RED HAT
LINUX
ADVANCED SERVER
Version 2.1 *with Red Hat Network*

redhat

MAN0044US
Printed in Canada.

86 A2 68EG 00

THE OFFICIAL RED HAT® LINUX® ADVANCED SERVER INSTALLATION GUIDE

8 1/4"

5 7/8"

Red Hat Linux Advanced Server 2.1

The Official Red Hat Linux Advanced Server Installation Guide

ISBN: N/A


Red Hat, Inc.

1801 Varsity Drive
Raleigh, NC 27606-2072 USA
+1 919 754 3700 (Voice)
+1 919 754 3701 (FAX)
888 733 4281 (Voice)
P.O. Box 13588
Research Triangle Park, NC 27709 USA

rhl-ig-as-x86(EN)-2.1-Print-RHI (2002-03-22T11:09-0400)

Contents

Red Hat Linux Advanced Server 2.1

Introduction

Welcome to the *Official Red Hat Linux Advanced Server Installation Guide*. This guide contains useful information to assist you during the installation of Red Hat Linux Advanced Server. From fundamental concepts such as installation preparation to the step-by-step installation procedure, this book will be a valuable resource as you install Red Hat Linux Advanced Server.

This manual will walk you through a typical installation using the Red Hat Linux Advanced Server CD-ROMs. Once you have completed the installation as outlined in this manual, you will have a fully functioning Red Hat Linux Advanced Server system.

Document Conventions

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

`command`

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

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

`filename`

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

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

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

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

application

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

Use Netscape Navigator to browse the Web.

[key]

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

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

[key]-[combination]

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

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

text found on a GUI interface

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

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

top level of a menu on a GUI screen or window

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

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

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

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

button on a GUI screen or window

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

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

`computer output`

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

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

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

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

`prompt`

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

```
$
```

```
#
```

```
[stephen@maturin stephen]$
```

```
leopard login:
```

`user input`

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

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

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

Note

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

Tip

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

Important

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

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

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

How to Use This Manual

This manual focuses on a CD-ROM based installation, so it is ideal for users (both new and old) who want a quick and simple installation solution. It will help you prepare your system, walk you through the installation, and assist you in the configuration of Red Hat Linux Advanced Server.

If you are an experienced user who wants to perform a Red Hat Linux Advanced Server CD-ROM installation, and you do not need a review of the basics, you can skip ahead to Chapter 3, *Installing Red Hat Linux Advanced Server* to begin the installation process.

Tip

Refer to the **Red Hat Frequently Asked Questions** for answers to questions and problems that may occur before, during, or after the installation. You will find the FAQ online at:

http://www.redhat.com/support/docs/faqs/rhl_general_faq/

We Need Feedback!

If you spot a typo in the *Official Red Hat Linux Advanced Server Installation Guide*, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla (http://bugzilla.redhat.com/bugzilla/) against the component *rhl-ig-as-x86*.

Be sure to mention the manual's identifier:

```
rhl-ig-as-x86(EN)-2.1-Print-RHI (2002-03-22T11:09-0400)
```

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

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

If you have a support question (for example, if you need help configuring X, or if you are not sure how to partition your hard drive[s]), please register your product at http://www.redhat.com/apps/activate/ and use the online support system for those type of requests.

Part I Installing Red Hat Linux Advanced Server

1 Steps to Get You Started

Before you install Red Hat Linux, you should perform the following steps:

1.1 Do You Have the Right Components?

If you have purchased an Official Red Hat Linux Advanced Server product, you are ready to go. However, mistakes occasionally happen, so now is a good time to double-check the contents of your product.

A black, red, and white Registration Information card is included with your product. A list of the contents of your boxed set version is on the back of the card. Please read over the list and check to make sure that you have all the CDs and manuals that are included with the version of Red Hat Linux Advanced Server that you purchased.

If you have purchased an Official Red Hat Linux Advanced Server product from Red Hat, Inc. (or one of its distributors), and you are missing one or more of the items listed, please let us know. Contact information is also available on the Registration Information card.

How to identify our official boxed set: The bottom of our box has an ISBN number next to one of the bar codes. That ISBN number should be in this form:

`1-58569-`*`x`*`-`*`y`*

(The *x* and *y* will be unique numbers.)

Red Hat partners with companies (international and domestic) so that we can makeRed Hat Linux Advanced Server available to you in the most convenient form. Because of these partnerships, you might find that your Red Hat Linux Advanced Server product may not have been actually produced by Red Hat.

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

1.1.1 Where to Find Other Manuals

If your particular product did not include all of the printed Red Hat Linux Advanced Server manuals, you can find them online or on the Red Hat Linux Advanced Server Documentation CD included with your official Red Hat Linux Advanced Server product.

To find the manuals in both HTML and PDF formats online, go to http://www.redhat.com/docs.

1.1.2 Registering Your Product

If you have purchased an Official Red Hat Linux Advanced Server product, you should register your product. Registration offers many useful services, such as installation support, access to Red Hat Network, and more.

To register your product, go to http://www.redhat.com/apps/activate/. You will find your *Product ID* on the Registration Information card in your Official Red Hat Linux Advanced Server boxed set. Once registered, you will have access to all the extras that Red Hat provides to its registered users.

For more information on registering and the scope of Red Hat's technical support offerings, see Appendix D, *Getting Technical Support*.

1.2 Is Your Hardware Compatible?

Hardware compatibility is particularly important if you have an older system or a system that you built yourself. Red Hat Linux Advanced Server 2.1 should be compatible with most hardware in systems that were factory built within the last two years. However, hardware specifications change almost daily, so it is hard to guarantee that your hardware will be 100% compatible.

The most recent list of supported hardware can be found at http://hardware.redhat.com/hcl/.

Refer to the *Red Hat Cluster Manager Installation and Administration Guide* for specific details on hardware installation and supplemental hardware information.

1.3 Do You Have Enough Disk Space?

Nearly every modern-day operating system (OS) uses **disk partitions**, and Red Hat Linux Advanced Server is no exception. When you install Red Hat Linux Advanced Server, you may have to work with disk partitions. If you have not worked with disk partitions before (or need a quick review of the basic concepts) read Appendix F, *An Introduction to Disk Partitions* before proceeding.

The disk space used by Red Hat Linux Advanced Server must be separate from the disk space used by other OSes you may have installed on your system, such as Windows, OS/2, or even a different version of Linux. At least two partitions (`/` and `swap`) must be dedicated to Red Hat Linux Advanced Server.

