to use, such as **http://server.example.com/kickstart/ks.cfg**. This file can be created with the **Kickstart Configurator**. Refer to Chapter 2, *Kickstart Configurator* for details.

Ignore the **Snapshot name** and **Ethernet** options. They are only used for diskless environments. For more information on configuring a diskless environment, refer to Chapter 4, *Diskless Environments* for details.

3.3.1. Command Line Configuration

If the network server is not running X, the **pxeboot** utility, a part of the **system-config-netboot** package, can be used to add hosts which are allowed to connect to the PXE server:

```
pxeboot -a -K <kickstart> -O <os-identifier> -r <value><host>
```

The following list explains the options:

- **-a** — Specifies that a host is to be added.
- **-K<kickstart>** — The location of the kickstart file, if available.
- **-O<os-identifier>** — Specifies the operating system identifier as defined in Section 3.2, "PXE Boot Configuration".
- **-r<value>** — Specifies the ram disk size.
- **<host>** — Specifies the IP address or hostname of the host to add.

For more information on command line options available for the **pxeboot** command, refer to the **pxeboot** man page.

3.4. Adding a Custom Boot Message

Optionally, modify **`/tftpboot/linux-install/msgs/boot.msg`** to use a custom boot message.

3.5. Performing the PXE Installation

For instructions on how to configure the network interface card with PXE support to boot from the network, consult the documentation for the NIC. It varies slightly per card.

After the system boots the installation program, refer to the *Installation Guide*.

Chapter 4. Diskless Environments

Some networks require multiple systems with the same configuration. They also require that these systems be easy to reboot, upgrade, and manage. One solution is to use a *diskless environment* in which most of the operating system, which can be read-only, is shared from a central server between the clients. The individual clients have their own directories on the central server for the rest of the operating system, which must be read/write. Each time the client boots, it mounts most of the OS from the NFS server as read-only and another directory as read-write. Each client has its own read-write directory so that one client can not affect the others.

The following steps are necessary to configure Red Hat Enterprise Linux to run on a diskless client:

1. Install Red Hat Enterprise Linux on a system so that the files can be copied to the NFS server. (Refer to the *Installation Guide* for details.) Any software to be used on the clients must be installed on this system and the **busybox-anaconda** package must be installed.

2. Create a directory on the NFS server to contain the diskless environment such as **/diskless/i386/RHEL4-AS/**. For example:

   ```
   mkdir -p /diskless/i386/RHEL4-AS/
   ```

   This directory is referred to as the **diskless directory**.

3. Create a subdirectory of this directory named **root/**:

   ```
   mkdir -p /diskless/i386/RHEL4-AS/root/
   ```

4. Copy Red Hat Enterprise Linux from the client system to the server using **rsync**. For example:

   ```
   rsync -a -e ssh installed-system.example.com:/ /diskless/i386/RHEL4-
   AS/root/
   ```

   The length of this operation depends on the network connection speed as well as the size of the file system on the installed system. Depending on these factors, this operation may take a while.

5. Start the **tftp** server

6. Configure the DHCP server

7. Finish creating the diskless environment as discussed in Section 4.2, "Finish Configuring the Diskless Environment".

8. Configure the diskless clients as discussed in Section 4.3, "Adding Hosts".

9. Configure each diskless client to boot via PXE and boot them.

4.1. Configuring the NFS Server

The shared read-only part of the operating system is shared via NFS.

Configure NFS to export the **root/** and **snapshot/** directories by adding them to **/etc/exports**. For example:

```
/diskless/i386/RHEL4-AS/root/     *(ro,sync,no_root_squash)
/diskless/i386/RHEL4-AS/snapshot/ *(rw,sync,no_root_squash)
```

Replace ***** with one of the hostname formats discussed in Section 21.3.2, "Hostname Formats". Make the hostname declaration as specific as possible, so unwanted systems can not access the NFS mount.

If the NFS service is not running, start it:

```
service nfs start
```

If the NFS service is already running, reload the configuration file:

```
service nfs reload
```

4.2. Finish Configuring the Diskless Environment

To use the graphical version of the **Network Booting Tool**, you must be running the X Window System, have root privileges, and have the **system-config-netboot** RPM package installed. To start the **Network Booting Tool** from the desktop, go to **Applications** (the main menu on the panel) => **System Settings** => **Server Settings** => **Network Booting Service**. Or, type the command **system-config-netboot** at a shell prompt (for example, in an **XTerm** or a **GNOME terminal**).

If starting the **Network Booting Tool** for the first time, select **Diskless** from the **First Time Druid**. Otherwise, select **Configure** => **Diskless** from the pull-down menu, and then click **Add**.

A wizard appears to step you through the process:

1. Click **Forward** on the first page.
2. On the **Diskless Identifier** page, enter a **Name** and **Description** for the diskless environment. Click **Forward**.
3. Enter the IP address or domain name of the NFS server configured in Section 4.1, "Configuring the NFS Server" as well as the directory exported as the diskless environment. Click **Forward**.
4. The kernel versions installed in the diskless environment are listed. Select the kernel version to boot on the diskless system.
5. Click **Apply** to finish the configuration.

After clicking **Apply**, the diskless kernel and image file are created based on the kernel selected. They are copied to the PXE boot directory **/tftpboot/linux-install/<os-identifier>/**. The directory **snapshot/** is created in the same directory as the **root/** directory (for example, **/diskless/i386/RHEL4-AS/snapshot/**) with a file called **files** in it. This file contains a list of files and directories that must be read/write for each diskless system. Do not modify this file. If additional entries must be added to the list, create a **files.custom** file in the same directory as the **files** file, and add each additional file or directory on a separate line.

4.3. Adding Hosts
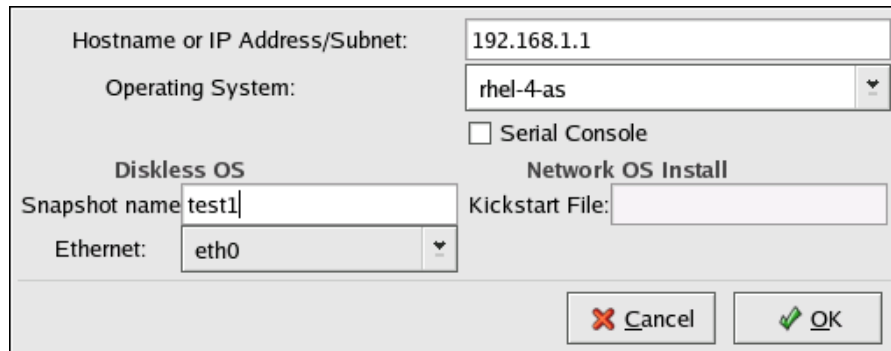
Each diskless client must have its own *snapshot* directory on the NFS server that is used as its read/write file system. The **Network Booting Tool** can be used to create these snapshot directories.

After completing the steps in Section 4.2, "Finish Configuring the Diskless Environment", a window appears to allow hosts to be added for the diskless environment. Click the **New** button. In the dialog shown in Figure 4.1, "Add Diskless Host", provide the following information:

- **Hostname or IP Address/Subnet** — Specify the hostname or IP address of a system to add it as a host for the diskless environment. Enter a subnet to specify a group of systems.
- **Operating System** — Select the diskless environment for the host or subnet of hosts.
- **Serial Console** — Select this checkbox to perform a serial installation.

- **Snapshot name** — Provide a subdirectory name to be used to store all of the read/write content for the host.
- **Ethernet** — Select the Ethernet device on the host to use to mount the diskless environment. If the host only has one Ethernet card, select **eth0**.

Ignore the **Kickstart File** option. It is only used for PXE installations.

Figure 4.1. Add Diskless Host

In the existing **snapshot/** directory in the diskless directory, a subdirectory is created with the **Snapshot name** specified as the file name. Then, all of the files listed in **snapshot/files** and **snapshot/files.custom** are copied copy from the **root/** directory to this new directory.

4.4. Booting the Hosts

Consult the documentation for your PXE card to configure the host to boot via PXE.

When the diskless client boots, it mounts the remote **root/** directory in the diskless directory as read-only. It also mounts its individual snapshot directory as read/write. Then it mounts all the files and directories in the **files** and **files.custom** files using the **mount -o bind** over the read-only diskless directory to allow applications to write to the root directory of the diskless environment if they need to.

Chapter 5. Basic System Recovery

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, single-user mode, and emergency mode, where you can use your own knowledge to repair the system.

5.1. Common Problems

You might need to boot into one of these recovery modes for any of the following reasons:

- You are unable to boot normally into Red Hat Enterprise Linux (runlevel 3 or 5).
- You are having hardware or software problems, and you want to get a few important files off of your system's hard drive.
- You forgot the root password.

5.1.1. Unable to Boot into Red Hat Enterprise Linux

This problem is often caused by the installation of another operating system after you have installed Red Hat Enterprise Linux. Some other operating systems assume that you have no other operating system(s) on your computer. They overwrite the Master Boot Record (MBR) that originally contained the GRUB boot loader. If the boot loader is overwritten in this manner, you cannot boot Red Hat Enterprise Linux unless you can get into rescue mode and reconfigure the boot loader.

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 might not be able to find it to mount the partition. To fix this problem, boot in rescue mode and modify the **/boot/grub/grub.conf** file.

For instructions on how to reinstall the GRUB boot loader from a rescue environment, refer to Section 5.2.1, "Reinstalling the Boot Loader".

5.1.2. Hardware/Software Problems

This category includes a wide variety of different situations. Two examples include failing hard drives and specifying an invalid root device or kernel in the boot loader configuration file. If either of these occur, you might not be able to reboot into Red Hat Enterprise Linux. However, if you boot into one of the system recovery modes, you might be able to resolve the problem or at least get copies of your most important files.

5.1.3. Root Password

What can you do if you forget your root password? To reset it to a different password, boot into rescue mode or single-user mode, and use the **passwd** command to reset the root password.

5.2. Booting into Rescue Mode

Rescue mode provides the ability to boot a small Red Hat Enterprise Linux environment entirely from CD-ROM, or some other boot method, instead of the system's hard drive.

As the name implies, rescue mode is provided to rescue you from something. During normal operation, your Red Hat Enterprise 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 Red Hat Enterprise Linux running completely enough to access 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 Red Hat Enterprise Linux from that hard drive.

To boot into rescue mode, you must be able to boot the system using one of the following methods [1]:

- By booting the system from an installation boot CD-ROM.
- By booting the system from other installation boot media, such as USB flash devices.
- By booting the system from the Red Hat Enterprise Linux CD-ROM #1.

Once you have booted using one of the described methods, add the keyword **rescue** as a kernel parameter. For example, for an x86 system, type the following command at the installation boot prompt:

```
linux rescue
```

You are prompted to answer a few basic questions, including which language to use. It also prompts you to select where a valid rescue image is located. Select from **Local CD-ROM**, **Hard Drive**, **NFS image**, **FTP**, or **HTTP**. The location selected must contain a valid installation tree, and the installation tree must be for the same version of Red Hat Enterprise Linux as the Red Hat Enterprise Linux disk from which you booted. If you used a boot CD-ROM or other media to start rescue mode, the installation tree must be from the same tree from which the media was created. For more information about how to setup an installation tree on a hard drive, NFS server, FTP server, or HTTP server, refer to the earlier section of this guide.

If you select a rescue image that does not require a network connection, you are asked whether or not you want to establish a network connection. A network connection is useful if you need to backup files to a different computer or install some RPM packages from a shared network location, for example.

The following message is displayed:

```
The rescue environment will now attempt to find your 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 file systems 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 attempts to mount your file system under the directory **/mnt/sysimage/**. If it fails to mount a partition, it notifies you. If you select **Read-Only**, it attempts to mount your file system under the directory **/mnt/sysimage/**, but in read-only mode. If you select **Skip**, your file system is not mounted. Choose **Skip** if you think your file system 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-3.00b#
```

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

Even if your file system is mounted, the default root partition while in rescue mode is a temporary root partition, not the root partition of the file system used during normal user mode (runlevel 3 or 5). If you selected to mount your file system and it mounted successfully, you can change the root partition of the rescue mode environment to the root partition of your file system by executing the following 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** to return to the prompt.

If you selected **Skip**, you can still try to mount a partition or LVM2 logical volume manually inside rescue mode by creating a directory such as */foo*, and typing the following command:

```
mount -t ext3 /dev/mapper/VolGroup00-LogVol02 /foo
```

In the above command, */foo* is a directory that you have created and */dev/mapper/VolGroup00-LogVol02* is the LVM2 logical volume you want to mount. If the partition is of type **ext2**, replace **ext3** with **ext2**.

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

```
fdisk -l
```

If you do not know the names of all LVM2 physical volumes, volume groups, or logical volumes, use the following commands to list them:

```
pvdisplay
```

```
vgdisplay
```

```
lvdisplay
```

From the prompt, you can run many useful commands, such as:

- **ssh**, **scp**, and **ping** if the network is started
- **dump** and **restore** for users with tape drives
- **parted** and **fdisk** for managing partitions
- **rpm** for installing or upgrading software
- **joe** for editing configuration files

> **Note**
>
> If you try to start other popular editors such as **emacs**, **pico**, or **vi**, the **joe** editor is started.

5.2.1. Reinstalling the Boot Loader

In many cases, the GRUB boot loader can mistakenly be deleted, corrupted, or replaced by other operating systems.

The following steps detail the process on how GRUB is reinstalled on the master boot record:

- Boot the system from an installation boot medium.
- Type **linux rescue** at the installation boot prompt to enter the rescue environment.
- Type **chroot /mnt/sysimage** to mount the root partition.

▷ Type **/sbin/grub-install /dev/hda** to reinstall the GRUB boot loader, where **/dev/hda** is the boot partition.

▷ Review the **/boot/grub/grub.conf** file, as additional entries may be needed for GRUB to control additional operating systems.

▷ Reboot the system.

5.3. Booting into Single-User Mode

One of the advantages of single-user mode is that you do not need a boot CD-ROM; however, it does not give you the option to mount the file systems as read-only or not mount them at all.

If your system boots, but does not allow you to log in when it has completed booting, try single-user mode.

In single-user mode, your computer boots to runlevel 1. Your local file systems are mounted, but your network is not activated. You have a usable system maintenance shell. Unlike rescue mode, single-user mode automatically tries to mount your file system. *Do not use single-user mode if your file system cannot be mounted successfully.* You cannot use single-user mode if the runlevel 1 configuration on your system is corrupted.

On an x86 system using GRUB, use the following steps to boot into single-user mode:

1. At the GRUB splash screen at boot time, press any key to enter the GRUB interactive menu.
2. Select **Red Hat Enterprise Linux** with the version of the kernel that you wish to boot and type **a** to append the line.
3. 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.

5.4. Booting into Emergency Mode

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

To boot into emergency mode, use the same method as described for single-user mode in Section 5.3, "Booting into Single-User Mode" with one exception, replace the keyword **single** with the keyword **emergency**.

---

[1] Refer to the earlier sections of this guide for more details.

Part II. File Systems

File system refers to the files and directories stored on a computer. A file system can have different formats called *file system types*. These formats determine how the information is stored as files and directories. Some file system types store redundant copies of the data, while some file system types make hard drive access faster. This part discusses the ext3, swap, RAID, and LVM file system types. It also discusses the `parted` utility to manage partitions and access control lists (ACLs) to customize file permissions.

Chapter 6. The ext3 File System

The default file system is the journaling *ext3* file system.

6.1. 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, the ext3 volumes are configured 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 migrate from ext2 to ext3 and gain the benefits of a robust journaling file system without reformatting. Refer to Section 6.3, "Converting to an ext3 File System" for more on how to perform this task.

The following sections walk you through the steps for creating and tuning ext3 partitions. For ext2 partitions, skip the partitioning and formating sections below and go directly to Section 6.3, "Converting to an ext3 File System".

6.2. 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 the system, you may want to 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 the **/etc/fstab** file.

6.3. 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 is visible as the file **.journal** in the root directory of the file system. If the file system is not mounted, the journal is hidden and does not appear in the file system at all.

> **Note**
>
> A default installation of Red Hat Enterprise Linux uses ext3 for all file systems.

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

```
/sbin/tune2fs -j <file_system>
```

where **<file_system>** is an appropriate LVM2 file system.

A valid LVM2 file system could be one of two types of entries:

» A mapped device — A logical volume in a volume group, for example, **/dev/mapper/VolGroup00-LogVol02**.

» A static device — A traditional storage volume, for example, **/dev/hdbX**, where *hdb* is a storage device name and *X* is the partition number.

Issue the **df** command to display mounted file systems. For more detailed information on the LVM file system, refer to Chapter 8, *LVM Configuration*.

For the remainder of this section, the sample commands use the following value:

```
/dev/mapper/VolGroup00-LogVol02
```

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

If you are transitioning your root file system, you must 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 configuration loads the **initrd**.

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

6.4. 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/mapper/VolGroup00-LogVol02
```

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

```
/sbin/tune2fs -O ^has_journal /dev/mapper/VolGroup00-LogVol02
```

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

```
/sbin/e2fsck -y /dev/mapper/VolGroup00-LogVol02
```

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

```
mount -t ext2 /dev/mapper/VolGroup00-LogVol02/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 want to permanently change the partition to ext2, remember to update the **/etc/fstab** file.

> **Tip**
>
> You can increase the size of an ext3 file system using **ext2online**. **ext2online** allows you to increase the size of an ext3 file system once it is mounted (online) and on a resizable logical volume. The root file system is set up by default on LVM2 logical volumes during installation. Note that **ext2online** will only work on ext3 file systems. For more information, refer to **man ext2online**.

Chapter 7. Logical Volume Manager (LVM)

7.1. What is LVM?

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

With LVM, a hard drive or set of hard drives is allocated to one or more *physical volumes*. A physical volume cannot 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 cannot be on a logical volume group because the boot loader cannot read it. If the root (**/**) partition is on a logical volume, create a separate **/boot/** partition which is not a part of a volume group.

Since a physical volume cannot span over multiple drives, to span over more than one drive, create one or more physical volumes per drive.
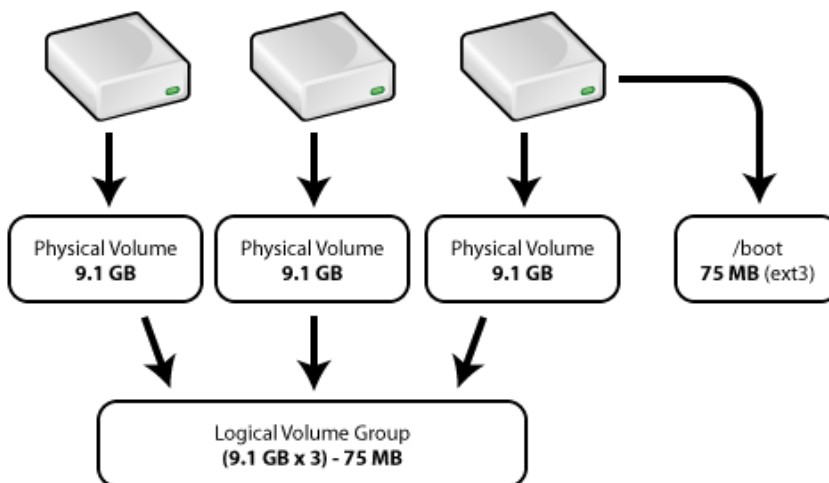
Figure 7.1. Logical Volume Group

The logical volume group is divided into *logical volumes*, which are assigned mount points, such as **/home** and **/** m and file system types, such as ext2 or 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 partitions that are logical volumes can be expanded.
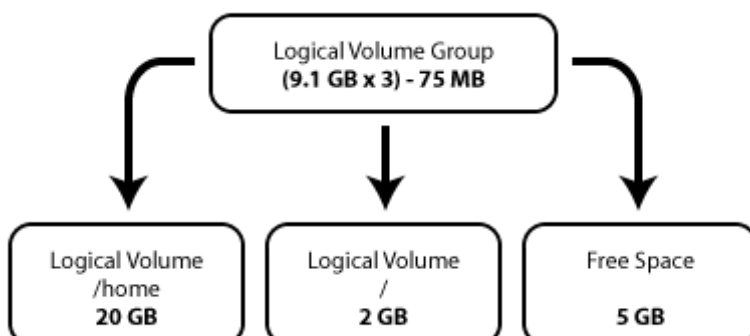
Figure 7.2. Logical Volumes

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, and the default Red Hat kernel is compiled with LVM support.

To learn how to configure LVM during the installation process, refer to Chapter 8, *LVM Configuration*.

7.2. What is LVM2?

LVM version 2, or LVM2, is the default for Red Hat Enterprise Linux, which uses the device mapper driver contained in the 2.6 kernel. LVM2, which is almost completely compatible with the earlier LVM1 version, can be upgraded from versions of Red Hat Enterprise Linux running the 2.4 kernel.

Although upgrading from LVM1 to LVM2 is usually seamless, refer to Section 7.3, "Additional Resources" for further details on more complex requirements and upgrading scenarios.

7.3. Additional Resources

Use these sources to learn more about LVM.

7.3.1. Installed Documentation

- `rpm -qd lvm` — This command shows all the documentation available from the `lvm` package, including man pages.
- `lvm help` — This command shows all LVM commands available.

7.3.2. Useful Websites

- http://sourceware.org/lvm2 — LVM2 webpage, which contains an overview, link to the mailing lists, and more.
- http://tldp.org/HOWTO/LVM-HOWTO/ — *LVM HOWTO* from the Linux Documentation Project.

Chapter 8. LVM Configuration

LVM can be configured during the graphical installation process, the text-based installation process, or during a kickstart installation. You can use the utilities from the `lvm` package to create your own LVM configuration post-installation, but these instructions focus on using **Disk Druid** during installation to complete this task.

Read Chapter 7, *Logical Volume Manager (LVM)* first to learn about LVM. An overview of the steps required to configure LVM include:

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

> **Note**
>
> Although the following steps are illustrated during a GUI installation, the same can be done during a text-based installation.

Two 9.1 GB SCSI drives (`/dev/sda` and `/dev/sdb`) are used in the following examples. They detail how to create a simple configuration using a single LVM volume group with associated logical volumes during installation.

8.1. Automatic Partitioning

On the `Disk Partitioning Setup` screen, select `Automatically partition`.

For Red Hat Enterprise Linux, LVM is the default method for disk partitioning. If you do not wish to have LVM implemented, or if you require RAID partitioning, manual disk partitioning through **Disk Druid** is required.

The following properties make up the automatically created configuration:

- The `/boot/` partition resides on its own non-LVM partition. In the following example, it is the first partition on the first drive (`/dev/sda1`). Bootable partitions *cannot* reside on LVM logical volumes.
- A single LVM volume group (`VolGroup00`) is created, which spans all selected drives and all remaining space available. In the following example, the remainder of the first drive (`/dev/sda2`), and the entire second drive (`/dev/sdb1`) are allocated to the volume group.
- Two LVM logical volumes (`LogVol00` and `LogVol01`) are created from the newly created spanned volume group. In the following example, the recommended swap space is automatically calculated and assigned to `LogVol01`, and the remainder is allocated to the root file system, `LogVol00`.
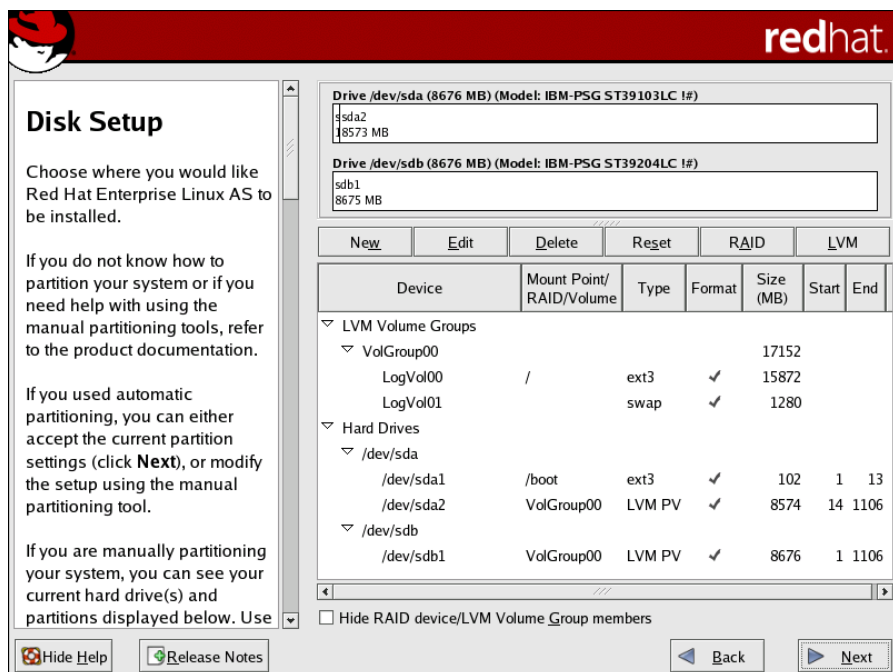
Figure 8.1. Automatic LVM Configuration With Two SCSI Drives

> **Note**
>
> If enabling quotas are of interest to you, it may be best to modify the automatic configuration to include other mount points, such as `/home/` or `/var/`, so that each file system has its own independent quota configuration limits.
> In most cases, the default automatic LVM partitioning is sufficient, but advanced implementations could warrant modification or manual configuration of the LVM partition tables.

> **Note**
>
> If you anticipate future memory upgrades, leaving some free space in the volume group would allow for easy future expansion of the swap space logical volume on the system; in which case, the automatic LVM configuration should be modified to leave available space for future growth.

8.2. Manual LVM Partitioning

The following section explains how to manually configure LVM for Red Hat Enterprise Linux. Because there are numerous ways to manually configure a system with LVM, the following example is similar to the default configuration done in Section 8.1, "Automatic Partitioning".

On the **Disk Partitioning Setup** screen, select **Manually partition with Disk Druid**.

8.2.1. Creating the `/boot/` Partition

In a typical situation, the disk drives are new, or formatted clean. The following figure, Figure 8.2, "Two Blank Drives, Ready For Configuration", shows both drives as raw devices with no partitioning configured.
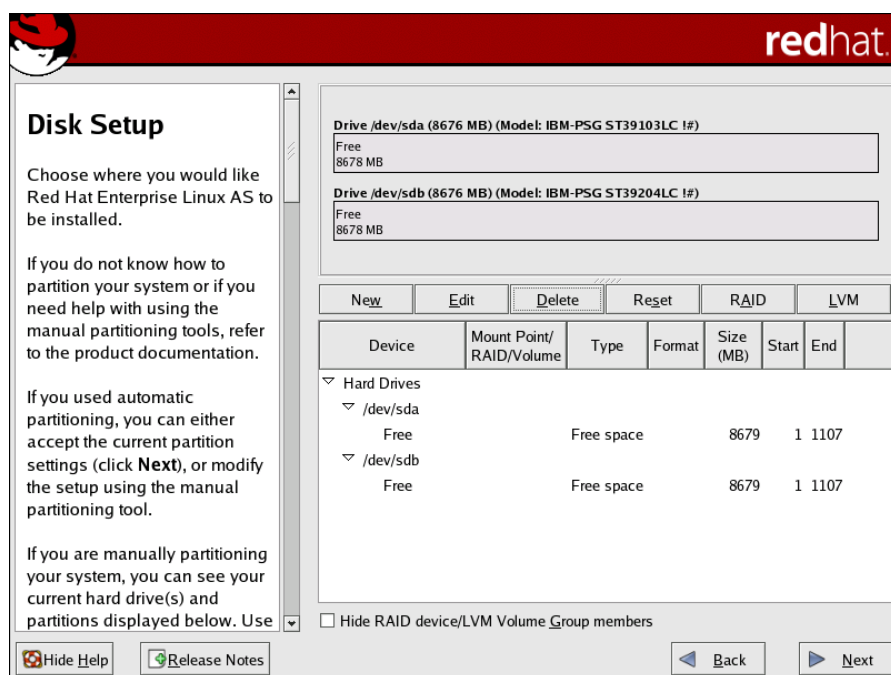
Figure 8.2. Two Blank Drives, Ready For Configuration

⚠️ **Warning**

The **/boot/** partition cannot reside on an LVM volume group because the GRUB boot loader cannot read it.

---

1. Select **New**.

2. Select **/boot** from the **Mount Point** pulldown menu.

3. Select **ext3** from the **File System Type** pulldown menu.

4. Select only the **sda** checkbox from the **Allowable Drives** area.

5. Leave **100** (the default) in the **Size (MB)** menu.

6. Leave the **Fixed size** (the default) radio button selected in the **Additional Size Options** area.

7. Select **Force to be a primary partition** to make the partition be a primary partition. A primary partition is one of the first four partitions on the hard drive. If unselected, the partition is created as a logical partition. If other operating systems are already on the system, unselecting this option should be considered. For more information on primary versus logical/extended partitions, refer to the appendix section of the *Installation Guide*.

Refer to Figure 8.3, "Creation of the Boot Partition" to verify your inputted values:
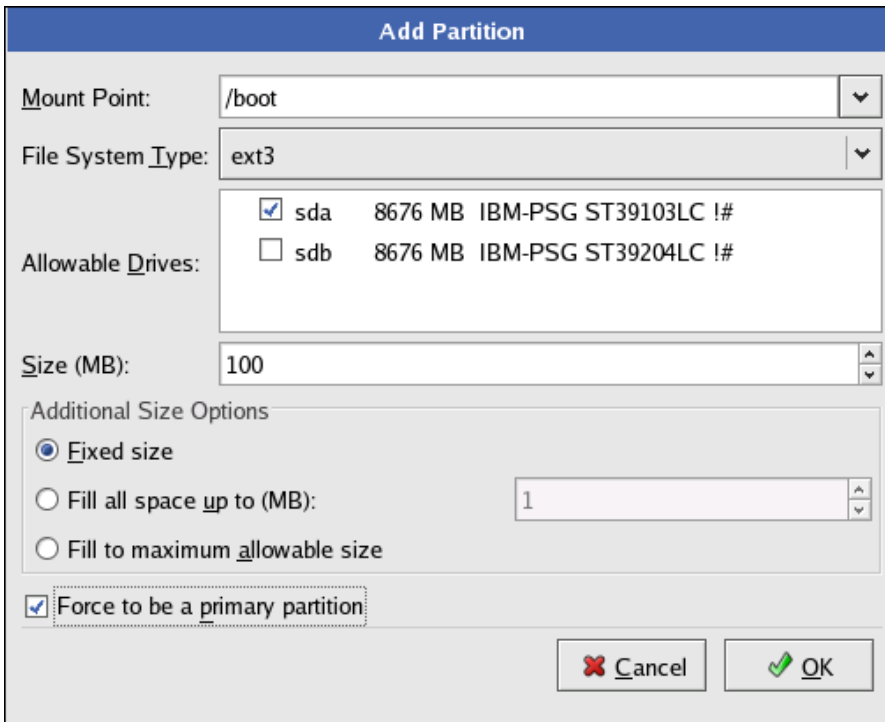
Figure 8.3. Creation of the Boot Partition

Click **OK** to return to the main screen. The following figure displays the boot partition correctly set:
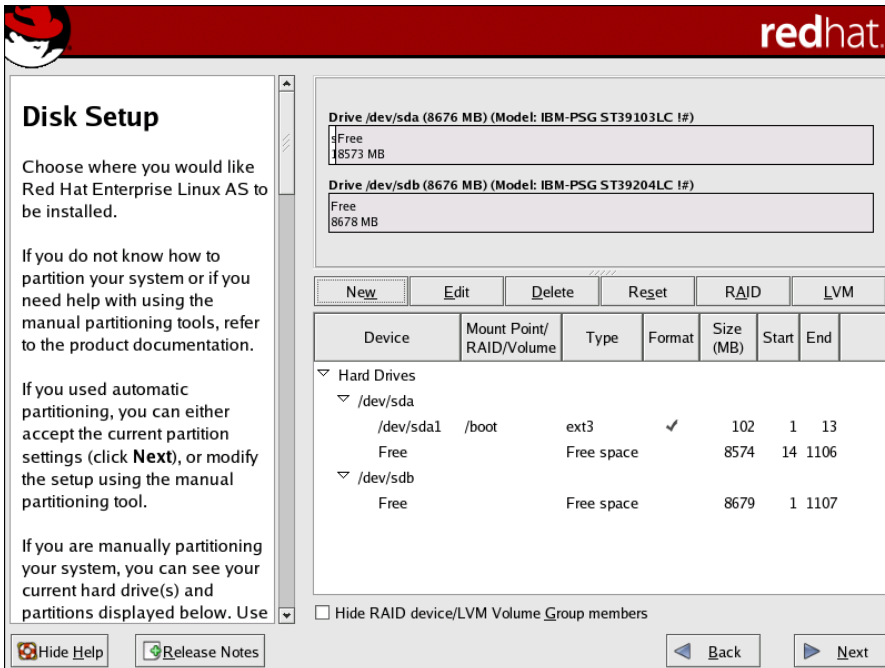
Figure 8.4. The /boot/ Partition Displayed

8.2.2. Creating the LVM Physical Volumes

Once the boot partition is created, the remainder of all disk space can be allocated to LVM partitions. The first step in creating a successful LVM implementation is the creation of the physical volume(s).

1. Select **New**.
2. Select **physical volume (LVM)** from the **File System Type** pulldown menu as shown in Figure 8.5, "Creating a Physical Volume".

Figure 8.5. Creating a Physical Volume

3. You cannot enter a mount point yet (you can once you have created all your physical volumes and then all volume groups).
4. A physical volume must be constrained to one drive. For **Allowable Drives**, select the drive on which the physical volume are created. If you have multiple drives, all drives are selected, and you must deselect all but one drive.
5. Enter the size that you want the physical volume to be.
6. 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 share the available free space on the disk.
7. Select **Force to be a primary partition** if you want the partition to be a primary partition.
8. Click **OK** to return to the main screen.

Repeat these steps 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. The following figure shows both drives completed after the repeated process:
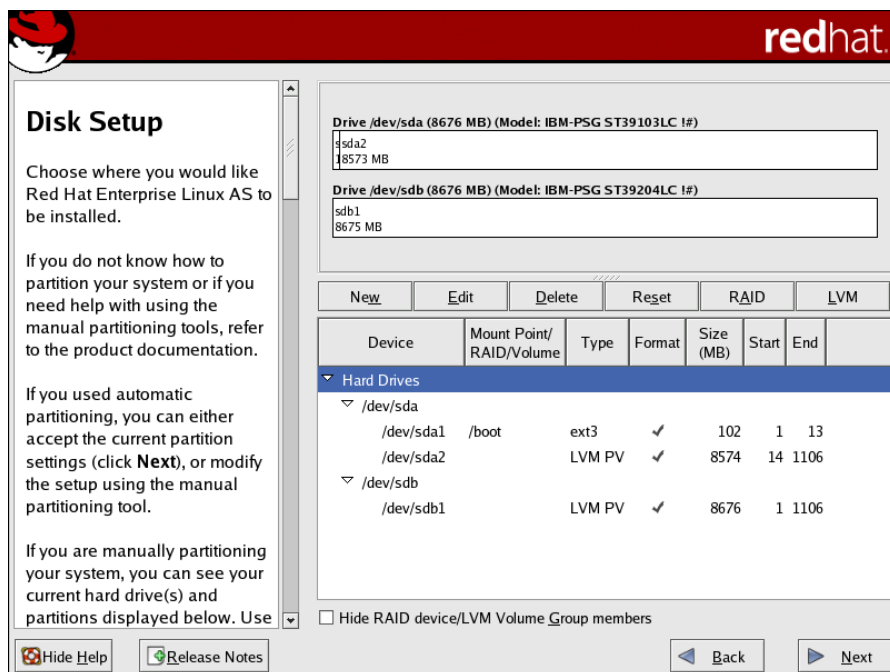
Figure 8.6. Two Physical Volumes Created

8.2.3. Creating the LVM Volume Groups

Once all the physical volumes are created, the volume groups can be created:

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 8.7. Creating an LVM Volume Group

2. Change the **Volume Group Name** if desired.
3. All logical volumes inside the volume group must be allocated in *physical extent* units. By default, the physical extent is set to 32 MB; thus, logical volume sizes must be divisible by 32 MBs. If you enter a size that is not a unit of 32 MBs, the installation program automatically selects the closest size in units of 32 MBs. It is not recommended that you change this setting.
4. Select which physical volumes to use for the volume group.

8.2.4. Creating the LVM Logical Volumes

Create logical volumes with mount points such as **/**, **/home/**, and swap space. Remember that **/boot** cannot 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 8.8, "Creating a Logical Volume" appears.

Figure 8.8. 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. The default automatic configuration does not do this, but this manual configuration example does — approximately 1 GB is left as free space for future expansion.

Figure 8.9. Pending Logical Volumes

Click **OK** to apply the volume group and all associated logical volumes.

The following figure shows the final manual configuration:
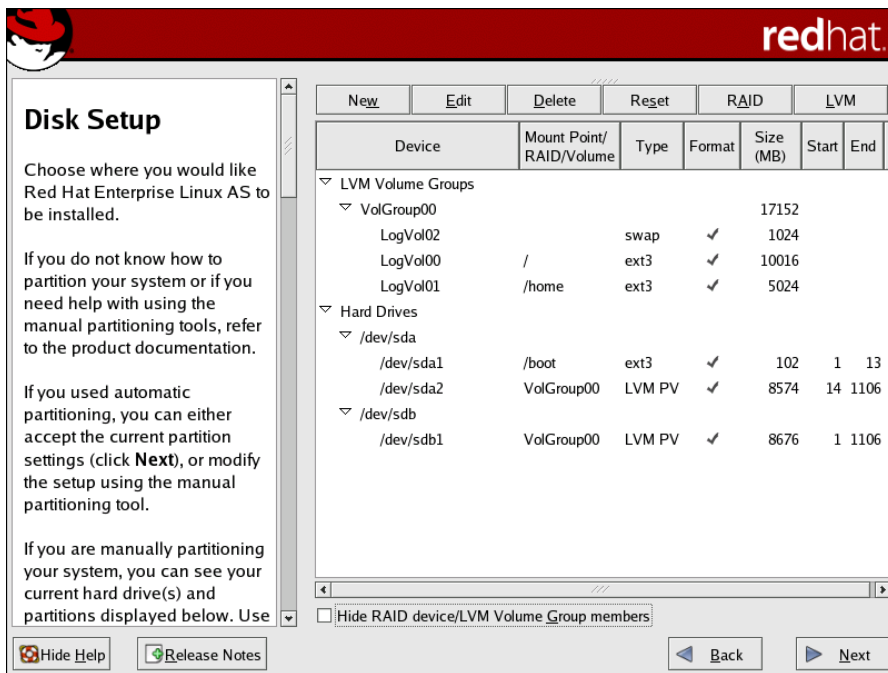
Figure 8.10. Final Manual Configuration

Chapter 9. Redundant Array of Independent Disks (RAID)

9.1. 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 appears to the computer as a single logical storage unit or drive.

RAID is a method in which information is spread across several disks. RAID uses 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 to increase bandwidth for reading or writing to disks, and to 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 the RAID array according to the RAID level used. When the data is to be read, the process is reversed, giving the illusion that the multiple drives in the array are actually one large drive.

9.2. Who Should Use RAID?

Those who need to keep large quantities of data on hand (such as system administrators) 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

9.3. Hardware RAID versus Software RAID

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

9.3.1. Hardware RAID

The hardware-based array 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.

9.3.2. 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 [2] 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 during installation, refer to the Chapter 10, *Software RAID Configuration*.

For those interested in learning more about what Software RAID has to offer, here are 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

9.4. 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. [3] 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.

- *Level 4* — Level 4 uses parity [4] 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 Enterprise Linux RAID installations. [5] 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.

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

[3] 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.

[4] 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.

[5] 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.

Chapter 10. Software RAID Configuration

Software RAID can be configured during the graphical installation process, the text-based installation process, or during a kickstart installation. This chapter discusses how to configure software RAID during installation, using the **Disk Druid** interface.

Read Chapter 9, *Redundant Array of Independent Disks (RAID)* first to learn about RAID, the differences between hardware and software RAID, and the differences between RAID 0, 1, and 5. An overview of the steps required to configure RAID include:

- Applying *software RAID partitions* to the physical hard drives.

  If you wish to have the boot partition (**/boot/**) reside on a RAID parition, it *must* be on a RAID 1 partition.
- Creating *RAID devices* from the software RAID partitions.
- *Optional:* Configuring *LVM* from the RAID devices. Refer to Chapter 8, *LVM Configuration* for more information on configuring LVM after first configuring RAID.
- Creating *file systems* from the RAID devices.

> **Note**
>
> Although the following steps are illustrated during a GUI installation, the same can be done during a text-based installation.
> Configuration of software RAID must be done manually in **Disk Druid** during the installation process.

Two 9.1 GB SCSI drives (**/dev/sda** and **/dev/sdb**) are used in the following examples. They detail how to create a simple RAID 1 configuration by implementing multiple RAID devices.

On the **Disk Partitioning Setup** screen, select **Manually partition with Disk Druid**.

10.1. Creating the RAID Partitions

In a typical situation, the disk drives are new or are formatted. Both drives are shown as raw devices with no partition configuration in Figure 10.1, "Two Blank Drives, Ready For Configuration".
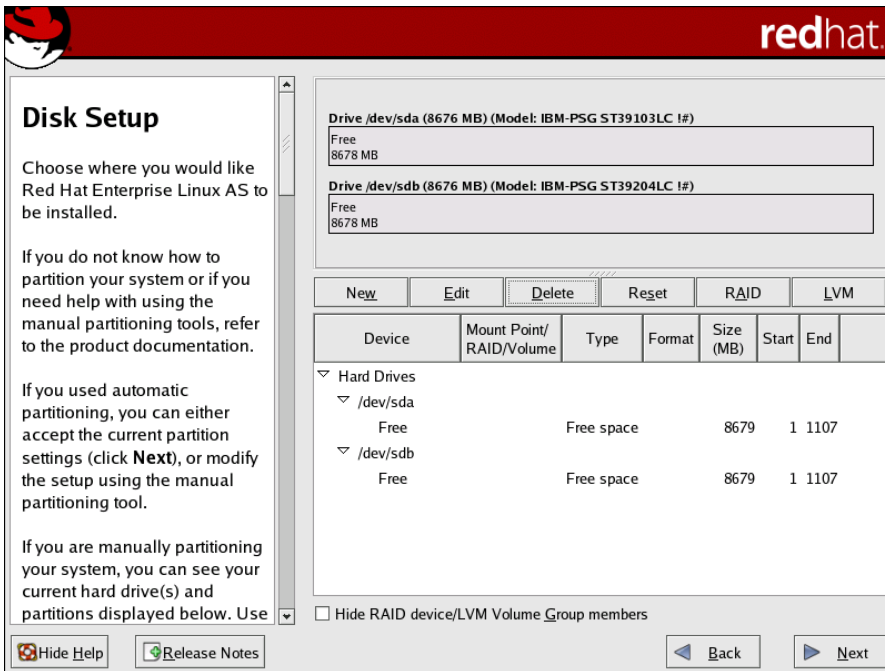
Figure 10.1. Two Blank Drives, Ready For Configuration

1. In **Disk Druid**, choose **RAID** to enter the software RAID creation screen.

2. Choose **Create a software RAID partition** to create a RAID partition as shown in Figure 10.2, "RAID Partition Options". Note that no other RAID options (such as entering a mount point) are available until RAID partitions, as well as RAID devices, are created.
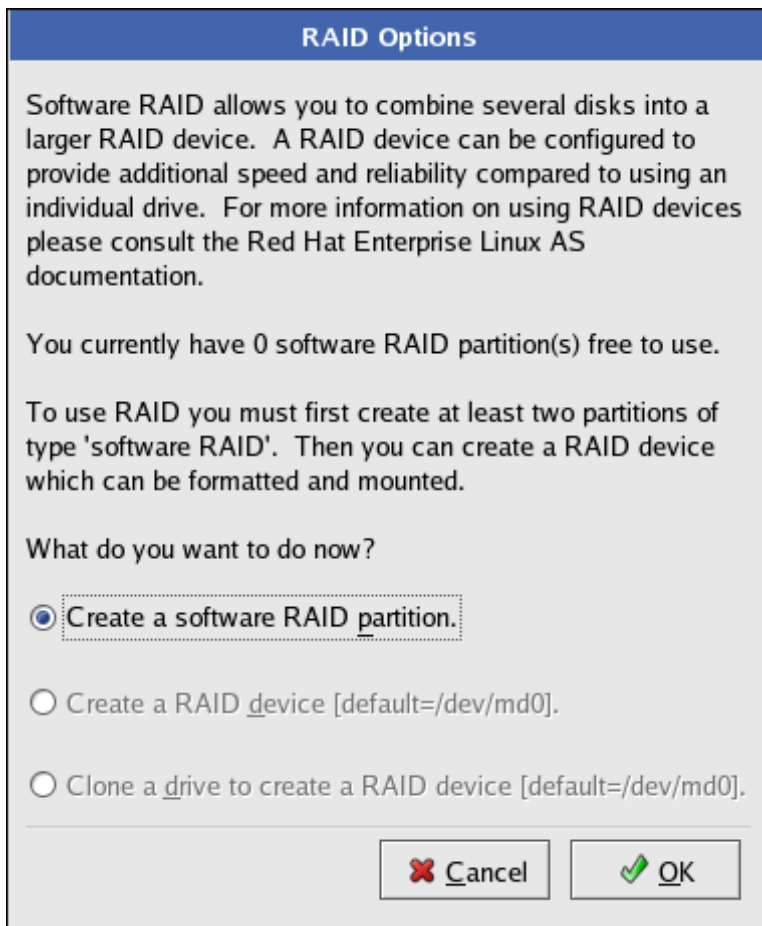
Figure 10.2. RAID Partition Options

3. A software RAID partition must be constrained to one drive. For **Allowable Drives**, select the drive on which RAID is to be created. If you have multiple drives, all drives are selected, and you must deselect all but one drive.

Figure 10.3. Adding a RAID Partition

4. Enter the size that you want the partition to be.

5. 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 share the available free space on the disk.

6. Select **Force to be a primary partition** if you want the partition to be a primary partition. A primary partition is one of the first four partitions on the hard drive. If unselected, the partition is created as a logical partition. If other operating systems are already on the system, unselecting this option should be considered. For more information on primary versus logical/extended partitions, refer to the appendix section of the *Installation Guide*.

7. 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 **/boot/** partition as a software RAID device, leaving the root partition (**/**), **/home/**, and swap as regular file systems. Figure 10.4, "RAID 1 Partitions Ready, Pre-Device and Mount Point Creation" shows successfully allocated space for the RAID 1 configuration (for **/boot/**), which is now ready for RAID device and mount point creation:
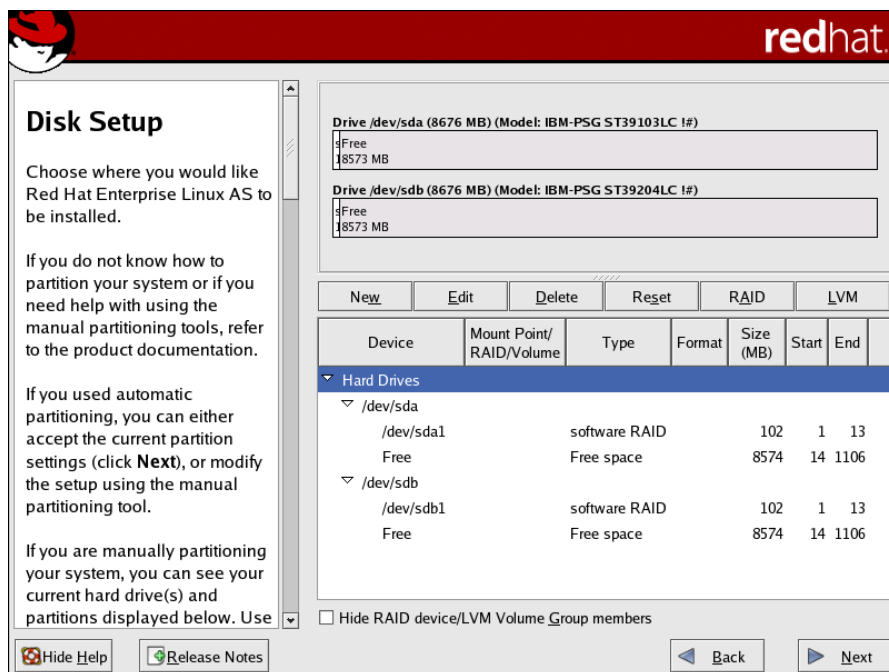
Figure 10.4. RAID 1 Partitions Ready, Pre-Device and Mount Point Creation

10.2. Creating the RAID Devices and Mount Points

Once you have all of your partitions created as `software RAID` partitions, the following steps create the RAID device and mount point:

1. Select the **RAID** button on the **Disk Druid** main partitioning screen (refer to Figure 10.5, "RAID Options").

2. Figure 10.5, "RAID Options" appears. Select `Create a RAID device`.
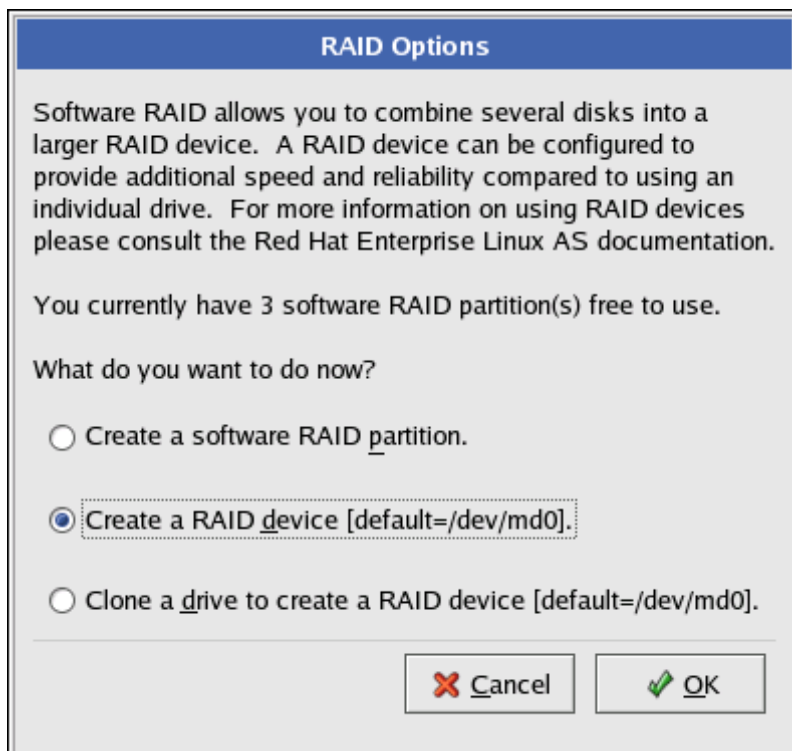
Figure 10.5. RAID Options

3. Next, Figure 10.6, "Making a RAID Device and Assigning a Mount Point" appears, where you can make a RAID device and assign a mount point.