Before you start the installation process, one of the following conditions must be met:

- Your computer must have enough *unpartitioned*[1] disk space for the installation of Red Hat Linux Advanced Server.

[1] Unpartitioned disk space means that the hard drive(s) you are installing to have not been divided into sections for data. When you partition a disk, each partition will behave like a separate disk drive.

- You must have one or more partitions that may be deleted, thereby freeing up enough disk space to install Red Hat Linux Advanced Server.

1.3.1 Installation Disk Space Requirements

Note

These recommendations are based on an installation that only installs one language (such as English). If you plan to install multiple languages to use on your system, you should increase the disk space requirements.

See Section 1.5, *Which Installation Class is Best For You?*, for further information regarding disk space requirements for your specific installation needs.

Advanced Server

An Advanced Server installation, choosing to install GNOME[2] or KDE[3], requires at least 1.0 GB of free space. Choosing both GNOME and KDE requires at least 1.3 GB of free disk space.

Custom

A Custom installation requires 350 MB for a minimal installation (without the Advanced Server packages) and at least 3.2 GB of free space if every package is selected.

If you are not sure that you meet these conditions, or if you want to know how to create free disk space for your Red Hat Linux Advanced Server installation, please refer to Appendix F, *An Introduction to Disk Partitions*.

1.4 Can You Install Using the CD-ROM?

There are several methods that can be used to install Red Hat Linux Advanced Server. This manual focuses on installing from the CD-ROM. For instructions on alternative installation methods, refer to Chapter 4, *Installing Red Hat Linux Advanced Server via Text Mode*.

[2] Acronym for GNU Network Object Model Environment. GNOME is part of the GNU project and part of the free software, or open source, movement. GNOME is a Windows-like desktop system that works on UNIX and UNIX-like systems. The main objective of GNOME is to provide a user-friendly suite of applications and an easy-to-use desktop.

[3] Acronym for K Desktop Environment. A network-transparent, contemporary desktop environment for UNIX workstations, KDE is part of the open source movement. It is free to anyone and its source code is available to anyone to modify.

Installing from a CD-ROM requires that you have purchased a Red Hat Linux Advanced Server 2.1 product , or you have a Red Hat Linux Advanced Server CD-ROM, and you have a CD-ROM drive. Most new computers will allow booting from the CD-ROM. If your system will support booting from the CD-ROM, it is an easy way to begin a local CD-ROM installation.

Your BIOS may need to be changed to allow booting from your CD-ROM drive. For more information about editing your BIOS, see Section 3.2.1, *Booting the Installation Program*.

1.4.1 Alternative Boot Methods

If you cannot boot from the CD-ROM drive, the following alternative boot method is available:

Local Boot Disk

If you need a **local boot disk**[4], you must create it. The local boot disk image file, `boot.img`, is located in the `images` directory on your Red Hat Linux Advanced Server CD-ROM. Refer to Section 1.4.2, *Making Installation Diskettes*, for more information on making a boot disk.

Note

USB Floppies — You can also boot the Red Hat Linux Advanced Server installation program using a USB floppy as a boot disk (if your system supports booting from a USB floppy).

Note

Although it is not required to boot your installation, you may occasionally find that a driver disk is needed to continue with the installation. Appendix G, *Driver Disks* explains why a driver disk may be necessary for your installation, and how to obtain one if needed.

1.4.2 Making Installation Diskettes

You may need to create a diskette from an **image file**; for example, you may need to use updated diskette images obtained from the Red Hat Linux errata page (http://www.redhat.com/apps/support/updates.html) or you may need to create a boot disk.

[4] A boot disk is a diskette you create during an installation that can later be used to boot the operating system. Normally, your computer boots from a hard disk, but if the hard disk is damaged, you can boot the computer from a bootable diskette.

An image file contains an exact copy (or image) of a diskette's contents. Since a diskette contains filesystem information in addition to the data contained in files, the image file is not usable until it has been written to a diskette.

To start, you will need a blank, formatted, high-density (1.44MB), 3.5-inch diskette. You will need access to a computer with a 3.5-inch diskette drive. The computer must be able to run either an MS-DOS program or the `dd` utility found on most Linux-like operating systems.

The `images` directory on your Red Hat Linux Advanced Server CD-ROM contains the boot images for Red Hat Linux/x86. Once you have selected the proper image (such as `boot.img` for a CD-ROM-based installation or `bootnet.img` for a network installation), transfer the image file onto a diskette.

Using the `rawrite` Utility

To make a diskette using MS-DOS, use the `rawrite` utility included on the Red Hat Linux Advanced Server CD-ROM in the `dosutils` directory. First, label a blank, formatted 3.5-inch diskette appropriately (such as "Boot Disk" or "Updates Disk"). Insert it into the diskette drive. Then, use the following commands (assuming your CD-ROM is drive `d:`):

```
C:\> d:
D:\> cd \dosutils
D:\dosutils> rawrite
Enter disk image source file name: ..\images\boot.img
Enter target diskette drive: a:
Please insert a formatted diskette into drive A: and
press --ENTER-- : [Enter]
D:\dosutils>
```

First, `rawrite` asks you for the filename of a diskette image; enter the directory and name of the image you wish to write (for example, `..\images\boot.img`). Then `rawrite` asks for a diskette drive to write the image to; enter `a:`. Finally, `rawrite` asks for confirmation that a formatted diskette is in the drive you have selected. After pressing [Enter] to confirm, `rawrite` copies the image file onto the diskette. If you need to make another diskette, label that diskette, and run `rawrite` again, specifying the appropriate image file.

Using the dd Command

To make a diskette under Linux (or any other Linux-like operating system), you must have permission to write to the device representing a 3.5-inch diskette drive (known as `/dev/fd0` under Linux).

First, label a blank, formatted diskette appropriately (such as "Boot Disk" or "Updates Disk"). Insert it into the diskette drive (but do not issue a `mount`[5] command). After mounting the Red Hat Linux

[5] When you mount a floppy or CD-ROM, you make that device's contents available to you. See the *Official Red Hat Linux Getting Started Guide* for more information.

Advanced Server CD-ROM, change to the directory containing the desired image file, and use the following command (changing the name of the image file and diskette device as appropriate):

```
# dd if=boot.img of=/dev/fd0 bs=1440k
```

To make another diskette, label that diskette, and run `dd` again, specifying the appropriate image file.

1.5 Which Installation Class is Best For You?

Usually, Red Hat Linux Advanced Server is installed on its own disk partition or set of partitions, or over another installation of Linux.

WARNING

Installing Red Hat Linux Advanced Server over another installation of Linux (including Red Hat Linux) does *not* preserve any information (files or data) from a prior installation. Make sure you save or back up any important files!

Red Hat Linux Advanced Server provides two different classes, or types, of installations:

Advanced Server

An Advanced Server installation is specifically targeted at server installations capable of high levels of availablity through load balancing and failover capacities. The Advanced Server configuration includes, at your option, the ability to install a default X Window System environment with management, as well as the necessary components for clustering together two or more systems to achieve higher levels of performance to meet the demands of high performance server environments.

Custom

A Custom installation allows you the greatest flexibility during your installation. You choose your boot loader, which packages you want, and more. Custom installations are most appropriate for those users more familiar with Red Hat Linux installations and for those afraid of losing complete flexibility.

These classes give you the option of simplifying the installation process (with some potential for loss of configuration flexibility), or retaining flexibility with a slightly more complex installation process. Next, take a detailed look at each class, so you can decide which one is right for you.

1.5.1 Advanced Server Installation

An Advanced Server installation is most appropriate for server installations capable of high levels of availablity through load balancing and failover capacities.

Below are the minimum recommended disk space requirements for an Advanced Server installation where only one language (such as English) will be installed.

Note

The *minimum recommended disk space requirements* as listed below, are just minimum recommendations for the installation itself. You should adjust these disk space requirements as appropriate for your specific computing needs (such as disk space for personal files, additional applications you may install at a later time, and so on).

- Advanced Server (minimum, no graphical interface): 800 MB
- Advanced Server (default, choosing GNOME or KDE): 1 GB
- Advanced Server (choosing GNOME and KDE): 1.3 GB
- Advanced Server (choosing everything, GNOME and KDE): 1.5 GB

If you plan to choose all group packages, as well as select additional individual packages, you may want to allow yourself 2.0 GB or more of disk space. This will provide space where additional data may be written.

What an Advanced Server Installation Will Do

An Advanced Server installation, with automatic partitioning, will create the following partitions:

Note

The partitions represented below were created on a system with 9 GB of hard drive space and 512 MB of RAM. Depending on the amount of hard drive space and memory you have available, these values may differ slightly.

- A 47 MB partition (mounted as `/boot`) in which the Linux kernel and related files are kept.
- A partition of at least 4877 MB (mounted as `/usr`).

- The size of the swap partition is determined by the amount of RAM in your system and the amount of space available on your hard drive. If you have 128 MB of RAM, then the swap partition created can be 128 MB – 256 MB (twice your RAM), depending on how much disk space is available.

  For this example, a 1020 MB swap partition (mounted as <swap>) is created.
- A 2738 MB ext partition.
- A partition of at least 2094 MB (mounted as `/home`).
- A 384 MB partition (mounted as `/`).
- A 259 MB partition (mounted as `/var`).

This partitioning scheme offers a reasonably flexible filesystem configuration for most server tasks.

1.5.2 Custom Installations

The *Custom installation* allows you the most flexibility during your installation. During a Custom installation, you have complete control over the packages that will be installed on your system.

The recommended disk space requirements for a Custom installation are as follows:

Note

The *minimum recommended disk space requirements* as listed below, are just minimum recommendations for the installation itself. You should adjust these disk space requirements as appropriate for your specific computing needs (such as disk space for personal files, additional applications you may install at a later time, and so on).

- Custom (minimum, without Advanced Server packages): 350 MB
- Custom (minimum, with Advanced Server packages): 500 MB
- Custom (default): 1 GB
- Custom (choosing everything): 3.2 GB

What a Custom Installation Will Do

As you might guess from the name, a custom installation puts the emphasis on flexibility. You have complete control over which packages will be installed on your system.

If you choose automatic partitioning, a Custom installation will create the following partitions:

Note

The partitions represented below were created on a system with 9 GB of hard drive space and 512 MB of RAM. Depending on the amount of hard drive space and memory you have available, these values may differ slightly.

- The size of the swap partition is determined by the amount of RAM in your system and the amount of space available on your hard drive. If you have 128 MB of RAM, then the swap partition created can be 128 MB – 256 MB (twice your RAM), depending on how much disk space is available.

  For this example, a 1020 MB swap partition (mounted as <swap>) is created.

- A 47 MB partition (mounted as `/boot`) in which the Linux kernel and related files reside.
- A 2609 MB root partition (mounted as `/`) in which all other files are stored (the exact size of this partition is dependent on your available disk space).

2 System Requirements Table

The most recent list of supported hardware can be found at http://hardware.redhat.com/hcl/.

Refer to the *Red Hat Cluster Manager Installation and Administration Guide* for specific details on hardware installation and supplemental hardware information.

This chapter provides you with a system requirements table, which will help you keep a record of your current system settings and requirements. Enter information about your system in the table provided as a handy reference to help make your Red Hat Linux Advanced Server installation go more smoothly.

Table 2–1 System Requirements Table

hard drive(s): type, label, size; ex: IDE hda=1.2 GB	
partitions: map of partitions and mount points; ex: `/dev/hda1=/home`, `/dev/hda2=/` (fill this in once you know where they will reside)	
memory: amount of RAM installed on your system; ex: 64 MB, 128 MB	
CD-ROM: interface type; ex: SCSI, IDE (ATAPI)	
SCSI adapter: if present, make and model number; ex: BusLogic SCSI Adapter, Adaptec 2940UW	
network card: if present, make and model number; ex: Tulip, 3COM 3C590	

mouse: type, protocol, and number of buttons; ex: generic 3 button PS/2 mouse, MouseMan 2 button serial mouse	
monitor: make, model, and manufacturer specifications; ex: Optiquest Q53, ViewSonic G773	
video card: make, model number and size of VRAM; ex: Creative Labs Graphics Blaster 3D, 8MB	
sound card: make, chipset and model number; ex: S3 SonicVibes, Sound Blaster 32/64 AWE	
IP, DHCP, and BOOTP addresses: four numbers, separated by dots; ex: 10.0.2.15	
netmask: four numbers, separated by dots; ex: 255.255.248.0	
gateway IP address: four numbers, separated by dots; ex: 10.0.2.245	
one or more name server IP addresses (DNS): one or more sets of dot-separated numbers; ex: 10.0.2.1	

domain name: the name given to your organization; ex: Red Hat's would be `redhat.com`	
hostname: the name of your computer; your personal choice of names; ex: `cookie`, `southpark`	

If any of these networking requirements or terms are unfamiliar to you, contact your network administrator for assistance.