Make RAID Device

Mount Point:	/boot
File System Type:	ext3
RAID Device:	md0
RAID Level:	RAID0
RAID Members:	☑ sda1 102 MB ☑ sdb1 102 MB
Number of spares:	0

✗ Cancel   ✓ OK

Figure 10.6. Making a RAID Device and Assigning a Mount Point

4. Enter a mount point.

5. Choose the file system type for the partition. At this point you can either configure a dynamic LVM file system or a traditional static ext2/ext3 file system. For more information on configuring LVM on a RAID device, select **physical volume (LVM)** and then refer to Chapter 8, *LVM Configuration*. If LVM is not required, continue on with the following instructions.

6. Select a device name such as **md0** for the RAID device.

7. Choose your RAID level. You can choose from **RAID 0**, **RAID 1**, and **RAID 5**. If you need assistance in determining which RAID level to implement, refer to Chapter 9, *Redundant Array of Independent Disks (RAID)*.

> **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 seperate RAID partition of **/boot/**, and you are making a RAID partition for the root file system (**/**), it must be RAID level 1 and must use one of the first two drives (IDE first, SCSI second).
>
>
>
> **Figure 10.7. The /boot/ Mount Error**

8. The RAID partitions created appear in the **RAID Members** list. Select which of these partitions should be used to create the RAID device.

9. If configuring RAID 1 or RAID 5, specify the number of spare partitions. If a software RAID partition fails, the spare is automatically 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). Select the partitions for the RAID device and the partition(s) for the spare(s).

10. After clicking **OK**, the RAID device appears in the **Drive Summary** list.

11. Repeat this chapter's entire process for configuring additional partitions, devices, and mount points, such as the root partition (**/**), **/home/**, or swap.

After completing the entire configuration, the figure as shown in Figure 10.8, "Final Sample RAID Configuration" resembles the default configuration, except for the use of RAID.
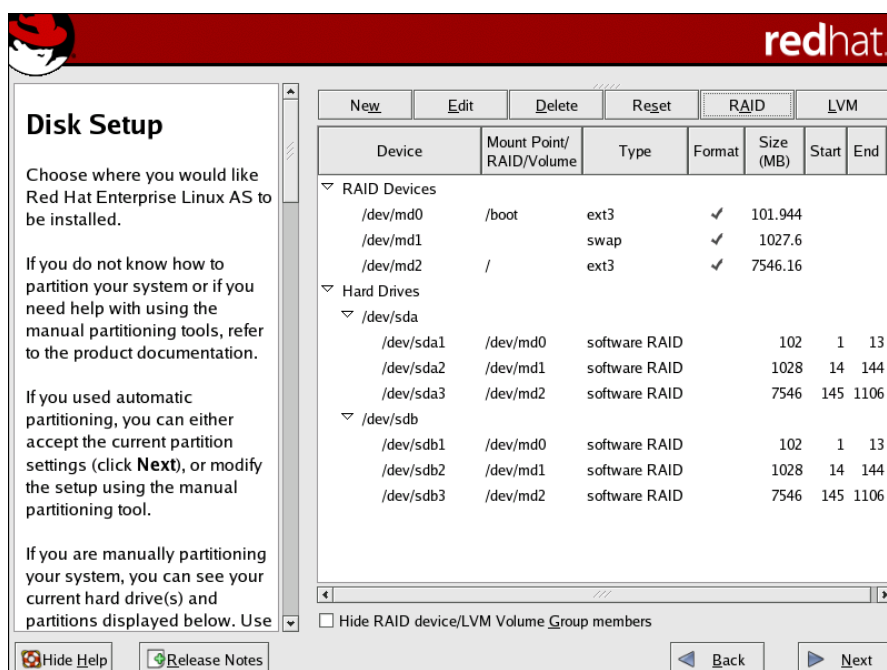
Figure 10.8. Final Sample RAID Configuration

The figure as shown in Figure 10.9, "Final Sample RAID With LVM Configuration" is an example of a RAID and LVM configuration.
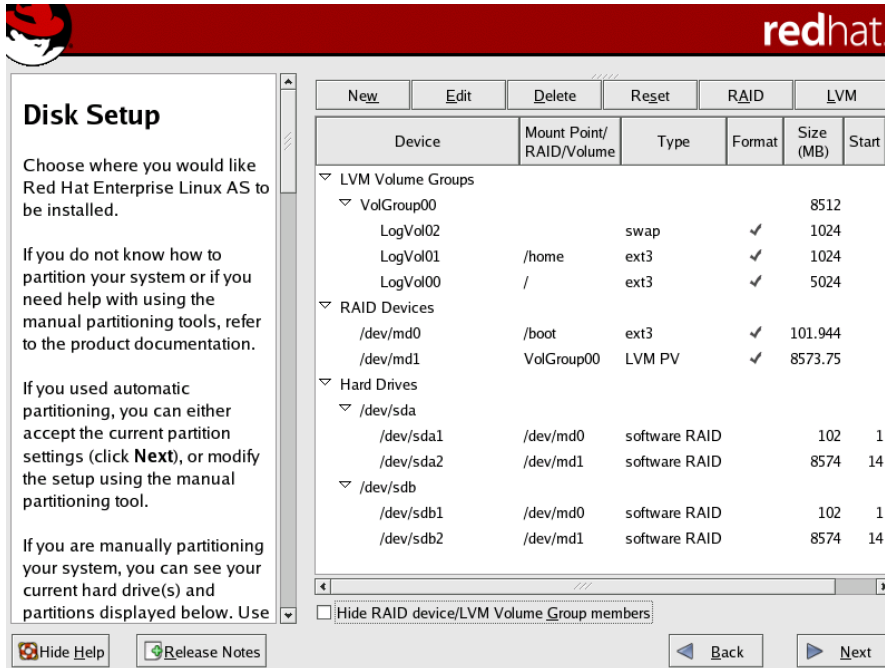
Figure 10.9. Final Sample RAID With LVM Configuration

You can continue with your installation process. Refer to the *Installation Guide* for further instructions.

Chapter 11. Swap Space

11.1. 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 RAM 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 should be equal to twice your computer's physical RAM for up to 2 GB of physical RAM. For physical RAM above 2 GB, the size of your swap should be equal to the amount of physical RAM above 2 GB. The size of your swap should never be less than 32 MB.

Using this basic formula, a system with 2 GB of physical RAM would have 4 GB of swap, while one with 3 GB of physical RAM would have 5 GB of swap.

> **Note**
>
> Unfortunately, deciding on the amount of swap to allocate to Red Hat Enterprise Linux is more of an art than a science, so hard rules are not possible. Each system's most used applications should be accounted for when determining swap size.

> **Important**
>
> File systems and LVM2 volumes assigned as swap space *cannot* be in use when being modified. For example, no system processes can be assigned the swap space, as well as no amount of swap should be allocated and used by the kernel. Use the `free` and `cat /proc/swaps` commands to verify how much and where swap is in use.
> The best way to achieve swap space modifications is to boot your system in rescue mode, and then follow the instructions (for each scenario) in the remainder of this chapter. Refer to Chapter 5, *Basic System Recovery* for instructions on booting into rescue mode. When prompted to mount the file system, select `Skip`.

11.2. 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 128 MB to 256 MB, but there is only 256 MB of swap space. It might be advantageous to increase the amount of swap space to 512 MB if you perform memory-intense operations or run applications that require a large amount of memory.

You have three options: create a new swap partition, create a new swap file, or extend swap on an existing LVM2 logical volume. It is recommended that you extend an existing logical volume.

11.2.1. Extending Swap on an LVM2 Logical Volume

To extend an LVM2 swap logical volume (assuming **/dev/VolGroup00/LogVol01** is the volume you

want to extend):

1. Disable swapping for the associated logical volume:

   ```
   # swapoff -v /dev/VolGroup00/LogVol01
   ```

2. Resize the LVM2 logical volume by 256 MB:

   ```
   # lvm lvresize /dev/VolGroup00/LogVol01 -L +256M
   ```

3. Format the new swap space:

   ```
   # mkswap /dev/VolGroup00/LogVol01
   ```

4. Enable the extended logical volume:

   ```
   # swapon -va
   ```

5. Test that the logical volume has been extended properly:

   ```
   # cat /proc/swaps # free
   ```

11.2.2. Creating an LVM2 Logical Volume for Swap

To add a swap volume group (assuming **/dev/VolGroup00/LogVol02** is the swap volume you want to add):

1. Create the LVM2 logical volume of size 256 MB:

   ```
   # lvm lvcreate VolGroup00 -n LogVol02 -L 256M
   ```

2. Format the new swap space:

   ```
   # mkswap /dev/VolGroup00/LogVol02
   ```

3. Add the following entry to the **/etc/fstab** file:

   ```
   /dev/VolGroup00/LogVol02 swap swap defaults 0 0
   ```

4. Enable the extended logical volume:

   ```
   # swapon -va
   ```

5. Test that the logical volume has been extended properly:

   ```
   # cat /proc/swaps # free
   ```

11.2.3. Creating a Swap File

To add a swap file:

1. Determine the size of the new swap file in megabytes and multiply by 1024 to determine the number of blocks. For example, the block size of a 64 MB swap file is 65536.

2. 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
```

3. Setup the swap file with the command:

```
mkswap /swapfile
```

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

```
swapon /swapfile
```

5. To enable it at boot time, edit **/etc/fstab** to include the following entry:

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

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

6. After adding the new swap file and enabling it, verify it is enabled by viewing the output of the command **cat /proc/swaps** or **free**.

11.3. Removing Swap Space

Sometimes it can be prudent to reduce swap space after installation. For example, say you downgraded the amount of RAM in your system from 1 GB to 512 MB, but there is 2 GB of swap space still assigned. It might be advantageous to reduce the amount of swap space to 1 GB, since the larger 2 GB could be wasting disk space.

You have three options: remove an entire LVM2 logical volume used for swap, remove a swap file, or reduce swap space on an existing LVM2 logical volume.

11.3.1. Reducing Swap on an LVM2 Logical Volume

To reduce an LVM2 swap logical volume (assuming **/dev/VolGroup00/LogVol01** is the volume you want to extend):

1. Disable swapping for the associated logical volume:

```
# swapoff -v /dev/VolGroup00/LogVol01
```

2. Reduce the LVM2 logical volume by 512 MB:

```
# lvm lvreduce /dev/VolGroup00/LogVol01 -L -512M
```

3. Format the new swap space:

```
# mkswap /dev/VolGroup00/LogVol01
```

4. Enable the extended logical volume:

```
# swapon -va
```

5. Test that the logical volume has been reduced properly:

```
# cat /proc/swaps # free
```

11.3.2. Removing an LVM2 Logical Volume for Swap

The swap logical volume cannot be in use (no system locks or processes on the volume). The easiest way to achieve this it to boot your system in rescue mode. Refer to Chapter 5, *Basic System Recovery* for instructions on booting into rescue mode. When prompted to mount the file system, select **Skip**.

To remove a swap volume group (assuming **/dev/VolGroup00/LogVol02** is the swap volume you want to remove):

1. Disable swapping for the associated logical volume:

   ```
   # swapoff -v /dev/VolGroup00/LogVol02
   ```

2. Remove the LVM2 logical volume of size 512 MB:

   ```
   # lvm lvremove /dev/VolGroup00/LogVol02
   ```

3. Remove the following entry from the **/etc/fstab** file:

   ```
   /dev/VolGroup00/LogVol02 swap swap defaults 0 0
   ```

4. Test that the logical volume has been extended properly:

   ```
   # cat /proc/swaps # free
   ```

11.3.3. Removing a Swap File

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 -v /swapfile
   ```

2. Remove its entry from the **/etc/fstab** file.
3. Remove the actual file:

   ```
   # rm /swapfile
   ```

11.4. 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.

Chapter 12. Managing Disk Storage

Introduction to different methods……..

12.1. Standard Partitions using `parted`

Many users need 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 users to perform these tasks. This chapter discusses how to use **`parted`** to perform file system tasks.

If you want to view the system's disk space usage or monitor the disk space usage, refer to Section 39.3, "File Systems".

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/sda`**, where **`/dev/sda`** is the device name for the drive you want to configure. The **`(parted)`** prompt is displayed. Type **`help`** to view a list of available commands.

If you want to create, remove, or resize a partition, the device cannot be in use (partitions cannot be mounted, and swap space cannot be enabled). The partition table should not be modified while in use because the kernel may not properly recognize the changes. Data could be overwritten by writing to the wrong partition because the partition table and partitions mounted do not match. The easiest way to achieve this it to boot your system in rescue mode. Refer to Chapter 5, *Basic System Recovery* 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 (system processes that use or lock the file system from being unmounted), you can unmount them with the **`umount`** command and turn off all the swap space on the hard drive with the **`swapoff`** command.

Table 12.1, "**`parted`** commands" contains a list of commonly used **`parted`** commands. The sections that follow explain some of them in more detail.

Table 12.1. `parted` commands

Command	Description
check *minor-num*	Perform a simple check of the file system
cp *fromto*	Copy file system from one partition to another; *from* and *to* are the minor numbers of the partitions
help	Display list of available commands
mklabel *label*	Create a disk label for the partition table
mkfs *minor-numfile-system-type*	Create a file system of type *file-system-type*
mkpart *part-typefs-typestart-mbend-mb*	Make a partition without creating a new file system
mkpartfs *part-typefs-typestart-mbend-mb*	Make a partition and create the specified file system
move *minor-numstart-mbend-mb*	Move the partition
name *minor-numname*	Name the partition for Mac and PC98 disklabels only
print	Display the partition table
quit	Quit **parted**
rescue *start-mbend-mb*	Rescue a lost partition from *start-mb* to *end-mb*
resize *minor-numstart-mbend-mb*	Resize the partition from *start-mb* to *end-mb*
rm *minor-num*	Remove the partition
select *device*	Select a different device to configure
set *minor-numflagstate*	Set the flag on a partition; *state* is either on or off

12.1.1. Viewing the Partition Table

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

```
print
```

A table similar to the following appears:

```
Disk geometry for /dev/sda: 0.000-8678.789 megabytes
Disk label type: msdos
Minor    Start        End      Type       Filesystem   Flags
1          0.031    101.975   primary    ext3          boot
2        101.975   5098.754   primary    ext3
3       5098.755   6361.677   primary     linux-swap
4       6361.677   8675.727   extended
5       6361.708   7357.895   logical    ext3


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                 lba
5        760.922   9758.232   logical    ext3
```

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/sda1**. 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, fat16, fat32, 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.

In this example, minor number 1 refers to the **/boot/** file system, minor number 2 refers to the root file system (**/**), minor number 3 refers to the swap, and minor number 5 refers to the **/home/** file system.

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

12.1.2. 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/*sda* is the device on which to create the partition:

```
parted /dev/sda
```

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 12.1.4, "Resizing a Partition" for details.

12.1.2.1. 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 *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 is 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 press **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.

12.1.2.2. Formating the Partition

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

```
/sbin/mkfs -t ext3 /dev/sda6
```

> **Warning**
>
> Formatting the partition permanently destroys any data that currently exists on the partition.

12.1.2.3. Labeling the Partition

Next, give the partition a label. For example, if the new partition is **/dev/sda6** and you want to label it **/work**:

```
e2label /dev/sda6 /work
```

By default, the 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.

12.1.2.4. Creating the Mount Point

As root, create the mount point:

```
mkdir /work
```

12.1.2.5. 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 system 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 is mounted at boot time. To mount the partition without rebooting, as root, type the command:

```
mount /work
```

12.1.3. 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/*sda* is the device on which to remove the partition:

```
parted /dev/sda
```

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 press **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.

12.1.4. 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/*sda* is the device on which to resize the partition:

```
parted /dev/sda
```

View the current partition table to determine the minor number of the partition to resize 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.

12.2. LVM Partition Management

The following commands can be found by issuing **lvm help** at a command prompt.

Table 12.2. LVM commands

Command	Description
dumpconfig	Dump the active configuration
formats	List the available metadata formats
help	Display the help commands
lvchange	Change the attributes of logical volume(s)
lvcreate	Create a logical volume
lvdisplay	Display information about a logical volume
lvextend	Add space to a logical volume
lvmchange	*Due to use of the device mapper, this command has been deprecated*
lvmdiskscan	List devices that may be used as physical volumes
lvmsadc	Collect activity data
lvmsar	Create activity report
lvreduce	Reduce the size of a logical volume
lvremove	Remove logical volume(s) from the system
lvrename	Rename a logical volume
lvresize	Resize a logical volume
lvs	Display information about logical volumes
lvscan	List all logical volumes in all volume groups
pvchange	Change attributes of physical volume(s)
pvcreate	Initialize physical volume(s) for use by LVM
pvdata	Display the on-disk metadata for physical volume(s)
pvdisplay	Display various attributes of physical volume(s)
pvmove	Move extents from one physical volume to another
pvremove	Remove LVM label(s) from physical volume(s)
pvresize	Resize a physical volume in use by a volume group
pvs	Display information about physical volumes
pvscan	List all physical volumes
segtypes	List available segment types
vgcfgbackup	Backup volume group configuration
vgcfgrestore	Restore volume group configuration
vgchange	Change volume group attributes
vgck	Check the consistency of a volume group
vgconvert	Change volume group metadata format
vgcreate	Create a volume group
vgdisplay	Display volume group information
vgexport	Unregister a volume group from the system

vgextend	Add physical volumes to a volume group
vgimport	Register exported volume group with system
vgmerge	Merge volume groups
vgmknodes	Create the special files for volume group devices in /dev/
vgreduce	Remove a physical volume from a volume group
vgremove	Remove a volume group
vgrename	Rename a volume group
vgs	Display information about volume groups
vgscan	Search for all volume groups
vgsplit	Move physical volumes into a new volume group
version	Display software and driver version information

Chapter 13. Implementing Disk Quotas

Disk space can be restricted by implementing disk quotas which alert a system administrator before a user consumes too much disk space or a partition becomes full.

Disk quotas can be configured for individual users as well as user groups. This kind of flexibility makes it possible to give each user a small quota to handle "personal" files (such as email and reports), while allowing the projects they work on to have more sizable quotas (assuming the projects are given their own groups).

In addition, quotas can be set not just to control the number of disk blocks consumed but to control the number of inodes (data structures that contain information about files in UNIX file systems). Because inodes are used to contain file-related information, this allows control over the number of files that can be created.

The **quota** RPM must be installed to implement disk quotas.

13.1. Configuring Disk Quotas

To implement disk quotas, use the following steps:

1. Enable quotas per file system by modifying the **/etc/fstab** file.
2. Remount the file system(s).
3. Create the quota database files and generate the disk usage table.
4. Assign quota policies.

Each of these steps is discussed in detail in the following sections.

13.1.1. Enabling Quotas

As root, using a text editor, edit the **/etc/fstab** file. Add the **usrquota** and/or **grpquota** options to the file systems that require quotas:

```
/dev/VolGroup00/LogVol00 /          ext3    defaults         1 1
LABEL=/boot              /boot      ext3    defaults         1 2
none                     /dev/pts   devpts  gid=5,mode=620   0 0
none                     /dev/shm   tmpfs   defaults         0 0
none                     /proc      proc    defaults         0 0
none                     /sys       sysfs   defaults         0 0
/dev/VolGroup00/LogVol02 /home      ext3    defaults,usrquota,grpquota  1 2
/dev/VolGroup00/LogVol01 swap       swap    defaults         0 0
.
.
.
```

In this example, the **/home** file system has both user and group quotas enabled.

> **Note**
>
> The following examples assume that a separate **/home** partition was created during the installation of Red Hat Enterprise Linux. Although not ideal, the root (**/**) partition (the installation default created partition) can be used for setting quota policies in the **/etc/fstab** file.

13.1.2. Remounting the File Systems

After adding the **usrquota** and/or **grpquota** options, remount each file system whose **fstab** entry has been modified. If the file system is not in use by any process, use one of the following methods:

- Issue the **umount** command followed by the **mount** command to remount the file system.
- Issue the **mount -o remount /home** command to remount the file system.

If the file system is currently in use, the easiest method for remounting the file system is to reboot the system.

13.1.3. Creating the Quota Database Files

After each quota-enabled file system is remounted, the system is capable of working with disk quotas. However, the file system itself is not yet ready to support quotas. The next step is to run the **quotacheck** command.

The **quotacheck** command examines quota-enabled file systems and builds a table of the current disk usage per file system. The table is then used to update the operating system's copy of disk usage. In addition, the file system's disk quota files are updated.

To create the quota files (**aquota.user** and **aquota.group**) on the file system, use the **-c** option of the **quotacheck** command. For example, if user and group quotas are enabled for the **/home** file system, create the files in the **/home** directory:

```
quotacheck -cug /home
```

The **-c** option specifies that the quota files should be created for each file system with quotas enabled, the **-u** option specifies to check for user quotas, and the **-g** option specifies to check for group quotas.

If neither the **-u** or **-g** options are specified, only the user quota file is created. If only **-g** is specified, only the group quota file is created.

After the files are created, run the following command to generate the table of current disk usage per file system with quotas enabled:

```
quotacheck -avug
```

The options used are as follows:

- **a** — Check all quota-enabled, locally-mounted file systems
- **v** — Display verbose status information as the quota check proceeds
- **u** — Check user disk quota information
- **g** — Check group disk quota information

After **quotacheck** has finished running, the quota files corresponding to the enabled quotas (user and/or group) are populated with data for each quota-enabled locally-mounted file system such as **/home**.

13.1.4. Assigning Quotas per User

The last step is assigning the disk quotas with the **edquota** command.

To configure the quota for a user, as root in a shell prompt, execute the command:

```
   edquota username
```

Perform this step for each user who needs a quota. For example, if a quota is enabled in **/etc/fstab** for the **/home** partition (**/dev/VolGroup00/LogVol02**) and the command **edquota testuser** is executed, the following is shown in the editor configured as the default for the system:

```
Disk quotas for user testuser (uid 501):
   Filesystem                blocks     soft      hard    inodes    soft    hard
   /dev/VolGroup00/LogVol02  440436        0        0     37418       0       0
```

> **Note**
>
> The text editor defined by the **EDITOR** environment variable is used by **edquota**. To change the editor, set the **EDITOR** environment variable in your **~/.bash_profile** file to the full path of the editor of your choice.

The first column is the name of the file system that has a quota enabled for it. The second column shows how many blocks the user is currently using. The next two columns are used to set soft and hard block limits for the user on the file system. The **inodes** column shows how many inodes the user is currently using. The last two columns are used to set the soft and hard inode limits for the user on the file system.

A hard limit is the absolute maximum amount of disk space that a user or group can use. Once this limit is reached, no further disk space can be used.

The soft limit defines the maximum amount of disk space that can be used. However, unlike the hard limit, the soft limit can be exceeded for a certain amount of time. That time is known as the *grace period*. The grace period can be expressed in seconds, minutes, hours, days, weeks, or months.

If any of the values are set to 0, that limit is not set. In the text editor, change the desired limits. For example:

```
Disk quotas for user testuser (uid 501):
   Filesystem                blocks     soft      hard    inodes    soft    hard
   /dev/VolGroup00/LogVol02  440436   500000   550000     37418       0       0
```

To verify that the quota for the user has been set, use the command:

```
   quota testuser
```

13.1.5. Assigning Quotas per Group

Quotas can also be assigned on a per-group basis. For example, to set a group quota for the **devel** group (the group must exist prior to setting the group quota), use the command:

```
   edquota -g devel
```

This command displays the existing quota for the group in the text editor:

Modify the limits, save the file, and then configure the quota.

To verify that the group quota has been set, use the command:

```
quota -g devel
```

13.1.6. Assigning Quotas per File System

To assign quotas based on each file system enabled for quotas, use the command:

```
edquota -t
```

Like the other **edquota** commands, this one opens the current quotas for the file system in the text editor:

```
Grace period before enforcing soft limits for users:
Time units may be: days, hours, minutes, or seconds
  Filesystem                        Block grace period  Inode grace period
  /dev/mapper/VolGroup00-LogVol02      7days                 7days
```

Change the block grace period or inode grace period, save the changes to the file, and exit the text editor.

13.2. Managing Disk Quotas

If quotas are implemented, they need some maintenance — mostly in the form of watching to see if the quotas are exceeded and making sure the quotas are accurate. Of course, if users repeatedly exceeds their quotas or consistently reaches their soft limits, a system administrator has a few choices to make depending on what type of users they are and how much disk space impacts their work. The administrator can either help the user determine how to use less disk space or increase the user's disk quota if needed.

13.2.1. Enabling and Disabling

It is possible to disable quotas without setting them to be 0. To turn all user and group quotas off, use the following command:

```
quotaoff -vaug
```

If neither the **-u** or **-g** options are specified, only the user quotas are disabled. If only **-g** is specified, only group quotas are disabled.

To enable quotas again, use the **quotaon** command with the same options.

For example, to enable user and group quotas for all file systems, use the following command:

```
quotaon -vaug
```

To enable quotas for a specific file system, such as **/home**, use the following command:

```
quotaon -vug /home
```

If neither the **-u** or **-g** options are specified, only the user quotas are enabled. If only **-g** is specified,

only group quotas are enabled.

13.2.2. Reporting on Disk Quotas

Creating a disk usage report entails running the **repquota** utility. For example, the command **repquota /home** produces this output:

```
*** Report for user quotas on device /dev/mapper/VolGroup00-LogVol02
Block grace time: 7days; Inode grace time: 7days
                        Block limits              File limits
User            used    soft    hard  grace    used  soft  hard  grace
----------------------------------------------------------------------
root      --      36       0       0              4     0     0
kristin   --     540       0       0            125     0     0
testuser  --  440400  500000  550000          37418     0     0
```

To view the disk usage report for all (option **-a**) quota-enabled file systems, use the command:

```
repquota -a
```

While the report is easy to read, a few points should be explained. The **--** displayed after each user is a quick way to determine whether the block or inode limits have been exceeded. If either soft limit is exceeded, a **+** appears in place of the corresponding **-**; the first **-** represents the block limit, and the second represents the inode limit.

The **grace** columns are normally blank. If a soft limit has been exceeded, the column contains a time specification equal to the amount of time remaining on the grace period. If the grace period has expired, **none** appears in its place.

13.2.3. Keeping Quotas Accurate

Whenever a file system is not unmounted cleanly (due to a system crash, for example), it is necessary to run **quotacheck**. However, **quotacheck** can be run on a regular basis, even if the system has not crashed. Running the following command periodically keeps the quotas more accurate (the options used have been described in Section 13.1.1, "Enabling Quotas"):

```
quotacheck -avug
```

The easiest way to run it periodically is to use **cron**. As root, either use the **crontab -e** command to schedule a periodic **quotacheck** or place a script that runs **quotacheck** in any one of the following directories (using whichever interval best matches your needs):

- **/etc/cron.hourly**
- **/etc/cron.daily**
- **/etc/cron.weekly**
- **/etc/cron.monthly**

The most accurate quota statistics can be obtained when the file system(s) analyzed are not in active use. Thus, the cron task should be schedule during a time where the file system(s) are used the least. If this time is various for different file systems with quotas, run **quotacheck** for each file system at different times with multiple cron tasks.

Refer to Chapter 34, *Automated Tasks* for more information about configuring **cron**.

13.3. Additional Resources

For more information on disk quotas, refer to the following resources.

13.3.1. Installed Documentation

- The **quotacheck**, **edquota**, **repquota**, **quota**, **quotaon**, and **quotaoff** man pages

13.3.2. Related Books

- *Introduction to System Administration* ; Red Hat, Inc — Available at http://www.redhat.com/docs/ and on the Documentation CD, this manual contains background information on storage management (including disk quotas) for new Red Hat Enterprise Linux system administrators.

Chapter 14. Access Control Lists

Files and directories have permission sets for the owner of the file, the group associated with the file, and all other users for the system. However, these permission sets have limitations. For example, different permissions cannot be configured for different users. Thus, *Access Control Lists* (ACLs) were implemented.

The Red Hat Enterprise Linux 4 kernel provides ACL support for the ext3 file system and NFS-exported file systems. ACLs are also recognized on ext3 file systems accessed via Samba.

Along with support in the kernel, the **acl** package is required to implement ACLs. It contains the utilities used to add, modify, remove, and retrieve ACL information.

The **cp** and **mv** commands copy or move any ACLs associated with files and directories.

14.1. Mounting File Systems

Before using ACLs for a file or directory, the partition for the file or directory must be mounted with ACL support. If it is a local ext3 file system, it can mounted with the following command:

```
mount -t ext3 -o acl <device-name><partition>
```

For example:

```
mount -t ext3 -o acl /dev/VolGroup00/LogVol02 /work
```

Alternatively, if the partition is listed in the **/etc/fstab** file, the entry for the partition can include the **acl** option:

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

If an ext3 file system is accessed via Samba and ACLs have been enabled for it, the ACLs are recognized because Samba has been compiled with the **--with-acl-support** option. No special flags are required when accessing or mounting a Samba share.

14.1.1. NFS

By default, if the file system being exported by an NFS server supports ACLs and the NFS client can read ACLs, ACLs are utilized by the client system.

To disable ACLs on NFS shares when configuring the server, include the **no_acl** option in the **/etc/exports** file. To disable ACLs on an NFS share when mounting it on a client, mount it with the **no_acl** option via the command line or the **/etc/fstab** file.

14.2. Setting Access ACLs

There are two types of ACLs: *access ACLs* and *default ACLs*. An access ACL is the access control list for a specific file or directory. A default ACL can only be associated with a directory; if a file within the directory does not have an access ACL, it uses the rules of the default ACL for the directory. Default ACLs are optional.

ACLs can be configured:

1. Per user

2. Per group

3. Via the effective rights mask

4. For users not in the user group for the file

The **setfacl** utility sets ACLs for files and directories. Use the **-m** option to add or modify the ACL of a file or directory:

```
setfacl -m <rules><files>
```

Rules (**<rules>**) must be specified in the following formats. Multiple rules can be specified in the same command if they are separated by commas.

u:<uid>:<perms>

Sets the access ACL for a user. The user name or UID may be specified. The user may be any valid user on the system.

g:<gid>:<perms>

Sets the access ACL for a group. The group name or GID may be specified. The group may be any valid group on the system.

m:<perms>

Sets the effective rights mask. The mask is the union of all permissions of the owning group and all of the user and group entries.

o:<perms>

Sets the access ACL for users other than the ones in the group for the file.

White space is ignored. Permissions (**<perms>**) must be a combination of the characters **r**, **w**, and **x** for read, write, and execute.

If a file or directory already has an ACL, and the **setfacl** command is used, the additional rules are added to the existing ACL or the existing rule is modified.

For example, to give read and write permissions to user andrius:

```
setfacl -m u:andrius:rw /project/somefile
```

To remove all the permissions for a user, group, or others, use the **-x** option and do not specify any permissions:

```
setfacl -x <rules><files>
```

For example, to remove all permissions from the user with UID 500:

```
setfacl -x u:500 /project/somefile
```

14.3. Setting Default ACLs

To set a default ACL, add **d:** before the rule and specify a directory instead of a file name.

For example, to set the default ACL for the **/share/** directory to read and execute for users not in the user group (an access ACL for an individual file can override it):

```
setfacl -m d:o:rx /share
```

14.4. Retrieving ACLs

To determine the existing ACLs for a file or directory, use the **getfacl** command:

```
getfacl <filename>
```

It returns output similar to the following:

```
# file: file
# owner: andrius
# group: andrius
user::rw-
user:smoore:r--
group::r--
mask::r--
other::r--
```

If a directory is specified, and it has a default ACL, the default ACL is also displayed such as:

```
# file: file
# owner: andrius
# group: andrius
user::rw-
user:smoore:r--
group::r--
mask::r--
other::r--
default:user::rwx
default:user:andrius:rwx
default:group::r-x
default:mask::rwx
default:other::r-x
```

14.5. Archiving File Systems With ACLs

> ⚠️ **Warning**
>
> The **tar** and **dump** commands do *not* backup ACLs.

The **star** utility is similar to the **tar** utility in that it can be used to generate archives of files; however, some of its options are different. Refer to Table 14.1, "Command Line Options for **star**" for a listing of more commonly used options. For all available options, refer to the **star** man page. The **star** package is required to use this utility.

Table 14.1. Command Line Options for `star`

Option	Description
-c	Creates an archive file.
-n	Do not extract the files; use in conjunction with **-x** to show what extracting the files does.
-r	Replaces files in the archive. The files are written to the end of the archive file, replacing any files with the same path and file name.
-t	Displays the contents of the archive file.
-u	Updates the archive file. The files are written to the end of the archive if they do not exist in the archive or if the files are newer than the files of the same name in the archive. This option only work if the archive is a file or an unblocked tape that may backspace.
-x	Extracts the files from the archive. If used with **-U** and a file in the archive is older than the corresponding file on the file system, the file is not extracted.
-help	Displays the most important options.
-xhelp	Displays the least important options.
-/	Do not strip leading slashes from file names when extracting the files from an archive. By default, they are striped when files are extracted.
-acl	When creating or extracting, archive or restore any ACLs associated with the files and directories.

14.6. Compatibility with Older Systems

If an ACL has been set on any file on a given file system, that file system has the **ext_attr** attribute. This attribute can be seen using the following command:

```
tune2fs -l <filesystem-device>
```

A file system that has acquired the **ext_attr** attribute can be mounted with older kernels, but those kernels do not enforce any ACLs which have been set.

Versions of the **e2fsck** utility included in version 1.22 and higher of the **e2fsprogs** package (including the versions in Red Hat Enterprise Linux 2.1 and 4) can check a file system with the **ext_attr** attribute. Older versions refuse to check it.

14.7. Additional Resources

Refer to the follow resources for more information.

14.7.1. Installed Documentation

- **acl** man page — Description of ACLs
- **getfacl** man page — Discusses how to get file access control lists

- **setfacl** man page — Explains how to set file access control lists
- **star** man page — Explains more about the **star** utility and its many options

14.7.2. Useful Websites

- http://acl.bestbits.at/ — Website for ACLs

Part III. Package Management

All software on a Red Hat Enterprise Linux system is divided into RPM packages which can be installed, upgraded, or removed. This part describes how to manage the RPM packages on a Red Hat Enterprise Linux system using graphical and command line tools.

Chapter 15. Package Management with RPM

The RPM Package Manager (RPM) is an open packaging system, available for anyone to use, which runs on Red Hat Enterprise 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 the **Package Management Tool** to perform many RPM commands.

During upgrades, RPM handles configuration files carefully, so that you never lose your customizations — something that you cannot 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 to install, remove, or upgrade an RPM package.

15.1. RPM Design Goals

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 Enterprise Linux), you do not 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 do not 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, verify the package. You are 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.

15.2. 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 refer to Section 15.5, "Additional Resources" for more information on RPM.

15.2.1. Finding RPM Packages

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

- The Red Hat Enterprise 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 — Refer to Chapter 16, *Red Hat Network* for more details on Red Hat Network

15.2.2. 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**). To install a package, log in as root and type the following command at a shell prompt:

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

If installation is successful, the following output is displayed:

```
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.

The signature of a package is checked automatically when installing or upgrading a package. The signature confirms that the package was signed by an authorized party. For example, if the verification of the signature fails, an error message such as the following is displayed:

```
error: V3 DSA signature: BAD, key ID 0352860f
```

If it is a new, header-only, signature, an error message such as the following is displayed:

```
error: Header V3 DSA signature: BAD, key ID 0352860f
```

If you do not have the appropriate key installed to verify the signature, the message contains the word **NOKEY** such as:

```
warning: V3 DSA signature: NOKEY, key ID 0352860f
```

Refer to Section 15.3, "Checking a Package's Signature" for more information on checking a package's signature.

> ⚠ **Warning**
>
> If you are installing a kernel package, you should use **rpm -ivh** instead. Refer to Chapter 36, *Manually Upgrading the Kernel* for details.

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

15.2.2.1. Package Already Installed

If the package of the same version is already installed, the following is displayed:

```
Preparing...                ########################################### [100%]
package foo-1.0-1 is already installed
```

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

```
rpm -ivh --replacepkgs 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.

15.2.2.2. 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, the following is displayed:

```
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
```

15.2.2.3. Unresolved Dependency

RPM packages can, essentially, depend on other packages, which means that they require other packages to be installed to run properly. If you try to install a package which has an unresolved dependency, output similar to the following is displayed:

```
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 a package from the Red Hat Enterprise Linux CD-ROM set, it usually suggest the package(s) needed to resolve the dependency. Find the suggested package(s) on the Red Hat Enterprise 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, output similar to the following is displayed:

```
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 option.

```
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 is displayed:

```
bar-2.0.20-3.i386.rpm
```

To force the installation anyway (which is not recommended since the package may not run correctly), use the **--nodeps** option.

15.2.3. 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, 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:

```
error: Failed dependencies:
        foo is needed by (installed) bar-2.0.20-3.i386.rpm
```

To cause RPM to ignore this error and uninstall the package anyway, which may break the package depending on it, use the **--nodeps** option.

15.2.4. 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
```

As part of upgrading a package, RPM automatically uninstalls any old versions of the **foo** package. In fact, you may want to always use **-U** to install packages which works even when there are no previous versions of the package installed.

> **Tip**
>
> You don't want to use the **-U** option for installing kernel packages because RPM replaces the previous kernel package. This does not affect a running system, but if the new kernel is unable to boot during your next restart, there would be no other kernel to boot instead.
> Using the **-i** option adds the kernel to your GRUB boot menu (**/etc/grub.conf**). Similarly, removing an old, unneeded kernel removes the kernel from GRUB.

Because 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, the output is similar to the following:

```
package foo-2.0-1 (which is newer than foo-1.0-1) is already installed
```

To force RPM to upgrade anyway, use the **--oldpackage** option:

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

15.2.5. 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 is upgraded to the newer version. However, RPM's freshen option does not install a package if no previously-installed package of the same name exists. This differs from RPM's upgrade option, as an upgrade *does* install packages, whether or not an older version of the package was already installed.

RPM's freshen option works for single packages or package groups. 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 does the job. If you use freshening, you do not have to delete any unwanted packages from the group that you downloaded before using RPM.

In this case, issue the following command:

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

RPM automatically upgrades only those packages that are already installed.

15.2.6. Querying

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

```
foo-2.0-1
```

> **Note**
>
> To query a package, replace **foo** with the actual package name.

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

- **-a** queries all currently installed packages.
- **-f** *<file>* queries the package which owns *<file>*. When specifying a file, you must specify the full path of the file (for example, **/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 Query 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, add **-v** to the command to display the lists in a familiar **ls -l** format.

15.2.7. 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 Verify 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 /usr/bin/vim
```

- 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 is no output. If there are any discrepancies, they are 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 period (**.**) means the test passed. The following characters denote failure of certain tests:

- **5** — MD5 checksum
- **S** — file size
- **L** — symbolic link
- **T** — file modification time
- **D** — device
- **U** — user
- **G** — group
- **M** — mode (includes permissions and file type)
- **?** — unreadable file

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

15.3. 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 file name of the RPM package):

```
rpm -K --nosignature <rpm-file>
```

The message **<rpm-file>: md5 OK** is displayed. 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 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 installation, 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 must import Red Hat's public key.

15.3.1. Importing Keys

To verify 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 includes:

```
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
```

15.3.2. 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, the following message is displayed: **md5 gpg OK**. That means that the signature of the package has been verified and that it is not corrupt.

15.4. 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. 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 and 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, enter:

```
rpm -qf /usr/bin/ggv
```

The output would look like the following:

```
ggv-2.6.0-2
```

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**. Enter the following command,

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

and the appropriate package is 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/free
```

The output would be similar to the following:

```
/usr/share/doc/procps-3.2.3/BUGS
/usr/share/doc/procps-3.2.3/FAQ
/usr/share/doc/procps-3.2.3/NEWS
/usr/share/doc/procps-3.2.3/TODO
/usr/share/man/man1/free.1.gz
/usr/share/man/man1/pgrep.1.gz
/usr/share/man/man1/pkill.1.gz
/usr/share/man/man1/pmap.1.gz
/usr/share/man/man1/ps.1.gz
/usr/share/man/man1/skill.1.gz
/usr/share/man/man1/slabtop.1.gz
/usr/share/man/man1/snice.1.gz
/usr/share/man/man1/tload.1.gz
/usr/share/man/man1/top.1.gz
/usr/share/man/man1/uptime.1.gz
/usr/share/man/man1/w.1.gz
/usr/share/man/man1/watch.1.gz
/usr/share/man/man5/sysctl.conf.5.gz
/usr/share/man/man8/sysctl.8.gz
/usr/share/man/man8/vmstat.8.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 crontabs-1.10-7.noarch.rpm
```

The output would be similar to the following:

```
Name        : crontabs                   Relocations: (not relocatable)
Version     : 1.10                              Vendor: Red Hat, Inc
Release     : 7                             Build Date: Mon 20 Sep 2004
05:58:10 PM EDT
Install Date: (not installed)             Build Host: tweety.build.redhat.com
Group       : System Environment/Base     Source RPM: crontabs-1.10-7.src.rpm
Size        : 1004                            License: Public Domain
Signature   : DSA/SHA1, Wed 05 Jan 2005 06:05:25 PM EST, Key ID
219180cddb42a60e
Packager    : Red Hat, Inc <http://bugzilla.redhat.com/bugzilla>
Summary     : Root crontab files used to schedule the execution of programs.
Description :
The crontabs package contains root crontab files. Crontab is the
program used to install, uninstall, or list the tables used to drive the
cron daemon. The cron daemon checks the crontab files to see when
particular commands are scheduled to be executed. If commands are
scheduled, then it executes them.
```

▷ Perhaps you now want to see what files the **crontabs** RPM installs. You would enter the following:

```
rpm -qlp crontabs-1.10-5.noarch.rpm
```

The output is similar to the following:

```
/etc/cron.daily
/etc/cron.hourly
/etc/cron.monthly
/etc/cron.weekly
/etc/crontab
/usr/bin/run-parts
```

These are just a few examples. As you use it, you will find many more uses for RPM.

15.5. 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.

15.5.1. Installed Documentation

▷ **rpm --help** — This command displays a quick reference of RPM parameters.

▷ **man rpm** — The RPM man page gives more detail about RPM parameters than the **rpm --help** command.

15.5.2. Useful Websites

▷ http://www.rpm.org/ — The RPM website.

▷ http://www.redhat.com/mailman/listinfo/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.

15.5.3. Related Books

▷ *Red Hat RPM Guide* by Eric Foster-Johnson; Wiley, John &Sons, Incorporated — This book is a comprehensive guide to RPM, from installing package to building RPMs.

Chapter 16. Red Hat Network

Red Hat Network is an Internet solution for managing one or more Red Hat Enterprise Linux systems. All Security Alerts, Bug Fix Alerts, and Enhancement Alerts (collectively known as Errata Alerts) can be downloaded directly from Red Hat using the **Package Updater** standalone application or through the RHN website available at https://rhn.redhat.com/.
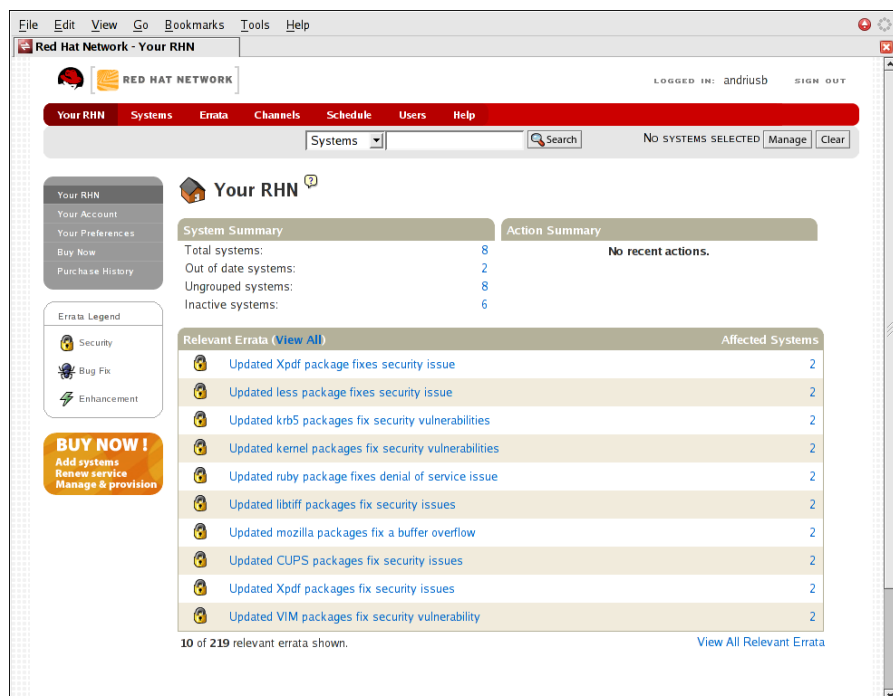
Figure 16.1. Your RHN

Red Hat Network saves you time because you receive email when updated packages are released. You do not have to search the Web for updated packages or security alerts. By default, Red Hat Network installs the packages as well. You do not have to learn how to use RPM or worry about resolving software package dependencies; RHN does it all.

Red Hat Network features include:

- Errata Alerts — learn when Security Alerts, Bug Fix Alerts, and Enhancement Alerts are issued for all the systems in your network

Figure 16.2. Relevant Errata

- Automatic email notifications — Receive an email notification when an Errata Alert is issued for your system(s)
- Scheduled Errata Updates — Schedule delivery of Errata Updates
- Package installation — Schedule package installation on one or more systems with the click of a button
- **Package Updater** — Use the **Package Updater** 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 through a secure Web browser connection from any computer

> ⚠️ **Warning**
>
> You must activate your Red Hat Enterprise Linux product before registering your system with Red Hat Network to make sure your system is entitled to the correct services. To activate your product, go to:
>
> http://www.redhat.com/apps/activate/

After activating your product, register it with Red Hat Network to receive Errata Updates. The registration process gathers information about the system that is required to notify you of updates. For example, a list of packages installed on the system is compiled so you are only notified about updates that are relevant to your system.

The first time the system is booted, the **Software Update Setup Assistant** prompts you to register. If you did not register then, select **Applications** (the main menu on the panel) => **System Tools** => **Package Updater** on your desktop to start the registration process. Alternately, execute the command

yum update from a shell prompt.
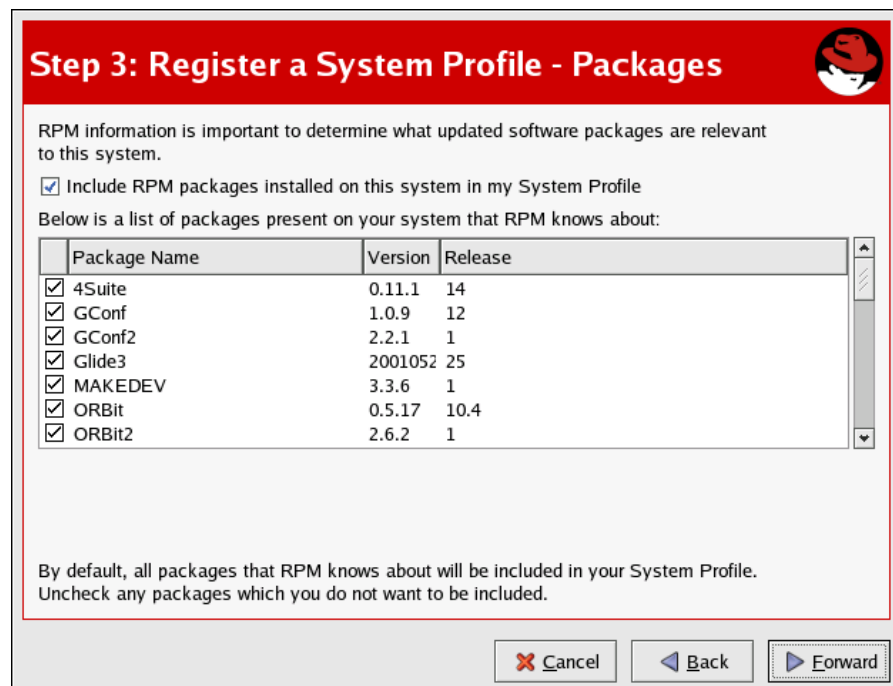
Figure 16.3. Registering with RHN

After registering, use one of the following methods to start receiving updates:

- Select **Applications** (the main menu on the panel) => **System Tools** => **Package Updater** on your desktop
- Execute the command **yum** from a shell prompt
- Use the RHN website at https://rhn.redhat.com/
- Click on the package icon when it appears in the panel to launch the **Package Updater**.

For more detailed instructions, refer to the documentation available at:

http://www.redhat.com/docs/manuals/RHNetwork/

> **Tip**
>
> Red Hat Enterprise Linux includes a convenient panel icon that displays visible alerts when there is an update for your Red Hat Enterprise Linux system. This panel icon is not present if no updates are available.