3 Installing Red Hat Linux Advanced Server

This chapter explains how to install Red Hat Linux Advanced Server from the CD-ROM using the graphical, mouse-based installation program. The following topics are discussed:

- Getting familiar with the installation program's user interface
- Starting the installation program
- Selecting an installation method
- Configuration steps during the installation (language, keyboard, mouse, etc.)
- Finishing the installation

3.1 The Installation Program User Interface

If you have used a **graphical user interface (GUI)** before, you will be familiar with this process; simply use your mouse to navigate the screens, "click" buttons, or enter text fields. You can also navigate through the installation using the [Tab] and [Enter] keys.

3.1.1 A Note about Virtual Consoles

The Red Hat Linux Advanced Server installation program offers more than the dialog boxes of the installation process. Several different kinds of diagnostic messages are available to you, in addition to providing a way to enter commands from a shell prompt. The installation program displays these messages on five **virtual consoles**, among which you can switch using a single keystroke combination.

These virtual consoles can be helpful if you encounter a problem while installing Red Hat Linux Advanced Server. Messages displayed on the installation or system consoles can help pinpoint a problem. Please see Table 3–1, *Console, Keystrokes, and Contents* for a listing of the virtual consoles, keystrokes used to switch to them, and their contents.

Table 3–1 Console, Keystrokes, and Contents

Console	Keystrokes	Contents
1	[Ctrl]-[Alt]-[F1]	installation dialog
2	[Ctrl]-[Alt]-[F2]	shell prompt
3	[Ctrl]-[Alt]-[F3]	install log (messages from installation program)

Console	Keystrokes	Contents
4	[Ctrl]-[Alt]-[F4]	system-related messages
5	[Ctrl]-[Alt]-[F5]	other messages
7	[Ctrl]-[Alt]-[F7]	X graphical display

Generally, there is no reason to leave the default console (virtual console #7) unless you are attempting to diagnose installation problems.

3.2 Starting the Installation Program

To start the installation, you must first boot the installation program. Please make sure you have all the resources you will need for the installation. If you have already read through Chapter 1, *Steps to Get You Started* , and followed the instructions, you should be ready to begin.

Note

Occasionally, some hardware components require a **driver disk** during the installation. A driver disk adds support for hardware that is not otherwise supported by the installation program. Refer to Appendix G, *Driver Disks* for more information.

3.2.1 Booting the Installation Program

Note

To create a boot disk, refer to Section 1.4.2, *Making Installation Diskettes*.

You can boot the Red Hat Linux Advanced Server installation program using any one of the following media (depending upon what your system can support):

- *Bootable CD-ROM* — Your machine supports a bootable CD-ROM drive and you want to perform a local CD-ROM installation.
- *Local boot disk* — Your machine will not support a bootable CD-ROM and you want to install from a local CD-ROM or a hard drive.
- *Network boot disk* — Use a network boot disk to install via NFS, FTP, and HTTP.

Insert the boot disk into your computer's first diskette drive and reboot (or boot using the CD-ROM, if your computer supports booting from it). Your BIOS settings may need to be changed to allow you to boot from the diskette or CD-ROM.

Tip

To change your BIOS settings, watch the instructions provided on your display when your computer first begins to boot. Often you will see a line of text telling you to press the [Del] or [F1] key to enter the BIOS settings.

Once you've entered your BIOS setup program, find the section where you can alter your boot sequence. The default is often C, A or A, C (depending on whether you boot from your hard drive [C] or a diskette drive [A]). Change this sequence so that the CD-ROM is first in your boot order and that C or A (whichever is your typical boot default) is second. This instructs the computer to first look at the CD-ROM drive for bootable media; if it does not find bootable media on the CD-ROM drive, it will then check your hard drive or diskette drive.

Save your changes before exiting the BIOS. For more information, please refer to the documentation that came with your system.

After a short delay, a screen containing the `boot:` prompt should appear. The screen contains information on a variety of boot options. Each boot option also has one or more help screens associated with it. To access a help screen, press the appropriate function key as listed in the line at the bottom of the screen.

As you boot the installation program, be aware of two issues:

- Once you see the `boot:` prompt, the installation program will automatically begin if you take no action within the first minute. To disable this feature, press one of the help screen function keys.
- If you press a help screen function key, there will be a slight delay while the help screen is read from the boot media.

Normally, you only need to press [Enter] to boot. Watch the boot messages to see if the Linux kernel detects your hardware. If your hardware is properly detected, please continue to the next section. If it does not properly detect your hardware, you may need to restart the installation in expert mode.

Additional Boot Options

If you do not wish to perform a CD-ROM GUI installation, you can start a text mode installation using the following boot command:

```
boot: text
```

For text mode installation instructions, please refer to Chapter 4, *Installing Red Hat Linux Advanced Server via Text Mode*.

If you are having trouble booting into the graphical installation program, you can try to boot using the no framebuffer (`nofb`) boot option.

At the boot command, enter the following:

```
boot: nofb
```

This option allows you to use the graphical installation program without using a framebuffer.

Enter expert mode using the following boot command:

```
boot: linux expert
```

If you need to perform the installation in **serial mode**, type the following command:

```
boot: linux console=<device>
```

In the above command, *<device>* should be the device you are using (such as ttyS0 or ttyS1). For example, `linux console=ttyS0,115200n8`.

Kernel Options

Options can also be passed to the kernel. For example, to instruct the kernel to use all the RAM in a system with 128 MB of RAM, enter:

```
boot: linux mem=128M
```

After entering any options, press [Enter] to boot using those options.

If you need to specify boot options to identify your hardware, please write them down. The boot options will be needed during the boot loader configuration portion of the installation (please see Section 3.14, *Boot Loader Installation* for more information).

Booting Without Diskettes

The Red Hat Linux Advanced Server CD-ROM can be booted by computers that support bootable CD-ROMs. Not all computers support this feature, so if your system cannot boot from the CD-ROM, there is one other way to start the installation without using a boot disk. The following method is specific to x86-based computers only.

If you have MS-DOS installed on your system, you can boot directly from the CD-ROM drive without using a boot disk. To do this (assuming your CD-ROM is drive `d:`), use the following commands:

```
C:\> d:
D:\> cd \dosutils
D:\dosutils> autoboot.bat
```

This method will not work if run in a DOS window — the `autoboot.bat` file must be executed with DOS as the only operating system. In other words, Windows cannot be running.

If your computer cannot boot directly from CD-ROM (and you cannot use a DOS-based autoboot), you will have to use a boot diskette to get things started.

3.3 Selecting an Installation Method

What type of installation method do you wish to use? The following installation methods are available:

CD-ROM

If you have a CD-ROM drive and the Red Hat Linux Advanced Server CD-ROM, you can use this method. You will need a boot disk or a bootable CD-ROM. PCMCIA boot and driver disks may also be used. Continue reading this chapter for further instructions.

Hard Drive

If you have copied the Red Hat Linux Advanced Server ISO images to a local hard drive, you can use this method. You will need a boot disk. PCMCIA boot and driver disks may also be used. Refer to Chapter 4, *Installing Red Hat Linux Advanced Server via Text Mode*, for hard drive installation instructions.

NFS Image

If you are installing from an NFS server which is exporting the Red Hat Linux Advanced Server CD-ROM(s) or a mirror image of Red Hat Linux, you can use this method. You will need a network boot disk. PCMCIA boot and driver disks may also be used. Refer to Chapter 4, *Installing Red Hat Linux Advanced Server via Text Mode*, for network installation instructions. Please note that NFS installations may also be performed in GUI mode.

FTP

If you are installing directly from an FTP server, use this method. You will need a network boot disk. PCMCIA boot and driver disks may also be used. Refer to Chapter 4, *Installing Red Hat Linux Advanced Server via Text Mode*, for FTP installation instructions.

HTTP

If you are installing directly from an HTTP (Web) server, use this method. You will need a network boot disk. PCMCIA boot and driver disks may also be used. Refer to Chapter 4, *Installing Red Hat Linux Advanced Server via Text Mode*, for HTTP installation instructions.

3.4 Beginning the Installation

If you are planning to install via CD-ROM using the graphical interface, please continue reading.

If you would rather perform a text mode installation, reboot your system and at the `boot:` prompt, type `text`. Refer to Chapter 4, *Installing Red Hat Linux Advanced Server via Text Mode* for a more information.

3.4.1 Installing from CD-ROM

To install Red Hat Linux Advanced Server from a CD-ROM, choose the **CD-ROM** option from the boot loader screen and select **OK**. When prompted, insert the Red Hat Linux Advanced Server CD into your CD-ROM drive (if you did not boot from the CD-ROM). Once the CD is in the CD-ROM drive, select **OK**, and press [Enter].

The installation program will then probe your system and attempt to identify your CD-ROM drive. It will start by looking for an IDE (also known as an ATAPI) CD-ROM drive. If found, you will continue to the next stage of the installation process (see Section 3.5, *Language Selection*).

Note

To abort the installation process at this time, reboot your machine and then eject the boot diskette or CD-ROM. You can safely cancel the installation at any point before the **About to Install** screen. See Section 3.23, *Preparing to Install* for more information.

If a CD-ROM drive is not detected, you will be asked what type of CD-ROM drive you have. Choose from the following types:

SCSI

Select this if your CD-ROM drive is attached to a supported SCSI adapter; the installation program will then ask you to choose a SCSI driver. Choose the driver that most closely resembles your adapter. You may specify options for the driver if necessary; however, most drivers will detect your SCSI adapter automatically.

Other

If your CD-ROM drive is neither an IDE nor a SCSI, it is an other. Sound cards with proprietary CD-ROM interfaces are good examples of this type of CD-ROM. The installation program will display a list of drivers for supported CD-ROM drives — choose a driver and, if necessary, specify any driver options.

Tip

A partial list of optional parameters for CD-ROM drives can be found in the *Official Red Hat Linux Reference Guide*, in the *General Parameters and Modules* appendix.

What If the IDE CD-ROM Was Not Found?

If you have an IDE (ATAPI) CD-ROM, but the installation program fails to find your IDE (ATAPI) CD-ROM and asks you what type of CD-ROM drive you have, try the following boot command. Restart the installation, and at the `boot:` prompt enter **`linux hdX=cdrom`**. Replace the ***X*** with one of the following letters, depending on the interface the unit is connected to, and whether it is configured as master or slave (also known as primary and secondary):

- `a` — first IDE controller, master
- `b` — first IDE controller, slave
- `c` — second IDE controller, master
- `d` — second IDE controller, slave

If you have a third and/or fourth controller, continue assigning letters in alphabetical order, going from controller to controller, and master to slave.

3.5 Language Selection

Using your mouse, select the language you would prefer to use for the installation and as the system default (see Figure 3–1, *Language Selection*).

Selecting the appropriate language will also help target your time zone configuration later in the installation. The installation program will try to define the appropriate time zone based on what you specify on this screen.

Once you select the appropriate language, click **Next** to continue.

Figure 3–1 Language Selection

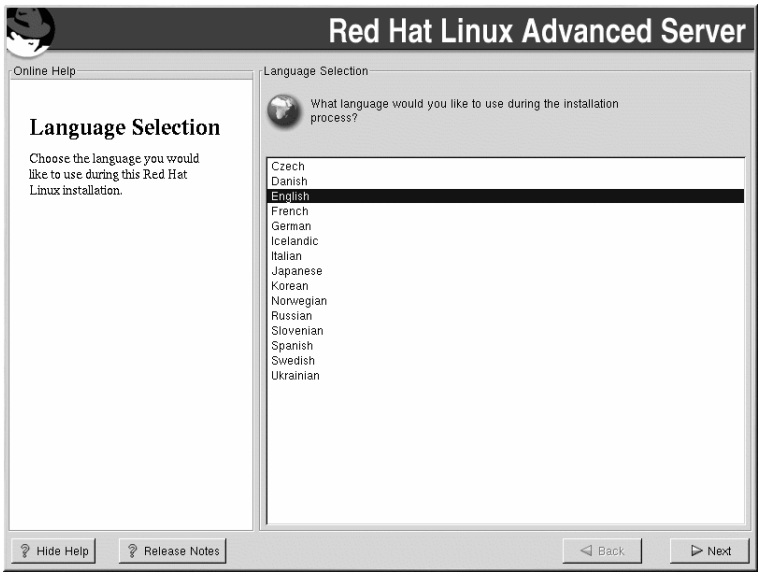

3.6 Keyboard Configuration

Choose the keyboard model that best fits your system (see Figure 3–2, *Keyboard Configuration*). If you cannot find an exact match, choose the best **Generic** match for your keyboard type (for example, **Generic 101-key PC**).