Part IV. Network-Related Configuration

After explaining how to configure the network, this part discusses topics related to networking such as how to allow remote logins, share files and directories over the network, and set up a Web server.

Chapter 17. Network Configuration

To communicate with each other, computers must have 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

It can also be used to configure IPsec connections, manage DNS settings, and manage the **/etc/hosts** file used to store additional hostnames and IP address combinations.

To use the **Network Administration Tool**, you must have root privileges. To start the application, go to the **Applications** (the main menu on the panel) => **System Settings** => **Network**, or type the command **system-config-network** at a shell prompt (for example, in an **XTerm** or a **GNOME terminal**). If you type the command, the graphical version is displayed if **X** is running; otherwise, the text-based version is displayed.

To use the command line version, execute the command **system-config-network-cmd --help** as root to view all of the options.
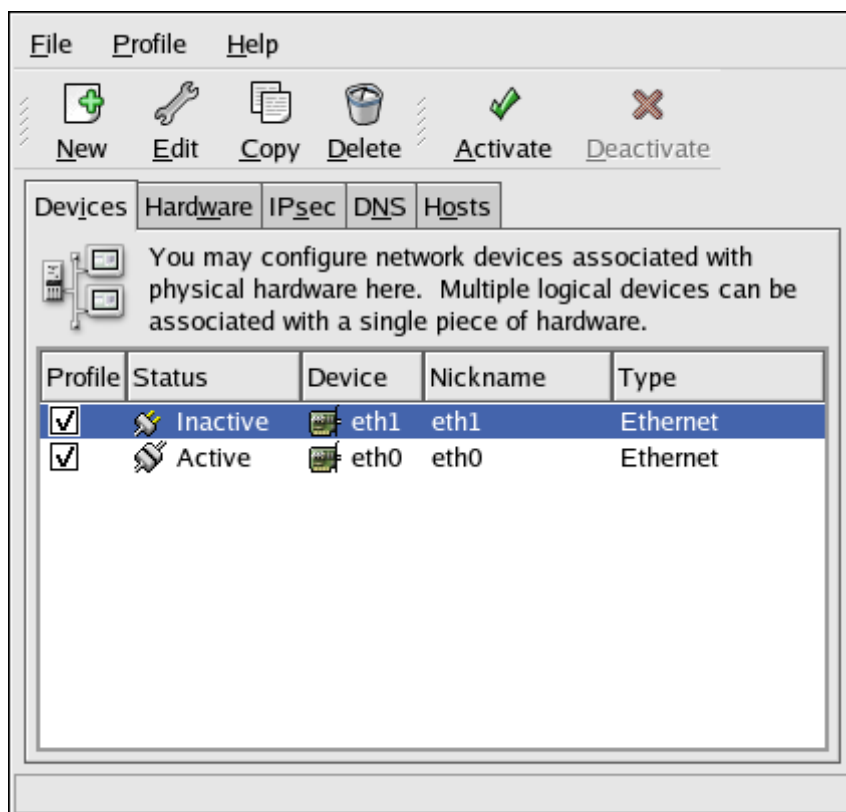
Figure 17.1. Network Administration Tool

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

17.1. Overview

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

1. Add a network device associated with the physical hardware device.
2. Add the physical hardware device to the hardware list, if it does not already exist.
3. Configure the hostname and DNS settings.
4. Configure any hosts that cannot be looked up through DNS.

This chapter discusses each of these steps for each type of network connection.

17.2. 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. Different networks are configured to use different network speeds; 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 **New** button on the toolbar.
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 detects supported Ethernet devices and prompts you to configure them. If you configured any Ethernet devices during the installation, they are displayed 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. In the **Configure Network Settings** window shown in Figure 17.2, "Ethernet Settings", 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.
7. Click **Apply** on the **Create Ethernet Device** page.

Figure 17.2. Ethernet Settings

After configuring the Ethernet device, it appears in the device list as shown in Figure 17.3, "Ethernet Device".
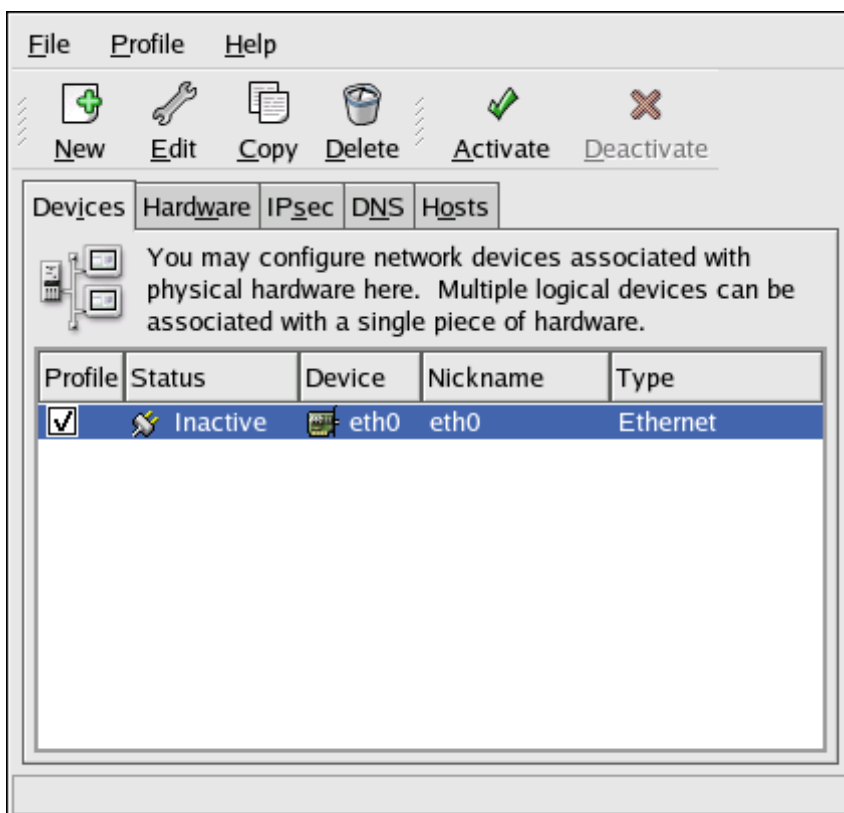
Figure 17.3. Ethernet Device

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. To change this setting, select to edit the device, modify the **Activate device when computer starts** value, and save the changes.

If you associate more than one device with an Ethernet card, the subsequent devices are *device aliases*. A device alias allows you to setup multiple virtual devices for one physical device, thus giving the one physical device more than one IP address. For example, you can configure an eth1 device and an eth1:1 device. For details, refer to Section 17.11, "Device Aliases".

17.3. 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 **New** button on the toolbar.
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.
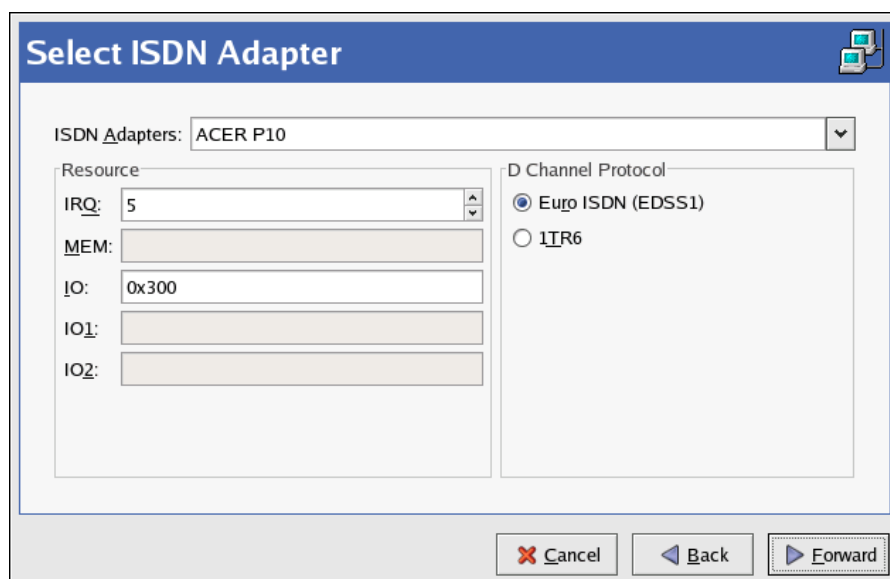
Figure 17.4. ISDN Settings

5. If your Internet Service Provider (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. In the **IP Settings** window, select the **Encapsulation Mode** and whether to obtain an IP address automatically or to set a static IP instead. Click **Forward** when finished.
7. On the **Create Dialup Connection** page, click **Apply**.

After configuring the ISDN device, it appears in the device list as a device with type **ISDN** as shown in Figure 17.5, "ISDN Device".

Be sure to select **File** => **Save** to save the changes.

After adding the ISDN 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 immediately, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.
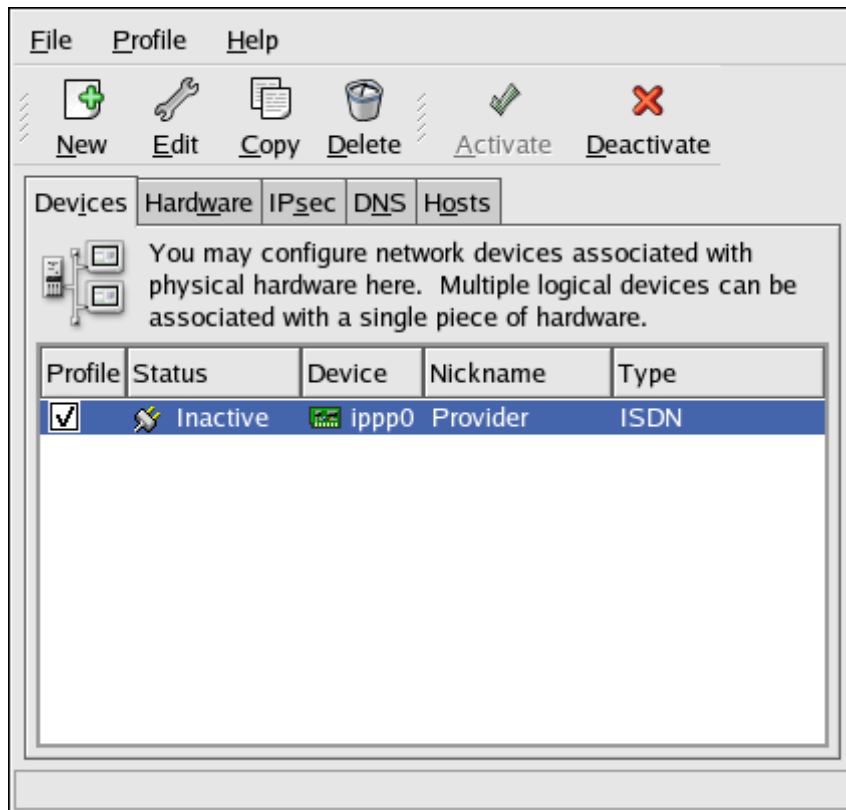
Figure 17.5. ISDN Device

17.4. 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 **New** button on the toolbar.
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), the **Network Administration Tool** assumes you want to use it to establish a modem connection. If there are no modems already configured, it tries to detect any modems in the system. This probe might take a while. If a modem is not found, a message is displayed to warn you that the settings shown are not values found from the probe.
5. After probing, the window in appears.

Figure 17.6. Modem Settings

6. Configure the modem device, baud rate, flow control, and modem volume. If you do not know these values, accept the defaults if the modem was probed successfully. If you do not have touch tone dialing, uncheck the corresponding checkbox. Click **Forward**.

7. 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 these values, contact your ISP. Click **Forward**.

8. On the **IP Settings** page, select whether to obtain an IP address automatically or whether to set one statically. Click **Forward** when finished.

9. On the **Create Dialup Connection** page, click **Apply**.

After configuring the modem device, it appears in the device list with the type **Modem** as shown in Figure 17.7, "Modem Device".
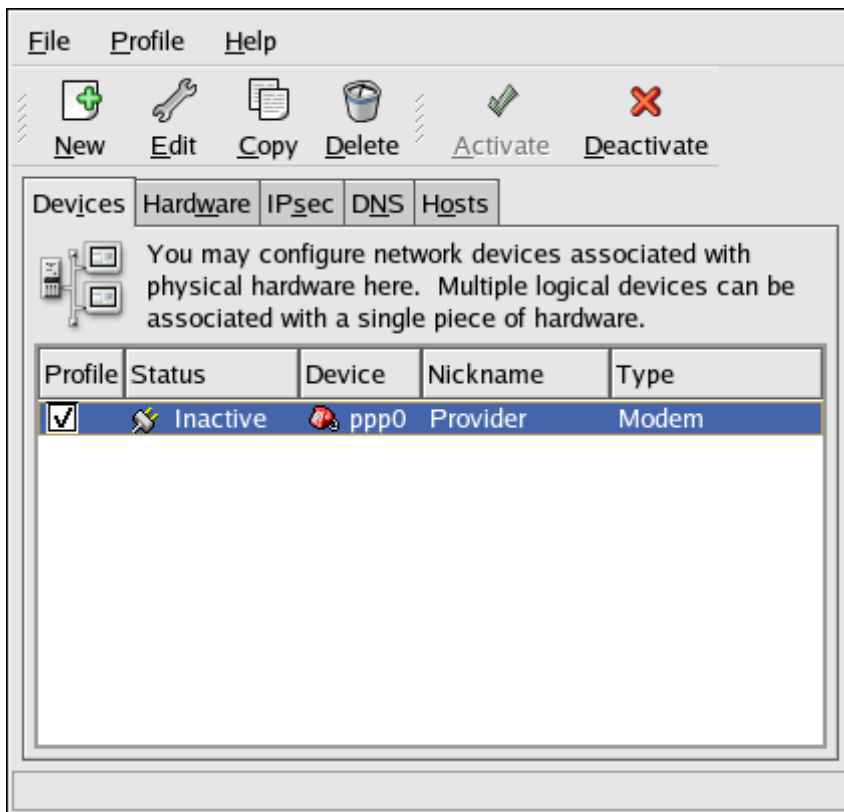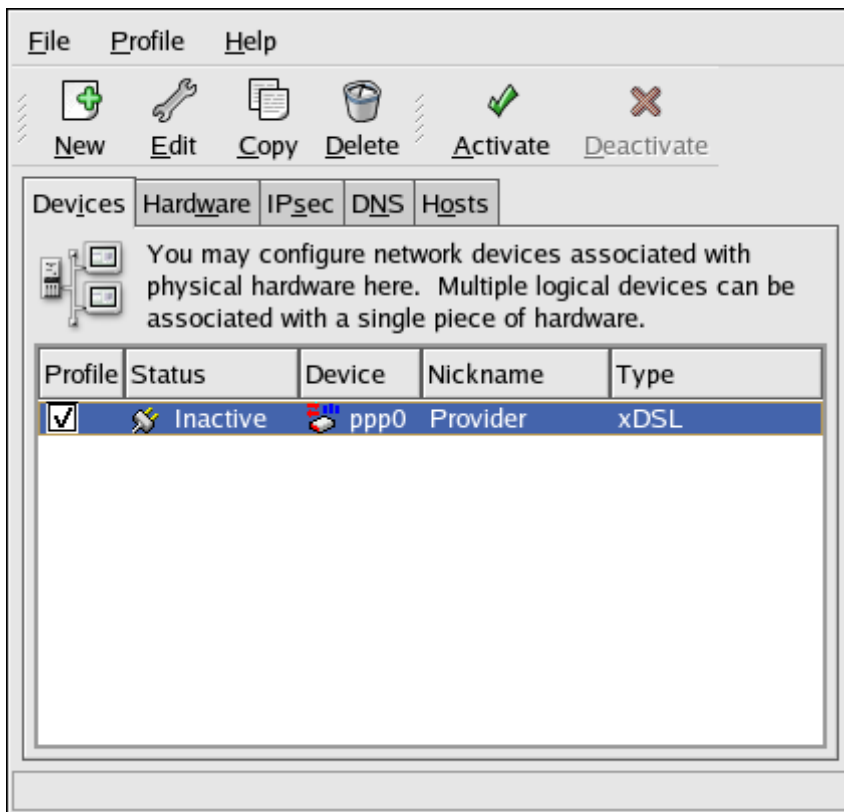
Figure 17.7. Modem Device

Be sure to select **File** => **Save** 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 immediately, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

17.5. Establishing an xDSL Connection

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

Some DSL providers require that the system is configured 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 17.2, "Establishing an Ethernet Connection" to configure your Ethernet card.

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

1. Click the **Devices** tab.
2. Click the **New** button.

3. Select **xDSL connection** from the **Device Type** list, and click **Forward**.

4. If your Ethernet card is in the hardware list, select the **Ethernet Device** from the pulldown menu from the page shown in Figure 17.8, "xDSL Settings". Otherwise, the **Select Ethernet Adapter** window appears.

> **Note**
>
> The installation program detects supported Ethernet devices and prompts you to configure them. If you configured any Ethernet devices during the installation, they are displayed in the hardware list on the **Hardware** tab.

Figure 17.8. 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**. If you have a T-Online account, instead of entering a **Login Name** and **Password** in the default window, click the **T-Online Account Setup** button and enter the required information. Click **Forward** to continue.

7. On the **Create DSL Connection** page, click **Apply**.

After configuring the DSL connection, it appears in the device list as shown in Figure 17.7, "Modem Device".
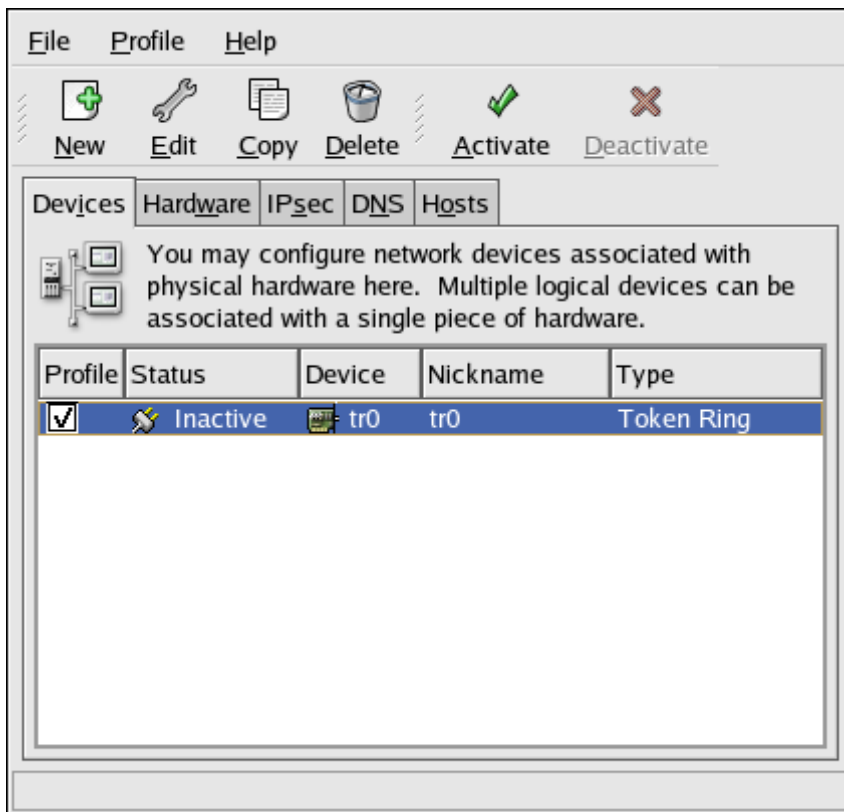
Figure 17.9. xDSL Device

Be sure to select **File** => **Save** 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 immediately, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

17.6. 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 rings 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 **New** button on the toolbar.

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 **Tokenring 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 17.10, "Token Ring Settings" appears. Select the manufacturer and model of the adapter. Select the device name. If this is the system's first token ring card, select **tr0**; 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.
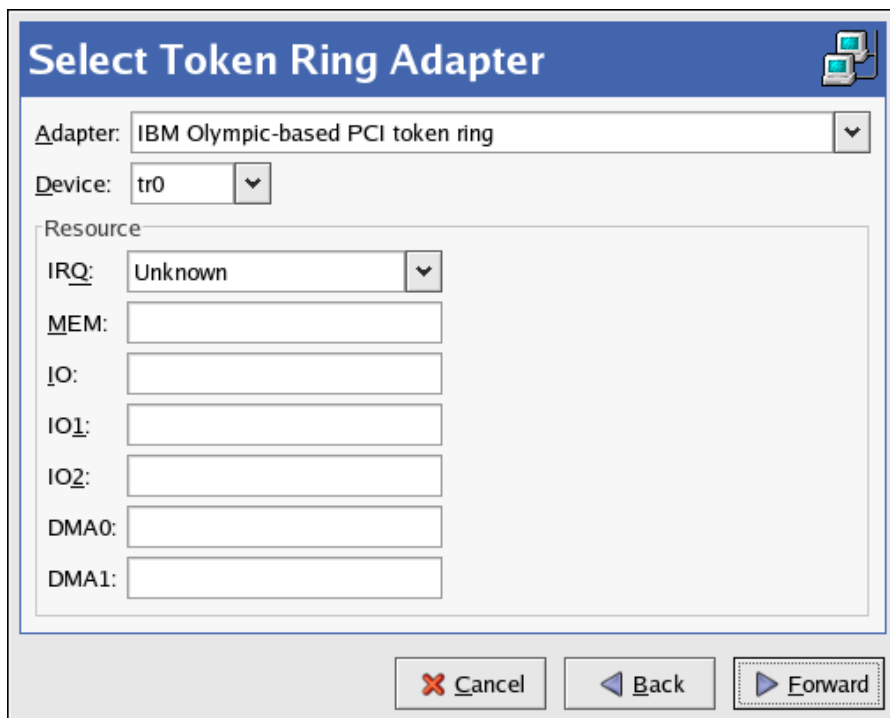
Figure 17.10. Token Ring Settings

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 Tokenring Device** page.

After configuring the token ring device, it appears in the device list as shown in Figure 17.11, "Token Ring Device".
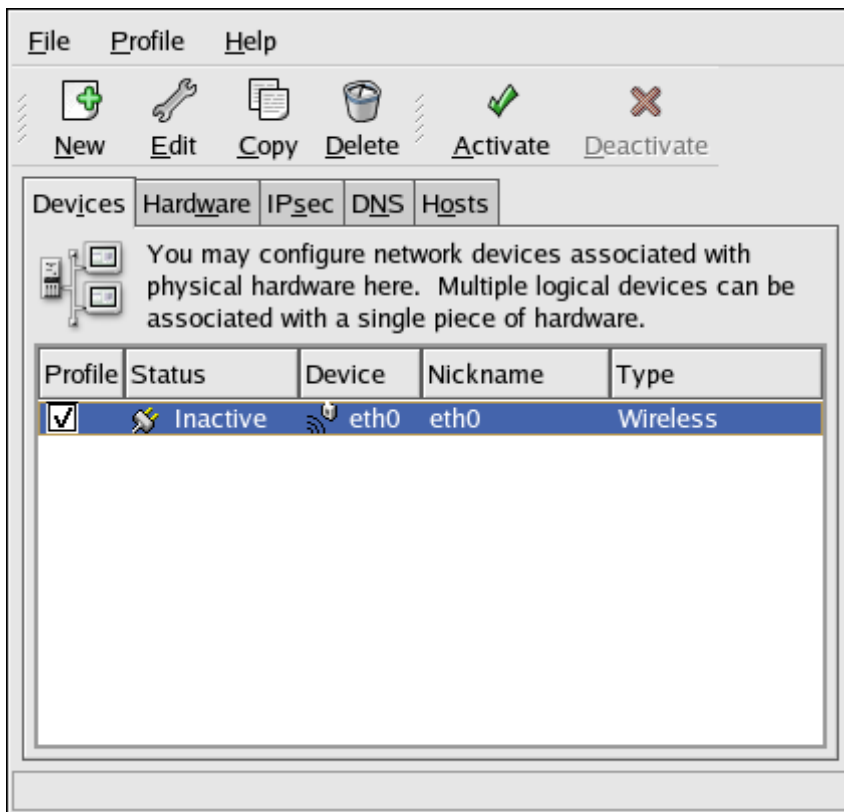
Figure 17.11. Token Ring Device

Be sure to select **File** => **Save** 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 immediately, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

17.7. 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 settings such as the SSID and key for the wireless device.

To add a wireless Ethernet connection, follow these steps:

1. Click the **Devices** tab.
2. Click the **New** button on the toolbar.
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 **Wireless card** list. Otherwise, select **Other Wireless 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, they are displayed in the hardware list on the **Hardware** tab.

5. If you selected **Other Wireless 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 17.12, "Wireless Settings", configure the settings for the wireless device.

Figure 17.12. 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 17.13, "Wireless Device".

Figure 17.13. Wireless Device

Be sure to select **File** => **Save** 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 immediately, as seen by its **Inactive** status. To activate the device, select it from the device list, and click the **Activate** button. If the system is configured to activate the device when the computer starts (the default), this step does not have to be performed again.

17.8. 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.

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

If the hostname is retrieved dynamically from DHCP or PPPoE (or retrieved from the ISP), do not change it.
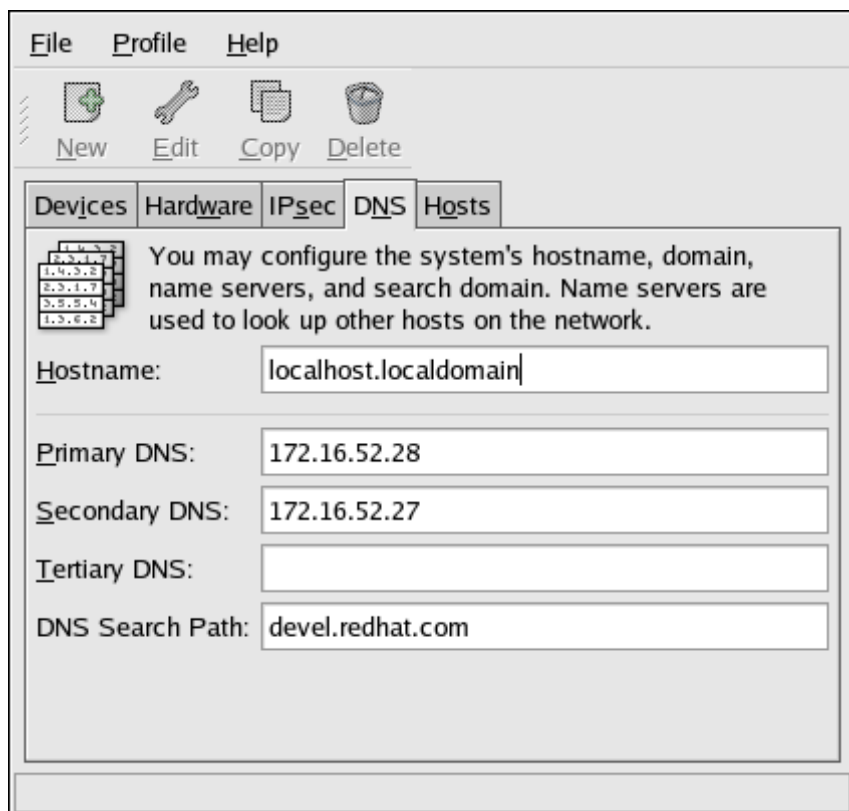
Figure 17.14. DNS Configuration

> **Note**
>
> The name servers section does not configure the system to be a name server. Instead, it configures which name servers to use when resolving IP addresses to hostnames and vice-versa.

> **Warning**
>
> If the hostname is changed and `system-config-network` is started on the local host, you may not be able to start another **X11** application. As such, you may have to re-login to a new desktop session.

17.9. 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 tries to 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 Enterprise 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, go to the **Hosts** tab, click the **New** button on the toolbar, provide the requested information, and click **OK**. Select **File** => **Save** or press **Ctrl**+**S** to save the changes to the **/etc/hosts** file. The network or network services do not need to be restarted since the current version of the file is referred to each time an address is resolved.

> ⚠️ **Warning**
>
> Do not remove the **localhost** entry. Even if the system does not have a network connection or have a network connection running constantly, some programs need to connect to the system via the localhost loopback interface.
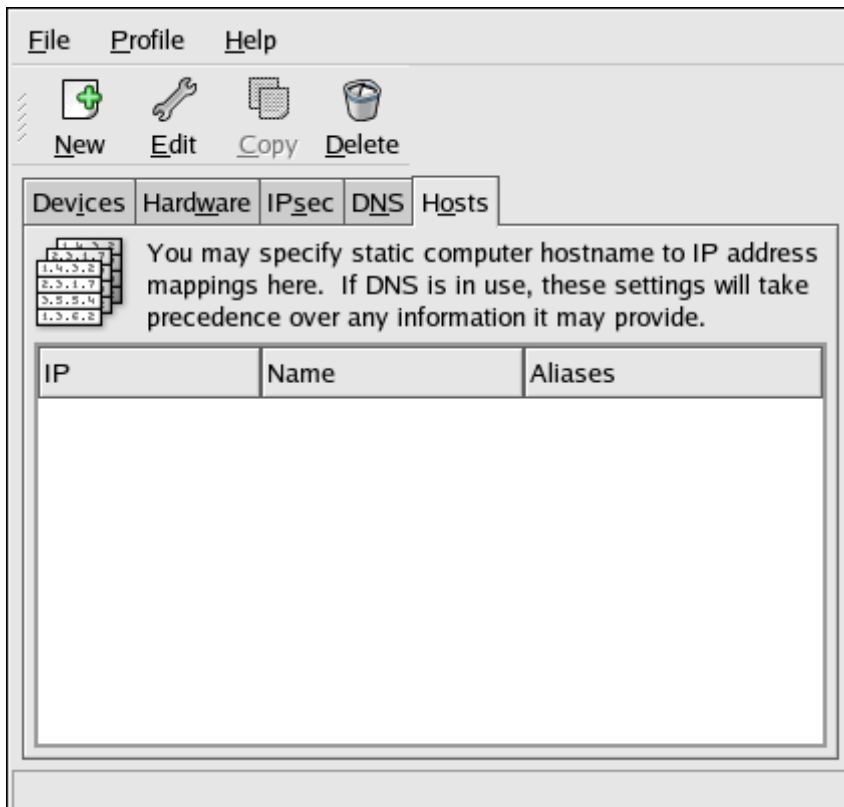
Figure 17.15. Hosts Configuration

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

17.10. 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 to be specifically 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 cannot 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 17.11, "Device Aliases" 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, select **Profile** => **New** from the pull-down menu, and enter a unique name for the profile.

You are now modifying the new profile as indicated by the status bar at the bottom of the main window.

Click on an existing device already in the list and click the **Copy** button to copy the existing device to a logical network device. If you use the **New** button, a network alias is created, which is incorrect. To change the properties of the logical device, select it from the list and click **Edit**. For example, the nickname can be changed to a more descriptive name, such as **eth0_office**, so that it can be recognized more easily.

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, if you create a logical device named **eth0_office** in a profile called **Office** and want to activate the logical device if the profile is selected, uncheck the **eth0** device and check the **eth0_office** device.

For example, Figure 17.16, "Office Profile" shows a profile called **Office** with the logical device **eth0_office**. It is configured to activate the first Ethernet card using DHCP.
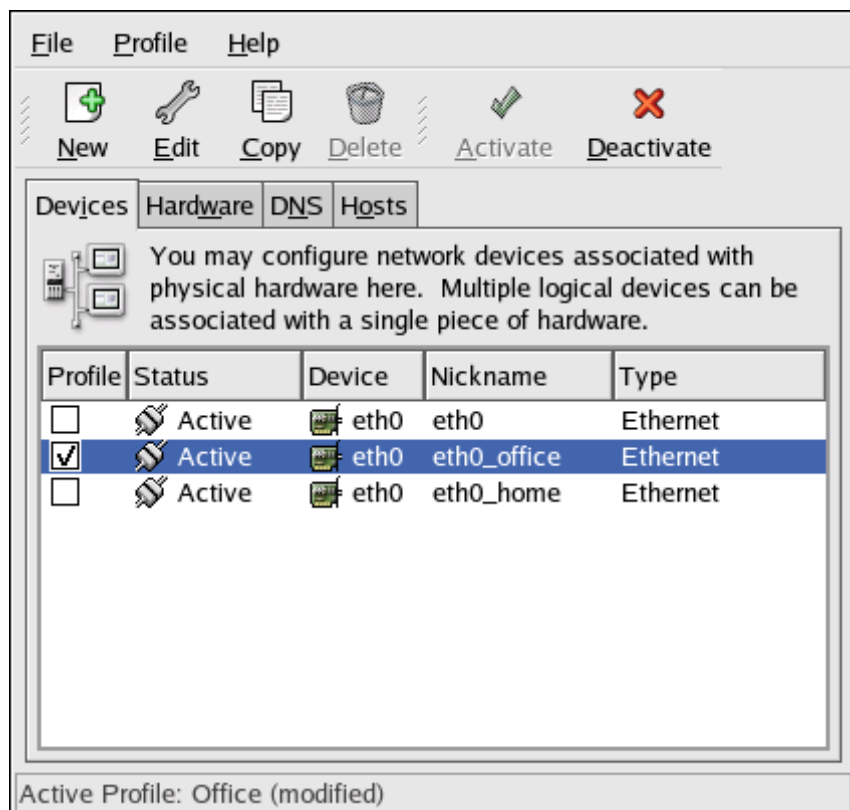
Figure 17.16. Office Profile

Notice that the **Home** profile as shown in Figure 17.17, "Home Profile" activates the **eth0_home** logical device, which is associated with **eth0**.
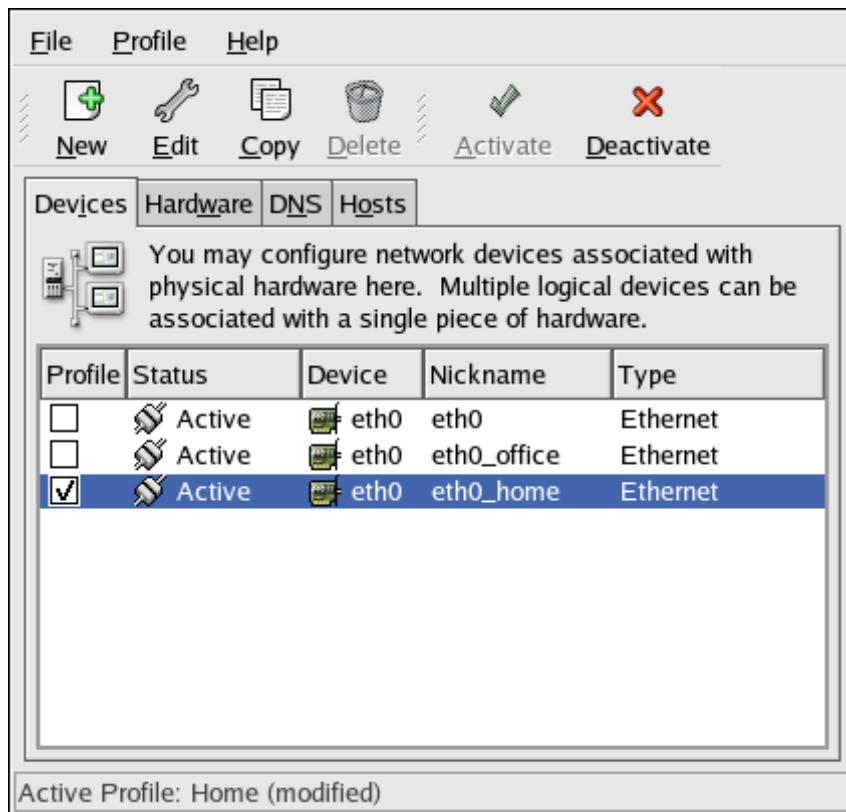
Figure 17.17. Home Profile

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

To activate a profile at boot time, modify the boot loader configuration file to include the **netprofile=<*profilename*>** option. For example, if the system uses GRUB as the boot loader and **/boot/grub/grub.conf** contains:

```
title Red Hat Enterprise Linux (2.6.9-5.EL)
        root (hd0,0)
   kernel /vmlinuz-2.6.9-5.EL ro root=/dev/VolGroup00/LogVol00 rhgb quiet
   initrd /initrd-2.6.9-5.EL.img
```

Modify it to the following (where **<*profilename*>** is the name of the profile to be activated at boot time):

```
title Red Hat Enterprise Linux (2.6.9-5.EL)
        root (hd0,0)
   kernel /vmlinuz-2.6.9-5.EL ro root=/dev/VolGroup00/LogVol00 \
    netprofile=<profilename> \      rhgb quiet
   initrd /initrd-2.6.9-5.EL.img
```

To switch profiles after the system has booted, go to **Applications** (the main menu on the panel) => **System Tools** => **Network Device Control** (or type the command `system-control-network`) to select a profile and activate it. The activate profile section only appears in the **Network Device Control** interface if more than the default **Common** interface exists.

Alternatively, execute the following command to enable a profile (replace ***<profilename>*** with the name of the profile):

```
system-config-network-cmd --profile <profilename> --activate
```

17.11. 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 multiple IP addresses for a system that only has one network card.

After configuring the Ethernet device —such as **eth0** —to use a static IP address (DHCP does not work with aliases), go to the **Devices** tab and click **New**. Select the Ethernet card to configure with an alias, set the static IP address for the alias, and click **Apply** to create it. Since a device already exists for the Ethernet card, the one just created is the alias, such as **eth0:1**.

> ⚠️ **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 17.18, "Network Device Alias Example" 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.
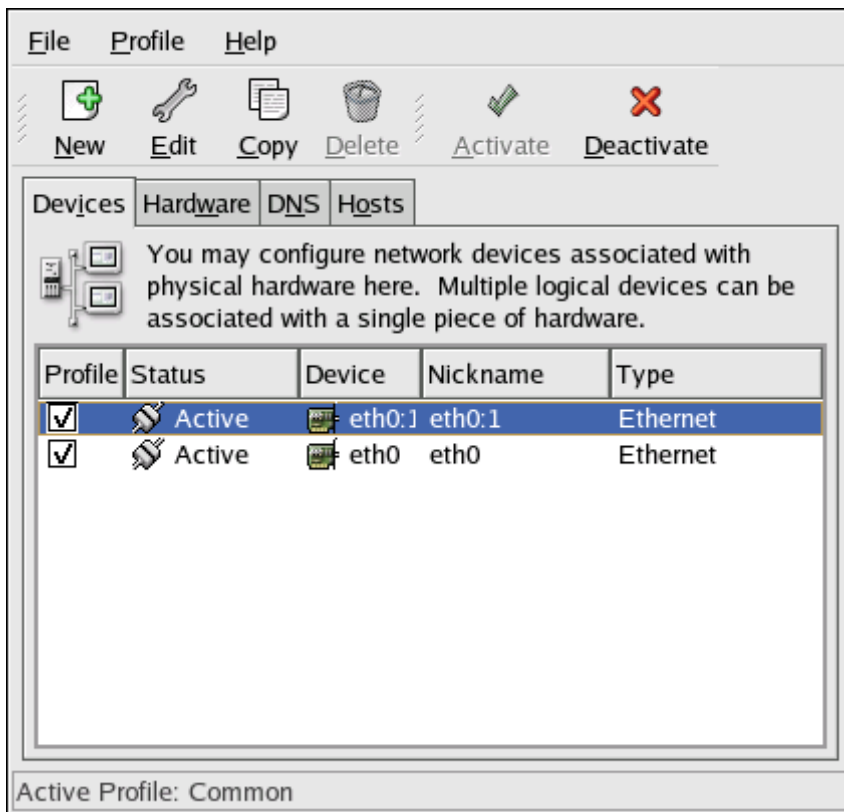
Figure 17.18. Network Device Alias Example

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 addresses:

```
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)
```

17.12. Saving and Restoring the Network Configuration

The command line version of **Network Administration Tool** can be used to save the system's network configuration to a file. This file can then be used to restore the network settings to a Red Hat Enterprise Linux system.

This feature can be used as part of an automated backup script, to save the configuration before upgrading or reinstalling, or to copy the configuration to a different Red Hat Enterprise Linux system.

To save, or *export*, the network configuration of a system to the file **/tmp/network-config**, execute the following command as root:

```
system-config-network-cmd -e > /tmp/network-config
```

To restore, or *import*, the network configuration from the file created from the previous command, execute the following command as root:

```
system-config-network-cmd -i -c -f /tmp/network-config
```

The **-i** option means to import the data, the **-c** option means to clear the existing configuration prior to importing, and the **-f** option specifies that the file to import is as follows.

Chapter 18. Firewalls

Information security is commonly thought of as a process and not a product. However, standard security implementations usually employ some form of dedicated mechanism to control access privileges and restrict network resources to users who are authorized, identifiable, and traceable. Red Hat Enterprise Linux includes several tools to assist administrators and security engineers with network-level access control issues.

Firewalls are one of the core components of a network security implementation. Several vendors market firewall solutions catering to all levels of the marketplace: from home users protecting one PC to data center solutions safeguarding vital enterprise information. Firewalls can be stand-alone hardware solutions, such as firewall appliances by Cisco, Nokia, and Sonicwall. Vendors such as Checkpoint, McAfee, and Symantec have also developed proprietary software firewall solutions for home and business markets.

Apart from the differences between hardware and software firewalls, there are also differences in the way firewalls function that separate one solution from another. Table 18.1, "Firewall Types" details three common types of firewalls and how they function:

Table 18.1. Firewall Types

Method	Description	Advantages	Disadvantages
NAT	*Network Address Translation* (NAT) places private IP subnetworks behind one or a small pool of public IP addresses, masquerading all requests to one source rather than several. The Linux kernel has built-in NAT functionality through the Netfilter kernel subsystem.	· Can be configured transparently to machines on a LAN · Protection of many machines and services behind one or more external IP addresses simplifies administration duties · Restriction of user access to and from the LAN can be configured by opening and closing ports on the NAT firewall/gateway	· Cannot prevent malicious activity once users connect to a service outside of the firewall
Packet Filter	A packet filtering firewall reads each data packet that passes through a LAN. It can read and process packets by header information and filters the packet based on sets of programmable rules implemented by the firewall administrator. The Linux kernel has built-in packet filtering functionality through the Netfilter kernel subsystem.	· Customizable through the `iptables` front-end utility · Does not require any customization on the client side, as all network activity is filtered at the router level rather than the application level · Since packets are not transmitted through a proxy, network performance is faster due to direct connection from client to remote host	· Cannot filter packets for content like proxy firewalls · Processes packets at the protocol layer, but cannot filter packets at an application layer · Complex network architectures can make establishing packet filtering rules difficult, especially if coupled with *IP masquerading* or local subnets and DMZ networks
Proxy	Proxy firewalls filter all requests of a certain protocol or type from LAN clients to a proxy machine, which then makes those requests to the Internet on behalf of the local client. A proxy machine acts as a buffer between malicious remote users and the internal network client machines.	· Gives administrators control over what applications and protocols function outside of the LAN · Some proxy servers can cache frequently-accessed data locally rather than having to use the Internet connection to request it. This helps to reduce bandwidth consumption · Proxy services can be logged and monitored closely, allowing tighter	· Proxies are often application-specific (HTTP, Telnet, etc.), or protocol-restricted (most proxies work with TCP-connected services only) · Application services cannot run behind a proxy, so your application servers must use a separate form of network security · Proxies can become a network bottleneck, as all requests and transmissions

	control over resource utilization on the network	are passed through one source rather than directly from a client to a remote service

18.1. Netfilter and IPTables

The Linux kernel features a powerful networking subsystem called *Netfilter*. The Netfilter subsystem provides stateful or stateless packet filtering as well as NAT and IP masquerading services. Netfilter also has the ability to *mangle* IP header information for advanced routing and connection state management. Netfilter is controlled using the `iptables` tool.

18.1.1. IPTables Overview

The power and flexibility of Netfilter is implemented using the `iptables` administration tool, a command line tool similar in syntax to its predecessor, `ipchains`.

A similar syntax does not mean similar implementation, however. `ipchains` requires intricate rule sets for: filtering source paths; filtering destination paths; and filtering both source and destination connection ports.

By contrast, `iptables` uses the Netfilter subsystem to enhance network connection, inspection, and processing. `iptables` features advanced logging, pre- and post-routing actions, network address translation, and port forwarding, all in one command line interface.

This section provides an overview of `iptables`.

18.2. Basic Firewall Configuration

Just as a firewall in a building attempts to prevent a fire from spreading, a computer firewall attempts to prevent malicious software from spreading to your computer. It also helps to prevent unauthorized users from accessing your computer.

In a default Red Hat Enterprise Linux installation, a firewall exists between your computer or network and any untrusted networks, for example the Internet. It determines which services on your computer remote users 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 Enterprise Linux system with an Internet connection.

18.2.1. Security Level Configuration Tool

During the `Firewall Configuration` screen of the Red Hat Enterprise Linux installation, you were given the option to enable a basic firewall as well as to allow specific devices, incoming services, and ports.

After installation, you can change this preference by using the **Security Level Configuration Tool**.

To start this application, use the following command:

```
[root@myServer ~] # system-config-selinux
```
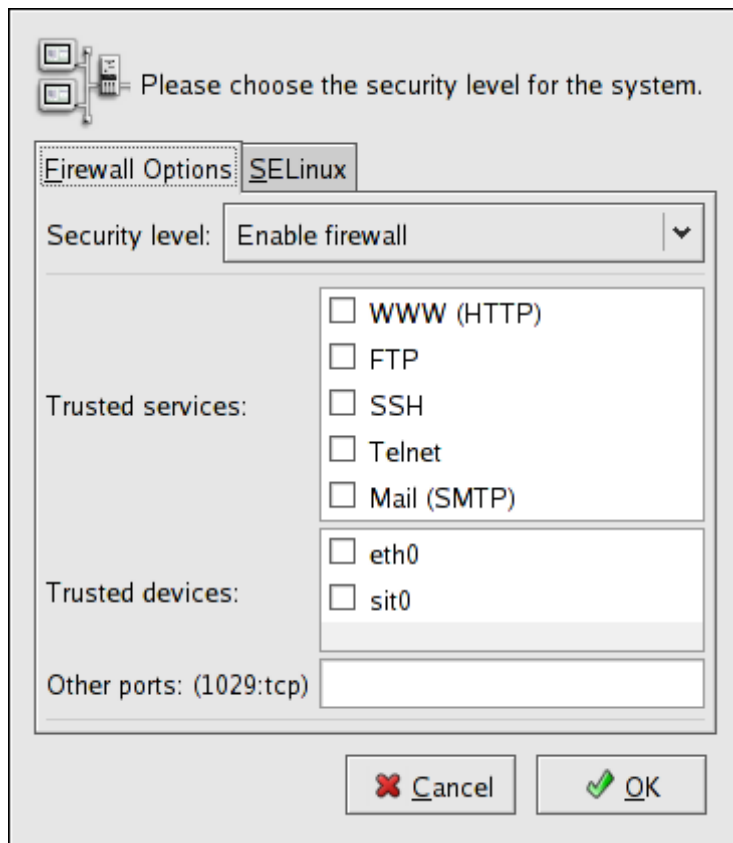
Figure 18.1. Security Level Configuration Tool

> **Note**
>
> The **Security Level Configuration Tool** only configures a basic firewall.

18.2.2. Enabling and Disabling the Firewall

Select one of the following options for the firewall:

▸ **Disabled** — Disabling the firewall provides complete access to your system and does no security checking. This should only be selected if you are running on a trusted network (not the Internet) or need to configure a custom firewall using the iptables command line tool.

> **Warning**
>
> Firewall configurations and any customized firewall rules are stored in the **/etc/sysconfig/iptables** file. If you choose **Disabled** and click **OK**, these configurations and firewall rules will be lost.

▸ **Enabled** — This option configures the system to reject incoming connections that are not in response to outbound requests, such as DNS replies or DHCP requests. If access to services running on this machine is needed, you can choose to allow specific services through the firewall.

If you are connecting your system to the Internet, but do not plan to run a server, this is the safest choice.

18.2.3. Trusted Services

Enabling options in the **`Trusted services`** list allows the specified service to pass through the firewall.

`WWW (HTTP)`

> The HTTP protocol is used by Apache (and by other Web servers) to serve web pages. If you plan on making your Web server publicly available, select this check box. This option is not required for viewing pages locally or for developing web pages. This service requires that the **`httpd`** package be installed.
>
> Enabling **`WWW (HTTP)`** will not open a port for HTTPS, the SSL version of HTTP. If this service is required, select the **`Secure WWW (HTTPS)`** check box.

`FTP`

> The FTP protocol is used to transfer files between machines on a network. If you plan on making your FTP server publicly available, select this check box. This service requires that the **`vsftpd`** package be installed.

`SSH`

> Secure Shell (SSH) is a suite of tools for logging into and executing commands on a remote machine. To allow remote access to the machine via ssh, select this check box. This service requires that the **`openssh-server`** package be installed.

`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. To allow remote access to the machine via telnet, select this check box. This service requires that the **`telnet-server`** package be installed.

`Mail (SMTP)`

> SMTP is a protocol that allows remote hosts to connect directly to your machine to deliver mail. You do not need to enable this service if you collect your mail from your ISP's server using POP3 or IMAP, or if you use a tool such as **`fetchmail`**. To allow delivery of mail to your machine, select this check box. Note that an improperly configured SMTP server can allow remote machines to use your server to send spam.

`NFS4`

> The Network File System (NFS) is a file sharing protocol commonly used on *NIX systems. Version 4 of this protocol is more secure than its predecessors. If you want to share files or directories on your system with other network users, select this check box.

`Samba`

> Samba is an implementation of Microsoft's proprietary SMB networking protocol. If you need to share files, directories, or locally-connected printers with Microsoft Windows machines, select this check box.

18.2.4. Other Ports

The **Security Level Configuration Tool** includes an **Other ports** section for specifying custom IP ports as being trusted by **iptables**. For example, to allow IRC and Internet printing protocol (IPP) to pass through the firewall, add the following to the **Other ports** section:

194:tcp,631:tcp

18.2.5. Saving the Settings

Click **OK** to save the changes and enable or disable the firewall. If **Enable firewall** was selected, the options selected are translated to **iptables** commands and written to the **/etc/sysconfig/iptables** file. The **iptables** service is also started so that the firewall is activated immediately after saving the selected options. If **Disable firewall** was selected, the **/etc/sysconfig/iptables** file is removed and the **iptables** service is stopped immediately.

The selected options are also written to the **/etc/sysconfig/system-config-selinux** file so that the settings can be restored the next time the application is started. Do not edit this file by hand.

Even though the firewall is activated immediately, the **iptables** service is not configured to start automatically at boot time. Refer to <u>Section 18.2.6, "Activating the IPTables Service"</u> for more information.

18.2.6. Activating the IPTables Service

The firewall rules are only active if the **iptables** service is running. To manually start the service, use the following command:

```
[root@myServer ~] # service iptables restart
```

To ensure that **iptables** starts when the system is booted, use the following command:

```
[root@myServer ~] # chkconfig --level 345 iptables on
```

The **ipchains** service is not included in Red Hat Enterprise Linux. However, if **ipchains** is installed (for example, an upgrade was performed and the system had **ipchains** previously installed), the **ipchains** and **iptables** services should not be activated simultaneously. To make sure the **ipchains** service is disabled and configured not to start at boot time, use the following two commands:

```
[root@myServer ~] # service ipchains stop
[root@myServer ~] # chkconfig --level 345 ipchains off
```

18.3. Using IPTables

The first step in using **iptables** is to start the **iptables** service. Use the following command to start the **iptables** service:

```
[root@myServer ~] # service iptables start
```

> **Note**
>
> The **ip6tables** service can be turned off if you intend to use the **iptables** service only. If you deactivate the **ip6tables** service, remember to deactivate the IPv6 network also. Never leave a network device active without the matching firewall.

To force **iptables** to start by default when the system is booted, use the following command:

```
[root@myServer ~] # chkconfig --level 345 iptables on
```

This forces **iptables** to start whenever the system is booted into runlevel 3, 4, or 5.

18.3.1. IPTables Command Syntax

The following sample **iptables** command illustrates the basic command syntax:

```
[root@myServer ~ ] # iptables -A <chain> -j <target>
```

The **-A** option specifies that the rule be appended to *<chain>*. Each chain is comprised of one or more *rules*, and is therefore also known as a *ruleset*.

The three built-in chains are INPUT, OUTPUT, and FORWARD. These chains are permanent and cannot be deleted. The chain specifies the point at which a packet is manipulated.

The **-j <target>** option specifies the target of the rule; i.e., what to do if the packet matches the rule. Examples of built-in targets are ACCEPT, DROP, and REJECT.

Refer to the **iptables** man page for more information on the available chains, options, and targets.

18.3.2. Basic Firewall Policies

Establishing basic firewall policies creates a foundation for building more detailed, user-defined rules.

Each **iptables** chain is comprised of a default policy, and zero or more rules which work in concert with the default policy to define the overall ruleset for the firewall.

The default policy for a chain can be either DROP or ACCEPT. Security-minded administrators typically implement a default policy of DROP, and only allow specific packets on a case-by-case basis. For example, the following policies block all incoming and outgoing packets on a network gateway:

```
[root@myServer ~ ] # iptables -P INPUT DROP
[root@myServer ~ ] # iptables -P OUTPUT DROP
```

It is also recommended that any *forwarded packets* — network traffic that is to be routed from the firewall to its destination node — be denied as well, to restrict internal clients from inadvertent exposure to the Internet. To do this, use the following rule:

```
[root@myServer ~ ] # iptables -P FORWARD DROP
```

When you have established the default policies for each chain, you can create and save further rules for your particular network and security requirements.

The following sections describe how to save iptables rules and outline some of the rules you might

implement in the course of building your iptables firewall.

18.3.3. Saving and Restoring IPTables Rules

Changes to **iptables** are transitory; if the system is rebooted or if the **iptables** service is restarted, the rules are automatically flushed and reset. To save the rules so that they are loaded when the **iptables** service is started, use the following command:

```
[root@myServer ~ ] # service iptables save
```

The rules are stored in the file **/etc/sysconfig/iptables** and are applied whenever the service is started or the machine is rebooted.

18.4. Common IPTables Filtering

Preventing remote attackers from accessing a LAN is one of the most important aspects of network security. The integrity of a LAN should be protected from malicious remote users through the use of stringent firewall rules.

However, with a default policy set to block all incoming, outgoing, and forwarded packets, it is impossible for the firewall/gateway and internal LAN users to communicate with each other or with external resources.

To allow users to perform network-related functions and to use networking applications, administrators must open certain ports for communication.

For example, to allow access to port 80 *on the firewall*, append the following rule:

```
[root@myServer ~ ] # iptables -A INPUT -p tcp -m tcp --dport 80 -j ACCEPT
```

This allows users to browse websites that communicate using the standard port 80. To allow access to secure websites (for example, https://www.example.com/), you also need to provide access to port 443, as follows:

```
[root@myServer ~ ] # iptables -A INPUT -p tcp -m tcp --dport 443 -j ACCEPT
```

> **Important**
>
> When creating an **iptables** ruleset, order is important.
> If a rule specifies that any packets from the 192.168.100.0/24 subnet be dropped, and this is followed by a rule that allows packets from 192.168.100.13 (which is within the dropped subnet), then the second rule is ignored.
> The rule to allow packets from 192.168.100.13 must precede the rule that drops the remainder of the subnet.
> To insert a rule in a specific location in an existing chain, use the **-I** option. For example:
>
> ```
> [root@myServer ~ ] # iptables -I INPUT 1 -i lo -p all -j ACCEPT
> ```
>
> This rule is inserted as the first rule in the INPUT chain to allow local loopback device traffic.

There may be times when you require remote access to the LAN. Secure services, for example SSH, can be used for encrypted remote connection to LAN services.

Administrators with PPP-based resources (such as modem banks or bulk ISP accounts), dial-up access can be used to securely circumvent firewall barriers. Because they are direct connections, modem connections are typically behind a firewall/gateway.

For remote users with broadband connections, however, special cases can be made. You can configure **iptables** to accept connections from remote SSH clients. For example, the following rules allow remote SSH access:

```
[root@myServer ~ ] # iptables -A INPUT -p tcp --dport 22 -j ACCEPT
[root@myServer ~ ] # iptables -A OUTPUT -p tcp --sport 22 -j ACCEPT
```

These rules allow incoming and outbound access for an individual system, such as a single PC directly connected to the Internet or a firewall/gateway. However, they do not allow nodes behind the firewall/gateway to access these services. To allow LAN access to these services, you can use *Network Address Translation* (NAT) with **iptables** filtering rules.

18.5. FORWARD and NAT Rules

Most ISPs provide only a limited number of publicly routable IP addresses to the organizations they serve.

Administrators must, therefore, find alternative ways to share access to Internet services without giving public IP addresses to every node on the LAN. Using private IP addresses is the most common way of allowing all nodes on a LAN to properly access internal and external network services.

Edge routers (such as firewalls) can receive incoming transmissions from the Internet and route the packets to the intended LAN node. At the same time, firewalls/gateways can also route outgoing requests from a LAN node to the remote Internet service.

This forwarding of network traffic can become dangerous at times, especially with the availability of modern cracking tools that can spoof *internal* IP addresses and make the remote attacker's machine act as a node on your LAN.

To prevent this, **iptables** provides routing and forwarding policies that can be implemented to prevent abnormal usage of network resources.

The **FORWARD** chain allows an administrator to control where packets can be routed within a LAN. For example, to allow forwarding for the entire LAN (assuming the firewall/gateway is assigned an internal IP address on eth1), use the following rules:

```
[root@myServer ~ ] # iptables -A FORWARD -i eth1 -j ACCEPT
[root@myServer ~ ] # iptables -A FORWARD -o eth1 -j ACCEPT
```

This rule gives systems behind the firewall/gateway access to the internal network. The gateway routes packets from one LAN node to its intended destination node, passing all packets through its **eth1** device.

> **Note**
>
> By default, the IPv4 policy in Red Hat Enterprise Linux kernels disables support for IP forwarding. This prevents machines that run Red Hat Enterprise Linux from functioning as dedicated edge routers. To enable IP forwarding, use the following command:
>
> ```
> [root@myServer ~ ] # sysctl -w net.ipv4.ip_forward=1
> ```
>
> This configuration change is only valid for the current session; it does not persist beyond a reboot or network service restart. To permanently set IP forwarding, edit the **/etc/sysctl.conf** file as follows:
> Locate the following line:
>
> ```
> net.ipv4.ip_forward = 0
> ```
>
> Edit it to read as follows:
>
> ```
> net.ipv4.ip_forward = 1
> ```
>
> Use the following command to enable the change to the **sysctl.conf** file:
>
> ```
> [root@myServer ~ ] # sysctl -p /etc/sysctl.conf
> ```

18.5.1. Postrouting and IP Masquerading

Accepting forwarded packets via the firewall's internal IP device allows LAN nodes to communicate with each other; however they still cannot communicate externally to the Internet.

To allow LAN nodes with private IP addresses to communicate with external public networks, configure the firewall for *IP masquerading*, which masks requests from LAN nodes with the IP address of the firewall's external device (in this case, eth0):

```
[root@myServer ~ ] # iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

This rule uses the NAT packet matching table (**-t nat**) and specifies the built-in POSTROUTING chain for NAT (**-A POSTROUTING**) on the firewall's external networking device (**-o eth0**).

POSTROUTING allows packets to be altered as they are leaving the firewall's external device.

The **-j MASQUERADE** target is specified to mask the private IP address of a node with the external IP address of the firewall/gateway.

18.5.2. Prerouting

If you have a server on your internal network that you want make available externally, you can use the **-j DNAT** target of the PREROUTING chain in NAT to specify a destination IP address and port where incoming packets requesting a connection to your internal service can be forwarded.

For example, if you want to forward incoming HTTP requests to your dedicated Apache HTTP Server at 172.31.0.23, use the following command:

```
[root@myServer ~ ] # iptables -t nat -A PREROUTING -i eth0 -p tcp --dport 80 -j
DNAT --to 172.31.0.23:80
```

This rule specifies that the nat table use the built-in PREROUTING chain to forward incoming HTTP requests exclusively to the listed destination IP address of 172.31.0.23.

> **Note**
>
> If you have a default policy of DROP in your FORWARD chain, you must append a rule to forward all incoming HTTP requests so that destination NAT routing is possible. To do this, use the following command:
>
> ```
> [root@myServer ~ ] # iptables -A FORWARD -i eth0 -p tcp --dport 80 -d
> 172.31.0.23 -j ACCEPT
> ```
>
> This rule forwards all incoming HTTP requests from the firewall to the intended destination; the Apache HTTP Server behind the firewall.

18.5.3. DMZs and IPTables

You can create **iptables** rules to route traffic to certain machines, such as a dedicated HTTP or FTP server, in a *demilitarized zone* (DMZ). A DMZ is a special local subnetwork dedicated to providing services on a public carrier, such as the Internet.

For example, to set a rule for routing incoming HTTP requests to a dedicated HTTP server at 10.0.4.2 (outside of the 192.168.1.0/24 range of the LAN), NAT uses the **PREROUTING** table to forward the packets to the appropriate destination:

```
[root@myServer ~ ] # iptables -t nat -A PREROUTING -i eth0 -p tcp --dport 80 -j
DNAT --to-destination 10.0.4.2:80
```

With this command, all HTTP connections to port 80 from outside of the LAN are routed to the HTTP server on a network separate from the rest of the internal network. This form of network segmentation can prove safer than allowing HTTP connections to a machine on the network.

If the HTTP server is configured to accept secure connections, then port 443 must be forwarded as well.

18.6. Malicious Software and Spoofed IP Addresses

More elaborate rules can be created that control access to specific subnets, or even specific nodes, within a LAN. You can also restrict certain dubious applications or programs such as trojans, worms, and other client/server viruses from contacting their server.

For example, some trojans scan networks for services on ports from 31337 to 31340 (called the *elite* ports in cracking terminology).

Since there are no legitimate services that communicate via these non-standard ports, blocking them can effectively diminish the chances that potentially infected nodes on your network independently communicate with their remote master servers.

The following rules drop all TCP traffic that attempts to use port 31337:

```
[root@myServer ~ ] # iptables -A OUTPUT -o eth0 -p tcp --dport 31337 --sport
31337 -j DROP
[root@myServer ~ ] # iptables -A FORWARD -o eth0 -p tcp --dport 31337 --sport
31337 -j DROP
```

You can also block outside connections that attempt to spoof private IP address ranges to infiltrate your LAN.

For example, if your LAN uses the 192.168.1.0/24 range, you can design a rule that instructs the Internet-facing network device (for example, eth0) to drop any packets to that device with an address in your LAN IP range.

Because it is recommended to reject forwarded packets as a default policy, any other spoofed IP address to the external-facing device (eth0) is rejected automatically.

```
[root@myServer ~ ] # iptables -A FORWARD -s 192.168.1.0/24 -i eth0 -j DROP
```

> **Note**
>
> There is a distinction between the **DROP** and **REJECT** targets when dealing with *appended* rules. The **REJECT** target denies access and returns a `connection refused` error to users who attempt to connect to the service. The **DROP** target, as the name implies, drops the packet without any warning.
> Administrators can use their own discretion when using these targets. However, to avoid user confusion and attempts to continue connecting, the **REJECT** target is recommended.

18.7. IPTables and Connection Tracking

You can inspect and restrict connections to services based on their *connection state.* A module within **iptables** uses a method called *connection tracking* to store information about incoming connections. You can allow or deny access based on the following connection states:

- **NEW** — A packet requesting a new connection, such as an HTTP request.
- **ESTABLISHED** — A packet that is part of an existing connection.
- **RELATED** — A packet that is requesting a new connection but is part of an existing connection. For example, FTP uses port 21 to establish a connection, but data is transferred on a different port (typically port 20).
- **INVALID** — A packet that is not part of any connections in the connection tracking table.

You can use the stateful functionality of **iptables** connection tracking with any network protocol, even if the protocol itself is stateless (such as UDP). The following example shows a rule that uses connection tracking to forward only the packets that are associated with an established connection:

```
[root@myServer ~ ] # iptables -A FORWARD -m state --state ESTABLISHED,RELATED -j
ACCEPT
```

18.8. IPv6

The introduction of the next-generation Internet Protocol, called IPv6, expands beyond the 32-bit address limit of IPv4 (or IP). IPv6 supports 128-bit addresses, and carrier networks that are IPv6 aware

are therefore able to address a larger number of routable addresses than IPv4.

Red Hat Enterprise Linux supports IPv6 firewall rules using the Netfilter 6 subsystem and the **ip6tables** command. In Red Hat Enterprise Linux 5, both IPv4 and IPv6 services are enabled by default.

The **ip6tables** command syntax is identical to **iptables** in every aspect except that it supports 128-bit addresses. For example, use the following command to enable SSH connections on an IPv6-aware network server:

```
[root@myServer ~ ] # ip6tables -A INPUT -i eth0 -p tcp -s 3ffe:ffff:100::1/128 --
dport 22 -j ACCEPT
```

For more information about IPv6 networking, refer to the IPv6 Information Page at http://www.ipv6.org/.

18.9. Additional Resources

There are several aspects to firewalls and the Linux Netfilter subsystem that could not be covered in this chapter. For more information, refer to the following resources.

18.9.1. Installed Documentation

- The **iptables** man page contains a brief summary of the various options.

18.9.2. Useful Websites

- http://www.netfilter.org/ — The official homepage of the Netfilter and **iptables** project.
- http://www.tldp.org/ — The Linux Documentation Project contains several useful guides relating to firewall creation and administration.
- http://www.iana.org/assignments/port-numbers — The official list of registered and common service ports as assigned by the Internet Assigned Numbers Authority.

18.9.3. Related Documentation

- *Red Hat Linux Firewalls*, by Bill McCarty; Red Hat Press — a comprehensive reference to building network and server firewalls using open source packet filtering technology such as Netfilter and **iptables**. It includes topics that cover analyzing firewall logs, developing firewall rules, and customizing your firewall using various graphical tools.
- *Linux Firewalls*, by Robert Ziegler; New Riders Press — contains a wealth of information on building firewalls using both 2.2 kernel **ipchains** as well as Netfilter and **iptables**. Additional security topics such as remote access issues and intrusion detection systems are also covered.

Chapter 19. Controlling Access to Services

Maintaining security on your system is extremely important, and one approach for this task is to manage access to system services carefully. 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 to minimize your exposure to possible bug exploits.

There are several different methods for managing access to system services. Decide which method of management 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 turn it off. Both the services managed by **xinetd** and the services in the **/etc/rc.d/init.d** hierarchy (also known as SysV services) 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 services to be started, stopped, and restarted.
- **ntsysv** — a text-based application that allows you to configure which services are started at boot time for each runlevel. Non-**xinetd** 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. 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 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. Refer to the *Reference Guide* and the *Security Guide* for more information about **iptables**.

Alternatively, if you are looking for a utility to set general access rules for your home machine, and/or if you are new to Linux, try the **Security Level Configuration Tool** (**system-config-securitylevel**), which allows you to select the security level for your system, similar to the **Firewall Configuration** screen in the installation program.

If you need more specific firewall rules, refer to the **iptables** chapter in the *Reference Guide*.

19.1. 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.

The following runlevels exist:

- 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 use a text login screen, you are operating in runlevel 3. If you use 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:5:initdefault:
```

Change the number in this line to the desired runlevel. The change does 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. The **telinit** command does not change the **/etc/inittab** file; it only changes the runlevel currently running. When the system is rebooted, it continues to boot the runlevel as specified in **/etc/inittab**.

19.2. 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. Refer to the *Reference Guide* and **hosts_access** in section 5 of the man pages (**man 5 hosts_access**) for details.

19.2.1. **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, and more

xinetd runs constantly and listens on all 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 the file only contains a few defaults and an instruction to include the **/etc/xinetd.d** directory. To enable or disable an **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 the **Services Configuration Tool**, **ntsysv**, or **chkconfig**. For a list of network services controlled by **xinetd**, review the contents of the **/etc/xinetd.d** directory with the command **ls /etc/xinetd.d**.

19.3. Services Configuration Tool

The **Services Configuration Tool** is a graphical application developed by Red Hat, Inc to configure which SysV services in the **/etc/rc.d/init.d** directory 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 the **Services Configuration Tool** from the desktop, go to the **Main Menu Button** (on the Panel) => **System Settings** => **Server Settings** => **Services** or type the command **system-config-services** at a shell prompt (for example, in an **XTerm** or a **GNOME terminal**).
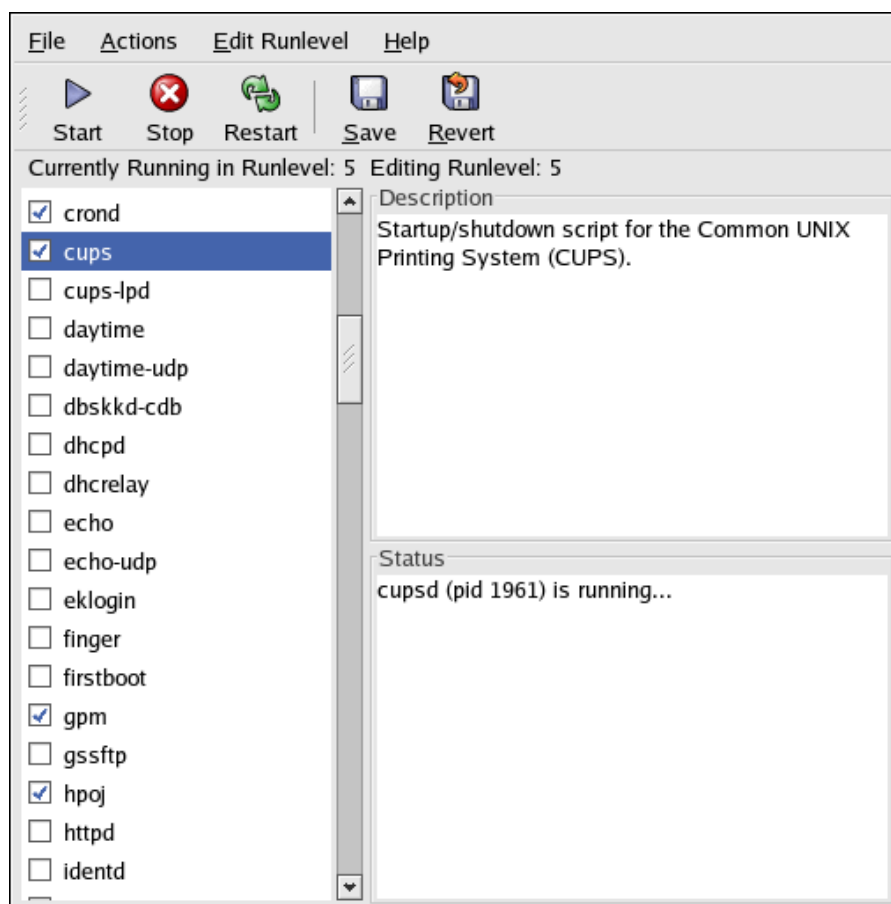
Figure 19.1. Services Configuration Tool

The **Services Configuration Tool** displays the current runlevel as well as the 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 19.1, "Runlevels" for a description of runlevels.

The **Services Configuration Tool** lists the services from the **/etc/rc.d/init.d** directory as well as the services controlled by **xinetd**. Click on the name of the service from the list on the left-hand side of the application to display a brief description of that service as well as the status of the service. If the service is not an **xinetd** service, the status window shows whether the service is currently running. If the service is controlled by **xinetd**, the status window displays the phrase **xinetd service**.

To start, stop, or restart a service immediately, select the service from the list and click the appropriate button on the toolbar (or choose the action from the **Actions** pulldown menu). If the service is an **xinetd** service, the action buttons are disabled because they can not be started or stopped individually.

If you enable/disable an **xinetd** service by checking or unchecking the checkbox next to the service name, you must select **File** => **Save Changes** from the pulldown menu to restart **xinetd** and immediately enable/disable the **xinetd** service that you changed. **xinetd** is also configured to remember the setting. You can enable/disable multiple **xinetd** services at a time and save the changes when you are finished.

For example, assume you check **rsync** to enable it in runlevel 3 and then save the changes. The **rsync** service is immediately enabled. The next time **xinetd** is started, **rsync** is still enabled.

> ⚠️ **Warning**
>
> When you save changes to **xinetd** services, **xinetd** is restarted, and the changes take place immediately. When you save changes to other services, the runlevel is reconfigured, but the changes do not take effect immediately.

To enable a non-**xinetd** service to start at boot time for the currently selected runlevel, check the checkbox beside the name of the service in the list. After configuring the runlevel, apply the changes by selecting **File** => **Save Changes** from the pulldown menu. The runlevel configuration is changed, but the runlevel is not restarted; thus, the changes do not take place immediately.

For example, assume you are configuring runlevel 3. If you change the value for the **httpd** service from checked to unchecked and then select **Save Changes**, the runlevel 3 configuration changes so that **httpd** is not started at boot time. However, runlevel 3 is not reinitialized, so **httpd** is still running. Select one of following options at this point:

1. Stop the **httpd** service — Stop the service by selecting it from the list and clicking the **Stop** button. A message appears stating that the service was stopped successfully.
2. Reinitialize the runlevel — Reinitialize the runlevel by going to a shell prompt 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 multiple services and want to activate the changes immediately.
3. Do nothing else — You do not have to stop the **httpd** service. You can wait until the system is rebooted for the service to stop. The next time the system is booted, the runlevel is initialized without the **httpd** service running.

To add a service to a runlevel, select the runlevel from the **Edit Runlevel** pulldown menu, and then select **Actions** => **Add Service**. To delete a service from a runlevel, select the runlevel from the **Edit Runlevel** pulldown menu, select the service to be deleted from the list on the left, and select **Actions** => **Delete Service**.

19.4. 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 configure runlevels. By default, only the current runlevel is configured. To configure a different runlevel, specify one or more runlevels with the **--level** option. For example, the command **ntsysv --level 345** configures runlevels 3, 4, and 5.

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. A ***** signifies that a service is set to on. Pressing the **F1** key displays a short description of the selected service.

> ⚠️ **Warning**
>
> Services managed by **xinetd** are immediately affected by **ntsysv**. For all other services, changes do not take effect immediately. 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.

19.5. `chkconfig`

The **chkconfig** command can also be used to activate and deactivate services. The **chkconfig --list** command displays a list of system services and whether they are started (**on**) or stopped (**off**) in runlevels 0-6. At the end of the list is a section for the services managed by **xinetd**.

If the **chkconfig --list** command is used to query a service managed by **xinetd**, it displays whether the **xinetd** service is enabled (**on**) or disabled (**off**). For example, the command **chkconfig --list finger** returns the following output:

```
finger          on
```

As shown, **finger** is enabled as an **xinetd** service. If **xinetd** is running, **finger** is enabled.

If you use **chkconfig --list** to query a service in **/etc/rc.d**, service's settings for each runlevel are displayed. For example, the command **chkconfig --list httpd** returns the following output:

```
httpd           0:off   1:off   2:on    3:on    4:on    5:on    6:off
```

chkconfig can also be used to configure 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 following command:

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

> ⚠️ **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.

19.6. Additional Resources

For more information, refer to the following resources.

19.6.1. Installed Documentation

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

19.6.2. Useful Websites

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

19.6.3. Related Books

- *Reference Guide* , Red Hat, Inc — This companion manual contains detailed information about how TCP wrappers and **xinetd** allow or deny access as well as how to configure network access using them. It also provides instructions for creating **iptables** firewall rules.
- *Security Guide* Red Hat, Inc — This manual discusses securing services with TCP wrappers and **xinetd** such as logging denied connection attempts.

Chapter 20. OpenSSH

OpenSSH is a free, open source implementation of the SSH (*S* ecure *SH* ell) protocols. It replaces **telnet**, **ftp**, **rlogin**, **rsh**, and **rcp** with secure, encrypted network connectivity tools. OpenSSH supports versions 1.3, 1.5, and 2 of the SSH protocol. Since OpenSSH version 2.9, the default protocol is version 2, which uses RSA keys as the default.

20.1. Why Use OpenSSH?

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

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

20.2. Configuring an OpenSSH Server

To run an OpenSSH server, you must first make sure that you have the proper RPM packages installed. The **openssh-server** package is required and depends on the **openssh** package.

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

To start the OpenSSH service, use the command **/sbin/service sshd start**. To stop the OpenSSH server, use the command **/sbin/service sshd stop**. If you want the daemon to start automatically at boot time, refer to Chapter 19, *Controlling Access to Services* for information on how to manage services.

If you reinstall, the reinstalled system creates a new set of identification keys. Any clients who had connected to the system with any of the OpenSSH tools before the reinstall will see the following message:

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@    WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!    @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that the RSA host key has just been changed.
```

If you want to keep the host keys generated for the system, backup the **/etc/ssh/ssh_host*key*** files and restore them after the reinstall. This process retains the system's identity, and when clients try to connect to the system after the reinstall, they will not receive the warning message.

20.3. 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.

20.3.1. 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 (**~/.ssh/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 *hostnamecommand***. 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 the remote directory **/usr/share/doc** will be displayed, and you will return to your local shell prompt.

20.3.2. 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:<remotefile>
```

The **<localfile>** specifies the source including path to the file, such as **/var/log/maillog**. The **<remotefile>** specifies the destination, which can be a new filename such as **/tmp/hostname-maillog**. For the remote system, if you do not have a preceding **/**, the path will be relative to the home directory of **username**, typically **/home/username/**.

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

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

This will transfer the local file **shadowman** to **/home/username/shadowman** on penguin.example.net. Alternately, you can leave off the final **shadowman** in the **scp** command.

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 including path, and **<newlocalfile>** specifies the destination including path.

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/
```

20.3.3. 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 those used by 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.

20.3.4. 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 Enterprise Linux 4 uses SSH Protocol 2 and RSA keys by default.

> **Tip**
>
> If you reinstall and 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.

20.3.4.1. Generating an RSA Key Pair for Version 2

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

1. To generate an 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 the `.ssh` directory using the following command:

   ```
   chmod 755 ~/.ssh
   ```

3. Copy the contents of **~/.ssh/id_rsa.pub** into the file **~/.ssh/authorized_keys** on the machine to which you want to connect. If the file **~/.ssh/authorized_keys** exist, append the contents of the file **~/.ssh/id_rsa.pub** to the file **~/.ssh/authorized_keys** on the other machine.

4. Change the permissions of the **authorized_keys** file using the following command:

   ```
   chmod 644 ~/.ssh/authorized_keys
   ```

5. If you are running GNOME, skip to Section 20.3.4.4, "Configuring **ssh-agent** with GNOME". If you are not running the X Window System, skip to Section 20.3.4.5, "Configuring **ssh-agent**".

20.3.4.2. 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 a single 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 the **.ssh** directory with the following command:

```
chmod 755 ~/.ssh
```

3. Copy the contents of **~/.ssh/id_dsa.pub** into the file **~/.ssh/authorized_keys** on the
machine to which you want to connect. If the file **~/.ssh/authorized_keys** exist, append the
contents of the file **~/.ssh/id_dsa.pub** to the file **~/.ssh/authorized_keys** on the other
machine.

4. Change the permissions of the **authorized_keys** file using the following command:

```
chmod 644 ~/.ssh/authorized_keys
```

5. If you are running GNOME, skip to Section 20.3.4.4, "Configuring **ssh-agent** with GNOME". If you
are not running the X Window System, skip to Section 20.3.4.5, "Configuring **ssh-agent**".

20.3.4.3. 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 systems that use DSA, 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** into 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 20.3.4.4, "Configuring **ssh-agent** with GNOME". If you
are not running GNOME, skip to Section 20.3.4.5, "Configuring **ssh-agent**".

20.3.4.4. 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** package contains the application 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 20.3.4.5, "Configuring **ssh-agent**".

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 Enterprise Linux CD-ROM set, from a Red Hat FTP mirror site, or using Red Hat Network.

2. Select **Main Menu Button** (on the Panel) => **Preferences** => **More 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.

3. 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**.

20.3.4.5. 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 (refer to Section 20.3.4.4, "Configuring **ssh-agent** with GNOME"), 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
```

2. 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.

3. 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.

20.4. Additional Resources

The OpenSSH and OpenSSL projects are in constant development, and the most up-to-date information for them is available from their websites. The man pages for OpenSSH and OpenSSL tools are also good sources of detailed information.

20.4.1. 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.

20.4.2. 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.

20.4.3. Related Books

⁕ *Reference Guide* — Learn the event sequence of an SSH connection, review a list of configuration files, and discover how SSH can be used for X forwarding.

Chapter 21. 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 Enterprise 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.

21.1. 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.

21.2. 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 the local machine (**/misc/local** in the above example) must exist before this command can be executed.

In this command, **shadowman.example.com** is the hostname of the NFS file server, **/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**.

21.2.1. 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,wsize=8192,timeo=14,intr
```

The mount point **/pub** must exist on the client machine before this command can be executed. 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** is mounted from the server.

21.2.2. Mounting NFS File Systems using autofs

A third option for mounting an NFS share is the use of the autofs service. 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 **/proj52** on the remote machine penguin.example.net at the mount point **/misc/myproject** on your machine, add the following line to **auto.master**:

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

Next, add the following line to **/etc/auto.misc**:

```
myproject   -rw,soft,intr,rsize=8192,wsize=8192 penguin.example.net:/proj52
```

The first field in **/etc/auto.misc** is the name of the **/misc** subdirectory. This subdirectory 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.

To start the autofs service, at a shell prompt, type the following command:

```
/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, and for information on managing services, refer to Chapter 19, *Controlling Access to Services*.

21.2.3. Using TCP

The default transport protocol for NFSv4 is TCP; however, the Red Hat Enterprise Linux 4 kernel includes support for NFS over UDP. To use NFS over UDP, include the **-o udp** option to **mount** when mounting the NFS-exported file system on the client system.

There are three ways to configure an NFS file system export. On demand via the command line (client side), automatically via the **/etc/fstab** file (client side), and automatically via autofs configuration files, such as **/etc/auto.master** and **/etc/auto.misc** (server side with NIS).

For example, on demand via the command line (client side):

```
mount -o udp shadowman.example.com:/misc/export /misc/local
```

When the NFS mount is specified in **/etc/fstab** (client side):

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

When the NFS mount is specified in an autofs configuration file for a NIS server, available for NIS enabled workstations:

```
myproject  -rw,soft,intr,rsize=8192,wsize=8192,udp penguin.example.net:/proj52
```

Since the default is TCP, if the **-o udp** option is not specified, the NFS-exported file system is accessed via TCP.

The advantages of using TCP include the following:

» Improved connection durability, thus less **NFS stale file handles** messages.
» Performance gain on heavily loaded networks because TCP acknowledges every packet, unlike UDP which only acknowledges completion.
» TCP has better congestion control than UDP (which has none). On a very congested network, UDP packets are the first packets that are dropped. This means that if NFS is writing data (in 8K chunks) all of that 8K must be retransmitted over UDP. Because of TCP's reliability, only parts of that 8K data are transmitted at a time.
» Error detection. When a TCP connection breaks (due to the server being unavailable) the client stops sending data and restarts the connection process once the server becomes available. With UDP, since it's connection-less, the client continues to pound the network with data until the server reestablishes a connection.

The main disadvantage is that there is a very small performance hit due to the overhead associated with the TCP protocol.

21.2.4. Preserving ACLs

The Red Hat Enterprise Linux 4 kernel provides ACL support for the ext3 file system and ext3 file systems mounted with the NFS or Samba protocols. Thus, if an ext3 file system has ACLs enabled for it and is NFS exported, and if the NFS client can read ACLs, they are used by the NFS client as well.

For more information about mounting NFS file systems with ACLs, refer to Chapter 14, *Access Control Lists*.

21.3. Exporting NFS File Systems

Sharing or serving 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, have root privileges, and have the **system-config-nfs** RPM package installed. To start the application, select

the **Main Menu Button** (on the Panel) => **System Settings** => **Server Settings** => **NFS**, or type the command **system-config-nfs**.
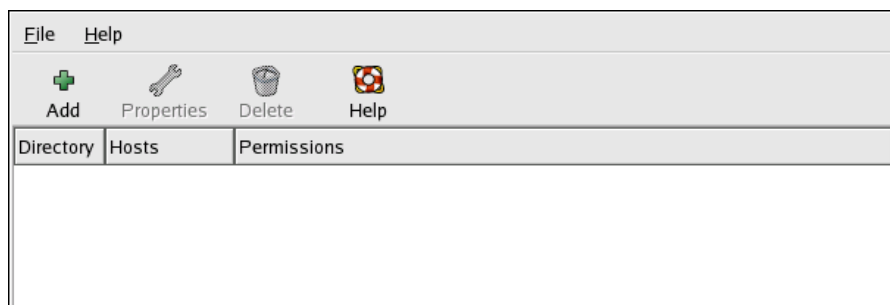
Figure 21.1. NFS Server Configuration Tool

To add an NFS share, click the **Add** button. The dialog box shown in Figure 21.2, "Add Share" appears.

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

- **Directory** — Specify the directory to share, such as **/tmp**.
- **Host(s)** — Specify the host(s) with which to share the directory. Refer to Section 21.3.2, "Hostname Formats" for an explanation of possible formats.
- **Basic permissions** — Specify whether the directory should have read-only or read/write permissions.
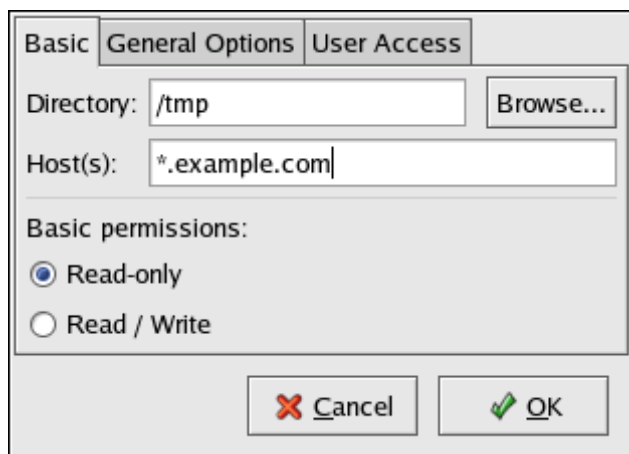
Figure 21.2. Add Share

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 clicking **OK** to add, edit, or delete an NFS share from the list, the changes take place immediately — the server daemon is restarted and the old configuration file is saved as **/etc/exports.bak**. The new configuration is written to **/etc/exports**.

The **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).

21.3.1. Command Line Configuration

If you prefer editing configuration files using a text editor or if you do not 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:

```
directoryhostname(options)
```

The only option that needs to be specified is one of **sync** or **async** (**sync** is recommended). If **sync** is specified, the server does not reply to requests before the changes made by the request are written to the disk.

For example,

```
/misc/export speedy.example.com(sync)
```

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

```
/misc/export speedy.example.com(rw,sync)
```

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

Refer to Section 21.3.2, "Hostname Formats" for an explanation of possible hostname formats.

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

> ⚠️ **Warning**
>
> 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.example.com(rw,sync)
> /misc/export speedy.example.com (rw,sync)
> ```
>
> The first line grants users from **speedy.example.com** read-write access and denies all other users. The second line grants users from **speedy.example.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
```

21.3.2. 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 wildcards — Use the * or ? character to specify a string match. Wildcards are not to be used with IP addresses; however, they may accidentally work if reverse DNS lookups fail. When specifying wildcards in fully qualified domain names, dots (.) are not included in the wildcard. 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.

21.3.3. 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 19, *Controlling Access to Services* for details.

21.4. Additional Resources

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

21.4.1. 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.

21.4.2. Useful Websites

- http://nfs.sourceforge.net/ — the NFS webpage, includes links to the mailing lists and FAQs.
- http://www.tldp.org/HOWTO/NFS-HOWTO/index.html — The *Linux NFS-HOWTO* from the Linux Documentation Project.

21.4.3. Related Books

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

Chapter 22. Samba

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

The Red Hat Enterprise Linux 4 kernel contains *Access Control List* (ACL) support for ext3 file systems. If the Samba server shares an ext3 file system with ACLs enabled for it, and the kernel on the client system contains support for reading ACLs from ext3 file systems, the client automatically recognizes and uses the ACLs. Refer to Chapter 14, *Access Control Lists* for more information on ACLs.

22.1. Why Use Samba?

Samba is useful if you have a network of both Windows and Linux machines. Samba allows files and printers to be shared by all the systems in a network. To share files between Linux machines only, use NFS as discussed in Chapter 21, *Network File System (NFS)*. To share printers between Linux machines only, you do not need to use Samba; refer to Chapter 33, *Printer Configuration*.

22.2. Configuring a Samba Server

The default configuration file (`/etc/samba/smb.conf`) allows users to view their home directories as a Samba share. It also shares all printers configured for the system as Samba shared printers. In other words, you can attach a printer to the system and print to it from the Windows machines on your network.

22.2.1. Graphical Configuration

To configure Samba using a graphical interface, use the **Samba Server Configuration Tool**. For command line configuration, skip to Section 22.2.2, "Command Line Configuration".

The **Samba Server Configuration Tool** is a graphical interface for managing Samba shares, users, and basic server settings. It modifies the configuration files in the `/etc/samba/` directory. Any changes to these files not made using the application are preserved.

To use this application, you must be running the X Window System, have root privileges, and have the `system-config-samba` RPM package installed. To start the **Samba Server Configuration Tool** from the desktop, go to the **Main Menu Button** (on the Panel) => **System Settings** => **Server Settings** => **Samba** or type the command `system-config-samba` at a shell prompt (for example, in an XTerm or a GNOME terminal).
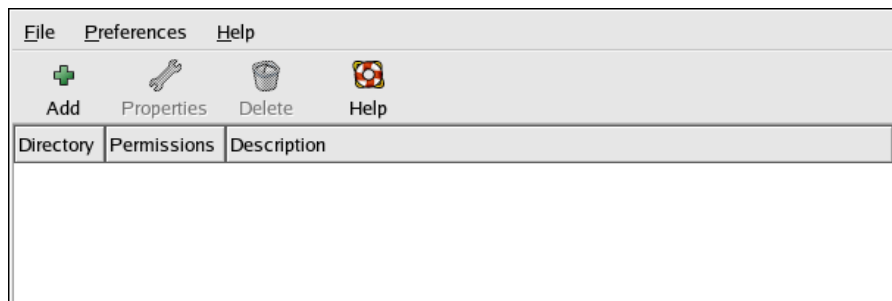
Figure 22.1. Samba Server Configuration Tool

> **Note**
>
> The **Samba Server Configuration Tool** does not display shared printers or the default stanza that allows users to view their own home directories on the Samba server.

22.2.1.1. Configuring Server Settings

The first step in configuring a Samba server is to configure the basic settings for the server and a few security options. After starting the application, select **Preferences** => **Server Settings** from the pulldown menu. The **Basic** tab is displayed as shown in Figure 22.2, "Configuring Basic Server Settings".

Figure 22.2. Configuring Basic Server Settings

On the **Basic** tab, specify which workgroup the computer should be in as well as a brief description of the computer. They correspond to the **workgroup** and **server string** options in **smb.conf**.

Figure 22.3. Configuring Security Server Settings

The **Security** tab contains the following options:

- **Authentication Mode** — This corresponds to the **security** option. Select one of the following types of authentication.
  - **ADS** — The Samba server acts as a domain member in an Active Directory Domain (ADS) realm.

For this option, Kerberos must be installed and configured on the server, and Samba must become a member of the ADS realm using the **net** utility, which is part of the **samba-client** package. Refer to the **net** man page for details. This option does not configure Samba to be an ADS Controller. Specify the realm of the Kerberos server in the **Kerberos Realm** field.

> **Note**
>
> The **Kerberos Realm** field must be supplied in all uppercase letters, such as **EXAMPLE.COM**.
> Use of your Samba server as a domain member in an ADS realm assumes proper configuration of Kerberos, including the **/etc/krb5.conf** file.

- **Domain** — The Samba server relies on a Windows NT Primary or Backup Domain Controller to verify the user. The server passes the username and password to the Controller and waits for it to return. Specify the NetBIOS name of the Primary or Backup Domain Controller in the **Authentication Server** field.

  The **Encrypted Passwords** option must be set to **Yes** if this is selected.

- **Server** — The Samba server tries to verify the username and password combination by passing them to another Samba server. If it can not, the server tries to verify using the user authentication mode. Specify the NetBIOS name of the other Samba server in the **Authentication Server** field.

- **Share** — Samba users do not have to enter a username and password combination on a per Samba server basis. They are not prompted for a username and password until they try to connect to a specific shared directory from a Samba server.

- **User** — (Default) Samba users must provide a valid username and password on a per Samba server basis. Select this option if you want the **Windows Username** option to work. Refer to Section 22.2.1.2, "Managing Samba Users" for details.

- **Encrypt Passwords** — This option must be enabled if the clients are connecting from a system with Windows 98, Windows NT 4.0 with Service Pack 3, or other more recent versions of Microsoft Windows. The passwords are transfered between the server and the client in an encrypted format instead of as a plain-text word that can be intercepted. This corresponds to the **encrypted passwords** option. Refer to Section 22.2.3, "Encrypted Passwords" for more information about encrypted Samba passwords.

- **Guest Account** — When users or guest users log into a Samba server, they must be mapped to a valid user on the server. Select one of the existing usernames on the system to be the guest Samba account. When guests log in to the Samba server, they have the same privileges as this user. This corresponds to the **guest account** option.

After clicking **OK**, the changes are written to the configuration file and the daemon is restart; thus, the changes take effect immediately.

22.2.1.2. Managing Samba Users

The **Samba Server Configuration Tool** requires that an existing user account be active on the system acting as the Samba server before a Samba user can be added. The Samba user is associated with the existing user account.
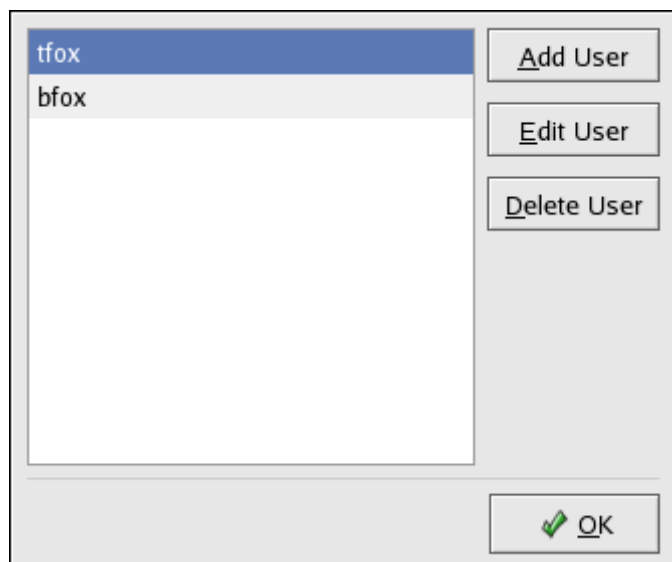
Figure 22.4. Managing Samba Users

To add a Samba user, select **Preferences** => **Samba Users** from the pulldown menu, and click the **Add User** button. In the **Create New Samba User** window select a **Unix Username** from the list of existing users on the local system.

If the user has a different username on a Windows machine and needs to log into the Samba server from the Windows machine, specify that Windows username in the **Windows Username** field. The **Authentication Mode** on the **Security** tab of the **Server Settings** preferences must be set to **User** for this option to work.

Also configure a **Samba Password** for the Samba User and confirm it by typing it again. Even if you select to use encrypted passwords for Samba, it is recommended that the Samba passwords for all users are different from their system passwords.

To edit an existing user, select the user from the list, and click **Edit User**. To delete an existing Samba user, select the user, and click the **Delete User** button. Deleting a Samba user does not delete the associated system user account.

The users are modified immediately after clicking the **OK** button.

22.2.1.3. Adding a Share

To create a Samba share, click the **Add** button from the main Samba configuration window.
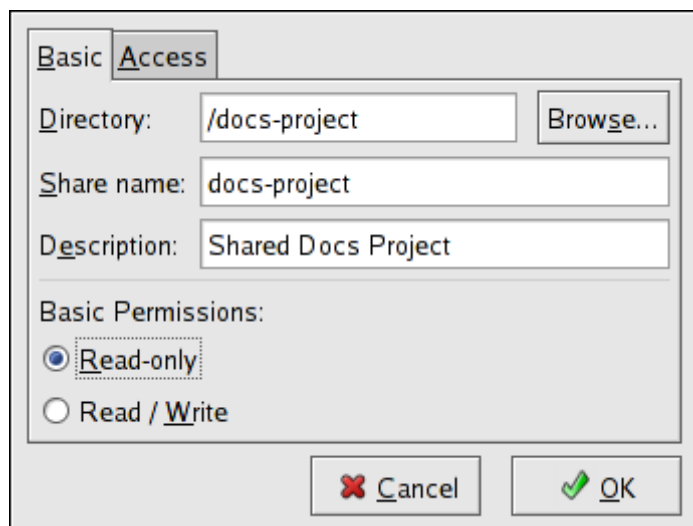
Figure 22.5. Adding a Share

The **Basic** tab configures the following options:

- **Directory** — The directory to share via Samba. The directory must exist before it can be entered here.
- **Share name** — The actual name of the share that is seen from remote machines. By default, it is the same value as **Directory**, but can be configured.
- **Descriptions** — A brief description of the share.
- **Basic Permissions** — Whether users should only be able to read the files in the shared directory or whether they should be able to read and write to the shared directory.

On the **Access** tab, select whether to allow only specified users to access the share or whether to allow all Samba users to access the share. If you select to allow access to specific users, select the users from the list of available Samba users.

The share is added immediately after clicking **OK**.

22.2.2. Command Line Configuration

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

To specify the Windows workgroup and a brief description of the Samba server, 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 is used as 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.

22.2.3. Encrypted Passwords

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 plain text passwords. However, Windows NT 4.0 with Service Pack 3 or higher, Windows 98, Windows 2000, Windows ME, and Windows XP require encrypted Samba passwords. To use Samba between a Linux system and a system running one of these Windows operating systems, 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 of your Windows machines — this is risky and may cause further conflicts. It is recommended that you use encrypted passwords for better security.

To configure Samba 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, and a Samba user account is not active until a password is set for it. For higher security, it is recommended that the user's Samba password be different from the user's system password. To set each Samba user's password, use the following command (replace *username* with each user's username):

   ```
   smbpasswd username
   ```

4. Encrypted passwords must be enabled. Since they are enabled by default, they do not have to be specifically enabled in the configuration file. However, they can not be disabled in the configuration file either. In the file **/etc/samba/smb.conf**, verify that the following line does not exist:

```
encrypt passwords = no
```

If it does exist but is commented out with a semi-colon (`;`) at the beginning of the line, then the line is ignored, and encrypted passwords are enabled. If this line exists but is not commented out, either remove it or comment it out.

To specifically enable encrypted passwords in the configuration file, add the following lines to **etc/samba/smb.conf**:

```
encrypt passwords = 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 the **Services Configuration Tool** to enable it at runtime. Refer to Chapter 19, *Controlling Access to Services* for details.