Next, choose the correct layout type for your keyboard (for example, U.S. English).

Creating special characters with multiple keystrokes (such as Ñ, Ô, and Ç) is done using "dead keys" (also known as compose key sequences). Dead keys are enabled by default. If you do not wish to use them, select **Disable dead keys**.

Tip

The following example will help you determine if you need dead keys enabled. An example of a dead key is the backspace (^H) key on a US English 101 Standard Keyboard. Dead keys are not exclusive to non-English keyboards.

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

Once you have made the appropriate selections, click **Next** to continue.

Figure 3–2 Keyboard Configuration

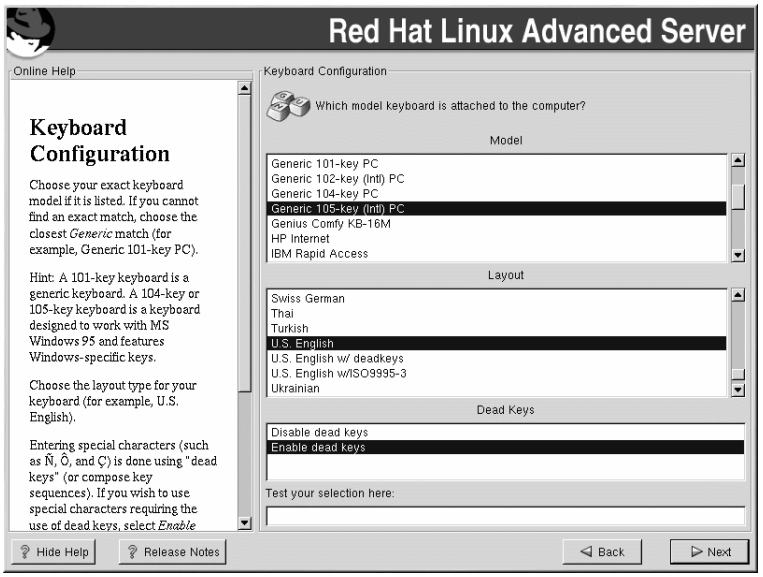

Tip

To change your keyboard type after you have installed Red Hat Linux Advanced Server, log in as root and use the `/usr/sbin/kbdconfig` command. Alternatively, you can type `setup` at the `root` prompt.

To become root, type `su -` at the shell prompt in a terminal window and then press [Enter]. Then, enter the root password.

Tip

To re-enable dead keys (assuming you chose to disable them during the installation), you will also need to comment out the line enabling dead keys in the `XF86Config-4` file (or, it you are using XFree86 version 3, the `XF86Config` file) in `/etc/X11`.

An **InputDevice** section, that would disable dead keys on a keyboard layout that used dead keys (for example, German), would look similar to the following:

```
Section "InputDevice"
        Identifier  "Keyboard0"
        Driver      "keyboard"
        Option      "XkbRules" "xfree86"
        Option      "XkbModel" "pc101"
        Option      "XkbLayout" "de"
        Option      "XkbVariant" "nodeadkeys"
EndSection
```

By default, the keyboard layout for various languages which use dead keys should have them enabled unless the `nodeadkeys` option is present.

3.7 Mouse Configuration

Choose the correct mouse type for your system. If you cannot find an exact match, choose a mouse type that you are sure is compatible with your system (see Figure 3–3, *Mouse Configuration*).

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

If you are not sure if your mouse is a serial or USB mouse, check the vendor documentation.

If you cannot find a mouse that you are sure is compatible with your system, select one of the **Generic** entries, based on your mouse's number of buttons, and its interface.

Figure 3–3 Mouse Configuration

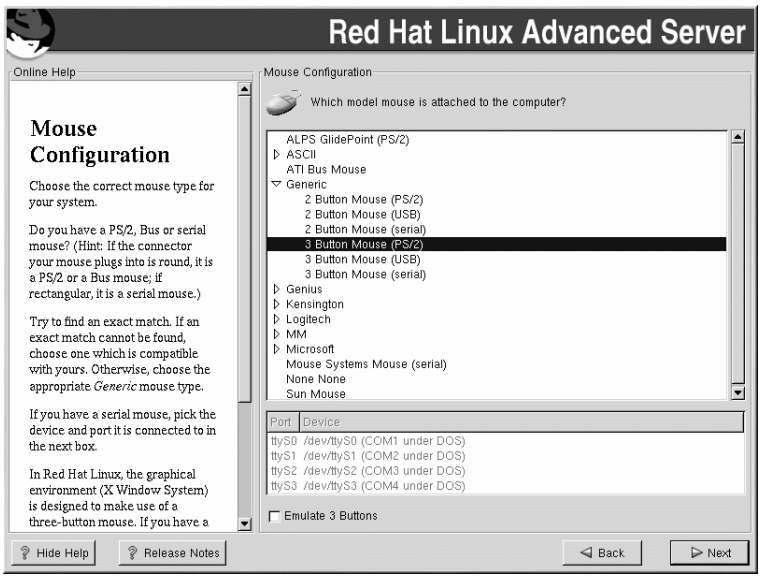

If you have a PS/2 or a bus mouse, you do not need to pick a port and device. If you have a serial mouse, you should choose the correct port and device that your serial mouse is on.

The **Emulate 3 Buttons** checkbox allows you to use a two-button mouse as if it had three buttons. In general, the X Window System is easier to use with a three-button mouse. If you select this checkbox, you can emulate a third, "middle" button by pressing both mouse buttons simultaneously.

Tip

To change your mouse configuration after you have completed the installation of Red Hat Linux Advanced Server, become root; then use the `/usr/sbin/mouseconfig` command from a shell prompt.

To configure your mouse to work as a left-handed mouse, reset the order of the mouse buttons. To do this, after you have booted your Red Hat Linux Advanced Server system, type `gpm -B 321` at the shell prompt.

3.8 Welcome to Red Hat Linux

The **Welcome** screen does not prompt you for any input. Please read over the help text in the left panel for additional instructions and information on where to register your Official Red Hat Linux Advanced Server product.

Please notice the **Hide Help** button at the bottom left corner of the screen. The help screen is open by default. If you do not want to view the help information, click on **Hide Help** to minimize the help portion of the screen.