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 Enterprise 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
```

22.2.4. Starting and Stopping the Server

On the server that is sharing directories via Samba, the **smb** service must be running.

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

```
/sbin/service smb status
```

Start the daemon with the following command:

```
/sbin/service smb start
```

Stop the daemon with the following command:

```
/sbin/service smb stop
```

To start the **smb** service at boot time, use the command:

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

You can also use **chkconfig**, **ntsysv**, or the **Services Configuration Tool** to configure which services start at boot time. Refer to Chapter 19, *Controlling Access to Services* for details.

> **Tip**
>
> To view active connections to the system, execute the command `smbstatus`.

22.3. Connecting to a Samba Share

You can use **Nautilus** to view available Samba shares on your network. Select **Main Menu Button** (on the Panel) => **Network Servers** to view a list of Samba workgroups on your network. You can also type `smb:` in the **Location:** bar of Nautilus to view the workgroups.

As shown in Figure 22.6, "SMB Workgroups in Nautilus", an icon appears for each available SMB workgroup on the network.
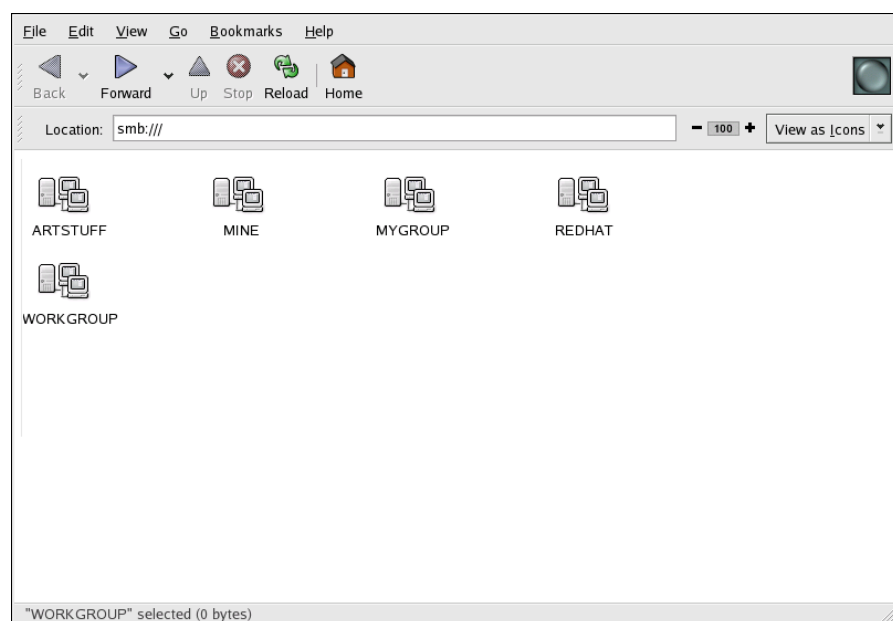
Figure 22.6. SMB Workgroups in Nautilus

Double-click one of the workgroup icons to view a list of computers within the workgroup.
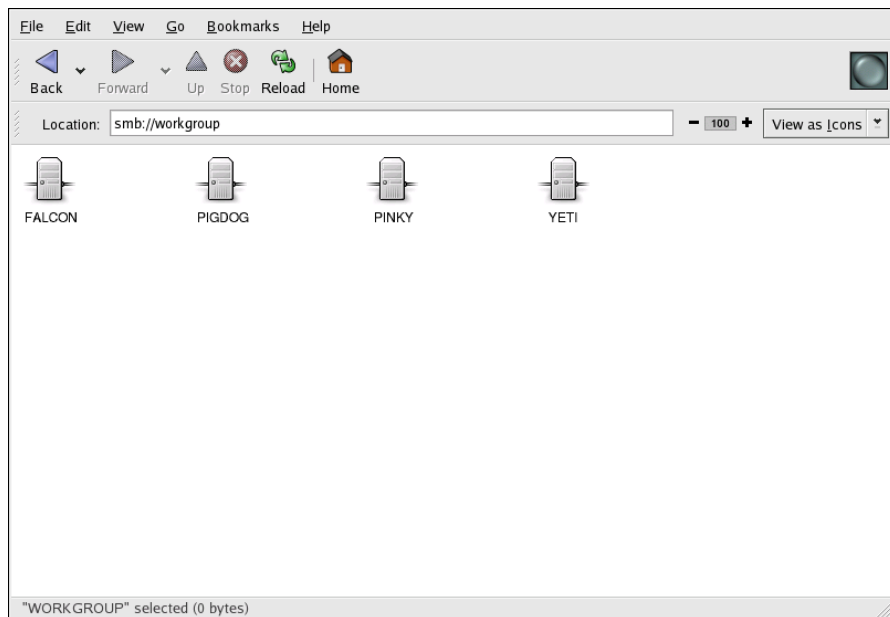
Figure 22.7. SMB Machines in Nautilus

As you can see from there is an icon for each machine within the workgroup. Double-click on an icon to view the Samba shares on the machine. If a username and password combination is required, you are prompted for them.

Alternately, you can also specify the Samba server and sharename in the **Location:** bar for **Nautilus** using the following syntax (replace **<servername>** and **<sharename>** with the appropriate values):

```
smb://<servername>/<sharename>/
```

22.3.1. Command Line

To query the network for Samba servers, use the **findsmb** command. For each server found, it displays its IP address, NetBIOS name, workgroup name, operating system, and SMB server version.

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

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

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.

22.3.2. Mounting the Share

Sometimes it is useful to mount a Samba share to a directory so that the files in the directory can be

treated as if they are part of the local file system.

To mount a Samba share to a directory, create the directory if it does not already exist, and execute the following command as root:

```
mount -t smbfs -o username=<username> //<servername>/<sharename>/mnt/point/
```

This command mounts **<sharename>** from **<servername>** in the local directory **/mnt/point/**.

22.4. Additional Resources

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

22.4.1. Installed Documentation

- **smb.conf** man page — explains how to configure the Samba configuration file
- **smbd** man page — describes how the Samba daemon works
- **smbclient** and **findsmb** man pages — learn more about these client tools
- **/usr/share/doc/samba-<version-number>/docs/** — help files included with the **samba** package

22.4.2. Useful Websites

- http://www.samba.org/ — The Samba webpage contains useful documentation, information about mailing lists, and a list of GUI interfaces.
- http://www.samba.org/samba/docs/using_samba/toc.html — an online version of *Using Samba, 2nd Edition* by Jay Ts, Robert Eckstein, and David Collier-Brown; O'Reilly &Associates

Chapter 23. Dynamic Host Configuration Protocol (DHCP)

Dynamic Host Configuration Protocol (DHCP) is a 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 the IP address, gateway, and DNS servers.

23.1. Why Use DHCP?

DHCP is useful for automatic configuration of client network interfaces. When configuring the client system, the administrator can choose DHCP and instead of entering 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 addresses. 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 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.

23.2. Configuring a DHCP Server

To configure a DHCP server, the **/etc/dhcpd.conf** configuration file must be created. A sample file can be found at **/usr/share/doc/dhcp-<*version*>/dhcpd.conf.sample**.

DHCP also uses the file **/var/lib/dhcp/dhcpd.leases** to store the client lease database. Refer to Section 23.2.2, "Lease Database" for more information.

23.2.1. 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, while other options can be declared for individual client systems.

The configuration file can contain extra tabs or blank lines for easier formatting. 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 Internet Engineering Task Force (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 3.0b2pl11 and previous versions used the ad-hoc mode; however, it has been deprecated. 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;
```

Refer to 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 the configuration file is changed, the changes do not take effect until the DHCP daemon is restarted with the command **service dhcpd restart**.

> **Tip**
>
> Instead of changing a DHCP configuration file and restarting the service each time, using the **omshell** command provides an interactive way to connect to, query, and change the configuration of a DHCP server. By using **omshell**, all changes can be made while the server is running. For more information on **omshell**, refer to the **omshell** man page.

In Example 23.1, "Subnet Declaration", the **routers**, **subnet-mask**, **domain-name**, **domain-name-servers**, and **time-offset** options are used for any **host** statements declared below it.

Additionally, a **subnet** can be declared, a **subnet** declaration must be included for every subnet in the network. If it is not, the DHCP server fails 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**.

Example 23.1. Subnet Declaration

```
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;
}
```

All subnets that share the same physical network should be declared within a **shared-network** declaration as shown in Example 23.2, "Shared-network Declaration". Parameters within the **shared-**

network, but outside the enclosed **subnet** declarations, are considered to be global parameters. The name of the **shared-network** should be a descriptive title for the network, such as using the title 'test-lab' to describe all the subnets in a test lab environment.

Example 23.2. Shared-network Declaration

```
shared-network name {
    option domain-name                  "test.redhat.com";
    option domain-name-servers      ns1.redhat.com, ns2.redhat.com;
    option routers                      192.168.0.254;
    more parameters for EXAMPLE shared-network
    subnet 192.168.1.0 netmask 255.255.252.0 {
        parameters for subnet
        range 192.168.1.1 192.168.1.254;
    }
    subnet 192.168.2.0 netmask 255.255.252.0 {
        parameters for subnet
        range 192.168.2.1 192.168.2.254;
    }
}
```

As demonstrated in Example 23.3, "Group Declaration", the **group** declaration can be used to apply global parameters to a group of declarations. For example, shared networks, subnets, and hosts can be grouped.

Example 23.3. Group Declaration

```
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;
    }
}
```

To configure a DHCP server that leases a dynamic IP address to a system within a subnet, modify Example 23.4, "Range Parameter" 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.

Example 23.4. Range Parameter

```
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;
}
```

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 23.5, "Static IP Address using DHCP", 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.

Note that the optional parameter **host-name** can also be used to assign a host name to the client.

Example 23.5. Static IP Address using DHCP

```
host apex {
    option host-name "apex.example.com";
    hardware ethernet 00:A0:78:8E:9E:AA;
    fixed-address 192.168.1.4;
}
```

> **Tip**
>
> The sample configuration file provided can be used as a starting point and custom configuration options can be added to it. To copy it to the proper location, use the following command:
>
> ```
>  cp /usr/share/doc/dhcp-<version-number>/dhcpd.conf.sample
> /etc/dhcpd.conf
> ```
>
> (where **<version-number>** is the DHCP version number).

For a complete list of option statements and what they do, refer to the **dhcp-options** man page.

23.2.2. 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, the **dhcpd.leases** file does not exist, but it is required to start the service. Do not create a new lease file. If you do, all old leases are lost which causes many problems. The correct solution is to rename the **dhcpd.leases~** backup file to **dhcpd.leases** and then start the daemon.

23.2.3. Starting and Stopping the Server

> **Important**
>
> When the DHCP server is started for the first time, it fails unless the **dhcpd.leases** file exists. Use the command **touch /var/lib/dhcp/dhcpd.leases** to create the file if it does not exist.
> If the same server is also running BIND as a DNS server, this step is not necessary, as starting the **named** service automatically checks for a **dhcpd.leases** file.

To start the DHCP service, use the command **/sbin/service dhcpd start**. To stop the DHCP server, use the command **/sbin/service dhcpd stop**.

By default, the DHCP service does not start at boot time. To configure the daemon to start automatically at boot time, refer to Chapter 19, *Controlling Access to Services* for information on how to manage services.

If more than one network interface is attached to the system, but the DHCP server should only be started on one of the interfaces, 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 for 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 the default port 67 is used, the server listens on port 67 for requests and responses to the client on port 68. If a port is specified here and the DHCP relay agent is used, the same port on which the DHCP relay agent should listen must be specified. Refer to Section 23.2.4, "DHCP Relay Agent" 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.

23.2.4. DHCP Relay Agent

The DHCP Relay Agent (`dhcrelay`) allows for the relay of 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`.

23.3. 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 a card is added after installation, **Kudzu** [6] should recognize it and prompt for the configuration of the corresponding kernel module for it. Be sure to check the 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 37, Kernel Modules for details on loading kernel modules.

To configure a DHCP client manually, 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
```

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
```

A configuration file is needed for each device to be configured to use DHCP.

Other options for the network script include:

- `DHCP_HOSTNAME` — Only use this option if the DHCP server requires the client to specify a hostname before receiving an IP address. (The DHCP server daemon in Red Hat Enterprise Linux

does not support this feature.)

- **PEERDNS=<*answer*>** , where **<*answer*>** is one of the following:
  - **yes** — Modify **/etc/resolv.conf** with information from the server. If using DHCP, then **yes** is the default.
  - **no** — Do not modify **/etc/resolv.conf**.
- **SRCADDR=<*address*>** , where **<*address*>** is the specified source IP address for outgoing packets.
- **USERCTL=<*answer*>** , where **<*answer*>** is one of the following:
  - **yes** — Non-root users are allowed to control this device.
  - **no** — Non-root users are not allowed to control this device.

If you prefer using a graphical interface, refer to Chapter 17, *Network Configuration* for details on using the **Network Administration Tool** to configure a network interface to use DHCP.

> **Tip**
>
> For advanced configurations of client DHCP options such as protocol timing, lease requirements and requests, dynamic DNS support, aliases, as well as a wide variety of values to override, prepend, or append to client-side configurations, refer to the **dhclient** and **dhclient.conf** man pages.

23.4. Additional Resources

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

23.4.1. 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.
- **/usr/share/doc/dhcp-<*version*>/** — Contains sample files, README files, and release notes for the specific version of the DHCP service.

---

[6] **Kudzu** is a hardware probing tool run at system boot time to determine what hardware has been added or removed from the system.

Chapter 24. Apache HTTP Server Configuration

Red Hat Enterprise Linux provides version 2.0 of the Apache HTTP Server. 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 *Reference Guide* for details.

If you configured the Apache HTTP Server with the **HTTP Configuration Tool** in previous versions of Red Hat Enterprise Linux and then performed an upgrade, you can use the **HTTP Configuration Tool** 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 `httpd` and `system-config-httpd` RPM packages need to be installed to use the **HTTP Configuration Tool**. It also requires the X Window System and root access. To start the application, go to the **Main Menu Button** => **System Settings** => **Server Settings** => **HTTP** or type the command `system-config-httpd` at a shell prompt (for example, in an XTerm or GNOME Terminal).

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 provided with Red Hat Enterprise Linux can be configured with the **HTTP Configuration Tool**. If additional modules are installed, they can not be configured using this tool.

> ⚠ **Warning**
>
> Do not edit the `/etc/httpd/conf/httpd.conf` configuration file by hand if you wish to use this tool. The **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 follows:

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. To serve more than one URL or virtual host, add any 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.
8. Exit the application and select to save your settings.

24.1. Basic Settings

Use the **Main** tab to configure the basic server settings.
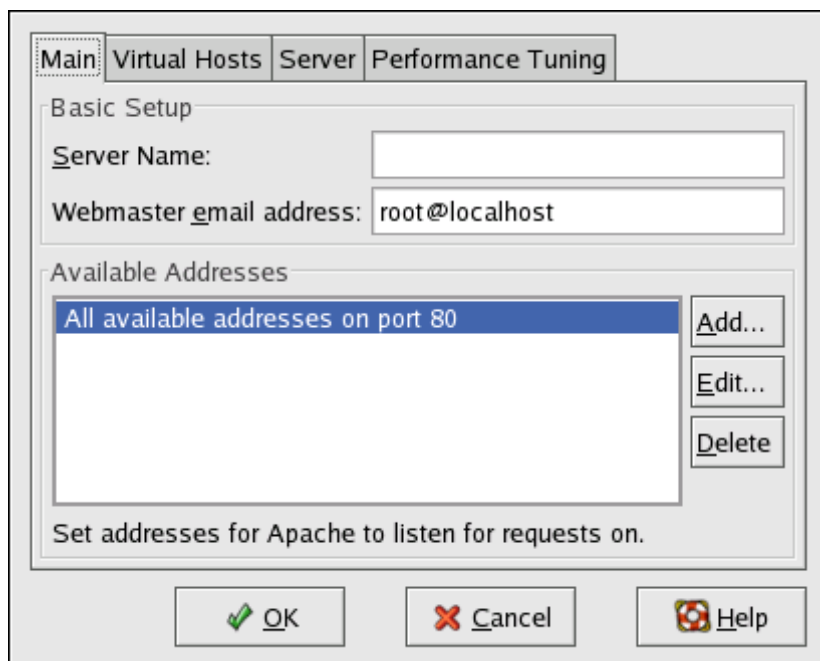
Figure 24.1. Basic Settings

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 set the server name to www.example.com while the server's real DNS name is foo.example.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 is used so that users can report a problem to the server's administrator. The default value is root@localhost.

Use the **Available Addresses** area to define the ports on which the server accepts 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 24.2, "Available Addresses" appears. 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 accepts connections in the **Address** field. Only specify one IP address per port number. 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 in the **Available Addresses** frame 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.

> **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.

Figure 24.2. Available Addresses

24.2. 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. A window as shown in Figure 24.3, "Site Configuration" appears. 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.

24.2.1. Site Configuration

The default values for the **Directory Page Search List** and **Error Pages** work for most servers. If you are unsure of these settings, do not modify them.
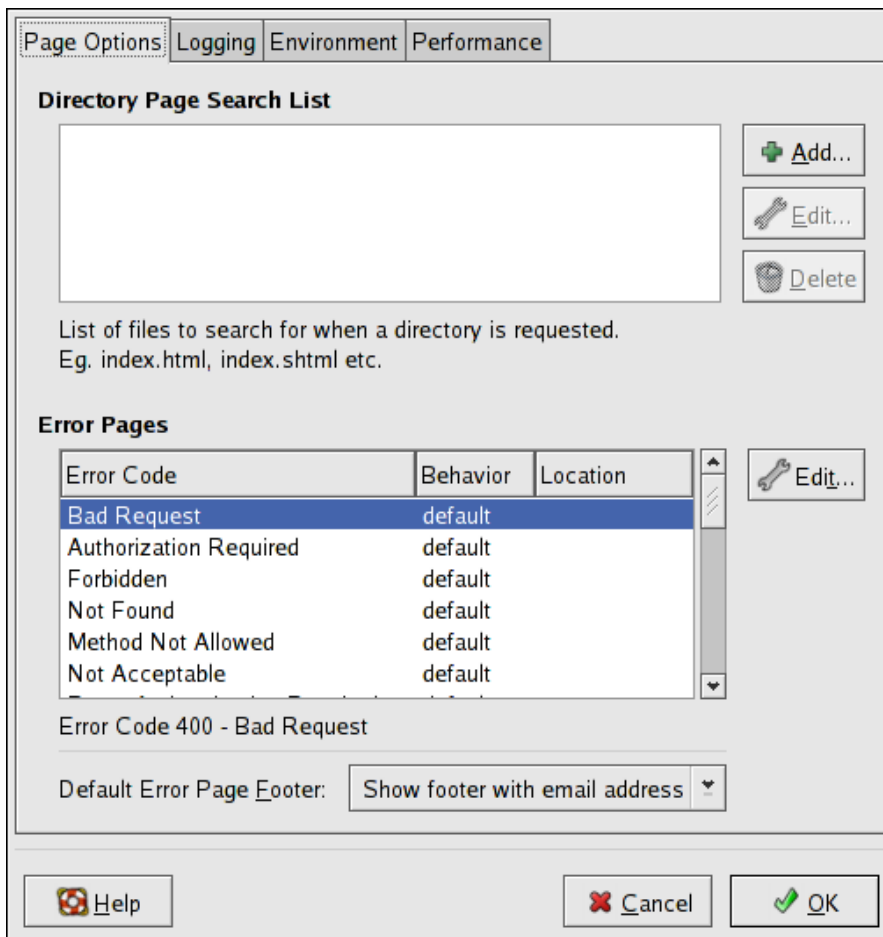
Figure 24.3. Site Configuration

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://www.example.com/this_directory/**, they are going to get either the **DirectoryIndex** page, if it exists, or a server-generated directory list. The server tries to find one of the files listed in the **DirectoryIndex** directive and returns 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 generates and returns 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 in 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 webpage that you created in a file called **404.html**, copy **404.html** to *DocumentRoot/../error/404.html*. In this case, *DocumentRoot* is the Document Root directory that you have defined (the default is **/var/www/html/**). If the Document Root is left as the default location, the file should be copied to **/var/www/error/404.html**. Then,

choose **File** as the Behavior for **404 - Not Found** error code and enter **/error/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 24.3.1.1, "General Options" 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.

24.2.2. Logging

Use the **Logging** tab to configure options for specific transfer and error logs.

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

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 file name do not start with a slash (/), the path is relative to the server root directory as configured. This option corresponds to the TransferLog directive.

Figure 24.4. Logging

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 file name do 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 the verbosity of the error messages in the error logs. 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 automatically resolves the IP address for each connection which requests a document from your Web server. Resolving the IP address means that your server makes 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 performs 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. Each individual connection made to look up each hostname adds 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**.

24.2.3. Environment Variables

Use the **Environment** tab to configure options for specific variables to set, pass, or unset for CGI scripts.

Sometimes it is necessary to modify environment variables for CGI scripts or server-side include (SSI) pages. 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.

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 Figure 24.5, "Environment Variables", and type **MAXNUM** in the **Environment Variable** text field and **50** in the **Value to set** text field. Click **OK** to add it to the list. 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 is 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** to add it to the list. The **Pass to CGI Scripts** section

configures the **PassEnv** directive.
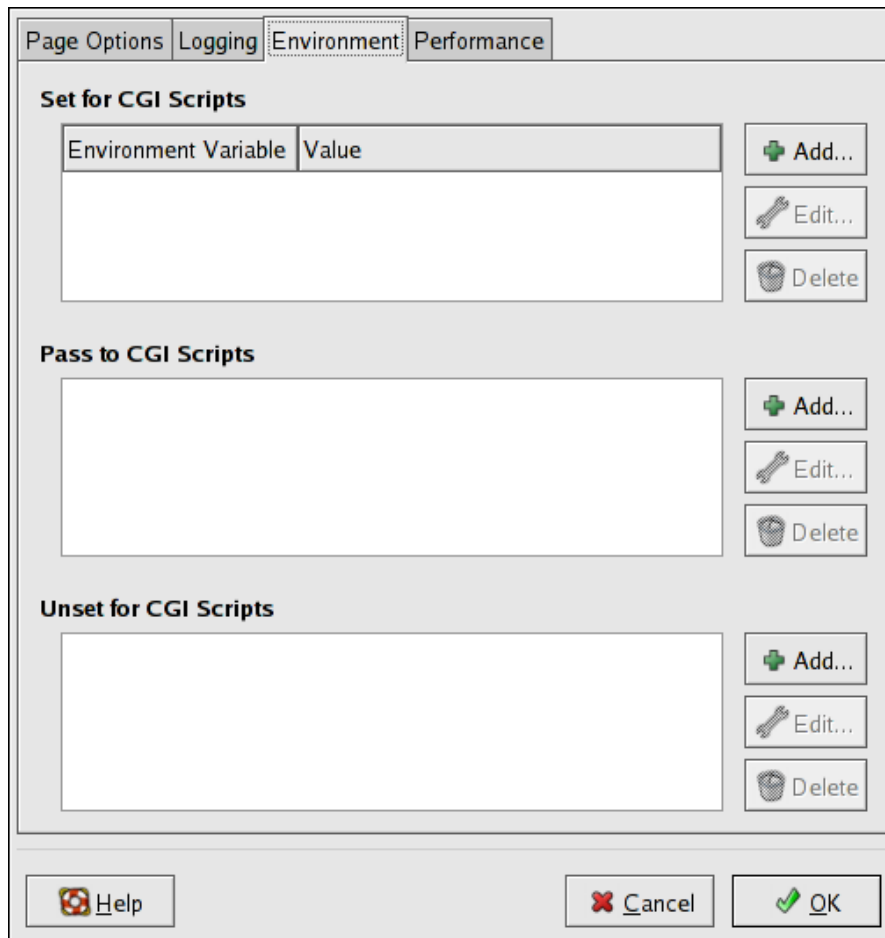
Figure 24.5. Environment Variables

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. Click **OK** to add it to the list. This corresponds to the **UnsetEnv** directive.

To edit any of these environment values, select it from the list and click the corresponding **Edit** button. To delete any entry from the list, select it and click the corresponding **Delete** button.

To learn more about environment variables in the Apache HTTP Server, refer to the following:

http://httpd.apache.org/docs-2.0/env.html

24.2.4. Directories

Use the **Directories** page in the **Performance** tab to configure options for specific directories. This corresponds to the **<Directory>** directive.
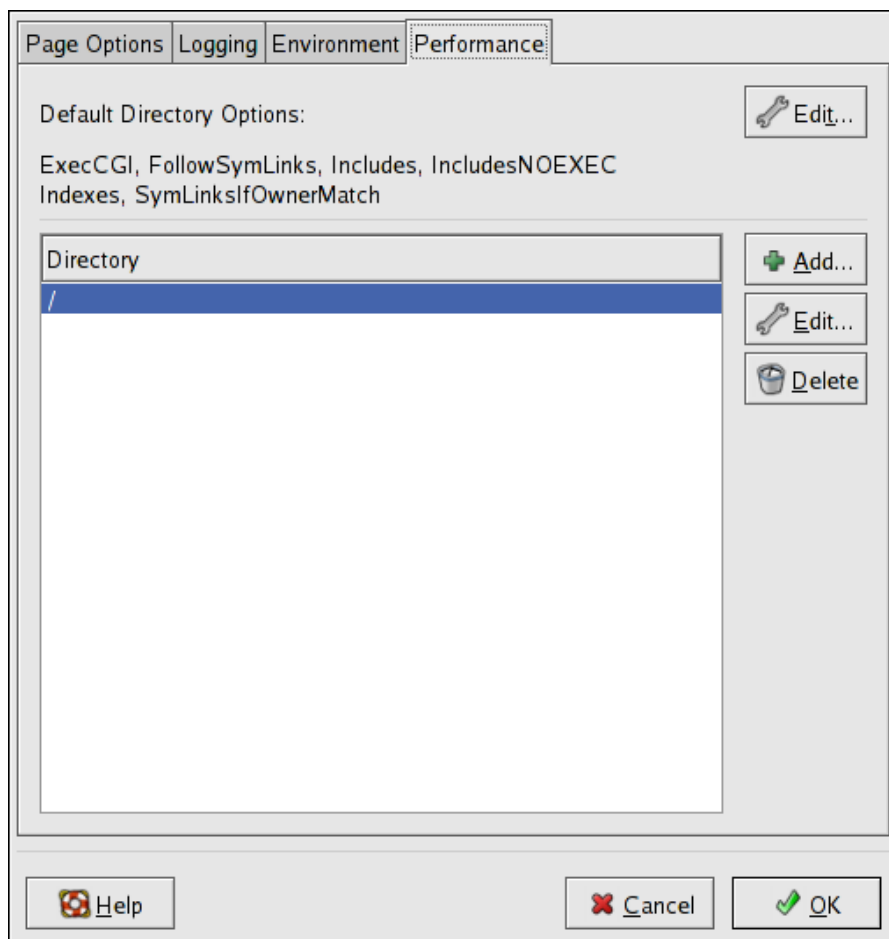
Figure 24.6. Directories

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. A window as shown in Figure 24.7, "Directory Settings" appears. Enter the directory to configure in the **Directory** text field at the bottom of the window. Select the options in the right-hand list and configure the **Order** directive with the left-hand side options. The **Order** directive controls the order in which allow and deny directives are evaluated. In the **Allow hosts from** and **Deny hosts from** text field, you can specify one of the following:

- Allow all hosts — Type **all** 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**
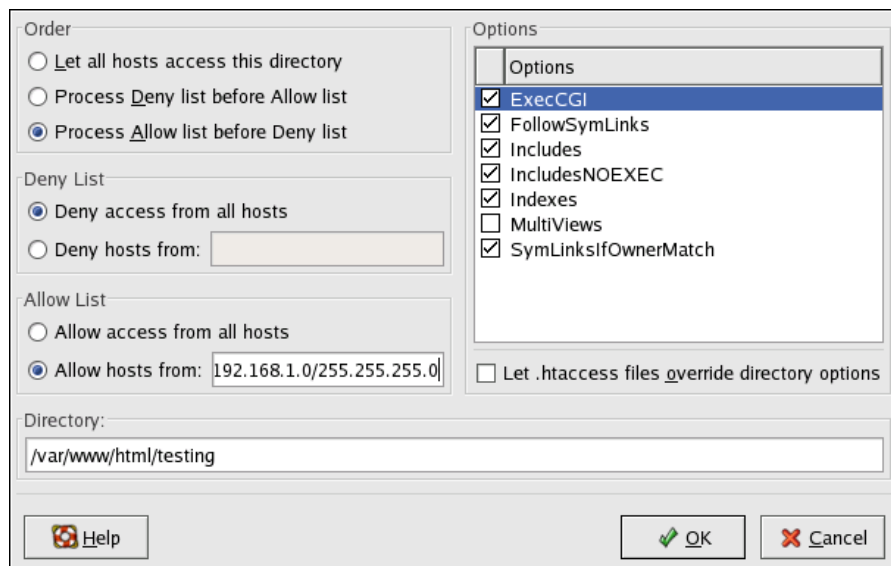
Figure 24.7. Directory Settings

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

24.3. Virtual Hosts Settings

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.example.com and http://www.anotherexample.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.

The **HTTP Configuration Tool** includes a default virtual host as shown in Figure 24.8, "Virtual Hosts".

Figure 24.8. Virtual Hosts
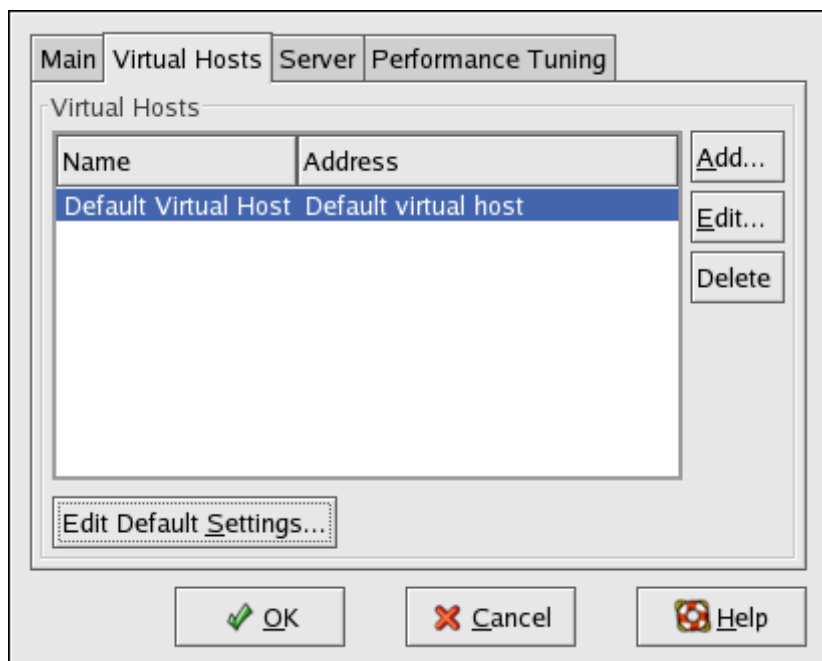
http://httpd.apache.org/docs-2.0/vhosts/ and the Apache HTTP Server documentation on your machine provide more information about virtual hosts.

24.3.1. 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 and clicking the **Edit** button.

24.3.1.1. 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. 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 multiple IP addresses, separate each IP address with spaces. To specify a port, use the syntax *IP Address:Port*. Use "colon, asterisk" (**:***) 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 multiple IP addresses, separate each IP address with spaces. To specify a port, use the syntax *IP Address:Port*. Use "colon, asterisk" (**:***) 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.

24.3.1.2. SSL

> **Note**
>
> You cannot 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 plan to use name-based virtual hosts, remember that they only work with your non-secure Web server.

Figure 24.9. SSL Support

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 the **HTTP Configuration Tool**, you must allow access through port 443 under the **Main** tab => **Available Addresses**. Refer to Section 24.1, "Basic Settings" for details. Then, select the virtual host name in the **Virtual Hosts** tab, click the **Edit** button, choose **SSL** from the left-hand menu, and check the **Enable SSL Support** option as shown in Figure 24.9, "SSL Support". The **SSL Configuration** section is pre-configured with the dummy digital certificate. The digital certificate provides authentication for your secure Web server and identifies the secure server to client Web browsers. You must purchase your own digital certificate. Do not use the dummy one provided for your website. For details on purchasing a CA-approved digital certificate, refer to the Chapter 25, *Apache HTTP Secure Server Configuration*.

24.3.1.3. 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 24.2, "Default Settings" for details on these options.

24.4. Server Settings

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

Figure 24.10. 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 executing a core dump. 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 are also 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** creates 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 answers requests. The default group is also apache.

24.5. 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.
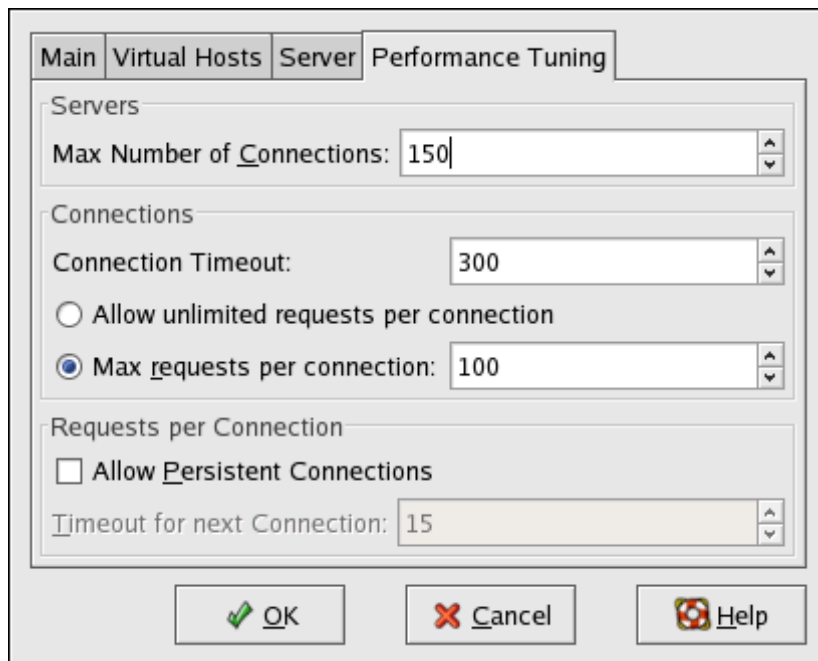
Figure 24.11. Performance Tuning

Set **Max Number of Connections** to the maximum number of simultaneous client requests that the server can handle. For each connection, a child **httpd** process is created. After this maximum number of processes is reached, no one else can 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 waits for receipts and transmissions during communications. Specifically, **Connection Timeout** defines how long your server waits to receive a GET request, how long it waits to receive TCP packets on a POST or PUT request, and how long it waits 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 is set 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 waits 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 the server to slow down, depending on how many users are trying to connect to it. The higher the number, the more server processes are waiting for another connection from the last client that connected to it.

24.6. 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 are prompted to confirm this decision. If you click **Yes** to confirm this choice, your settings are not 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 appears. If you answer **Yes**, your settings are saved in `/etc/httpd/conf/httpd.conf`. Remember that your original configuration file is overwritten with your new settings.

If this is the first time that you have used the **HTTP Configuration Tool**, a dialog window appears, 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 saves 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.

24.7. Additional Resources

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

24.7.1. Installed Documentation

- `/usr/share/docs/httpd-<version>/migration.html` — 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 migration the configuration file manually.

24.7.2. Useful Websites

- http://www.apache.org/ — *The Apache Software Foundation*.
- http://httpd.apache.org/docs-2.0/ — The Apache Software Foundation's documentation on Apache HTTP Server version 2.0, including the *Apache HTTP Server Version 2.0 User's Guide*.

24.7.3. Related Books

- *Apache: The Definitive Guide* by Ben Laurie and Peter Laurie; O'Reilly & Associates, Inc.
- *Reference Guide* ; Red Hat, Inc — This companion manual includes instructions for migrating from Apache HTTP Server version 1.3 to Apache HTTP Server version 2.0 manually, more details about the Apache HTTP Server directives, and instructions for adding modules to the Apache HTTP Server.

Chapter 25. Apache HTTP Secure Server Configuration

25.1. 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 are 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, regular HTTP communications between a browser and a Web server are sent in plain text, 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 points to appropriate places where you can find more in-depth documentation on particular subjects.

This chapter shows you how to install these programs. You can 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 server.

The **mod_ssl** configuration file is located at **/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 the **/etc/httpd/conf/httpd.conf** file. This statement is included by default in the default Apache HTTP Server configuration file.

25.2. An Overview of Security-Related Packages

To enable the secure server, you must 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 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 need all of these if you intend to load any extra modules, other than the modules provided with this product. Refer to the *Reference Guide* for more information on loading modules onto your secure server using Apache's dynamic shared object (DSO) functionality.

If you do not intend to load other modules onto your Apache HTTP Server, you do not need to install this package.

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 securely copying 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.

The **openssh-askpass-gnome** package can be used in conjunction with the GNOME desktop environment to display a graphical dialog window 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 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**; **sftp**, a secure replacement for **ftp** (for transferring files between machines); 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 20, *OpenSSH*, the *Reference Guide*, 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. It provides encryption for non-SSL aware daemons and protocols (such as POP, IMAP, and LDAP) without requiring any changes to the daemon's code.

> ### Note
>
> Newer implementations of various daemons now provide their services natively over SSL, such as **dovecot** or OpenLDAP's **slapd** server, which may be more desirable than using **stunnel**.
> For example, use of **stunnel** only provides wrapping of protocols, while the native support in OpenLDAP's **slapd** can also handle in-band upgrades for using encryption in response to a **StartTLS** client request.

Table 25.1, "Security Packages" displays a summary of the secure server packages and whether each package is optional for the installation of a secure server.

Table 25.1. Security Packages

Package Name	Optional?
`httpd`	no
`mod_ssl`	no
`openssl`	no
`httpd-devel`	yes
`openssh`	yes
`openssh-askpass`	yes
`openssh-askpass-gnome`	yes
`openssh-clients`	yes
`openssh-server`	yes
`openssl-devel`	yes
`stunnel`	yes

25.3. An Overview of Certificates and Security

Your secure 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 as well as the mutual authentication between browsers and your secure server. The CA-approved digital certificate provides authentication for your secure server (the CA puts its reputation behind its certification of your organization's identity). When your browser is communicating using SSL encryption, the **https://** prefix is used 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, use public cryptography to create a public and private key pair. In most cases, you send your certificate request (including your public key), proof of your company's identity, and

payment to a CA. The CA verifies the certificate request and your identity, and then sends back a certificate for your secure 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 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 are not automatically accepted by a user's browser — users are prompted by the browser to accept the certificate and create the secure connection. Refer to Section 25.5, "Types of Certificates" 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 must install it on your secure server.

25.4. Using Pre-Existing Keys and Certificates

If you already have an existing key and certificate (for example, if you are installing the secure server to replace another company's secure server product), you can probably use your existing key and certificate with the secure server. The following two situations provide instances where you are not able to use your existing 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 must 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 server. However, you are not be allowed to because VeriSign issues certificates for one specific server software and IP address/domain name combination.

  If you change either of those parameters (for example, if you previously used a different secure server product), the VeriSign certificate you obtained to use with the previous configuration will not work with the new configuration. You must obtain a new certificate.

If you have an existing key and certificate that you can use, you do 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 25.9, "Testing The Certificate".

If you are upgrading from the Red Hat Secure Web Server, your old key (**httpsd.key**) and certificate (**httpsd.crt**) are located in **/etc/httpd/conf/**. Move and rename your key and certificate so that the secure 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 server with the command:

```
/sbin/service httpd start
```

You are prompted to enter your passphrase. After you type it in and press **Enter**, the server starts.

25.5. Types of Certificates

If you installed your secure server from the RPM package provided by Red Hat, a random key and a test certificate are generated and put into the appropriate directories. Before you begin using your secure server, however, you must generate your own key and obtain a certificate which correctly identifies your server.

You need a key and a certificate to operate your secure 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 (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 webpages to the browser.

If your secure server is being accessed by the public at large, your secure 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 automatically accept. If a browser encounters a certificate whose authorizing CA is not in the list, the browser asks the user to either accept or decline the connection.

You can generate a self-signed certificate for your secure server, but be aware that a self-signed certificate does not provide the same functionality as a CA-signed certificate. A self-signed certificate is not automatically recognized by most Web browsers and 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 is to be used in a production environment, a CA-signed certificate is recommended.

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. Red Hat does not make recommendations on which certificate authority to choose. Your decision may be based on your past experiences, on the experiences of your friends or colleagues, or purely on monetary factors.

    Once you have decided upon a CA, you 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 provide you with a digital certificate.

5. Install this certificate on your secure 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. Refer to Section 25.6, "Generating a Key" for instructions.

25.6. Generating a Key

You must be root to generate a key.

First, use the **cd** command to change 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.keyrm ssl.crt/server.crt
```

Next, create your own random key. Change to the **/usr/share/ssl/certs/** directory and type in the following command:

```
make genkey
```

Your system displays 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 pass phrase:
```

You now must enter in a passphrase. For security reason, it should contain at least eight characters, include numbers and/or punctuation, and it should not be a word in a dictionary. Also, remember that your passphrase is case sensitive.

Note

You are required to remember and enter this passphrase every time you start your secure server. If you forget this passphrase, the key must be completely re-generated.

Re-type the passphrase to verify that it is correct. Once you have typed it in correctly, **/etc/httpd/conf/ssl.key/server.key**, the file containing your key, is created.

Note that if you do not want to type in a passphrase every time you start your secure server, you must use the following two commands instead of **make genkey** to create the key.

Use the following command to create your key:

```
/usr/bin/openssl genrsa 1024 > /etc/httpd/conf/ssl.key/server.key
```

Then, use the following command to make sure the permissions are set correctly for the file:

```
chmod go-rwx /etc/httpd/conf/ssl.key/server.key
```

After you use the above commands to create your key, you do not need to use a passphrase to start your secure server.

> ⚠️ **Warning**
>
> Disabling the passphrase feature for your secure server is a security risk. It is *not* recommended that you disable the passphrase feature for secure server.

Problems associated with not using a passphrase 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 webpages that appear to be from your secure 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), secure server's passphrase may seem unnecessary. However, since your secure server should not need to be re-booted very often, the extra security provided by entering a passphrase 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 no longer works and the CA is not able to help you. Your only option is to request (and pay for) a new certificate.

If you are going to purchase a certificate from a CA, continue to Section 25.7, "Generating a Certificate Request to Send to a CA". If you are generating your own self-signed certificate, continue to Section 25.8, "Creating a Self-Signed Certificate".

25.7. 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 need to send to the CA of your choice. Make sure you are in the **/usr/share/ssl/certs/** directory, and type the following command:

```
make certreq
```

Your system displays the following output and asks you for your passphrase (unless you disabled the passphrase option):

```
umask 77 ; \
/usr/bin/openssl req -new -key -set_serial num /etc/httpd/conf/ssl.key/server.key
-out /etc/httpd/conf/ssl.csr/server.csr
Using configuration from /usr/share/ssl/openssl.cnf
Enter pass phrase:
```

Type in the passphrase that you chose when you were generating your key unless you don't need to. Next, your system displays some instructions and then ask for a series of responses from you. Your inputs are incorporated into the certificate request. The display, with example responses, looks similar to the following:

```
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]:USState or Province Name (full name)
[Berkshire]:North CarolinaLocality Name (eg, city) [Newbury]:RaleighOrganization
Name (eg, company) [My Company Ltd]:Test CompanyOrganizational Unit Name (eg,
section) []:TestingCommon Name (your name or server's hostname)
[]:test.example.comEmail Address []:admin@example.comPlease 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 is to be used, shown like the following:

```
Country Name (2 letter code) [GB]:
```

The default input, in brackets, is **GB**. Accept the default by pressing **Enter** or fill in your country's two letter code.

You have to type in the rest of the values. All of these should be self-explanatory, but you must 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 rejects CSRs which include information they perceive as invalid.

» For **Common Name**, make sure you type in the *real* name of your secure 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 special characters like @, #, & !, and etc. Some CAs reject a certificate request which contains a special character. 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.

The file **/etc/httpd/conf/ssl.csr/server.csr** is created when you have finished entering your information. 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 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 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**. Be sure to keep a backup of this file.

25.8. Creating a Self-Signed Certificate

You can create your own self-signed certificate. Note that a self-signed certificate does not provide the security guarantees of a CA-signed certificate. Refer to Section 25.5, "Types of Certificates" for more details about certificates.

To make your own self-signed certificate, first create a random key using the instructions provided in Section 25.6, "Generating a Key". Once you have a key, make sure you are in the **/usr/share/ssl/certs/** directory, and type the following command:

```
make testcert
```

The following output is shown and you are prompted for your passphrase (unless you generated a key without a passphrase):

```
umask 77 ; \
/usr/bin/openssl req -new -key -set_serial num /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 pass phrase:
```

Next, you are asked for more information. The computer's output and a set of inputs looks like the following (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 CarolinaLocality Name (eg,
city) [Newbury]:RaleighOrganization Name (eg, company) [My Company Ltd]:My
Company, Inc.Organizational Unit Name (eg, section) []:DocumentationCommon Name
(your name or server's hostname) []:myhost.example.comEmail Address
[]:myemail@example.com
```

After you provide the correct information, a self-signed certificate is created in **/etc/httpd/conf/ssl.crt/server.crt**. Restart the secure server after generating the certificate with following the command:

```
/sbin/service httpd restart
```

25.9. Testing The Certificate

To test the test certificate installed by default, either a CA-signed certificate, or a self-signed certificate, point your Web browser to the following home page (replacing *server.example.com* with your domain name):

```
https://server.example.com
```

> **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 probably automatically accepts the certificate (without prompting you for input) and creates the secure connection. Your browser does 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.

Once your browser accepts the certificate, your secure server displays a default home page.

25.10. Accessing The Server

To access your secure server, use a URL similar to the following:

```
https://server.example.com
```

Your non-secure server can be accessed using an URL similar to the following:

```
http://server.example.com
```

The standard port for secure Web communications is port 443. The standard port for non-secure Web communications is port 80. The secure server default configuration listens on both of the two standard ports. Therefore, do 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 (for example, anything other than 80 or 443), you must 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 attempts to connect to a non-secure server listening on port 12331:

```
http://server.example.com:12331
```

25.11. Additional Resources

Refer to Section 24.7, "Additional Resources" for more information about the Apache HTTP Server.

25.11.1. Useful Websites

▹ 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/.

25.11.2. Related Books

- *Apache: The Definitive Guide* , 3rd edition, by Ben Laurie and Peter Laurie, O'Reilly & Associates, Inc.

Chapter 26. Authentication Configuration

When a user logs in to a Red Hat Enterprise Linux system, the username and password combination must be verified, or *authenticated*, as a valid and active user. Sometimes the information to verify the user is located on the local system, and other times the system defers the authentication to a user database on a remote system.

The **Authentication Configuration Tool** provides a graphical interface for configuring NIS, LDAP, and Hesiod to retrieve user information as well as for configuring LDAP, Kerberos, and SMB as authentication protocols.

> **Note**
>
> If you configured a medium or high security level during installation or with the **Security Level Configuration Tool**, network authentication methods, including NIS and LDAP, are not allowed through the firewall.

This chapter does not explain each of the different authentication types in detail. Instead, it explains how to use the **Authentication Configuration Tool** to configure them. For more information about the specific authentication types, refer to the *Reference Guide*.

To start the graphical version of the **Authentication Configuration Tool** from the desktop, select the **Main Menu Button** (on the Panel) => **System Settings** => **Authentication** or type the command **system-config-authentication** at a shell prompt (for example, in an **XTerm** or a **GNOME terminal**). To start the text-based version, type the command **authconfig** as root at a shell prompt.

> **Important**
>
> After exiting the authentication program, the changes made take effect immediately.

26.1. User Information

The **User Information** tab has several options. To enable an option, click the empty checkbox beside it. To disable an option, click the checkbox beside it to clear the checkbox. Click **OK** to exit the program and apply the changes.
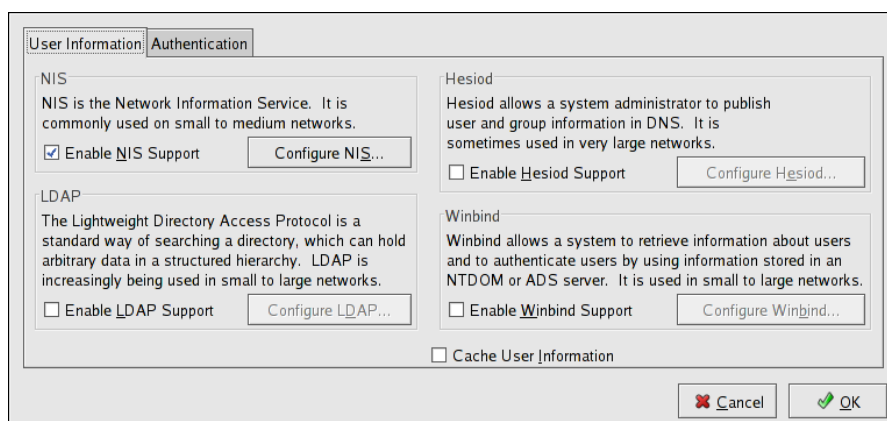
Figure 26.1. User Information

The following list explains what each option configures:

- **Enable NIS Support** — Select this option to configure the system as an NIS client which connects to an NIS server for user and password authentication. Click the **Configure NIS** button to specify the NIS domain and NIS server. If the NIS server is not specified, the daemon attempts to find it via broadcast.

  The **ypbind** package must be installed for this option to work. If NIS support is enabled, the **portmap** and **ypbind** services are started and are also enabled to start at boot time.

- **Enable LDAP Support** — Select this option to configure the system to retrieve user information via LDAP. Click the **Configure LDAP** button to specify the **LDAP Search Base DN** and **LDAP Server**. If **Use TLS to encrypt connections** is selected, Transport Layer Security is used to encrypt passwords sent to the LDAP server.

  The **openldap-clients** package must be installed for this option to work.

  For more information about LDAP, refer to the *Reference Guide*.

- **Enable Hesiod Support** — Select this option to configure the system to retrieve information from a remote Hesiod database, including user information.

  The **hesiod** package must be installed.

- **Winbind** — Select this option to configure the system to connect to a Windows Active Directory or a Windows domain controller. User information can be accessed, as well as server authentication options can be configured.

- **Cache User Information** — Select this option to enable the name service cache daemon (**nscd**) and configure it to start at boot time.

  The **nscd** package must be installed for this option to work.

26.2. Authentication

The **Authentication** tab allows for the configuration of network authentication methods. To enable an option, click the empty checkbox beside it. To disable an option, click the checkbox beside it to clear the checkbox.
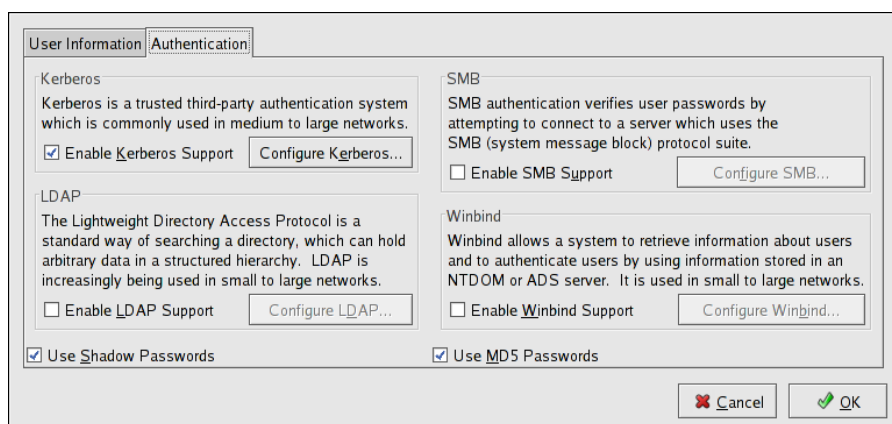
Figure 26.2. Authentication

The following explains what each option configures:

- **Enable Kerberos Support** — Select this option to enable Kerberos authentication. Click the **Configure Kerberos** button to configure:

- **`Realm`** — Configure the realm for the Kerberos server. The realm is the network that uses Kerberos, composed of one or more KDCs and a potentially large number of clients.

- **`KDC`** — Define the Key Distribution Center (KDC), which is the server that issues Kerberos tickets.

- **`Admin Servers`** — Specify the administration server(s) running **`kadmind`**.

The **`krb5-libs`** and **`krb5-workstation`** packages must be installed for this option to work. Refer to the *Reference Guide* for more information on Kerberos.

- **`Enable LDAP Support`** — Select this option to have standard PAM-enabled applications use LDAP for authentication. Click the **`Configure LDAP`** button to specify the following:

  - **`Use TLS to encrypt connections`** — Use Transport Layer Security to encrypt passwords sent to the LDAP server.

  - **`LDAP Search Base DN`** — Retrieve user information by its Distinguished Name (DN).

  - **`LDAP Server`** — Specify the IP address of the LDAP server.

The **`openldap-clients`** package must be installed for this option to work. Refer to the *Reference Guide* for more information about LDAP.

- **`Use Shadow Passwords`** — Select this option to store passwords in shadow password format in the **`/etc/shadow`** file instead of **`/etc/passwd`**. Shadow passwords are enabled by default during installation and are highly recommended to increase the security of the system.

The **`shadow-utils`** package must be installed for this option to work. For more information about shadow passwords, refer to the *Users and Groups* chapter in the *Reference Guide*.

- **`Enable SMB Support`** — This option configures PAM to use an SMB server to authenticate users. Click the **`Configure SMB`** button to specify:

  - **`Workgroup`** — Specify the SMB workgroup to use.

  - **`Domain Controllers`** — Specify the SMB domain controllers to use.

- **`Winbind`** — Select this option to configure the system to connect to a Windows Active Directory or a Windows domain controller. User information can be accessed, as well as server authentication options can be configured.

- **`Use MD5 Passwords`** — Select this option to enable MD5 passwords, which allows passwords to be up to 256 characters instead of eight characters or less. It is selected by default during installation and is highly recommended for increased security.

26.3. Command Line Version

The **Authentication Configuration Tool** can also be run as a command line tool with no interface. The command line version can be used in a configuration script or a kickstart script. The authentication options are summarized in Table 26.1, "Command Line Options".

> ### Tip
>
> These options can also be found in the **`authconfig`** man page or by typing **`authconfig --help`** at a shell prompt.

Table 26.1. Command Line Options

Option	Description
--enableshadow	Enable shadow passwords
--disableshadow	Disable shadow passwords
--enablemd5	Enable MD5 passwords
--disablemd5	Disable MD5 passwords
--enablenis	Enable NIS
--disablenis	Disable NIS
--nisdomain=<*domain*>	Specify NIS domain
--nisserver=<*server*>	Specify NIS server
--enableldap	Enable LDAP for user information
--disableldap	Disable LDAP for user information
--enableldaptls	Enable use of TLS with LDAP
--disableldaptls	Disable use of TLS with LDAP
--enableldapauth	Enable LDAP for authentication
--disableldapauth	Disable LDAP for authentication
--ldapserver=<*server*>	Specify LDAP server
--ldapbasedn=<*dn*>	Specify LDAP base DN
--enablekrb5	Enable Kerberos
--disablekrb5	Disable Kerberos
--krb5kdc=<*kdc*>	Specify Kerberos KDC
--krb5adminserver=<*server*>	Specify Kerberos administration server
--krb5realm=<*realm*>	Specify Kerberos realm
--enablekrb5kdcdns	Enable use of DNS to find Kerberos KDCs
--disablekrb5kdcdns	Disable use of DNS to find Kerberos KDCs
--enablekrb5realmdns	Enable use of DNS to find Kerberos realms
--disablekrb5realmdns	Disable use of DNS to find Kerberos realms
--enablesmbauth	Enable SMB
--disablesmbauth	Disable SMB
--smbworkgroup=<*workgroup*>	Specify SMB workgroup
--smbservers=<*server*>	Specify SMB servers
--enablewinbind	Enable winbind for user information by default
--disablewinbind	Disable winbind for user information by default
--enablewinbindauth	Enable winbindauth for authentication by default
--disablewinbindauth	Disable winbindauth for authentication by default

--smbsecurity=<*user|server|domain|ads*>	Security mode to use for Samba and winbind
--smbrealm=<*STRING*>	Default realm for Samba and winbind when **security=ads**
--smbidmapuid=<*lowest-highest*>	UID range winbind assigns to domain or ADS users
--smbidmapgid=<*lowest-highest*>	GID range winbind assigns to domain or ADS users
--winbindseparator=<\>	Character used to separate the domain and user part of winbind usernames if **winbindusedefaultdomain** is not enabled
--winbindtemplatehomedir=</home/%D/%U>	Directory that winbind users have as their home
--winbindtemplateprimarygroup=<*nobody*>	Group that winbind users have as their primary group
--winbindtemplateshell=</bin/false>	Shell that winbind users have as their default login shell
--enablewinbindusedefaultdomain	Configures winbind to assume that users with no domain in their usernames are domain users
--disablewinbindusedefaultdomain	Configures winbind to assume that users with no domain in their usernames are not domain users
--winbindjoin=<*Administrator*>	Joins the winbind domain or ADS realm now as this administrator
--enablewins	Enable WINS for hostname resolution
--disablewins	Disable WINS for hostname resolution
--enablehesiod	Enable Hesiod
--disablehesiod	Disable Hesiod
--hesiodlhs=<*lhs*>	Specify Hesiod LHS
--hesiodrhs=<*rhs*>	Specify Hesiod RHS
--enablecache	Enable **nscd**
--disablecache	Disable **nscd**
--nostart	Do not start or stop the **portmap**, **ypbind**, or **nscd** services even if they are configured
--kickstart	Do not display the user interface
--probe	Probe and display network defaults

Part V. System Configuration

Part of a system administrator's job is configuring the system for various tasks, types of users, and hardware configurations. This section explains how to configure a Red Hat Enterprise Linux system.

Chapter 27. 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 essentially 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 owns the files.

In contrast, *every* user who logs in at the console is 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**.

27.1. Disabling Shutdown Via `Ctrl+Alt+Del`

By default, **/etc/inittab** specifies that your system is set to shutdown and reboot in response to a **Ctrl**+**Alt**+**Del** key combination used at the console. To completely disable this ability, 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 want to allow certain non-root users the right to shutdown or reboot the system from the console using **Ctrl**+**Alt**+**Del**. You can restrict this privilege to certain users, by taking the following steps:

1. Add the **-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.

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 **shutdown.allow** file is a list of usernames, one per line, like the following:

   ```
   stephen
   jack
   sophie
   ```

According to this example **shutdown.allow** file, the users 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** command 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 continues; if not, an error message is written to the system console instead.

For more information on **shutdown.allow**, refer to the **shutdown** man page.

27.2. Disabling Console Program Access

To disable access by users to console programs, run the following 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/poweroffrm -f
/etc/security/console.apps/haltrm -f /etc/security/console.apps/reboot
```

27.3. 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]* vc/[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]* vc/[0-9][0-9]* :[0-9]\.[0-9] :[0-9] /dev/ttyS1
```

27.4. 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/* /mnt/floppy*
<sound>=/dev/dsp* /dev/audio* /dev/midi* \
       /dev/mixer* /dev/sequencer \
       /dev/sound/* /dev/beep \
/dev/snd/*
<cdrom>=/dev/cdrom* /dev/cdroms/* /dev/cdwriter* /mnt/cdrom*
```

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 /dev/usb/scanner*
```

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

That is 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 are given ownership of the **/dev/scanner** device with the permissions of 0600 (readable and writable by you only). When you log out, the device is owned by root and still has the permissions 0600 (now readable and writable by root only).

27.5. Enabling Console Access for Other Applications

To make other applications accessible to console users, a bit more work is required.

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/binln -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 **/usr/bin/foo** is executed, **consolehelper** is called, which authenticates the user with the help of **/usr/sbin/userhelper**. To authenticate the user, **consolehelper** asks for the user's password if **/etc/pam.d/foo** is a copy of **/etc/pam.d/halt** (otherwise, it does precisely what is specified in **/etc/pam.d/foo**) and then runs **/usr/sbin/foo** with root permissions.

In the PAM configuration file, an application can be configured to use the *pam_timestamp* module to remember (or 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 or KDE 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.

27.6. The `floppy` Group

If, for whatever reason, console access is not appropriate for you and your non-root users are required access to your system's diskette drive, this can be done using the **floppy** group. Add the user(s) to the **floppy** group using the tool of your choice. For example, the **gpasswd** command can be used to add user fred to the **floppy** group:

```
gpasswd -a fred floppy
```

Now, user fred is able to access the system's diskette drive from the console.

Chapter 28. Date and Time Configuration

The **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 use the tool. There are three ways to start the application:

- From the desktop, go to **Applications** (the main menu on the panel) => **System Settings** => **Date & Time**
- From the desktop, right-click on the time in the toolbar and select **Adjust Date and Time**.
- Type the command `system-config-date`, `system-config-time`, or `dateconfig` at a shell prompt (for example, in an **XTerm** or a **GNOME** terminal).

28.1. Time and Date Properties

As shown in , the first tabbed window that appears is for configuring the system date and time.
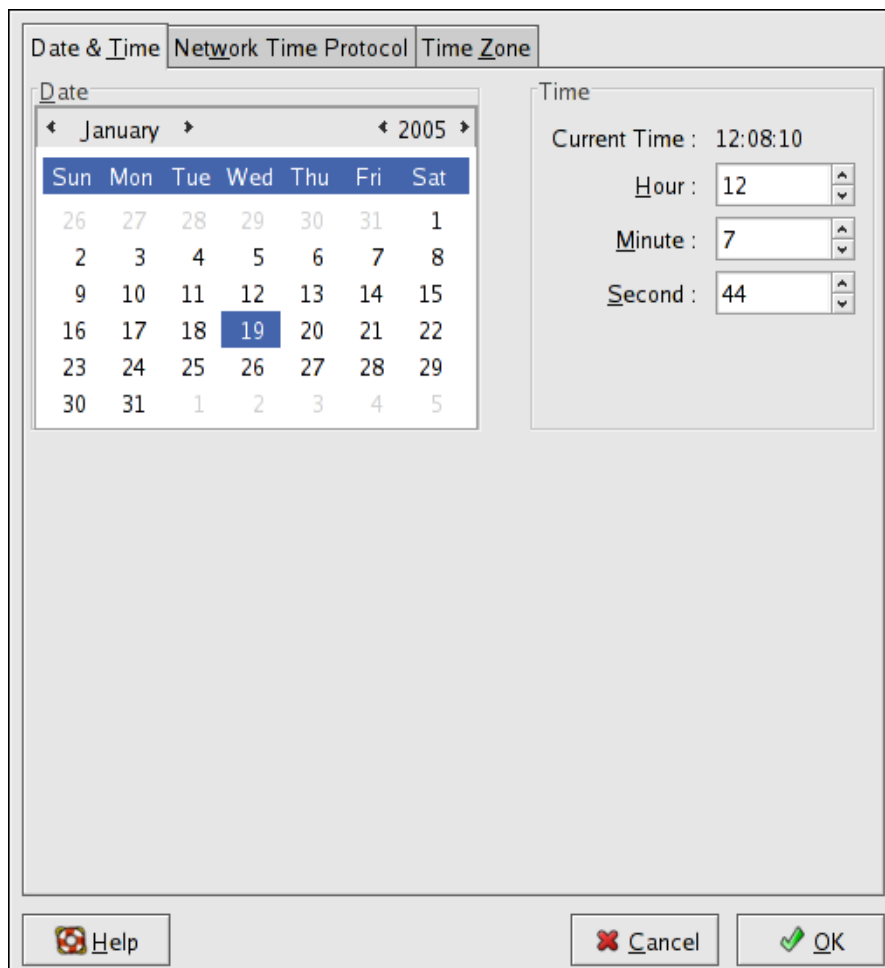
Figure 28.1. Time and Date Properties

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.

To change the time, use the up and down arrow buttons beside the **Hour**, **Minute**, and **Second** in the **Time** section.

Clicking the **OK** button applies any changes made to the date and time, the NTP daemon settings, and the time zone settings. It also exits the program.

28.2. Network Time Protocol (NTP) Properties

As shown in Figure 28.2, "NTP Properties", the second tabbed window that appears is for configuring NTP.
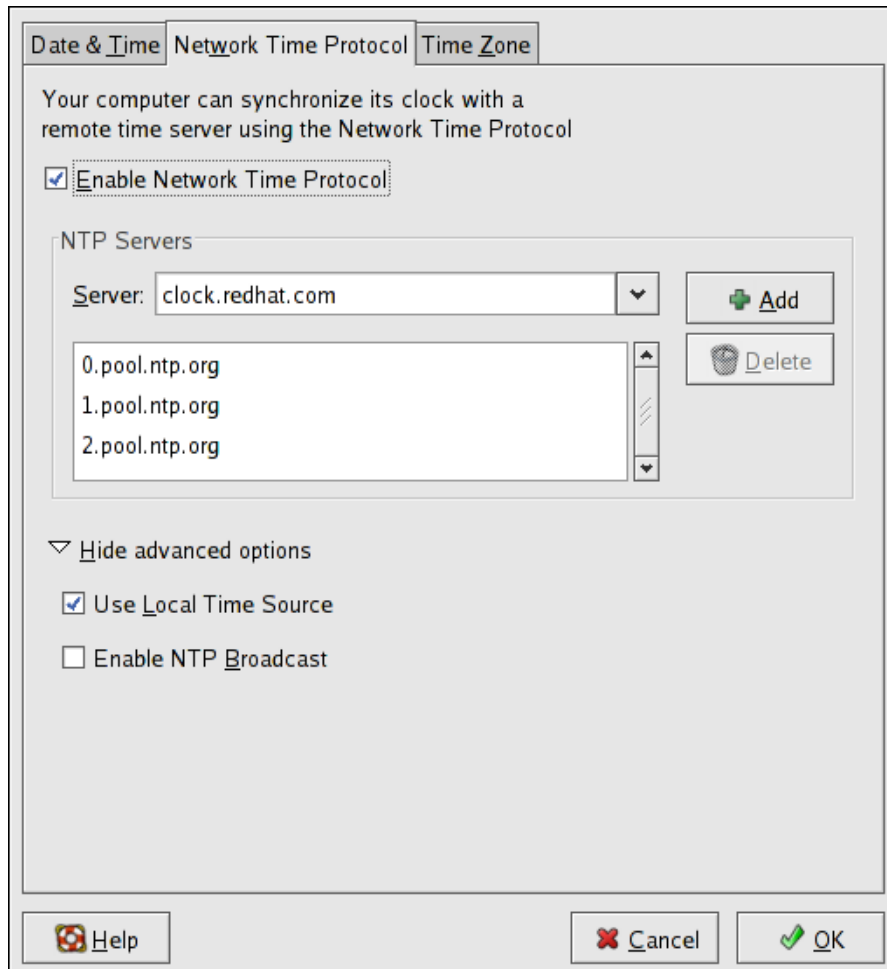
Figure 28.2. NTP Properties

The Network Time Protocol (NTP) daemon synchronizes the system clock with a remote time server or time source. The application allows you to configure an NTP daemon to synchronize your system clock with a remote server. To enable this feature, select **Enable Network Time Protocol**. This enables the **NTP Servers** list and other options. You can choose one of the predefined servers, edit a predefined server by clicking the **Edit** or add a new server name by clicking **Add**. Your system does not start synchronizing with the NTP server until you click **OK**. After clicking **OK**, the configuration is saved and the NTP daemon is started (or restarted if it is already running).

Clicking the **OK** button applies any changes made to the date and time, the NTP daemon settings, and the time zone settings. It also exits the program.

28.3. Time Zone Configuration

As shown in Figure 28.3, "Timezone Properties", the third tabbed window that appears is for configuring the system time zone.

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 desired region. The map zooms into the region selected, after which you may choose the city specific to your time zone. A red **X** appears and the time zone selection changes in the list below the map.

Alternatively, you can also use the list below the map. In the same way that the map lets you choose a region before choosing a city, the list of time zones is now a treelist, with cities and countries grouped within their specific continents. Non-geographic time zones have also been added to address needs in the scientific community.

Click **OK** to apply the changes and exit the program.
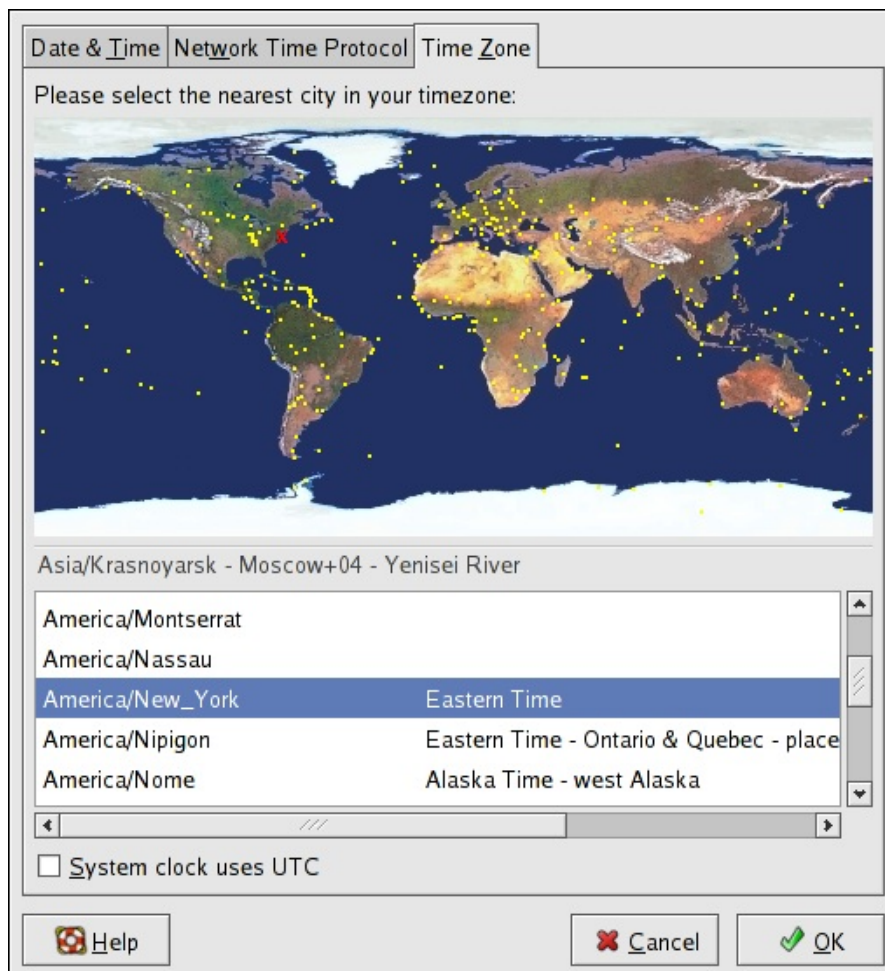
Figure 28.3. Timezone Properties

If your system clock is set to use UTC, select the **System clock uses UTC** option. UTC stands for the *Universal Time, Coordinated*, also known as Greenwich Mean Time (GMT). Other time zones are determined by adding or subtracting from the UTC time.

Chapter 29. Keyboard Configuration

The installation program allows users to configure a keyboard layout for their systems. To configure a different keyboard layout after installation, use the **Keyboard Configuration Tool**.

To start the **Keyboard Configuration Tool**, select **Applications** (the main menu on the panel) => **System Settings** => **Keyboard**, or type the command `system-config-keyboard` at a shell prompt.

Figure 29.1. Keyboard Configuration

Select a keyboard layout from the list (for example, `U.S. English`) and click `OK`. For changes to take effect, you should log out of your graphical desktop session and log back in.

Chapter 30. Mouse Configuration

The installation program allows users to select the type of mouse connected to the system. To configure a different mouse type for the system, use the **Mouse Configuration Tool**.

To start the **Mouse Configuration Tool**, type the command **system-config-mouse** at a shell prompt (for example, in an XTerm or GNOME terminal). If the X Window System is not running, the text-based version of the tool is started.

Figure 30.1. Mouse Configuration

Select the new mouse type for the system. If the specific mouse model is not listed, select one of the **Generic** entries, based on the mouse's number of buttons and its interface. If there is not an exact match, select the generic match that is most compatible with the system and the mouse.

> **Tip**
>
> Select the **Generic - Wheel Mouse** entry, with the proper mouse port, to enable the scroll button on the mouse.

The scroll button on a wheel mouse can be used as the middle mouse button for cutting text, pasting text, and other middle mouse button functions. If the mouse only has two buttons, select **Emulate 3 buttons** to use a two-button mouse as a three-button mouse. When this option enabled, clicking the two mouse buttons simultaneously emulates a middle mouse button click.

If a serial port mouse is selected, click the **Serial devices** button to configure the correct serial device number, such as **/dev/ttyS0** for the mouse.

Click **OK** to save the new mouse type. The selection is written to the file **/etc/sysconfig/mouse**, and the console mouse service, **gpm** is restarted. The changes are also written to the X Window System configuration file **/etc/X11/xorg.conf**; however, the mouse type change is not automatically applied to the current X session. To enable the new mouse type, log out of the graphical desktop and log back in.

> **Tip**
>
> To reset the order of the mouse buttons for a left-handed user, go to **Applications** (the main menu on the panel) => **Preferences** => **Mouse**, and select **Left-handed mouse** for the mouse orientation.

Chapter 31. X Window System Configuration

During installation, the system's monitor, video card, and display settings are configured. To change any of these settings after installation, use the **X Configuration Tool**.

To start the **X Configuration Tool**, go to **Applications** (the main menu on the panel) => **Administration** => **Display**, or type the command `system-config-display` at a shell prompt (for example, in an XTerm or GNOME terminal). If the X Window System is not running, a small version of X is started to run the program.

After changing any of the settings, log out of the graphical desktop and log back in to enable the changes.

31.1. Display Settings

The `Settings` tab allows users to change the *resolution* and *color depth*. The display of a monitor consists of tiny dots called *pixels*. The number of pixels displayed at one time is called the resolution. For example, the resolution 1024x768 means that 1024 horizontal pixels and 768 vertical pixels are used. The higher the resolution values, the more images the monitor can display at one time.

The color depth of the display determines how many possible colors are displayed. A higher color depth means more contrast between colors.
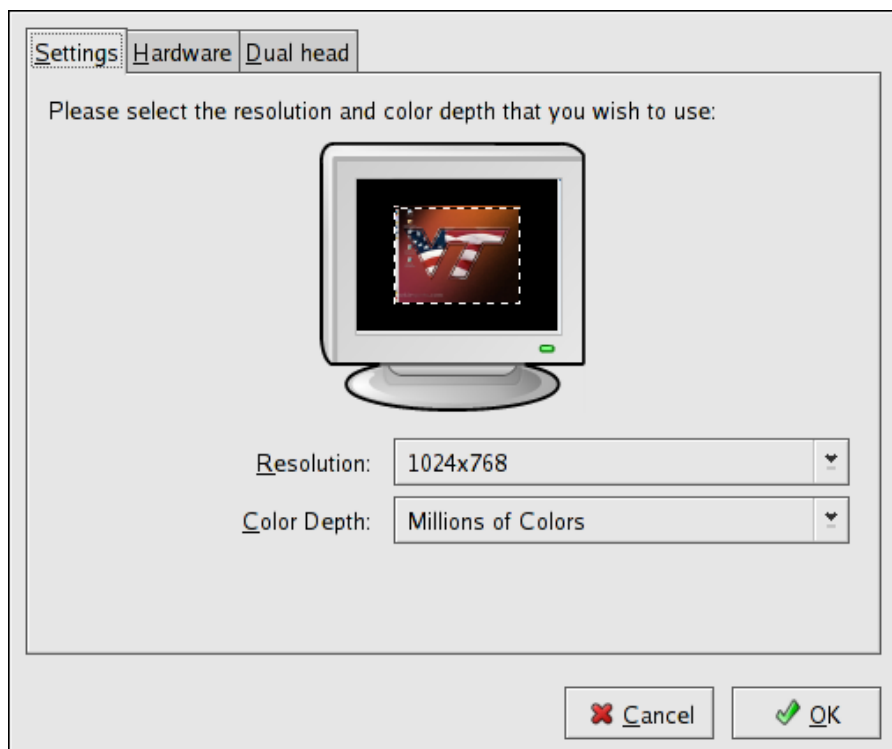
Figure 31.1. Display Settings

31.2. Display Hardware Settings

When the **X Configuration Tool** is started, it probes the monitor and video card. If the hardware is probed properly, the information for it is shown on the `Hardware` tab as shown in Figure 31.2, "Display Hardware Settings".
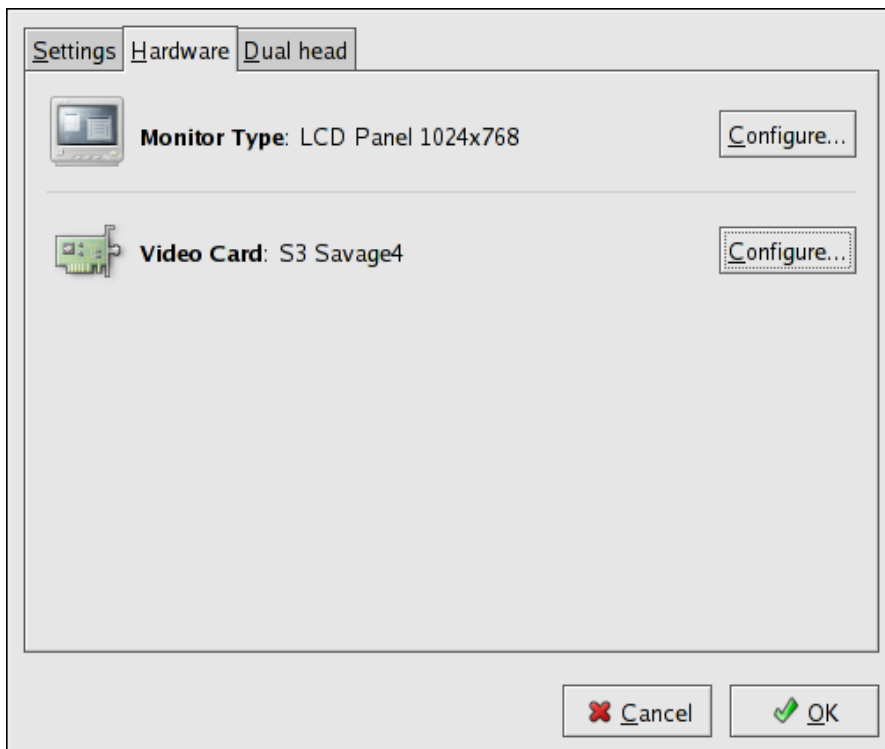
Figure 31.2. Display `Hardware Settings`

To change the monitor type or any of its settings, click the corresponding **`Configure`** button. To change the video card type or any of its settings, click the **`Configure`** button beside its settings.

31.3. Dual Head Display Settings

If multiple video cards are installed on the system, dual head monitor support is available and is configured via the **`Dual head`** tab, as shown in Figure 31.3, "Dual Head Display Settings".
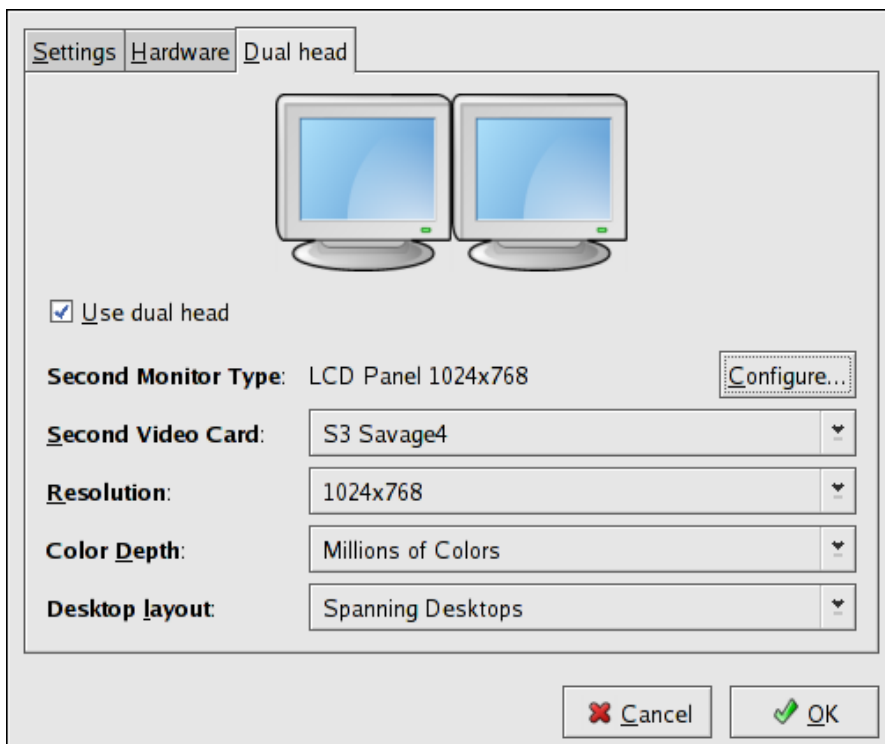
Figure 31.3. Dual Head Display Settings

To enable use of Dual head, check the **Use dual head** checkbox.

To configure the second monitor type, click the corresponding **Configure** button. You can also configure the other Dual head settings by using the corresponding drop-down list.

For the **Desktop layout** option, selecting **Spanning Desktops** allows both monitors to use an enlarged usable workspace. Selecting **Individual Desktops** shares the mouse and keyboard among the displays, but restricts windows to a single display.

Chapter 32. Users and Groups

The control of *users* and *groups* is a core element of Red Hat Enterprise Linux system administration.

Users can be either people (meaning accounts tied to physical users) or accounts which exist for specific applications to use.

Groups are logical expressions of organization, tying users together for a common purpose. Users within a group can read, write, or execute files owned by that group.

Each user and group has a unique numerical identification number called a *userid* (*UID*) and a *groupid* (*GID*), respectively.

A user who creates a file is also the owner and group owner of that file. The file is assigned separate read, write, and execute permissions for the owner, the group, and everyone else. The file owner can be changed only by the root user, and access permissions can be changed by both the root user and file owner.

32.1. User and Group Configuration

The **User Manager** allows you to view, modify, add, and delete local users and groups.

To use the **User Manager**, you must be running the X Window System, have root privileges, and have the **system-config-users** RPM package installed. To start the **User Manager** from the desktop, go to **Applications** (the main menu on the panel) => **Administration** => **Users & Groups**. You can also type the command **system-config-users** at a shell prompt (for example, in an XTerm or a GNOME terminal).
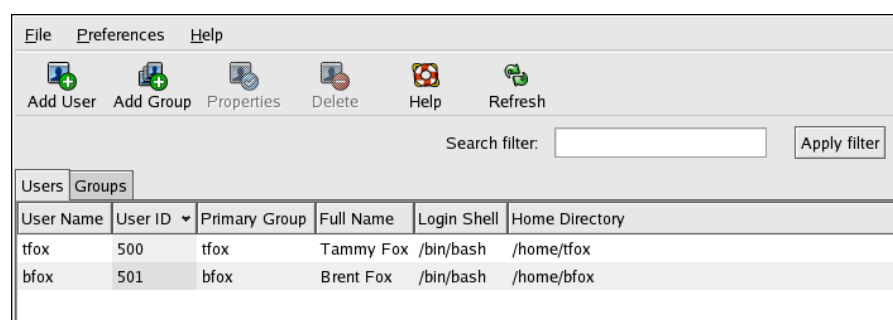
Figure 32.1. User Manager

To view a list of local users on the system, click the **Users** tab. To view a list of local groups on the system, click the **Groups** tab.

To find a specific user or group, type the first few letters of the name in the **Search filter** field. Press **Enter** or click the **Apply filter** button. The filtered list is displayed.

To sort the users or groups, click on the column name. The users or groups are sorted according to the value of that column.

Red Hat Enterprise Linux reserves user IDs below 500 for system users. By default, **User Manager** does not display system users. To view all users, including the system users, go to **Edit** => **Preferences** and uncheck **Hide system users and groups** from the dialog box.

32.1.1. Adding a New User

To add a new user, click the **Add User** button. A window as shown in appears. 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.

> **Tip**
>
> It is advisable to use a much longer password, as this makes it more difficult for an intruder to guess it and access the account without permission. It is also recommended that the password not be based on a dictionary term; use 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**.

If you select to create the home directory, default configuration files are copied from the **/etc/skel/** directory into the new home directory.

Red Hat Enterprise 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 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 above 500 is assigned to the new user. Because Red Hat Enterprise Linux reserves user IDs below 500 for system users, it is not advisable to manually assign user IDs 1-499.

Click **OK** to create the user.

Figure 32.2. New User

To configure more advanced user properties, such as password expiration, modify the user's properties after adding the user. Refer to Section 32.1.2, "Modifying User Properties" for more information.

32.1.2. 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 menu (or choose **File** => **Properties** from the pulldown menu). A window similar to Figure 32.3, "User Properties" appears.

Figure 32.3. User Properties

The **User Properties** window is divided into multiple tabbed pages:

- **User Data** — Shows the 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 **Local password is locked** to lock the user account and prevent the user from logging into the system.
- **Password Info** — Displays the date that the user's password last changed. To force the user to change passwords after a certain number of days, select **Enable password expiration** and enter a desired value in the **Days before change required:** field. The number of days before the user's password expires, the number of days before the user is warned to change passwords, and days before the account becomes inactive can also be changed.
- **Groups** — Allows you to view and configure the Primary Group of the user, as well as other groups that you want the user to be a member of.

32.1.3. Adding a New Group

To add a new user group, click the **Add Group** button. A window similar to Figure 32.4, "New Group" appears. 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. Note that Red Hat Enterprise Linux also reserves group IDs lower than 500 for system groups.
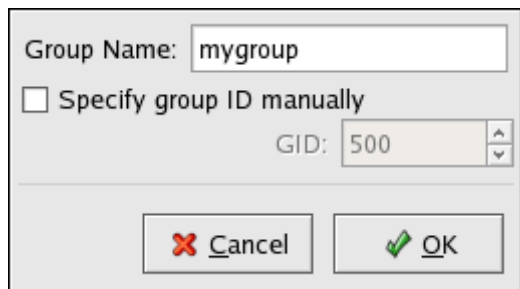
Figure 32.4. New Group

Click **OK** to create the group. The new group appears in the group list.

32.1.4. Modifying Group Properties

To view the properties of an existing group, select the group from the group list and click **Properties** from the menu (or choose **File** => **Properties** from the pulldown menu). A window similar to Figure 32.5, "Group Properties" appears.
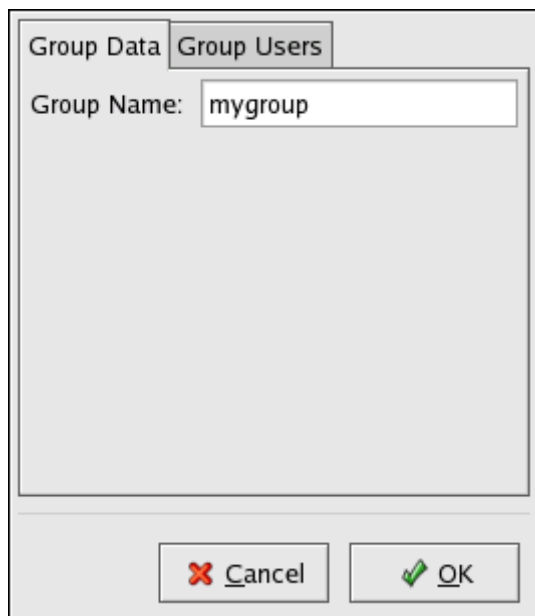
Figure 32.5. Group Properties

The **Group Users** tab displays which users are members of the group. Use this tab to add or remove users from the group. Click **OK** to save your changes.

32.2. User and Group Management Tools

Managing users and groups can be a tedious task; this is why Red Hat Enterprise Linux provides tools and conventions to make them easier to manage.

The easiest way to manage users and groups is through the graphical application, **User Manager** (**system-config-users**). For more information on **User Manager**, refer to Section 32.1, "User and Group Configuration".

The following command line tools can also be used to manage users and groups:

- **useradd**, **usermod**, and **userdel** — Industry-standard methods of adding, deleting and modifying user accounts

- **groupadd**, **groupmod**, and **groupdel** — Industry-standard methods of adding, deleting, and modifying user groups

- **gpasswd** — Industry-standard method of administering the **/etc/group** file

- **pwck**, **grpck** — Tools used for the verification of the password, group, and associated shadow files

- **pwconv**, **pwunconv** — Tools used for the conversion of passwords to shadow passwords and back to standard passwords

32.2.1. Command Line Configuration

If you prefer command line tools or do not have the X Window System installed, use this section to configure users and groups.

32.2.2. Adding a User

To add a user to the system:

1. Issue the **useradd** command to create a locked user account:

   > useradd *<username>*

2. Unlock the account by issuing the **passwd** command to assign a password and set password aging guidelines:

   > passwd *<username>*

Command line options for **useradd** are detailed in Table 32.1, "**useradd** Command Line Options".

Table 32.1. `useradd` Command Line Options

Option	Description
-c '*<comment>*'	*<comment>* can be replaced with any string. This option is generally used to specify the full name of a user.
-d *<home-dir>*	Home directory to be used instead of default **/home/*<username>*/**
-e *<date>*	Date for the account to be disabled in the format YYYY-MM-DD
-f *<days>*	Number of days after the password expires until the account is disabled. If **0** is specified, the account is disabled immediately after the password expires. If **-1** is specified, the account is not be disabled after the password expires.
-g *<group-name>*	Group name or group number for the user's default group. The group must exist prior to being specified here.
-G *<group-list>*	List of additional (other than default) group names or group numbers, separated by commas, of which the user is a member. The groups must exist prior to being specified here.
-m	Create the home directory if it does not exist.
-M	Do not create the home directory.
-n	Do not create a user private group for the user.
-r	Create a system account with a UID less than 500 and without a home directory
-p *<password>*	The password encrypted with **crypt**
-s	User's login shell, which defaults to **/bin/bash**
-u *<uid>*	User ID for the user, which must be unique and greater than 499

32.2.3. Adding a Group

To add a group to the system, use the command **groupadd**:

```
groupadd <group-name>
```

Command line options for **groupadd** are detailed in Table 32.2, "**groupadd** Command Line Options".

Table 32.2. `groupadd` Command Line Options

Option	Description
-g *<gid>*	Group ID for the group, which must be unique and greater than 499
-r	Create a system group with a GID less than 500
-f	When used with **-g** *<gid>* and *<gid>* already exists, **groupadd** will choose another unique *<gid>* for the group.

32.2.4. Password Aging

For security reasons, it is advisable to require users to change their passwords periodically. This can be done when adding or editing a user on the **Password Info** tab of the **User Manager**.

To configure password expiration for a user from a shell prompt, use the **chage** command, followed by

an option from Table 32.3, "**chage** Command Line Options", followed by the username of the user.

> ⭐ **Important**
>
> Shadow passwords must be enabled to use the **chage** command.

Table 32.3. chage Command Line Options

Option	Description
-m *<days>*	Specifies the minimum number of days between which the user must change passwords. If the value is 0, the password does not expire.
-M *<days>*	Specifies the maximum number of days for which the password is valid. When the number of days specified by this option plus the number of days specified with the **-d** option is less than the current day, the user must change passwords before using the account.
-d *<days>*	Specifies the number of days since January 1, 1970 the password was changed
-I *<days>*	Specifies the number of inactive days after the password expiration before locking the account. If the value is 0, the account is not locked after the password expires.
-E *<date>*	Specifies the date on which the account is locked, in the format YYYY-MM-DD. Instead of the date, the number of days since January 1, 1970 can also be used.
-W *<days>*	Specifies the number of days before the password expiration date to warn the user.

> 💬 **Tip**
>
> If the **chage** command is followed directly by a username (with no options), it displays the current password aging values and allows them to be changed.

You can configure a password to expire the first time a user logs in. This forces users to change passwords the first time they log in.

> 💬 **Note**
>
> This process will not work if the user logs in using the SSH protocol.

1. *Lock the user password* — If the user does not exist, use the **useradd** command to create the user account, but do not give it a password so that it remains locked.

   If the password is already enabled, lock it with the command:

   ```
   usermod -L username
   ```

2. *Force immediate password expiration* — Type the following command:

```
chage -d 0 username
```

This command sets the value for the date the password was last changed to the epoch (January 1, 1970). This value forces immediate password expiration no matter what password aging policy, if any, is in place.

3. *Unlock the account* — There are two common approaches to this step. The administrator can assign an initial password or assign a null password.

> ⚠️ **Warning**
>
> Do not use the **passwd** command to set the password as it disables the immediate password expiration just configured.