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

3.9 Install Options

Choose the type of installation you would like to perform (see Figure 3–4, *Choosing Your Installation Type*). Red Hat Linux Advanced Server allows you to choose the installation type that best fits your needs. Your options are **Advanced Server** and **Custom**.

Figure 3–4 Choosing Your Installation Type

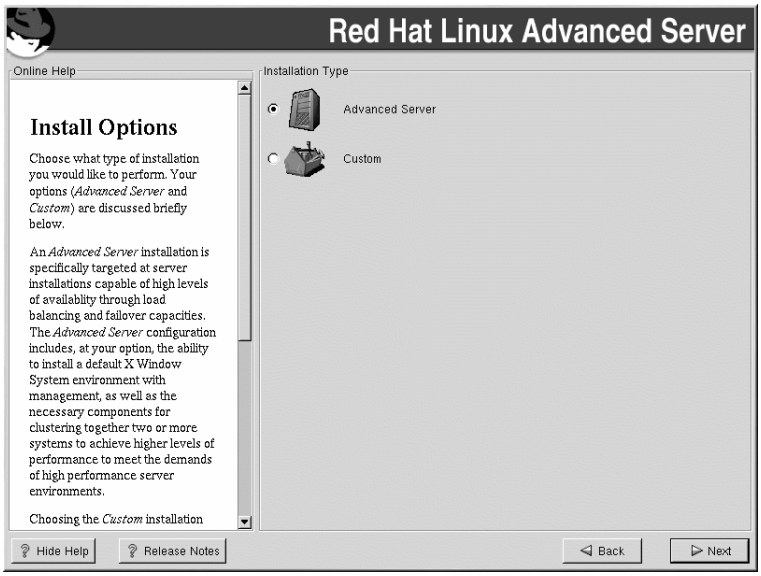

For more information about the different installation classes, please refer to Section 1.5, *Which Installation Class is Best For You?*.

3.10 Disk Partitioning Setup

Partitioning allows you to divide your hard drive into isolated sections, where each section behaves as its own hard drive. Partitioning is particularly useful if you run more than one operating system. If you are not sure how you want your system to be partitioned, read Appendix F, *An Introduction to Disk Partitions* for more information.

On this screen, you can choose to perform automatic partitioning, or manual partitioning using Disk Druid or fdisk (see Figure 3–5, *Disk Partitioning Setup*).

Automatic partitioning allows you to perform an installation without having to partition your drive(s) yourself. If you do not feel comfortable with partitioning your system, it is recommended that you *do not* choose to partition manually and instead let the installation program partition for you.

To partition manually, choose either the Disk Druid or fdisk (recommended for experts only) partitioning tool.

Figure 3–5 Disk Partitioning Setup

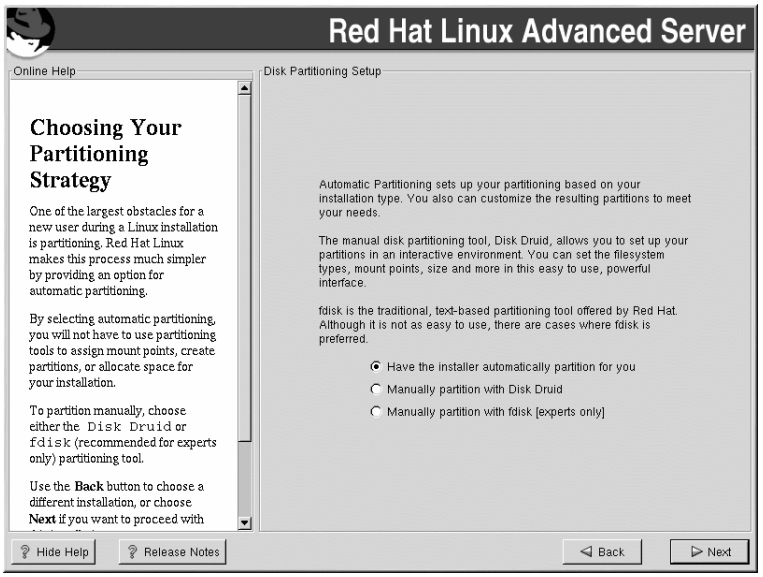

If you chose to manually partition using Disk Druid, refer to Section 3.12, *Partitioning Your System*.

If you chose to manually partition using fdisk, refer to Section 3.13, *Partitioning with fdisk*.

3.11 Automatic Partitioning

Figure 3–6 Automatic Partitioning

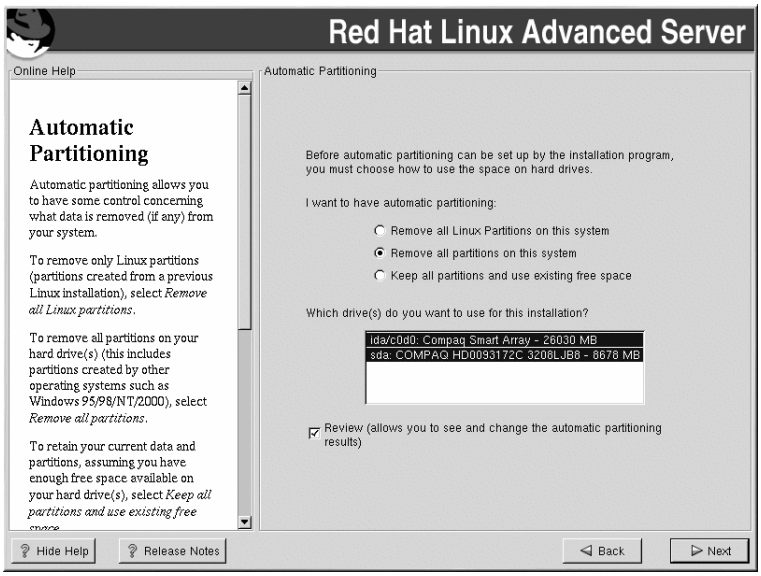

Automatic partitioning allows you to have some control concerning what data is removed (if any) from your system. Your options are:

- **Remove all Linux partitions on this system** — select this option to remove only Linux partitions (partitions created from a previous Linux installation). This will not remove other partitions you may have on your hard drive(s).
- **Remove all partitions on this system** — select this option to remove all partitions on your hard drive(s) (this includes partitions created by other operating systems such as Windows 95/98/NT/2000).