To assign an initial password, use the following steps:

- ▷ Start the command line Python interpreter with the **python** command. It displays the following:

```
Python 2.4.3 (#1, Jul 21 2006, 08:46:09)
[GCC 4.1.1 20060718 (Red Hat 4.1.1-9)] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

- ▷ At the prompt, type the following commands. Replace *<password>* with the password to encrypt and *<salt>* with a random combination of at least 2 of the following: any alphanumeric character, the slash (/) character or a dot (.):

```
import crypt; print crypt.crypt("<password>","<salt>")
```

The output is the encrypted password, similar to **'12CsGd8FRcMSM'**.

- ▷ Press **Ctrl-D** to exit the Python interpreter.
- ▷ At the shell, enter the following command (replacing *<encrypted-password>* with the encrypted output of the Python interpreter):

```
usermod -p "<encrypted-password>" <username>
```

Alternatively, you can assign a null password instead of an initial password. To do this, use the following command:

```
usermod -p "" username
```

> ⚠️ **Warning**
>
> Using a null password, while convenient, is a highly unsecure practice, as any third party can log in first an access the system using the unsecure username. Always make sure that the user is ready to log in before unlocking an account with a null password.

In either case, upon initial log in, the user is prompted for a new password.

32.2.5. Explaining the Process

The following steps illustrate what happens if the command **useradd juan** is issued on a system that

has shadow passwords enabled:

1. A new line for **juan** is created in **/etc/passwd**. The line has the following characteristics:
   - It begins with the username **juan**.
   - There is an **x** for the password field indicating that the system is using shadow passwords.
   - A UID greater than 499 is created. (Under Red Hat Enterprise Linux, UIDs and GIDs below 500 are reserved for system use.)
   - A GID greater than 499 is created.
   - The optional GECOS information is left blank.
   - The home directory for **juan** is set to **/home/juan/**.
   - The default shell is set to **/bin/bash**.

2. A new line for **juan** is created in **/etc/shadow**. The line has the following characteristics:
   - It begins with the username **juan**.
   - Two exclamation points (**!!**) appear in the password field of the **/etc/shadow** file, which locks the account.

   > **Note**
   >
   > If an encrypted password is passed using the **-p** flag, it is placed in the **/etc/shadow** file on the new line for the user.

   - The password is set to never expire.

3. A new line for a group named **juan** is created in **/etc/group**. A group with the same name as a user is called a *user private group*. For more information on user private groups, refer to Section 32.1.1, "Adding a New User".

   The line created in **/etc/group** has the following characteristics:
   - It begins with the group name **juan**.
   - An **x** appears in the password field indicating that the system is using shadow group passwords.
   - The GID matches the one listed for user **juan** in **/etc/passwd**.

4. A new line for a group named **juan** is created in **/etc/gshadow**. The line has the following characteristics:
   - It begins with the group name **juan**.
   - An exclamation point (**!**) appears in the password field of the **/etc/gshadow** file, which locks the group.
   - All other fields are blank.

5. A directory for user **juan** is created in the **/home/** directory. This directory is owned by user **juan** and group **juan**. However, it has read, write, and execute privileges *only* for the user **juan**. All other permissions are denied.

6. The files within the **/etc/skel/** directory (which contain default user settings) are copied into the new **/home/juan/** directory.

At this point, a locked account called **juan** exists on the system. To activate it, the administrator must next assign a password to the account using the **passwd** command and, optionally, set password aging guidelines.

32.3. Standard Users

Table 32.4, "Standard Users" lists the standard users configured in the `/etc/passwd` file by an `Everything` installation. The groupid (GID) in this table is the *primary group* for the user. See Section 32.4, "Standard Groups" for a listing of standard groups.

Table 32.4. Standard Users

User	UID	GID	Home Directory	Shell
root	0	0	**/root**	**/bin/bash**
bin	1	1	**/bin**	**/sbin/nologin**
daemon	2	2	**/sbin**	**/sbin/nologin**
adm	3	4	**/var/adm**	**/sbin/nologin**
lp	4	7	**/var/spool/lpd**	**/sbin/nologin**
sync	5	0	**/sbin**	**/bin/sync**
shutdown	6	0	**/sbin**	**/sbin/shutdown**
halt	7	0	**/sbin**	**/sbin/halt**
mail	8	12	**/var/spool/mail**	**/sbin/nologin**
news	9	13	**/etc/news**	
uucp	10	14	**/var/spool/uucp**	**/sbin/nologin**
operator	11	0	**/root**	**/sbin/nologin**
games	12	100	**/usr/games**	**/sbin/nologin**
gopher	13	30	**/var/gopher**	**/sbin/nologin**
ftp	14	50	**/var/ftp**	**/sbin/nologin**
nobody	99	99	**/**	**/sbin/nologin**
rpm	37	37	**/var/lib/rpm**	**/sbin/nologin**
vcsa	69	69	**/dev**	**/sbin/nologin**
dbus	81	81	/	**/sbin/nologin**
ntp	38	38	**/etc/ntp**	**/sbin/nologin**
canna	39	39	**/var/lib/canna**	**/sbin/nologin**
nscd	28	28	/	**/sbin/nologin**
rpc	32	32	/	**/sbin/nologin**
postfix	89	89	**/var/spool/postfix**	**/sbin/nologin**
mailman	41	41	**/var/mailman**	**/sbin/nologin**
named	25	25	**/var/named**	**/bin/false**
amanda	33	6	**var/lib/amanda/**	**/bin/bash**
postgres	26	26	**/var/lib/pgsql**	**/bin/bash**
exim	93	93	**/var/spool/exim**	**/sbin/nologin**
sshd	74	74	**/var/empty/sshd**	**/sbin/nologin**
rpcuser	29	29	**/var/lib/nfs**	**/sbin/nologin**
nsfnobody	65534	65534	**/var/lib/nfs**	**/sbin/nologin**
pvm	24	24	**/usr/share/pvm3**	**/bin/bash**
apache	48	48	**/var/www**	**/sbin/nologin**
xfs	43	43	**/etc/X11/fs**	**/sbin/nologin**
gdm	42	42	**/var/gdm**	**/sbin/nologin**
htt	100	101	**/usr/lib/im**	**/sbin/nologin**
mysql	27	27	**/var/lib/mysql**	**/bin/bash**
webalizer	67	67	**/var/www/usage**	**/sbin/nologin**

mailnull	47	47	**/var/spool/mqueue**	**/sbin/nologin**
smmsp	51	51	**/var/spool/mqueue**	**/sbin/nologin**
squid	23	23	**/var/spool/squid**	**/sbin/nologin**
ldap	55	55	**/var/lib/ldap**	**/bin/false**
netdump	34	34	**/var/crash**	**/bin/bash**
pcap	77	77	**/var/arpwatch**	**/sbin/nologin**
radiusd	95	95	**/**	**/bin/false**
radvd	75	75	**/**	**/sbin/nologin**
quagga	92	92	**/var/run/quagga**	**/sbin/login**
wnn	49	49	**/var/lib/wnn**	**/sbin/nologin**
dovecot	97	97	**/usr/libexec/dovecot**	**/sbin/nologin**

32.4. Standard Groups

Table 32.5, "Standard Groups" lists the standard groups configured by an **Everything** installation. Groups are stored in the **/etc/group** file.

Table 32.5. Standard Groups

Group	GID	Members
root	0	root
bin	1	root, bin, daemon
daemon	2	root, bin, daemon
sys	3	root, bin, adm
adm	4	root, adm, daemon
tty	5	
disk	6	root
lp	7	daemon, lp
mem	8	
kmem	9	
wheel	10	root
mail	12	mail, postfix, exim
news	13	news
uucp	14	uucp
man	15	
games	20	
gopher	30	
dip	40	
ftp	50	
lock	54	
nobody	99	
users	100	
rpm	37	
utmp	22	
floppy	19	
vcsa	69	
dbus	81	
ntp	38	
canna	39	
nscd	28	
rpc	32	
postdrop	90	
postfix	89	
mailman	41	
exim	93	
named	25	
postgres	26	
sshd	74	
rpcuser	29	
nfsnobody	65534	
pvm	24	

apache	48
xfs	43
gdm	42
htt	101
mysql	27
webalizer	67
mailnull	47
smmsp	51
squid	23
ldap	55
netdump	34
pcap	77
quaggavt	102
quagga	92
radvd	75
slocate	21
wnn	49
dovecot	97
radiusd	95

32.5. User Private Groups

Red Hat Enterprise Linux uses a *user private group* (*UPG*) scheme, which makes UNIX groups easier to manage.

A UPG is created whenever a new user is added to the system. A UPG has the same name as the user for which it was created and that user is the only member of the UPG.

UPGs make it safe to set default permissions for a newly created file or directory, allowing both the user and *the group of that user* to make modifications to the file or directory.

The setting which determines what permissions are applied to a newly created file or directory is called a *umask* and is configured in the **/etc/bashrc** file. Traditionally on UNIX systems, the **umask** is set to **022**, which allows only the user who created the file or directory to make modifications. Under this scheme, all other users, *including members of the creator's group*, are not allowed to make any modifications. However, under the UPG scheme, this "group protection" is not necessary since every user has their own private group.

32.5.1. Group Directories

Many IT organizations like to create a group for each major project and then assign people to the group if they need to access that project's files. Using this traditional scheme, managing files has been difficult; when someone creates a file, it is associated with the primary group to which they belong. When a single person works on multiple projects, it is difficult to associate the right files with the right group. Using the UPG scheme, however, groups are automatically assigned to files created within a directory with the *setgid* bit set. The setgid bit makes managing group projects that share a common directory very simple because any files a user creates within the directory are owned by the group which owns the directory.

Let us say, for example, that a group of people need to work on files in the **/usr/share/emacs/site-**

lisp/ directory. Some people are trusted to modify the directory, but certainly not everyone is trusted. First create an **emacs** group, as in the following command:

```
/usr/sbin/groupadd emacs
```

To associate the contents of the directory with the **emacs** group, type:

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

Now, it is possible to add the proper users to the group with the **gpasswd** command:

```
/usr/bin/gpasswd -a <username> emacs
```

To allow users to create files within the directory, use the following command:

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

When a user creates a new file, it is assigned the group of the user's default private group. Next, set the setgid bit, which assigns everything created in the directory the same group permission as the directory itself (**emacs**). Use the following command:

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

At this point, because the default umask of each user is 002, all members of the **emacs** group can create and edit files in the **/usr/share/emacs/site-lisp/** directory without the administrator having to change file permissions every time users write new files.

32.6. Shadow Passwords

In multiuser environments it is very important to use *shadow passwords* (provided by the **shadow-utils** package). Doing so enhances the security of system authentication files. For this reason, the installation program enables shadow passwords by default.

The following lists the advantages pf shadow passwords have over the traditional way of storing passwords on UNIX-based systems:

- Improves system security by moving encrypted password hashes from the world-readable **/etc/passwd** file to **/etc/shadow**, which is readable only by the root user.
- Stores information about password aging.
- Allows the use the **/etc/login.defs** file to enforce security policies.

Most utilities provided by the **shadow-utils** package work properly whether or not shadow passwords are enabled. However, since password aging information is stored exclusively in the **/etc/shadow** file, any commands which create or modify password aging information do not work.

The following is a list of commands which do not work without first enabling shadow passwords:

- **chage**
- **gpasswd**
- **/usr/sbin/usermod -e** or **-f** options
- **/usr/sbin/useradd -e** or **-f** options

32.7. Additional Resources

For more information about users and groups, and tools to manage them, refer to the following resources.

32.7.1. Installed Documentation

▷ Related man pages — There are a number of man pages for the various applications and configuration files involved with managing users and groups. Some of the more important man pages have been listed here:

User and Group Administrative Applications

- `man chage` — A command to modify password aging policies and account expiration.
- `man gpasswd` — A command to administer the `/etc/group` file.
- `man groupadd` — A command to add groups.
- `man grpck` — A command to verify the `/etc/group` file.
- `man groupdel` — A command to remove groups.
- `man groupmod` — A command to modify group membership.
- `man pwck` — A command to verify the `/etc/passwd` and `/etc/shadow` files.
- `man pwconv` — A tool to convert standard passwords to shadow passwords.
- `man pwunconv` — A tool to convert shadow passwords to standard passwords.
- `man useradd` — A command to add users.
- `man userdel` — A command to remove users.
- `man usermod` — A command to modify users.

Configuration Files

- `man 5 group` — The file containing group information for the system.
- `man 5 passwd` — The file containing user information for the system.
- `man 5 shadow` — The file containing passwords and account expiration information for the system.

Chapter 33. Printer Configuration

Printer Configuration Tool allows users to configure a printer. This tool helps maintain the printer configuration file, print spool directories, print filters, and printer classes.

Red Hat Enterprise Linux 4 uses the Common Unix Printing System (CUPS). If a system was upgraded from a previous Red Hat Enterprise Linux version that used CUPS, the upgrade process preserves the configured queues.

Using **Printer Configuration Tool** requires root privileges. To start the application, select **Applications** (the main menu on the panel) => **Administration** => **Printing**, or type the command **system-config-printer** at a shell prompt.
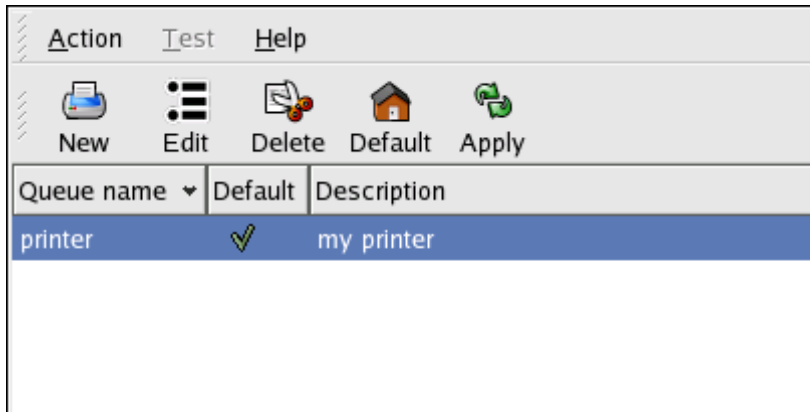
Figure 33.1. Printer Configuration Tool

The following types of print queues can be configured:

- **AppSocket/HP JetDirect** — a printer connected directly to the network through HP JetDirect or Appsocket interface instead of a computer.
- **Internet Printing Protocol (IPP)** — a printer that can be accessed over a TCP/IP network via the Internet Printing Protocol (for example, a printer attached to another Red Hat Enterprise Linux system running CUPS on the network).
- **LPD/LPR Host or Printer** — 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 Enterprise Linux system running LPD on the network).
- **Networked Windows (SMB)** — a printer attached to a different system which is sharing a printer over an SMB network (for example, a printer attached to a Microsoft Windows™ machine).
- **Networked JetDirect** — a printer connected directly to the network through HP JetDirect instead of a computer.

> **Important**
>
> If you add a new print queue or modify an existing one, you must apply the changes for them to take effect.

Clicking the **Apply** button prompts the printer daemon to restart with the changes you have configured.

Clicking the **Revert** button discards unapplied changes.

33.1. Adding a Local Printer

To add a local printer, such as one attached through a parallel port or USB port on your computer, click the **New Printer** button in the main **Printer Configuration Tool** window to display the window in Figure 33.2, "Adding a Printer".
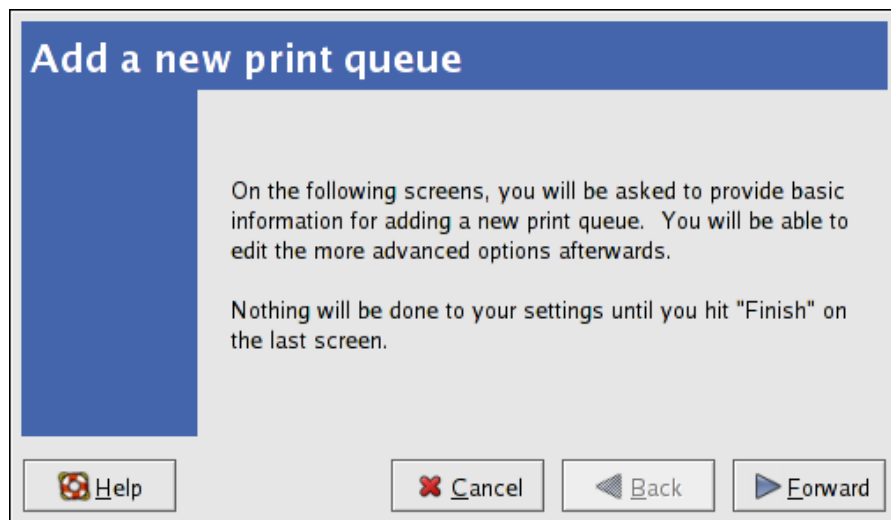
Figure 33.2. Adding a Printer

Click **Forward** to proceed.

Enter a unique name for the printer in the **Printer Name** field. The printer name can contain letters, numbers, dashes (-), and underscores (_); it *must not* contain any spaces.

You can also use the **Description** and **Location** fields to further distinguish this printer from others that may be configured on your system. Both of these fields are optional, and may contain spaces.

Click **Forward** to open the **New Printer** dialogue (refer to Figure 33.3, "Adding a Local Printer"). If the printer has been automatically detected, the printer model appears in **Select Connection**. Select the printer model and click **Forward** to continue.

If the device does not automatically appear, select the device to which the printer is connected (such as **LPT #1** or **Serial Port #1**) in **Select Connection**.
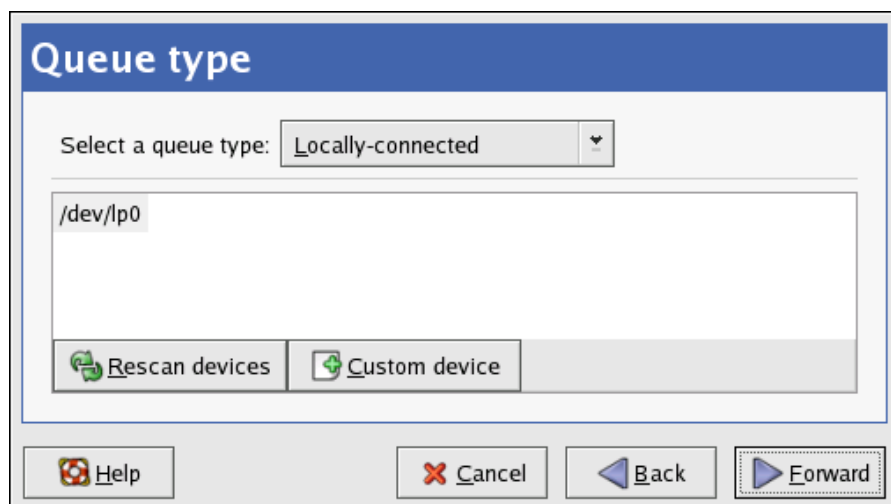
Figure 33.3. Adding a Local Printer

Next, select the printer type. Refer to Section 33.5, "Selecting the Printer Model and Finishing" for details.

33.2. Adding an IPP Printer

An IPP printer is a printer attached to a different system on the same TCP/IP network. The system this printer is attached to may either be running CUPS or simply configured to use IPP.

If a firewall is enabled on the printer server, then the firewall should be configured to allow send / receive connections on the incoming UDP port 631. If a firewall is enabled on on the client (the system sending the print request) then the firewall must be allowed to accept and create connections through port 631.

You can add a networked IPP printer by clicking the **New Printer** button in the main **Printer Configuration Tool** window to display the window in Figure 33.2, "Adding a Printer". Enter the **Printer Name** (printer names cannot contain spaces and may contain letters, numbers, dashes (-), and underscores (_)), **Description**, and **Location** to distinguish this printer from others that you may configure on your system. Click **Forward** to proceed.

In the window shown in Figure 33.4, "Adding an IPP Printer", enter the hostname of the IPP printer in the **Hostname** field as well as a unique name for the printer in the **Printername** field.
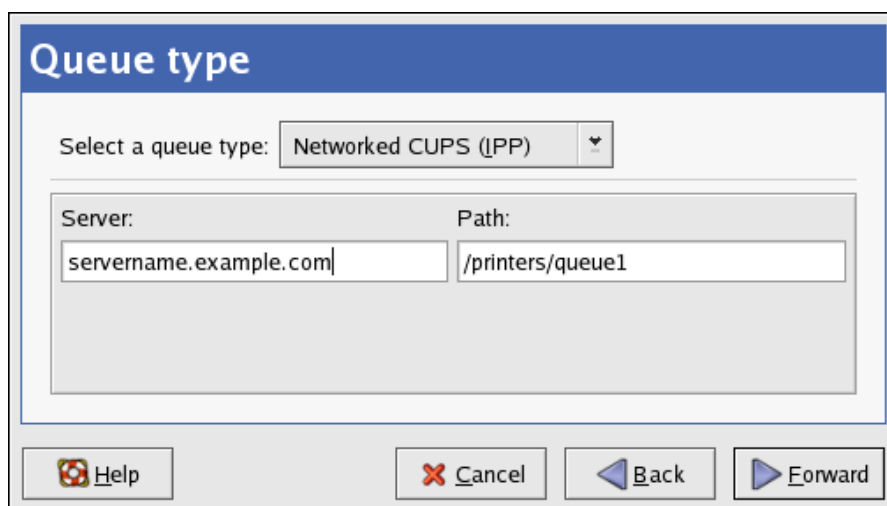
Figure 33.4. Adding an IPP Printer

Click **Forward** to continue.

Next, select the printer type. Refer to Section 33.5, "Selecting the Printer Model and Finishing" for details.

33.3. Adding a Samba (SMB) Printer

You can add a Samba (SMB) based printer share by clicking the **New Printer** button in the main **Printer Configuration Tool** window to display the window in Figure 33.2, "Adding a Printer". Enter a unique name for the printer in the **Printer Name** field. The printer name can contain letters, numbers, dashes (-), and underscores (_); it *must not* contain any spaces.

You can also use the **Description** and **Location** fields to further distinguish this printer from others that may be configured on your system. Both of these fields are optional, and may contain spaces.
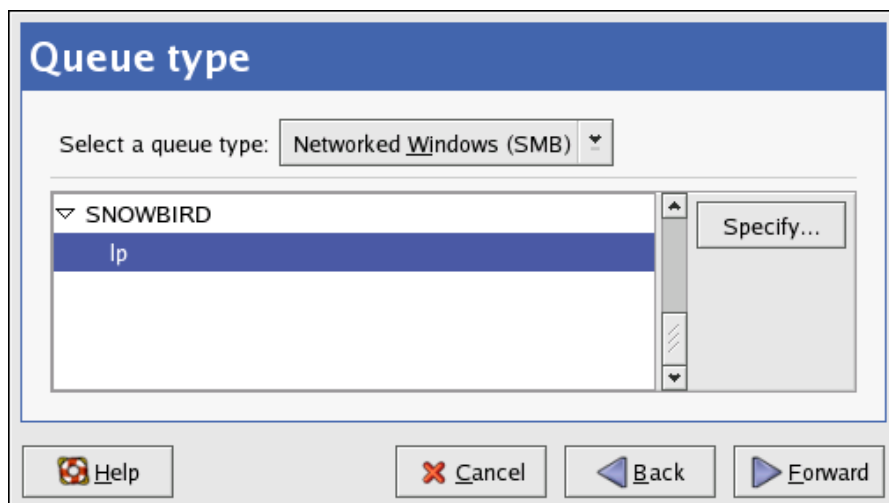
Figure 33.5. Adding a SMB Printer

As shown in Figure 33.5, "Adding a SMB Printer", available SMB shares are automatically detected and listed in the **Share** column. Click the arrow ( ) beside a Workgroup to expand it. From the expanded list, select a printer.

If the printer you are looking for does not appear in the list, enter the SMB address in the **smb://** field. Use the format **computer name/printer share**. In Figure 33.5, "Adding a SMB Printer", the **computer name** is **dellbox**, while the **printer share** is **r2**.

In the **Username** field, enter the username to access the printer. This user must exist on the SMB 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.

Enter the **Password** (if required) for the user specified in the **Username** field.

You can then test the connection by clicking **Verify**. Upon successful verification, a dialog box appears confirming printer share accessibility.

Next, select the printer type. Refer to Section 33.5, "Selecting the Printer Model and Finishing" for details.

> ⚠️ **Warning**
>
> Samba printer usernames and passwords are stored in the printer server as unencrypted files readable by root and lpd. Thus, other users that have root access to the printer server can view the username and password you use to access the Samba printer.
>
> As such, when you choose a username and password to access a Samba printer, it is advisable that you choose a password that is different from what you use to access your local Red Hat Enterprise Linux system.
>
> If there are files shared on the Samba print server, it is recommended that they also use a password different from what is used by the print queue.

33.4. Adding a JetDirect Printer

To add a JetDirect or AppSocket connected printer share, click the **New Printer** button in the main **Printer Configuration Tool** window to display the window in Figure 33.2, "Adding a Printer". Enter a

unique name for the printer in the **Printer Name** field. The printer name can contain letters, numbers, dashes (-), and underscores (_); it *must not* contain any spaces.

You can also use the **Description** and **Location** fields to further distinguish this printer from others that may be configured on your system. Both of these fields are optional, and may contain spaces.
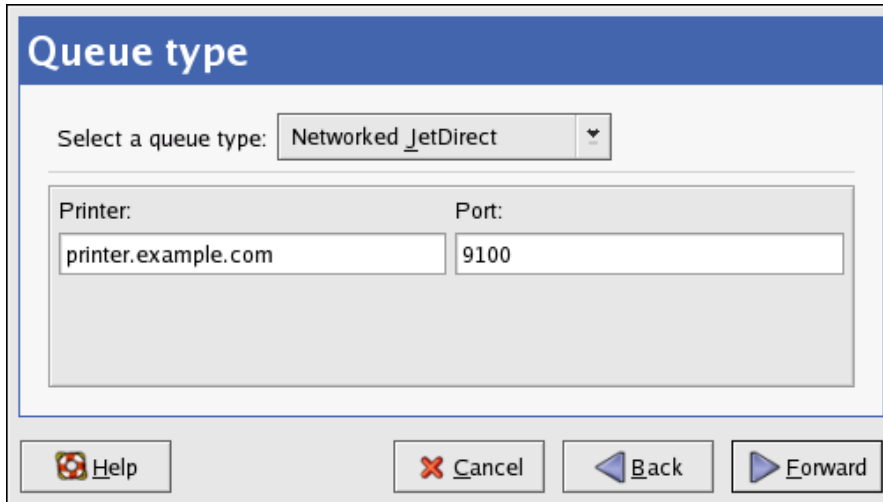
Figure 33.6. Adding a JetDirect Printer

Click **Forward** to continue.

Text fields for the following options appear:

- **Hostname** — The hostname or IP address of the JetDirect printer.
- **Port Number** — The port on the JetDirect printer that is listening for print jobs. The default port is 9100.

Next, select the printer type. Refer to for details.

33.5. Selecting the Printer Model and Finishing

Once you have properly selected a printer queue type, you can choose either option:

- Select a Printer from database - If you select this option, choose the make of your printer from the list of **Makes**. If your printer make is not listed, choose **Generic**.
- Provide PPD file - A PostScript Printer Description (PPD) file may also be provided with your printer. This file is normally provided by the manufacturer. If you are provided with a PPD file, you can choose this option and use the browser bar below the option description to select the PPD file.

Refer to .
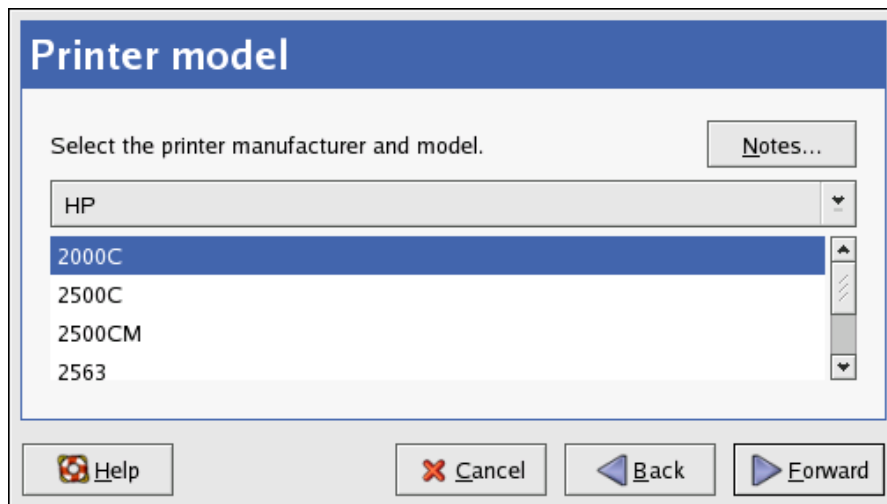
Figure 33.7. Selecting a Printer Model

After choosing an option, click **Forward** to continue. Figure 33.7, "Selecting a Printer Model" appears. You now have to choose the corresponding model and driver for the printer.

The recommended printed driver is automatically selected based on the printer model you chose. 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 a printer driver to process the data that is sent to the printer.

If you have a PPD file for the device (usually provided by the manufacturer), you can select it by choosing **Provide PPD file**. You can then browse the filesystem for the PPD file by clicking **Browse**.

33.5.1. 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.

After applying the changes, print a test page to ensure the configuration is correct. Refer to Section 33.6, "Printing a Test Page" for details.

33.6. Printing a Test Page

After you have configured your printer, you should print a test page to make sure the printer is functioning properly. To print a test page, select the printer that you want to try out from the printer list, then click **Print Test Page** from the printer's **Settings** tab.

If you change the print driver or modify the driver options, you should print a test page to test the different configuration.

33.7. Modifying Existing Printers

To delete an existing printer, select the printer and click the **Delete** button on the toolbar. The printer is removed from the printer list once you confirm deletion of the printer configuration.

To set the default printer, select the printer from the printer list and click the **Make Default Printer** button in the **Settings** tab.

33.7.1. The `Settings` Tab

To change printer driver configuration, click the corresponding name in the **Printer** list and click the **Settings** tab.

You can modify printer settings such as make and model, make a printer the default, print a test page, change the device location (URI), and more.

Figure 33.8. Settings Tab

33.7.2. The `Policies` Tab

To change settings in print output, click the **Policies** tab.

For example, to create a *banner page* (a page that describes aspects of the print job such as the originating printer, the username from the which the job originated, and the security status of the document being printed) click the **Starting Banner** or **Ending Banner** drop-menu and choose the option that best describes the nature of the print jobs (such as **topsecret**, **classified**, or **confidential**).

Figure 33.9. Policies Tab

You can also configure the **Error Policy** of the printer, by choosing an option from the drop-down menu. You can choose to abort the print job, retry, or stop it.

33.7.3. The `Access Control` Tab

You can change user-level access to the configured printer by clicking the **Access Control** tab.

Add users using the text box and click the **Add** button beside it. You can then choose to only allow use of the printer to that subset of users or deny use to those users.

Figure 33.10. Access Control Tab

33.7.4. The `Printer and Job Options` Tab

The **Printer Options** tab contains various configuration options for the printer media and output.

Figure 33.11. Printer Options Tab

- **Page Size** — Allows the paper size to be selected. The options include US Letter, US Legal, A3, and A4

- **`Media Source`** — set to **`Automatic`** by default. Change this option to use paper from a different tray.
- **`Media Type`** — Allows you to change paper type. Options include: Plain, thick, bond, and transparency.
- **`Resolution`** — Configure the quality and detail of the printout (default is 300 dots per inch (dpi).
- **`Toner Saving`** — Choose whether the printer uses less toner to conserve resources.

You can also configure printer job options using the **`Job Options`** tab. Use the drop-menu and choose the job options you wish to use, such as **`Landscape`** modes (horizontal or vertical printout), **`copies`**, or **`scaling`** (increase or decrease the size of the printable area, which can be used to fit an oversize print area onto a smaller physical sheet of print medium).

33.8. Managing Print Jobs

When you send a print job to the printer daemon, such as printing a 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 the job number, and more.

During the printing process, the Printer Status icon appears in the **Notification Area** on the panel. To check the status of a print job, double click the Printer Status, which displays a window similar to Figure 33.12, "GNOME Print Status".

Figure 33.12. GNOME Print Status

To cancel a specific print job listed in the **GNOME Print Status**, select it from the list and select **Edit** => **`Cancel Documents`** from the pulldown menu.

To view the list of print jobs in the print spool from a shell prompt, type the command **`lpq`**. The last few lines look similar to the following:

Example 33.1. Example of `lpq` output

```
Rank    Owner/ID                Class   Job Files        Size Time
active  user@localhost+902      A       902 sample.txt   2050 01:20:46
```

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 33.1, "Example of `lpq` output". 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`** prints the text file **`sample.txt`**. The print filter determines what type of file it is and converts it into a format the printer can understand.

33.9. Additional Resources

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

33.9.1. Installed Documentation

- **map lpr** — The manual page for the **lpr** command that allows you to print files from the command line.
- **man lprm** — The manual page for the command line utility to remove print jobs from the print queue.
- **man mpage** — The manual page for 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.

33.9.2. Useful Websites

- http://www.linuxprinting.org — *GNU/Linux Printing* contains a large amount of information about printing in Linux.
- http://www.cups.org/ — Documentation, FAQs, and newsgroups about CUPS.

Chapter 34. 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 Enterprise Linux is pre-configured 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 Enterprise Linux comes with several automated tasks utilities: **cron**, **at**, and **batch**.

34.1. 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 schedule one-time tasks, refer to Section 34.2, "At and Batch".

To use the cron service, the **vixie-cron** RPM package must be 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**.

34.1.1. 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 **SHELL** variable tells the system which shell environment to use (in this example the bash shell), while 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 is not 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 following 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 or feb)
- **dayofweek** — any integer from 0 to 7, where 0 or 7 represents Sunday (or the short name of the week such as sun or mon)

- **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)

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 shown in the **/etc/crontab** file, the **run-parts** script executes 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 directories should be shell scripts.

If a cron task is required 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 for examples.

Example 34.1. Crontab 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
```

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.

34.1.2. 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 usernames 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, users listed in **cron.deny** are not allowed to use cron.

34.1.3. 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 19, *Controlling Access to Services* for details on starting the cron service automatically at boot time.

34.2. At and Batch

While cron is used to schedule recurring tasks, the **at** command is used to schedule a one-time task at a specific time and 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**, the **at** RPM package must be 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**.

34.2.1. 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 4:00 a.m. If the time is already past, it is executed at the specified time the next day.
- midnight — Specifies 12:00 a.m.
- noon — Specifies 12:00 p.m.
- teatime — Specifies 4:00 p.m.
- 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 five days from now.

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**. Multiple commands 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 <u>Section 34.2.3, "Viewing Pending Jobs"</u> for more information.

Usage of the **at** command can be restricted. For more information, refer to <u>Section 34.2.5, "Controlling Access to At and Batch"</u> for details.

34.2.2. 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**. Multiple commands 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 <u>Section 34.2.3, "Viewing Pending Jobs"</u> for more information.

Usage of the **batch** command can be restricted. For more information, refer to <u>Section 34.2.5, "Controlling Access to At and Batch"</u> for details.

34.2.3. Viewing Pending Jobs

To view pending **at** and **batch** jobs, use the **atq** command. The **atq** command displays a list of pending jobs, with each job on a line. Each line follows the job number, date, hour, job class, and username format. Users can only view their own jobs. If the root user executes the **atq** command, all jobs for all users are displayed.

34.2.4. Additional Command Line Options

Additional command line options for **at** and **batch** include:

Table 34.1. at and batch Command Line Options

Option	Description
-f	Read the commands or shell script from a file instead of specifying them at the prompt.
-m	Send email to the user when the job has been completed.
-v	Display the time that the job is executed.

34.2.5. 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, users listed in **at.deny** are not allowed to use **at** or **batch**.

34.2.6. 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 19, *Controlling Access to Services* for details on starting the cron service automatically at boot time.

34.3. Additional Resources

To learn more about configuring automated tasks, refer to the following resources.

34.3.1. 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.
- **at** man page — description of **at** and **batch** and their command line options.

Chapter 35. 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 when trying to troubleshoot a problem with the system such as trying to load a kernel driver or when 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.

35.1. 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.

You may notice 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.

35.2. 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, root privileges are required 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 **Applications** (the main menu on the panel) => **System Tools** => **System Logs**, or type the command **system-logviewer** at a shell prompt.

The application only displays log files that exist; thus, the list might differ from the one shown in Figure 35.1, "Log Viewer".

To filter the contents of the log file for keywords, type the keyword(s) in the **Filter for** text field, and click **Filter**. Click **Reset** to reset the contents.
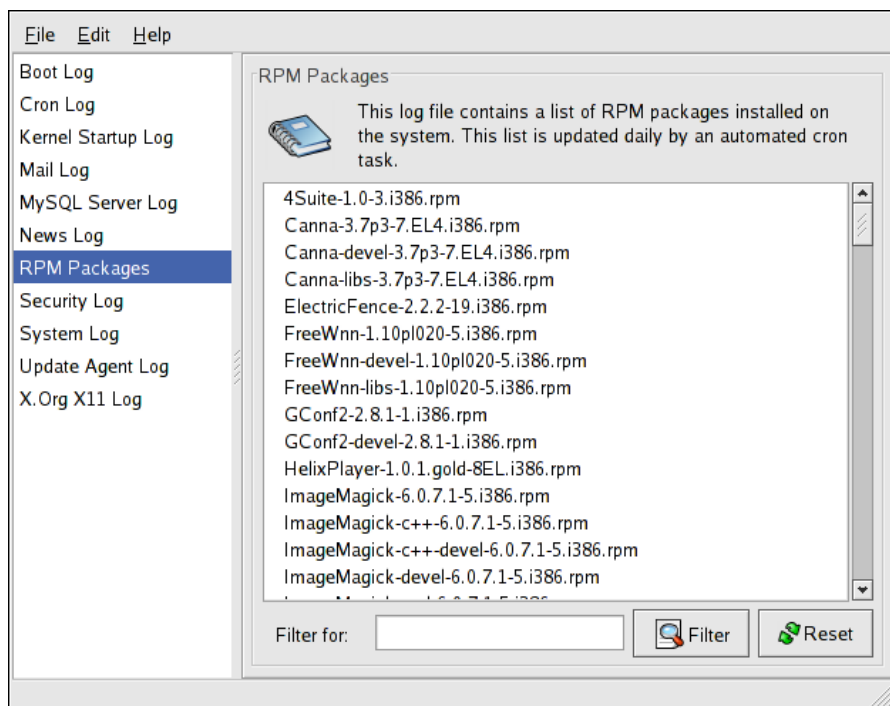
Figure 35.1. Log Viewer

By default, the currently viewable log file is refreshed every 30 seconds. To change the refresh rate, select **Edit** => `Preferences` from the pulldown menu. The window shown in Figure 35.2, "Log File Locations" appears. 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`.

On the `Log Files` tab in the Preferences, the log file locations can be modified. Select the log file from the list, and click the `Edit` 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.
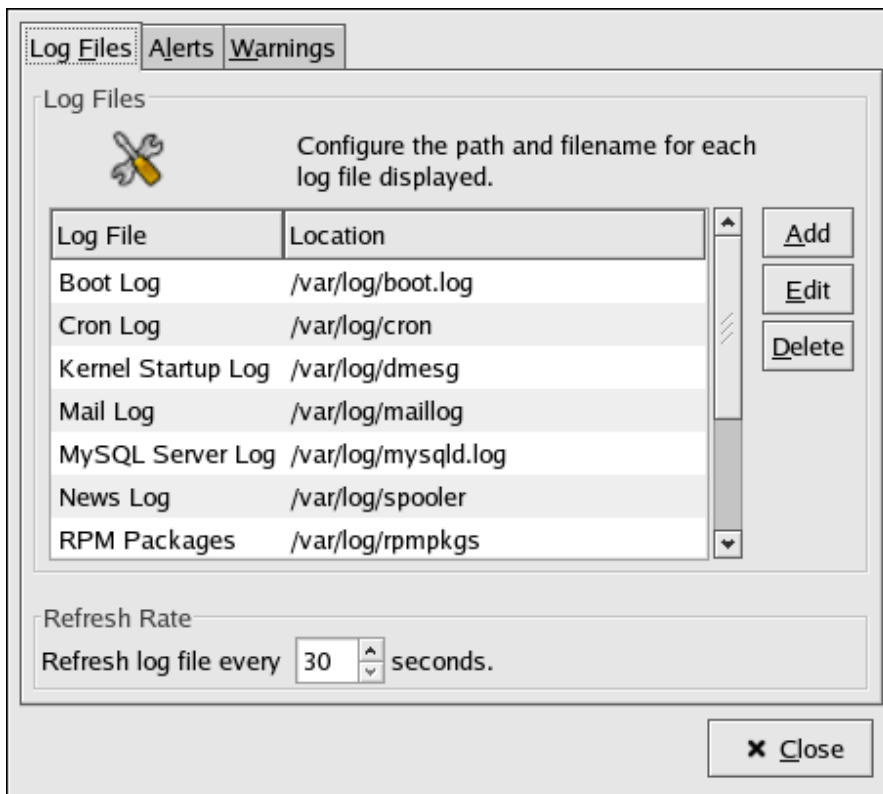
Figure 35.2. Log File Locations

35.3. Adding a Log File

To add a log file to the list, select **Edit** => **Preferences**, and click the **Add** button in the **Log Files** tab.
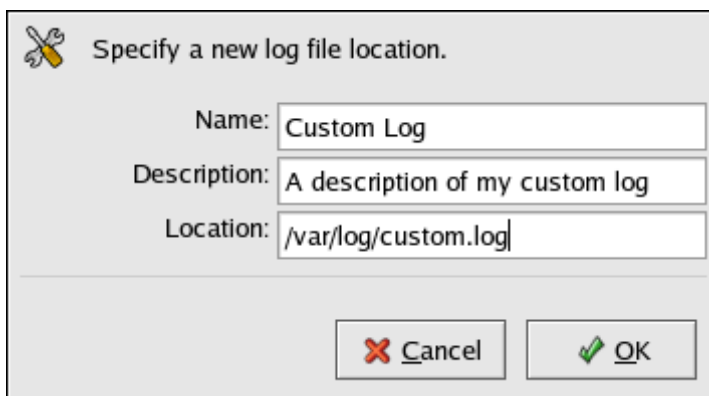
Figure 35.3. Adding a Log File

Provide a name, description, and the location of the log file to add. After clicking **OK**, the file is immediately added to the viewing area, if the file exists.

35.4. Examining Log Files

Log Viewer can be configured to display an alert icon beside lines that contain key alert words and a warning icon beside lines that contain key warning words.

To add alerts words, select **Edit** => **Preferences** from the pulldown 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**.

The alert icon is displayed to the left of the lines that contains any of the alert words.
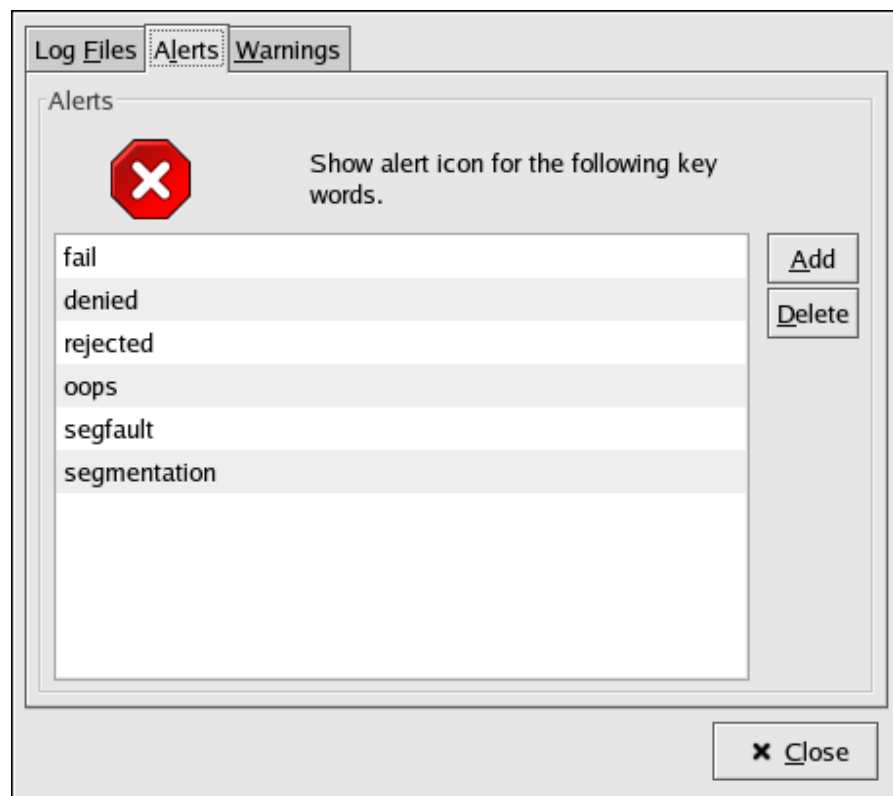
Figure 35.4. Alerts

To add warning words, select **Edit** => **Preferences** from the pull-down menu, and click on the **Warnings** tab. Click the **Add** button to add a warning word. To delete a warning word, select the word from the list, and click **Delete**.

The warning icon is displayed to the left of the lines that contains any of the warning words.
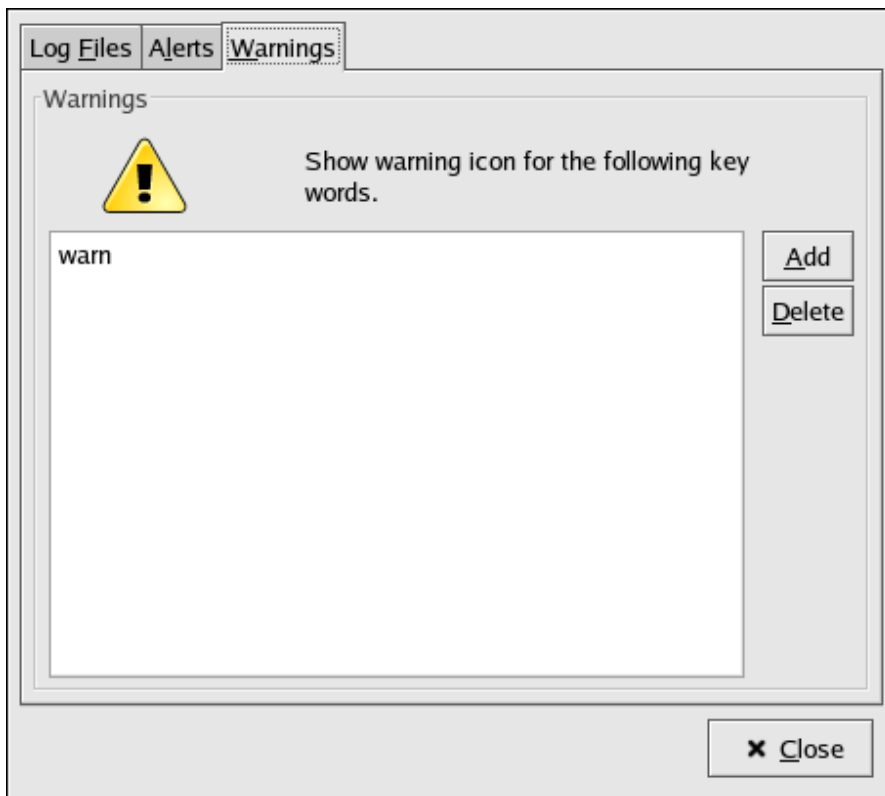
Figure 35.5. Warning

Chapter 36. Manually Upgrading the Kernel

The Red Hat Enterprise 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 first pass a rigorous set of quality assurance tests.

Red Hat Enterprise Linux kernels are packaged in RPM format so that they are easy to upgrade and verify using the **Red Hat User Agent**, or the **up2date** command. The **Red Hat User Agent** automatically queries the Red Hat Network servers and determines which packages need to be updated on your machine, including the kernel. This chapter is *only* useful for those individuals that require manual updating of kernel packages, without using the **up2date** command.

> ⚠️ **Warning**
>
> Please note, that building a custom kernel is not supported by the Red Hat Global Services Support team, and therefore is not explored in this manual.

> 💬 **Tip**
>
> Use of **up2date** is *highly* recommended by Red Hat for installing upgraded kernels.

For more information on Red Hat Network, the **Red Hat User Agent**, and **up2date**, refer to *Chapter 16, Red Hat Network*.

36.1. Overview of Kernel Packages

Red Hat Enterprise Linux contains the following kernel packages (some may not apply to your architecture):

- **kernel** — Contains the kernel and the following key features:
    - Uniprocessor support for x86 and Athlon systems (can be run on a multi-processor system, but only one processor is utilized)
    - Multi-processor support for all other architectures
    - For x86 systems, only the first 4 GB of RAM is used; use the **kernel-hugemem** package for x86 systems with over 4 GB of RAM
- **kernel-devel** — Contains the kernel headers and makefiles sufficient to build modules against the **kernel** package.
- **kernel-hugemem** — (only for i686 systems) In addition to the options enabled for the **kernel** package, the key configuration options are as follows:
    - Support for more than 4 GB of RAM (up to 64 GB for x86)

    > 💬 **Note**
    >
    > **kernel-hugemem** is required for memory configurations higher than 16 GB.

    - PAE (Physical Address Extension) or 3 level paging on x86 processors that support PAE
    - Support for multiple processors

- 4GB/4GB split — 4GB of virtual address space for the kernel and almost 4GB for each user process on x86 systems
- **kernel-hugemem-devel** — Contains the kernel headers and makefiles sufficient to build modules against the **kernel-hugemem** package.
- **kernel-smp** — Contains the kernel for multi-processor systems. The following are the key features:
  - Multi-processor support
  - Support for more than 4 GB of RAM (up to 16 GB for x86)
  - PAE (Physical Address Extension) or 3 level paging on x86 processors that support PAE
- **kernel-smp-devel** — Contains the kernel headers and makefiles sufficient to build modules against the **kernel-smp** package.
- **kernel-utils** — Contains utilities that can be used to control the kernel or system hardware.
- **kernel-doc** — Contains documentation files from the kernel source. Various portions of the Linux kernel and the device drivers shipped with it are documented in these files. Installation of this package provides a reference to the options that can be passed to Linux kernel modules at load time.

  By default, these files are placed in the **/usr/share/doc/kernel-doc-<version>/** directory.

> **Note**
>
> The **kernel-source** package has been removed and replaced with an RPM that can only be retrieved from Red Hat Network. This ***.src.rpm** must then be rebuilt locally using the **rpmbuild** command. Refer to the latest distribution Release Notes, including all updates, at https://www.redhat.com/docs/manuals/enterprise/ for more information on obtaining and installing the kernel source package.

36.2. Preparing to Upgrade

Before upgrading the kernel, take a few precautionary steps. The first step is to make sure working boot media exists for the system in case a problem occurs. If the boot loader is not configured properly to boot the new kernel, the system cannot be booted into Red Hat Enterprise Linux without working boot media.

For example, to create a boot diskette, login as root, and type the following command at a shell prompt:

```
/sbin/mkbootdisk `uname -r`
```

> **Tip**
>
> Refer to the **mkbootdisk** man page for more options. Creating bootable media via CD-Rs, CD-RWs, and USB flash drives are also supported given the system BIOS also supports it.

Reboot the machine with the boot media and verify that it works before continuing.

Hopefully, the media is not needed, but store it in a safe place just in case.

To determine which kernel packages are installed, execute the following command at a shell prompt:

```
rpm -qa | grep kernel
```

The output contains some or all of the following packages, depending on the system's architecture (the version numbers and packages may differ):

```
kernel-2.6.9-5.EL kernel-devel-2.6.9-5.EL kernel-utils-2.6.9-5.EL kernel-doc-
2.6.9-5.EL kernel-smp-2.6.9-5.EL kernel-smp-devel-2.6.9-5.EL kernel-hugemem-
devel-2.6.9-5.EL
```

From the output, determine which packages need to be download for the kernel upgrade. For a single processor system, the only required package is the **kernel** package. Refer to Section 36.1, "Overview of Kernel Packages" for descriptions of the different packages.

In the file name, each kernel package contains the architecture for which the package was built. The format is kernel-**<variant>**-**<version>**.**<arch>**.rpm, where **<variant>** is **smp**, **utils**, or so forth. The **<arch>** is one of the following:

- **x86_64** for the AMD64 architecture
- **ia64** for the Intel® Itanium™ architecture
- **ppc64** for the IBM® eServer™ pSeries™ architecture
- **ppc64** for the IBM® eServer™ iSeries™ architecture
- **s390** for the IBM® S/390® architecture
- **s390x** for the IBM® eServer™ zSeries® architecture
- x86 variant: The x86 kernels are optimized for different x86 versions. The options are as follows:
  - **i686** for Intel® Pentium® II, Intel® Pentium® III, Intel® Pentium® 4, AMD Athlon®, and AMD Duron® systems

36.3. Downloading the Upgraded Kernel

There are several ways to determine if an updated kernel is available for the system.

- Security Errata — Go to the following location for information on security errata, including kernel upgrades that fix security issues:

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

- Via Quarterly Updates — Refer to the following location for details:

  http://www.redhat.com/apps/support/errata/rhlas_errata_policy.html

- Via Red Hat Network — Download and install the kernel RPM packages. Red Hat Network can download the latest kernel, upgrade the kernel on the system, create an initial RAM disk image if needed, and configure the boot loader to boot the new kernel. For more information, refer to http://www.redhat.com/docs/manuals/RHNetwork/.

If Red Hat Network was used to download and install the updated kernel, follow the instructions in Section 36.5, "Verifying the Initial RAM Disk Image" and Section 36.6, "Verifying the Boot Loader", only *do not* change the kernel to boot by default. Red Hat Network automatically changes the default kernel to the latest version. To install the kernel manually, continue to Section 36.4, "Performing the Upgrade".

36.4. Performing the Upgrade

After retrieving all of the necessary packages, it is time to upgrade the 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 the old kernel is kept in case there are problems with the new kernel.

Use the **-i** argument with the **rpm** command to keep the old kernel. Do *not* use the **-U** option, since it overwrites the currently installed kernel, which creates boot loader problems. Issue the following command (the kernel version may vary):

```
rpm -ivh kernel-2.6.9-5.EL.<arch>.rpm
```

If the system is a multi-processor system, install the **kernel-smp** packages as well (the kernel version may vary):

```
rpm -ivh kernel-smp-2.6.9-5.EL.<arch>.rpm
```

If the system is **i686**-based and contains more than 4 GB of RAM, install the **kernel-hugemem** package built for the **i686** architecture as well (the kernel version might vary):

```
rpm -ivh kernel-hugemem-2.6.9-5.EL.i686.rpm
```

The next step is to verify that the initial RAM disk image has been created. Refer to Section 36.5, "Verifying the Initial RAM Disk Image" for details.

36.5. Verifying the Initial RAM Disk Image

If the system uses the ext3 file system, a SCSI controller, or uses labels to reference partitions in **/etc/fstab**, an initial RAM disk is needed. The initial RAM disk allows 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.

On the Red Hat Enterprise Linux architectures other than IBM eServer iSeries, the initial RAM disk can be created with the **mkinitrd** command. However, this step is performed automatically if the kernel and its associated packages are installed or upgraded from the RPM packages distributed by Red Hat, Inc; thus, it does not need to be executed manually. To verify that it was created, use the command **ls -l /boot** to make sure the **initrd-<version>.img** file was created (the version should match the version of the kernel just installed).

On iSeries systems, the initial RAM disk file and **vmlinux** file are combined into one file, which is created with the **addRamDisk** command. This step is performed automatically if the kernel and its associated packages are installed or upgraded from the RPM packages distributed by Red Hat, Inc; thus, it does not need to be executed manually. To verify that it was created, use the command **ls -l /boot** to make sure the **/boot/vmlinitrd-<kernel-version>** file was created (the version should match the version of the kernel just installed).

The next step is to verify that the boot loader has been configured to boot the new kernel. Refer to

for details.

36.6. Verifying the Boot Loader

The `kernel` RPM package configures the boot loader to boot the newly installed kernel (except for IBM eServer iSeries systems). 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, the system does not boot into Red Hat Enterprise Linux properly. If this happens, boot the system with the boot media created earlier and try configuring the boot loader again.

36.6.1. x86 Systems

All x86 systems use GRUB as the boot loader, which includes all AMD64 systems.

36.6.1.1. GRUB

Confirm that the file `/boot/grub/grub.conf` contains a `title` section with the same version as the `kernel` package just installed (if the `kernel-smp` or `kernel-hugemem` package was installed, a section exists for it as well):

```
# Note that you do not have to rerun grub after making changes to this file
# NOTICE:  You have a /boot partition.   This means that
#          all kernel and initrd paths are relative to /boot/, eg.
#          root (hd0,0)
#          kernel /vmlinuz-version ro root=/dev/hda2
#          initrd /initrd-version.img
#boot=/dev/hda
default=1
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz
title Red Hat Enterprise Linux (2.6.9-5.EL)
        root (hd0,0)
        kernel /vmlinuz-2.6.9-5.EL ro root=LABEL=/
        initrd /initrd-2.6.9-5.EL.img
title Red Hat Enterprise Linux (2.6.9-1.906_EL)
        root (hd0,0)
        kernel /vmlinuz-2.6.9-1.906_EL ro root=LABEL=/
        initrd /initrd-2.6.9-1.906_EL.img
```

If a separate `/boot/` partition was created, the paths to the kernel and initrd image are relative to `/boot/`.

Notice that the default is not set to the new kernel. 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 first title section, set `default` to `0`.

Begin testing the new kernel by rebooting the computer and watching the messages to ensure that the hardware is detected properly.

36.6.2. Itanium Systems

Itanium systems use ELILO as the boot loader, which uses `/boot/efi/EFI/redhat/elilo.conf` as the configuration file. Confirm that this file contains an `image` section with the same version as the

kernel package just installed:

```
prompt
timeout=50
default=old

image=vmlinuz-2.6.9-5.EL
        label=linux
        initrd=initrd-2.6.9-5.EL.img
        read-only
        append="root=LABEL=/"
image=vmlinuz-2.6.9-1.906_EL
        label=old
        initrd=initrd-2.6.9-1.906.img
        read-only
        append="root=LABEL=/"
```

Notice that the default is not set to the new kernel. To configure ELILO to boot the new kernel, change the value of the **default** variable to the value of the **label** for the **image** section that contains the new kernel.

Begin testing the new kernel by rebooting the computer and watching the messages to ensure that the hardware is detected properly.

36.6.3. IBM S/390 and IBM eServer zSeries Systems

The IBM S/390 and IBM eServer zSeries systems use z/IPL as the boot loader, which uses **/etc/zipl.conf** as the configuration file. Confirm that the file contains a section with the same version as the kernel package just installed:

```
[defaultboot]
default=old
target=/boot/
[linux]
        image=/boot/vmlinuz-2.6.9-5.EL
        ramdisk=/boot/initrd-2.6.9-5.EL.img
        parameters="root=LABEL=/"
[old]
        image=/boot/vmlinuz-2.6.9-1.906_EL
        ramdisk=/boot/initrd-2.6.9-1.906_EL.img
        parameters="root=LABEL=/"
```

Notice that the default is not set to the new kernel. To configure z/IPL to boot the new kernel by default change the value of the **default** variable to the name of the section that contains the new kernel. The first line of each section contains the name in brackets.

After modifying the configuration file, run the following command as root to enable the changes:

```
/sbin/zipl
```

Begin testing the new kernel by rebooting the computer and watching the messages to ensure that the hardware is detected properly.

36.6.4. IBM eServer iSeries Systems

The **/boot/vmlinitrd-<kernel-version>** file is installed when you upgrade the kernel. However, you must use the **dd** command to configure the system to boot the new kernel:

1. As root, issue the command `cat /proc/iSeries/mf/side` to determine the default side (either A, B, or C).

2. As root, issue the following command, where **<kernel-version>** is the version of the new kernel and **<side>** is the side from the previous command:

```
  dd if=/boot/vmlinitrd-<kernel-version>
of=/proc/iSeries/mf/<side>/vmlinux bs=8k
```

Begin testing the new kernel by rebooting the computer and watching the messages to ensure that the hardware is detected properly.

36.6.5. IBM eServer pSeries Systems

IBM eServer pSeries systems use YABOOT as the boot loader, which uses **/etc/aboot.conf** as the configuration file. Confirm that the file contains an **image** section with the same version as the **kernel** package just installed:

```
boot=/dev/sda1
init-message=Welcome to Red Hat Enterprise Linux!
Hit <TAB> for boot options

partition=2
timeout=30
install=/usr/lib/yaboot/yaboot
delay=10
nonvram

image=/vmlinux--2.6.9-5.EL
        label=old
        read-only
        initrd=/initrd--2.6.9-5.EL.img
        append="root=LABEL=/"

image=/vmlinux-2.6.9-5.EL
        label=linux
        read-only
        initrd=/initrd-2.6.9-5.EL.img
        append="root=LABEL=/"
```

Notice that the default is not set to the new kernel. The kernel in the first image is booted by default. To change the default kernel to boot either move its image stanza so that it is the first one listed or add the directive **default** and set it to the **label** of the image stanza that contains the new kernel.

Begin testing the new kernel by rebooting the computer and watching the messages to ensure that the hardware is detected properly.

Chapter 37. 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*, sometimes referred to as a *driver*, is dynamically loaded into memory.

During installation, the hardware on the system is probed. Based on this probing and the information provided by the user, 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 new hardware is added after installation and the hardware requires a kernel module, the system must be configured to load the proper kernel module for the new hardware. When the system is booted with the new hardware, the **Kudzu** program runs, detects the new hardware if it is supported, and configures the module for it. The module can also be specified manually by editing the module configuration file, `/etc/modprobe.conf`.

> **Note**
>
> Video card modules used to display the X Window System interface are part of the `xorg-X11` packages, not the kernel; thus, this chapter does not apply to them.

For example, if a system included an SMC EtherPower 10 PCI network adapter, the module configuration file contains the following line:

```
alias eth0 tulip
```

If a second network card is added to the system and is identical to the first card, add the following line to `/etc/modprobe.conf`:

```
alias eth1 tulip
```

Refer to the *Reference Guide* for an alphabetical list of kernel modules and supported hardware for those modules.

37.1. Kernel Module Utilities

A group of commands for managing kernel modules is available if the `module-init-tools` package is installed. Use these commands to determine if a module has been loaded successfully or when trying different modules for a piece of new hardware.

The command `/sbin/lsmod` displays a list of currently loaded modules. For example:

```
Module                  Size   Used by
nfs                    218437  1
lockd                   63977  2 nfs
parport_pc              24705  1
lp                      12077  0
parport                 37129  2 parport_pc,lp
autofs4                 23237  2
i2c_dev                 11329  0
i2c_core                22081  1 i2c_dev
sunrpc                 157093  5 nfs,lockd
button                   6481  0
battery                  8901  0
ac                       4805  0
md5                      4033  1
ipv6                   232833  16
ohci_hcd                21713  0
e100                    39493  0
mii                      4673  1 e100
floppy                  58481  0
sg                      33377  0
dm_snapshot             17029  0
dm_zero                  2369  0
dm_mirror               22957  2
ext3                   116809  2
jbd                     71257  1 ext3
dm_mod                  54741   6 dm_snapshot,dm_zero,dm_mirror
ips                     46173  2
aic7xxx                148121  0
sd_mod                  17217  3
scsi_mod               121421  4 sg,ips,aic7xxx,sd_mod
```

For each line, the first column is the name of the module, the second column is the size of the module, and the third column is the use count.

The **/sbin/lsmod** output is less verbose and easier to read than the output from viewing **/proc/modules**.

To load a kernel module, use the **/sbin/modprobe** command followed by the kernel module name. By default, **modprobe** attempts 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, meaning that other modules must be loaded first for it to load. The **/sbin/modprobe** command checks for these dependencies and loads the module dependencies before loading the specified module.

For example, the command

```
/sbin/modprobe e100
```

loads any module dependencies and then the **e100** module.

To print to the screen all commands as **/sbin/modprobe** executes them, use the **-v** option. For example:

```
/sbin/modprobe -v e100
```

Output similar to the following is displayed:

```
 /sbin/insmod /lib/modules/2.6.9-5.EL/kernel/drivers/net/e100.ko Using
 /lib/modules/2.6.9-5.EL/kernel/drivers/net/e100.ko Symbol version prefix
 'smp_'
```

The **/sbin/insmod** command also exists to load kernel modules; however, it does not resolve dependencies. Thus, it is recommended that the **/sbin/modprobe** command be used.

To unload kernel modules, use the **/sbin/rmmod** command followed by the module name. The **rmmod** utility only unloads modules that are not in use and that are not a dependency of other modules in use.

For example, the command

```
 /sbin/rmmod e100
```

unloads the **e100** kernel module.

Another useful kernel module utility is **modinfo**. Use the command **/sbin/modinfo** to display information about a kernel module. The general syntax is:

```
 /sbin/modinfo [options]<module>
```

Options include **-d**, which displays a brief description of the module, and **-p**, which lists the parameters the module supports. For a complete list of options, refer to the **modinfo** man page (**man modinfo**).

37.2. Persistent Module Loading

Kernel modules are usually loaded directly by the facility that requires them, which is given correct settings in the **/etc/modprobe.conf** file. However, it is sometimes necessary to explicitly force the loading of a module at boot time.

Red Hat Enterprise Linux checks for the existence of the **/etc/rc.modules** file at boot time, which contains various commands to load modules. The **rc.modules** should be used, and *not***rc.local** because **rc.modules** is executed earlier in the boot process.

For example, the following commands configure loading of the **foo** module at boot time (as root):

```
# echo modprobe foo >> /etc/rc.modules
# chmod +x /etc/rc.modules
```

> **Tip**
>
> This approach is not necessary for network and SCSI interfaces because they have their own specific mechanisms.

37.3. Additional Resources

For more information on kernel modules and their utilities, refer to the following resources.

37.3.1. 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/share/doc/kernel-doc-<version>/Documentation/kbuild/modules.txt`** — how to compile and use kernel modules.

37.3.2. Useful Websites

» http://www.redhat.com/mirrors/LDP/HOWTO/Module-HOWTO/index.html — *Linux Loadable Kernel Module HOWTO* from the Linux Documentation Project.

Chapter 38. Mail Transport Agent (MTA) Configuration

A *Mail Transport Agent* (MTA) is essential for sending email. A *Mail User Agent* (MUA) such as **Evolution**, **Mozilla Mail**, **Thunderbird**, and **Mutt**, is used to read and compose email. When a user sends an email from an MUA, the message is handed off to the MTA, which sends the message through 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 Enterprise Linux 4 provides three MTAs: Sendmail, Postfix, and Exim. If all three are installed, `sendmail` is the default MTA. The **Mail Transport Agent Switcher** allows for the selection of either `sendmail`, `postfix`, or `exim` as the default MTA for the system.

The `system-switch-mail` RPM package must be installed to use the text-based version of the **Mail Transport Agent Switcher** program. If you want to use the graphical version, the `system-switch-mail-gnome` package must also be installed.

To start the **Mail Transport Agent Switcher**, select **Applications** (the main menu on the panel) => **Preferences** => **More Preferences** => **Mail Transport Agent Switcher**, or type the command `system-switch-mail` 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 38.1, "Mail Transport Agent Switcher". If X is not detected, it starts in text-mode. To force **Mail Transport Agent Switcher** to run in text-mode, use the command `system-switch-mail-nox`.

Figure 38.1. Mail Transport Agent Switcher

If you select **OK** to change the MTA, the selected mail daemon is enabled to start at boot time, and the unselected mail daemons are disabled so that they do not start at boot time. The selected mail daemon is started, and any other mail daemon is stopped; thus making the changes take place immediately.

For more information about email protocols and MTAs, refer to the *Reference Guide.*

Part VI. System Monitoring

System administrators also monitor system performance. Red Hat Enterprise Linux contains tools to assist administrators with these tasks.

Chapter 39. 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 Enterprise Linux system using simple commands and a few simple programs.

39.1. System Processes

The **ps ax** command displays a list of current system processes, including processes owned by other users. To display the owner alongside each process, use the **ps aux** command. This list is a static list; in other words, it is a snapshot of what was 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 output from the **top** command is provided as follows:

```
top - 15:02:46 up 35 min,  4 users,  load average: 0.17, 0.65, 1.00
Tasks: 110 total,   1 running, 107 sleeping,   0 stopped,   2 zombie
Cpu(s): 41.1% us,  2.0% sy,  0.0% ni, 56.6% id,  0.0% wa,  0.3% hi,  0.0% si
Mem:    775024k total,   772028k used,     2996k free,    68468k buffers
Swap:  1048568k total,      176k used,  1048392k free,   441172k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
 4624 root      15   0 40192  18m 7228 S 28.4  2.4  1:23.21 X
 4926 mhideo    15   0 55564  33m 9784 S 13.5  4.4  0:25.96 gnome-terminal
 6475 mhideo    16   0  3612  968  760 R  0.7  0.1  0:00.11 top
 4920 mhideo    15   0 20872  10m 7808 S  0.3  1.4  0:01.61 wnck-applet
    1 root      16   0  1732  548  472 S  0.0  0.1  0:00.23 init
    2 root      34  19     0    0    0 S  0.0  0.0  0:00.00 ksoftirqd/0
    3 root       5 -10     0    0    0 S  0.0  0.0  0:00.03 events/0
    4 root       6 -10     0    0    0 S  0.0  0.0  0:00.02 khelper
    5 root       5 -10     0    0    0 S  0.0  0.0  0:00.00 kacpid
   29 root       5 -10     0    0    0 S  0.0  0.0  0:00.00 kblockd/0
   47 root      16   0     0    0    0 S  0.0  0.0  0:01.74 pdflush
   50 root      11 -10     0    0    0 S  0.0  0.0  0:00.00 aio/0
   30 root      15   0     0    0    0 S  0.0  0.0  0:00.05 khubd
   49 root      16   0     0    0    0 S  0.0  0.0  0:01.44 kswapd0
```

To exit **top**, press the **q** key.

contains useful interactive commands that you can use with **top**. For more information, refer to the **top**(1) manual page.

Table 39.1. Interactive `top` commands

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

If you prefer a graphical interface for **top**, you can use the **GNOME System Monitor**. To start it from the desktop, select **System** => **Administration** => **System Monitor** or type **gnome-system-monitor** at a shell prompt (such as an XTerm). Select the **Process Listing** tab.

The **GNOME System Monitor** allows you to search for a process in the list of running processes. Using the Gnome System Monitor, you can also view all processes, your processes, or active processes.

The **Edit** menu item allows you to:

» Stop a process.
» Continue or start a process.
» End a processes.
» Kill a process.
» Change the priority of a selected process.
» Edit the System Monitor preferences. These include changing the interval seconds to refresh the list and selecting process fields to display in the System Monitor window.

The **View** menu item allows you to:

» View only active processes.
» View all processes.
» View my processes.
» View process dependencies.
» Hide a process.
» View hidden processes.
» View memory maps.
» View the files opened by the selected process.

To stop a process, select it and click **End Process**. Alternatively you can also stop a process by selecting it, clicking **Edit** on your menu and selecting **Stop Process**.

To sort the information by a specific column, click on the name of the column. This sorts the information by the selected column in ascending order. Click on the name of the column again to toggle the sort between ascending and descending order.
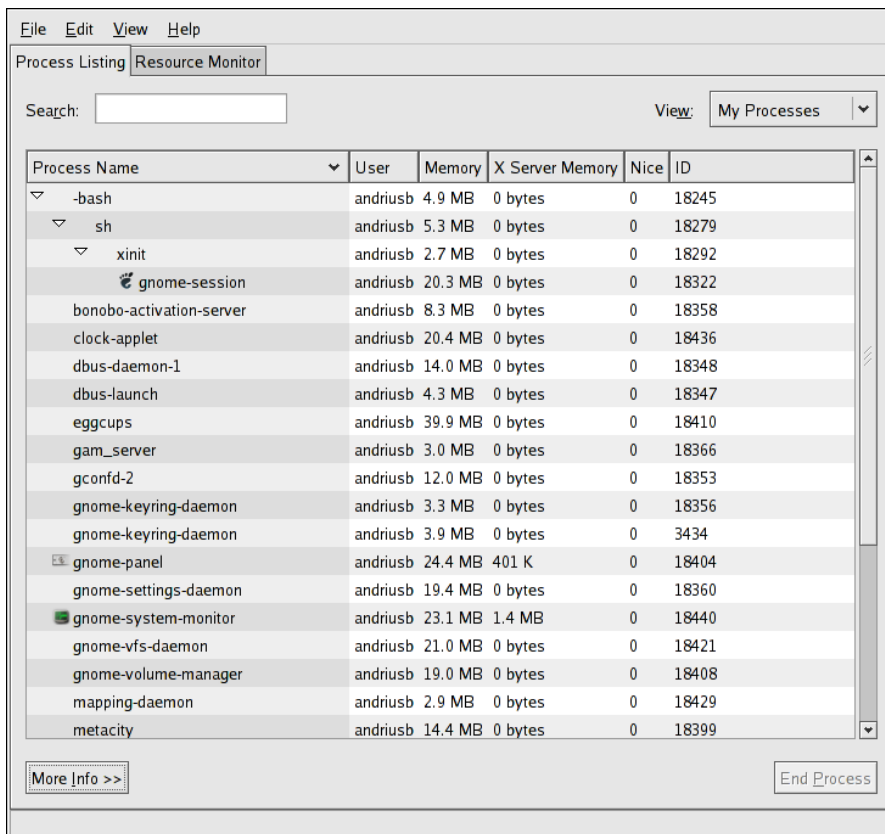
Figure 39.1. GNOME System Monitor

39.2. 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:       645712      549720       95992           0      176248      224452
 -/+ buffers/cache:      149020      496692
  Swap:     1310712           0     1310712
```

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

```
              total        used        free      shared     buffers      cached
  Mem:          630         536          93           0         172         219
 -/+ buffers/cache:         145         485
  Swap:        1279           0        1279
```

If you prefer a graphical interface for `free`, you can use the **GNOME System Monitor**. To start it from the desktop, go to **System** => **Administration** => **System Monitor** or type `gnome-system-monitor` at a shell prompt (such as an XTerm). Click on the **Resources** tab.
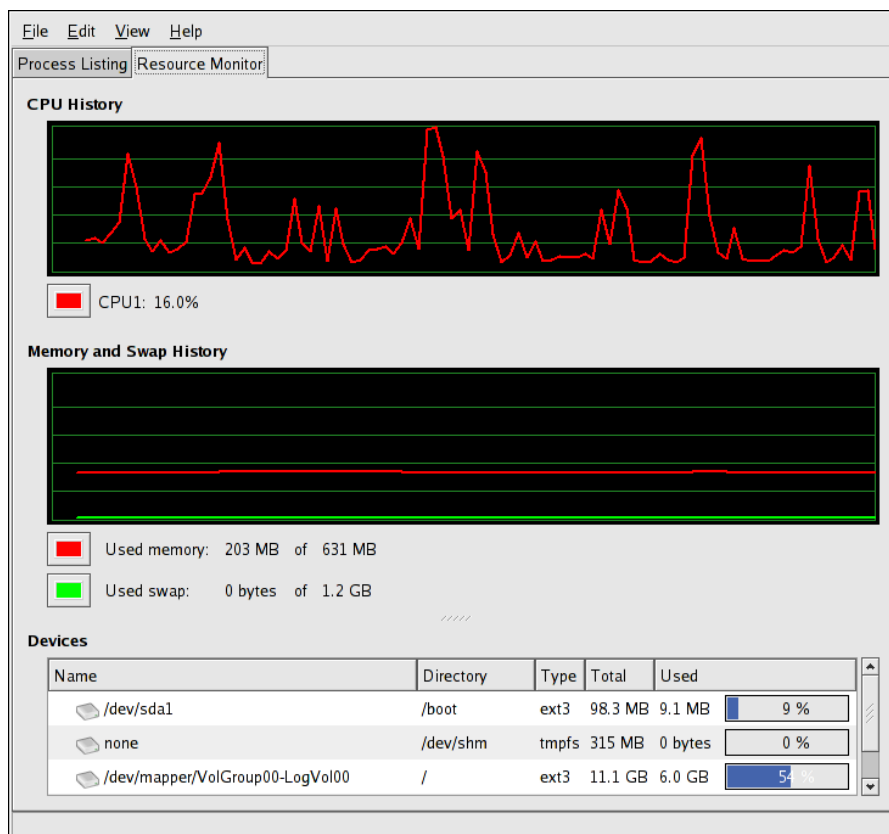
Figure 39.2. GNOME System Monitor - Resources tab

39.3. 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/mapper/VolGroup00-LogVol00
                     11675568    6272120    4810348   57% / /dev/sda1
                     100691        9281      86211   10% /boot
none                  322856           0     322856    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/mapper/VolGroup00-LogVol00
                      12G  6.0G  4.6G  57% / /dev/sda1
   99M  9.1M   85M  10% /boot
none    316M     0  316M   0% /dev/shm
```

In the list of mounted 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 is displayed in a list. The grand total

for the current directory and subdirectories are also 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 **Gnome System Monitor** by clicking on **System** => **Administration** => **System Monitor** or type **gnome-system-monitor** at a shell prompt (such as an XTerm). Select the File Systems tab to view the system's partitions. The figure below illustrates the File Systems tab.

Figure 39.3. GNOME System Monitor - File Systems

39.4. 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 **System** (the main menu on the panel) => **Administration** => **Hardware** or type **hwbrowser** at a shell prompt. As shown in Figure 39.4, "Hardware Browser", it displays your CD-ROM devices, diskette drives, 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 is displayed.
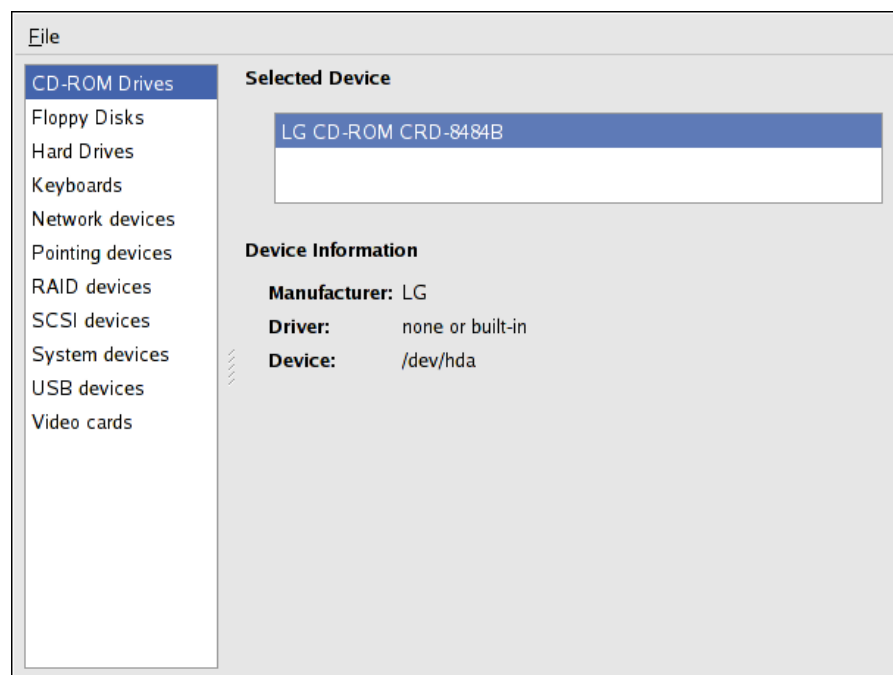
Figure 39.4. Hardware Browser

The **Device Manager** application can also be used to display your system hardware. This application can be started by selecting **System** (the main menu on the panel) => **Administration** => **Hardware** like the **Hardware Browser**. To start the application from a terminal, type **hal-device-manager**. Depending on your installation preferences, the graphical menu above may start this application or the **Hardware Browser** when clicked. The figure below illustrates the **Device Manager** window.

Figure 39.5. Device Manager

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:

```
00:00.0 Host bridge: ServerWorks CNB20LE Host Bridge (rev 06)
00:00.1 Host bridge: ServerWorks CNB20LE Host Bridge (rev 06)
00:01.0 VGA compatible controller: S3 Inc. Savage 4 (rev 04)
00:02.0 Ethernet controller: Intel Corp. 82557/8/9 [Ethernet Pro 100] (rev 08)
00:0f.0 ISA bridge: ServerWorks OSB4 South Bridge (rev 50)
00:0f.1 IDE interface: ServerWorks OSB4 IDE Controller
00:0f.2 USB Controller: ServerWorks OSB4/CSB5 OHCI USB Controller (rev 04)
01:03.0 SCSI storage controller: Adaptec AIC-7892P U160/m (rev 02)
01:05.0 RAID bus controller: IBM ServeRAID Controller
```

The **lspci** is also useful to determine the network card in your system if you do not know the manufacturer or model number.

39.5. Additional Resources

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

39.5.1. 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.
- **lspci** manual page — Type **man lspci** to learn more about the **lspci** command and its many options.
- **/proc/** directory — The contents of the **/proc/** directory can also be used to gather more detailed system information.

Chapter 40. OProfile

OProfile is a low overhead, system-wide performance monitoring tool. It uses the performance monitoring hardware on the processor to retrieve information about the kernel and executables on the system, such as when memory is referenced, the number of L2 cache requests, and the number of hardware interrupts received. On a Red Hat Enterprise Linux system, the **oprofile** RPM package must be installed to use this tool.

Many processors include dedicated performance monitoring hardware. This hardware makes it possible to detect when certain events happen (such as the requested data not being in cache). The hardware normally takes the form of one or more *counters* that are incremented each time an event takes place. When the counter value, essentially rolls over, an interrupt is generated, making it possible to control the amount of detail (and therefore, overhead) produced by performance monitoring.

OProfile uses this hardware (or a timer-based substitute in cases where performance monitoring hardware is not present) to collect *samples* of performance-related data each time a counter generates an interrupt. These samples are periodically written out to disk; later, the data contained in these samples can then be used to generate reports on system-level and application-level performance.

OProfile is a useful tool, but be aware of some limitations when using it:

- *Use of shared libraries* — Samples for code in shared libraries are not attributed to the particular application unless the **--separate=library** option is used.
- *Performance monitoring samples are inexact* — When a performance monitoring register triggers a sample, the interrupt handling is not precise like a divide by zero exception. Due to the out-of-order execution of instructions by the processor, the sample may be recorded on a nearby instruction.
- ***opreport** does not associate samples for inline functions' properly* — **opreport** uses a simple address range mechanism to determine which function an address is in. Inline function samples are not attributed to the inline function but rather to the function the inline function was inserted into.
- *OProfile accumulates data from multiple runs* — OProfile is a system-wide profiler and expects processes to start up and shut down multiple times. Thus, samples from multiple runs accumulate. Use the command **opcontrol --reset** to clear out the samples from previous runs.
- *Non-CPU-limited performance problems* — OProfile is oriented to finding problems with CPU-limited processes. OProfile does not identify processes that are asleep because they are waiting on locks or for some other event to occur (for example an I/O device to finish an operation).

40.1. Overview of Tools

Table 40.1, "OProfile Commands" provides a brief overview of the tools provided with the **oprofile** package.

Table 40.1. OProfile Commands

Command	Description
`op_help`	Displays available events for the system's processor along with a brief description of each.
`op_import`	Converts sample database files from a foreign binary format to the native format for the system. Only use this option when analyzing a sample database from a different architecture.
`opannotate`	Creates annotated source for an executable if the application was compiled with debugging symbols. Refer to Section 40.5.3, "Using `opannotate`" for details.
`opcontrol`	Configures what data is collected. Refer to Section 40.2, "Configuring OProfile" for details.
`opreport`	Retrieves profile data. Refer to Section 40.5.1, "Using `opreport`" for details.
`oprofiled`	Runs as a daemon to periodically write sample data to disk.

40.2. Configuring OProfile

Before OProfile can be run, it must be configured. At a minimum, selecting to monitor the kernel (or selecting not to monitor the kernel) is required. The following sections describe how to use the **opcontrol** utility to configure OProfile. As the **opcontrol** commands are executed, the setup options are saved to the **/root/.oprofile/daemonrc** file.

40.2.1. Specifying the Kernel

First, configure whether OProfile should monitor the kernel. This is the only configuration option that is required before starting OProfile. All others are optional.

To monitor the kernel, execute the following command as root:

```
opcontrol --setup --vmlinux=/usr/lib/debug/lib/modules/`uname -r`/vmlinux
```

> **Note**
>
> The **debuginfo** package must be installed (which contains the uncompressed kernel) in order to monitor the kernel.

To configure OProfile not to monitor the kernel, execute the following command as root:

```
opcontrol --setup --no-vmlinux
```

This command also loads the **oprofile** kernel module, if it is not already loaded, and creates the **/dev/oprofile/** directory, if it does not already exist. Refer to Section 40.6, "Understanding /dev/oprofile/" for details about this directory.

> **Note**
>
> Even if OProfile is configured not to profile the kernel, the SMP kernel still must be running so that the `oprofile` module can be loaded from it.

Setting whether samples should be collected within the kernel only changes what data is collected, not how or where the collected data is stored. To generate different sample files for the kernel and application libraries, refer to Section 40.2.3, "Separating Kernel and User-space Profiles".

40.2.2. Setting Events to Monitor

Most processors contain *counters*, which are used by OProfile to monitor specific events. As shown in Table 40.2, "OProfile Processors and Counters", the number of counters available depends on the processor.

Table 40.2. OProfile Processors and Counters

Processor	cpu_type	Number of Counters
Pentium Pro	i386/ppro	2
Pentium II	i386/pii	2
Pentium III	i386/piii	2
Pentium 4 (non-hyper-threaded)	i386/p4	8
Pentium 4 (hyper-threaded)	i386/p4-ht	4
Athlon	i386/athlon	4
AMD64	x86-64/hammer	4
Itanium	ia64/itanium	4
Itanium 2	ia64/itanium2	4
TIMER_INT	timer	1
IBM eServer iSeries and pSeries	timer	1
	ppc64/power4	8
	ppc64/power5	6
	ppc64/970	8
IBM eServer S/390 and S/390x	timer	1
IBM eServer zSeries	timer	1

Use Table 40.2, "OProfile Processors and Counters" to verify that the correct processor type was detected and to determine the number of events that can be monitored simultaneously. `timer` is used as the processor type if the processor does not have supported performance monitoring hardware.

If `timer` is used, events cannot be set for any processor because the hardware does not have support for hardware performance counters. Instead, the timer interrupt is used for profiling.

If `timer` is not used as the processor type, the events monitored can be changed, and counter 0 for the processor is set to a time-based event by default. If more than one counter exists on the processor, the counters other than counter 0 are not set to an event by default. The default events monitored are shown in Table 40.3, "Default Events".

Table 40.3. Default Events

Processor	Default Event for Counter	Description
Pentium Pro, Pentium II, Pentium III, Athlon, AMD64	CPU_CLK_UNHALTED	The processor's clock is not halted
Pentium 4 (HT and non-HT)	GLOBAL_POWER_EVENTS	The time during which the processor is not stopped
Itanium 2	CPU_CYCLES	CPU Cycles
TIMER_INT	(none)	Sample for each timer interrupt
ppc64/power4	CYCLES	Processor Cycles
ppc64/power5	CYCLES	Processor Cycles
ppc64/970	CYCLES	Processor Cycles

The number of events that can be monitored at one time is determined by the number of counters for the processor. However, it is not a one-to-one correlation; on some processors, certain events must be mapped to specific counters. To determine the number of counters available, execute the following command:

```
cat /dev/oprofile/cpu_type
```

The events available vary depending on the processor type. To determine the events available for profiling, execute the following command as root (the list is specific to the system's processor type):

```
op_help
```

The events for each counter can be configured via the command line or with a graphical interface. For more information on the graphical interface, refer to Section 40.8, "Graphical Interface". If the counter cannot be set to a specific event, an error message is displayed.

To set the event for each configurable counter via the command line, use **opcontrol**:

```
opcontrol --event=<event-name>:<sample-rate>
```

Replace **<event-name>** with the exact name of the event from **op_help**, and replace **<sample-rate>** with the number of events between samples.

40.2.2.1. Sampling Rate

By default, a time-based event set is selected. It creates a sample every 100,000 clock cycles per processor. If the timer interrupt is used, the timer is set to whatever the jiffy rate is and is not user-settable. If the **cpu_type** is not **timer**, each event can have a *sampling rate* set for it. The sampling rate is the number of events between each sample snapshot.

When setting the event for the counter, a sample rate can also be specified:

```
opcontrol --event=<event-name>:<sample-rate>
```

Replace **<sample-rate>** with the number of events to wait before sampling again. The smaller the count, the more frequent the samples. For events that do not happen frequently, a lower count may be needed to capture the event instances.

> **Warning**
>
> Be extremely careful when setting sampling rates. Sampling too frequently can overload the system, causing the system to appear as if it is frozen or causing the system to actually freeze.

40.2.2.2. Unit Masks

If the **cpu_type** is not **timer**, *unit masks* may also be required to further define the event.

Unit masks for each event are listed with the **op_help** command. The values for each unit mask are listed in hexadecimal format. To specify more than one unit mask, the hexadecimal values must be combined using a bitwise *or* operation.

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>
```

40.2.3. Separating Kernel and User-space Profiles

By default, kernel mode and user mode information is gathered for each event. To configure OProfile not to count events in kernel mode for a specific counter, execute the following command:

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>:0
```

Execute the following command to start profiling kernel mode for the counter again:

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>:1
```

To configure OProfile not to count events in user mode for a specific counter, execute the following command:

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>:<kernel>:0
```

Execute the following command to start profiling user mode for the counter again:

```
opcontrol --event=<event-name>:<sample-rate>:<unit-mask>:<kernel>:1
```

When the OProfile daemon writes the profile data to sample files, it can separate the kernel and library profile data into separate sample files. To configure how the daemon writes to sample files, execute the following command as root:

```
opcontrol --separate=<choice>
```

<choice> can be one of the following:

- **none** — do not separate the profiles (default)
- **library** — generate per-application profiles for libraries
- **kernel** — generate per-application profiles for the kernel and kernel modules
- **all** — generate per-application profiles for libraries and per-application profiles for the kernel and kernel modules

If **--separate=library** is used, the sample file name includes the name of the executable as well as the name of the library.

40.3. Starting and Stopping OProfile

To start monitoring the system with OProfile, execute the following command as root:

```
opcontrol --start
```

Output similar to the following is displayed:

```
Using log file /var/lib/oprofile/oprofiled.log
Daemon started.
Profiler running.
```

The settings in **/root/.oprofile/daemonrc** are used.

The OProfile daemon, **oprofiled**, is started; it periodically writes the sample data to the **/var/lib/oprofile/samples/** directory. The log file for the daemon is located at **/var/lib/oprofile/oprofiled.log**.

To stop the profiler, execute the following command as root:

```
opcontrol --shutdown
```

40.4. Saving Data

Sometimes it is useful to save samples at a specific time. For example, when profiling an executable, it may be useful to gather different samples based on different input data sets. If the number of events to be monitored exceeds the number of counters available for the processor, multiple runs of OProfile can be used to collect data, saving the sample data to different files each time.

To save the current set of sample files, execute the following command, replacing *<name>* with a unique descriptive name for the current session.

```
opcontrol --save=<name>
```

The directory **/var/lib/oprofile/samples/*name*/** is created and the current sample files are copied to it.

40.5. Analyzing the Data

Periodically, the OProfile daemon, **oprofiled**, collects the samples and writes them to the **/var/lib/oprofile/samples/** directory. Before reading the data, make sure all data has been written to this directory by executing the following command as root:

```
opcontrol --dump
```

Each sample file name is based on the name of the executable. For example, the samples for the default event on a Pentium III processor for **/bin/bash** becomes:

```
\{root\}/bin/bash/\{dep\}/\{root\}/bin/bash/CPU_CLK_UNHALTED.100000
```

The following tools are available to profile the sample data once it has been collected:

- **opreport**

- **opannotate**

Use these tools, along with the binaries profiled, to generate reports that can be further analyzed.

> ⚠️ **Warning**
>
> The executable being profiled must be used with these tools to analyze the data. If it must change after the data is collected, backup the executable used to create the samples as well as the sample files.

Samples for each executable are written to a single sample file. Samples from each dynamically linked library are also written to a single sample file. While OProfile is running, if the executable being monitored changes and a sample file for the executable exists, the existing sample file is automatically deleted. Thus, if the existing sample file is needed, it must be backed up, along with the executable used to create it before replacing the executable with a new version. Refer to Section 40.4, "Saving Data" for details on how to backup the sample file.

40.5.1. Using opreport

The **opreport** tool provides an overview of all the executables being profiled.

The following is part of an example output:

```
Profiling through timer interrupt
        TIMER:0|
  samples|      %|
------------------
    25926 97.5212 no-vmlinux
      359  1.3504 pi
       65  0.2445 Xorg
       62  0.2332 libvte.so.4.4.0
       56  0.2106 libc-2.3.4.so
       34  0.1279 libglib-2.0.so.0.400.7
       19  0.0715 libXft.so.2.1.2
       17  0.0639 bash
        8  0.0301 ld-2.3.4.so
        8  0.0301 libgdk-x11-2.0.so.0.400.13
        6  0.0226 libgobject-2.0.so.0.400.7
        5  0.0188 oprofiled
        4  0.0150 libpthread-2.3.4.so
        4  0.0150 libgtk-x11-2.0.so.0.400.13
        3  0.0113 libXrender.so.1.2.2
        3  0.0113 du
        1  0.0038 libcrypto.so.0.9.7a
        1  0.0038 libpam.so.0.77
        1  0.0038 libtermcap.so.2.0.8
        1  0.0038 libX11.so.6.2
        1  0.0038 libgthread-2.0.so.0.400.7
        1  0.0038 libwnck-1.so.4.9.0
```

Each executable is listed on its own line. The first column is the number of samples recorded for the executable. The second column is the percentage of samples relative to the total number of samples. The third column is the name of the executable.

Refer to the **opreport** man page for a list of available command line options, such as the **-r** option used to sort the output from the executable with the smallest number of samples to the one with the largest number of samples.

40.5.2. Using opreport on a Single Executable

To retrieve more detailed profiled information about a specific executable, use **opreport**:

```
opreport <mode><executable>
```

<executable> must be the full path to the executable to be analyzed. *<mode>* must be one of the following:

-l

> List sample data by symbols. For example, the following is part of the output from running the command **opreport -l /lib/tls/libc-<version>.so**:

```
samples   %         symbol name
12        21.4286   __gconv_transform_utf8_internal
5          8.9286   _int_malloc
4          7.1429   malloc
3          5.3571   __i686.get_pc_thunk.bx
3          5.3571   _dl_mcount_wrapper_check
3          5.3571   mbrtowc
3          5.3571   memcpy
2          3.5714   _int_realloc
2          3.5714   _nl_intern_locale_data
2          3.5714   free
2          3.5714   strcmp
1          1.7857   __ctype_get_mb_cur_max
1          1.7857   __unregister_atfork
1          1.7857   __write_nocancel
1          1.7857   _dl_addr
1          1.7857   _int_free
1          1.7857   _itoa_word
1          1.7857   calc_eclosure_iter
1          1.7857   fopen@@GLIBC_2.1
1          1.7857   getpid
1          1.7857   memmove
1          1.7857   msort_with_tmp
1          1.7857   strcpy
1          1.7857   strlen
1          1.7857   vfprintf
1          1.7857   write
```

> The first column is the number of samples for the symbol, the second column is the percentage of samples for this symbol relative to the overall samples for the executable, and the third column is the symbol name.

> To sort the output from the largest number of samples to the smallest (reverse order), use **-r** in conjunction with the **-l** option.

-i *<symbol-name>*

> List sample data specific to a symbol name. For example, the following output is from the command **opreport -l -i __gconv_transform_utf8_internal**

/lib/tls/libc-<*version*>.so:

```
samples  %         symbol name
12       100.000   __gconv_transform_utf8_internal
```

The first line is a summary for the symbol/executable combination.

The first column is the number of samples for the memory symbol. The second column is the percentage of samples for the memory address relative to the total number of samples for the symbol. The third column is the symbol name.

-d

List sample data by symbols with more detail than **-l**. For example, the following output is from the command **opreport -l -d __gconv_transform_utf8_internal /lib/tls/libc-<*version*>.so**:

```
vma       samples  %         symbol name
00a98640 12        100.000   __gconv_transform_utf8_internal
00a98640 1          8.3333
00a9868c 2         16.6667
00a9869a 1          8.3333
00a986c1 1          8.3333
00a98720 1          8.3333
00a98749 1          8.3333
00a98753 1          8.3333
00a98789 1          8.3333
00a98864 1          8.3333
00a98869 1          8.3333
00a98b08 1          8.3333
```

The data is the same as the **-l** option except that for each symbol, each virtual memory address used is shown. For each virtual memory address, the number of samples and percentage of samples relative to the number of samples for the symbol is displayed.

-x<*symbol-name*>

Exclude the comma-separated list of symbols from the output.

session:<*name*>

Specify the full path to the session or a directory relative to the **/var/lib/oprofile/samples/** directory.

40.5.3. Using **opannotate**

The **opannotate** tool tries to match the samples for particular instructions to the corresponding lines in the source code. The resulting files generated should have the samples for the lines at the left. It also puts in a comment at the beginning of each function listing the total samples for the function.

For this utility to work, the executable must be compiled with GCC's **-g** option. By default, Red Hat Enterprise Linux packages are not compiled with this option.

The general syntax for **opannotate** is as follows:

```
opannotate --search-dirs <src-dir> --source <executable>
```

The directory containing the source code and the executable to be analyzed must be specified. Refer to the **opannotate** man page for a list of additional command line options.

40.6. Understanding /dev/oprofile/

The **/dev/oprofile/** directory contains the file system for OProfile. Use the **cat** command to display the values of the virtual files in this file system. For example, the following command displays the type of processor OProfile detected:

```
cat /dev/oprofile/cpu_type
```

A directory exists in **/dev/oprofile/** for each counter. For example, if there are 2 counters, the directories **/dev/oprofile/0/** and **dev/oprofile/1/** exist.

Each directory for a counter contains the following files:

- **count** — The interval between samples.
- **enabled** — If 0, the counter is off and no samples are collected for it; if 1, the counter is on and samples are being collected for it.
- **event** — The event to monitor.
- **kernel** — If 0, samples are not collected for this counter event when the processor is in kernel-space; if 1, samples are collected even if the processor is in kernel-space.
- **unit_mask** — Defines which unit masks are enabled for the counter.
- **user** — If 0, samples are not collected for the counter event when the processor is in user-space; if 1, samples are collected even if the processor is in user-space.

The values of these files can be retrieved with the **cat** command. For example:

```
cat /dev/oprofile/0/count
```

40.7. Example Usage

While OProfile can be used by developers to analyze application performance, it can also be used by system administrators to perform system analysis. For example:

- *Determine which applications and services are used the most on a system* — **opreport** can be used to determine how much processor time an application or service uses. If the system is used for multiple services but is under performing, the services consuming the most processor time can be moved to dedicated systems.
- *Determine processor usage* — The **CPU_CLK_UNHALTED** event can be monitored to determine the processor load over a given period of time. This data can then be used to determine if additional processors or a faster processor might improve system performance.

40.8. Graphical Interface

Some OProfile preferences can be set with a graphical interface. To start it, execute the **oprof_start** command as root at a shell prompt.

After changing any of the options, save them by clicking the **Save and quit** button. The preferences are written to **/root/.oprofile/daemonrc**, and the application exits. Exiting the application does not stop OProfile from sampling.

On the **Setup** tab, to set events for the processor counters as discussed in Section 40.2.2, "Setting Events to Monitor", select the counter from the pulldown menu and select the event from the list. A brief description of the event appears in the text box below the list. Only events available for the specific counter and the specific architecture are displayed. The interface also displays whether the profiler is running and some brief statistics about it.
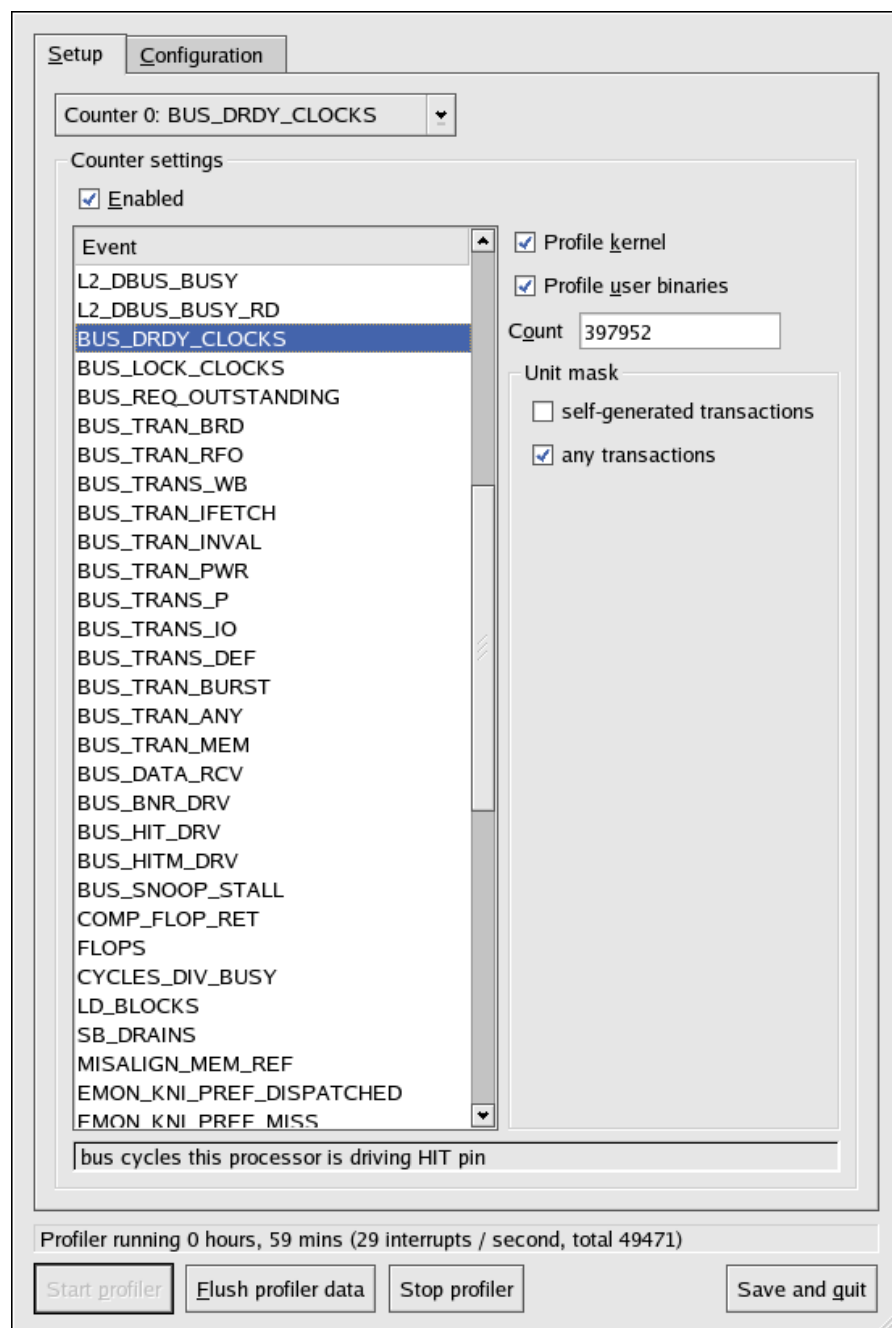
Figure 40.1. OProfile Setup

On the right side of the tab, select the **Profile kernel** option to count events in kernel mode for the currently selected event, as discussed in Section 40.2.3, "Separating Kernel and User-space Profiles". If this option is unselected, no samples are collected for the kernel.

Select the **Profile user binaries** option to count events in user mode for the currently selected event, as discussed in Section 40.2.3, "Separating Kernel and User-space Profiles". If this option is unselected, no samples are collected for user applications.

Use the **Count** text field to set the sampling rate for the currently selected event as discussed in Section 40.2.2.1, "Sampling Rate".

If any unit masks are available for the currently selected event, as discussed in Section 40.2.2.2, "Unit Masks", they are displayed in the **Unit Masks** area on the right side of the **Setup** tab. Select the checkbox beside the unit mask to enable it for the event.

On the **Configuration** tab, to profile the kernel, enter the name and location of the **vmlinux** file for the kernel to monitor in the **Kernel image file** text field. To configure OProfile not to monitor the kernel, select **No kernel image**.

Figure 40.2. OProfile Configuration

If the **Verbose** option is selected, the **oprofiled** daemon log includes more information.

If **Per-application kernel samples files** is selected, OProfile generates per-application profiles for the kernel and kernel modules as discussed in Section 40.2.3, "Separating Kernel and User-space Profiles". This is equivalent to the **opcontrol --separate=kernel** command. If **Per-application shared libs samples files** is selected, OProfile generates per-application profiles for libraries. This is equivalent to the **opcontrol --separate=library** command.

To force data to be written to samples files as discussed in Section 40.5, "Analyzing the Data", click the **Flush profiler data** button. This is equivalent to the **opcontrol --dump** command.

To start OProfile from the graphical interface, click **Start profiler**. To stop the profiler, click **Stop profiler**. Exiting the application does not stop OProfile from sampling.

40.9. Additional Resources

This chapter only highlights OProfile and how to configure and use it. To learn more, refer to the following resources.

40.9.1. Installed Docs

- **/usr/share/doc/oprofile-<version>/oprofile.html** — *OProfile Manual*
- **oprofile** man page — Discusses **opcontrol**, **opreport**, **opannotate**, and **op_help**

40.9.2. Useful Websites

http://oprofile.sourceforge.net/ — Contains the latest documentation, mailing lists, IRC channels, and more.

Part VII. Appendix

Revision History

Revision 1.0-1.33.402	Fri Oct 25 2013	Rüdiger Landmann

Rebuild with Publican 4.0.0

Revision 1.0-1.33	July 24 2012	Ruediger Landmann

Rebuild for Publican 3.0

Revision 1.0-1	Thu Sep 18 2008	Don Domingo

migrated to new automated build system

Index

Symbols

A

D

H

M

S

U

V

ypbind, User Information