If you select this option, all data on the selected hard drive(s) will be removed by the installation program. Do not select this option if you have information that you want to keep on the hard drive(s) where you are installing Red Hat Linux Advanced Server.

- **Keep all partitions and use existing free space** — select this option to retain your current data and partitions, assuming you have enough free space available on your hard drive(s).

Using your mouse, choose the hard drive(s) on which you want Red Hat Linux Advanced Server to be installed. If you have two or more hard drives, you can choose which hard drive(s) should contain this installation. Unselected hard drives, and any data on them, will not be touched.

If you have a RAID card, be aware that some BIOSes do not support booting from the RAID card. In cases such as these, the `/boot` partition must be created on a partition outside of the RAID array, such as on a separate hard drive. An internal hard drive is necessary to use for partition creation with problematic RAID cards.

If you have chosen to automatically partition your system, you should select **Review** and manually edit your `/boot` partition.

To review and make any necessary changes to the partitions created by automatic partitioning, select the **Review** option. After selecting **Review** and clicking **Next** to move forward, you will see the partitions created for you in Disk Druid. You will also be able to make modifications to these partitions if they do not meet your needs.

Click **Next** once you have made your selections to proceed.

3.12 Partitioning Your System

If you chose automatic partitioning and did not select **Review**, please skip ahead to Section 3.16, *Network Configuration*.

If you chose automatic partitioning and selected **Review**, you can either accept the current partition settings (click **Next**), or modify the setup using Disk Druid, the manual partitioning tool.

If you chose **Manually partition with fdisk**, please skip ahead to Section 3.13, *Partitioning with fdisk*.

At this point, you must tell the installation program where to install Red Hat Linux Advanced Server. This is done by defining mount points for one or more disk partitions in which Red Hat Linux Advanced Server will be installed. You may also need to create and/or delete partitions at this time (refer to Figure 3–7, *Partitioning with Disk Druid*).

Note

If you have not yet planned how you will set up your partitions, refer to Appendix F, *An Introduction to Disk Partitions*. At a bare minimum, you need an appropriately-sized root partition, and a swap partition equal to twice the amount of RAM you have on the system.

Figure 3–7 Partitioning with Disk Druid

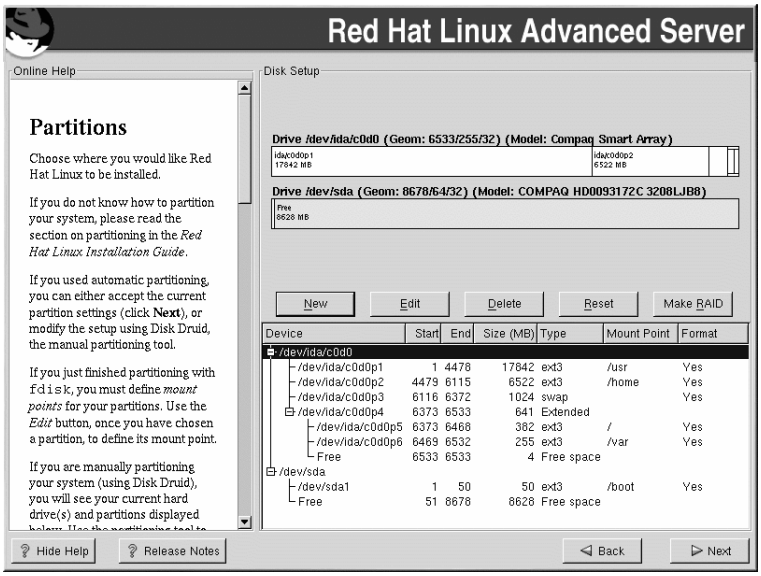

The partitioning tool used in Red Hat Linux Advanced Server 2.1 is Disk Druid. With the exception of certain esoteric situations, Disk Druid can handle the partitioning requirements for a typical Red Hat Linux Advanced Server installation.

3.12.1 Graphical Display of Hard Drive(s)

Disk Druid offers a graphical representation of your hard drive(s).

Using your mouse, click once to highlight a particular field in the graphical display. Double-click to edit an existing partition or to create a partition out of existing free space.

Above the display, you will see the **drive** name (such as /dev/hda), the **geom** (which shows the hard disk's geometry and consists of three numbers representing the number of cylinders, heads, and sectors as reported by the hard disk), and the **model** of the hard drive as detected by the installation program.

3.12.2 Disk Druid's Buttons

These buttons control Disk Druid's actions. They are used to change the attributes of a partition (for example the filesystem type and mount point) and also to create RAID devices. Buttons on this screen are also used to accept the changes you have made, or to exit Disk Druid. For further explanation, take a look at each button in order:

- **New**: Used to request a new partition. When selected, a dialog box appears containing fields (such as mount point and size) that must be filled in.
- **Edit**: Used to modify attributes of the partition currently selected in the **Partitions** section. Selecting **Edit** opens a dialog box. Some or all of the fields can be edited, depending on whether the partition information has already been written to disk.

  You can also edit free space as represented in the graphical display to create a new partition within that space. Either highlight the free space and then select the **Edit** button, or double-click on the free space to edit it.
- **Delete**: Used to remove the partition currently highlighted in the **Current Disk Partitions** section. You will be asked to confirm the deletion of any partition.
- **Reset**: Used to restore Disk Druid to its original state. All changes made will be lost if you **Reset** the partitions.
- **Make RAID**: **Make RAID** can be used if you want to provide software RAID redundancy to any or all disk partitions. *It should only be used if you have experience using RAID.* To read more about RAID, please refer to *RAID (Redundant Array of Independent Disks)* in the *Official Red Hat Linux Customization Guide*.

  To make a RAID device, you must first create software RAID partitions. Once you have created two or more software RAID partitions, select **Make RAID** to join the software RAID partitions into a RAID device.

3.12.3 Partition Fields

Above the partition hierarchy are labels which present information about the partitions you are creating. The labels are defined as follows:

- **Device**: This field displays the partition's device name.
- **Start**: This field shows the sector on your hard drive where the partition begins.
- **End**: This field shows the sector on your hard drive where the partition ends.
- **Size**: This field shows the partition's size (in MB).
- **Type**: This field shows the partition's type (for example, ext2, ext3, or vfat).
- **Mount Point**: A mount point is the location within the directory hierarchy at which a volume exists; the volume is "mounted" at this location. This field indicates where the partition will be mounted. If a partition exists, but is not set, then you need to define its mount point. Double-click on the partition or click the **Edit** button.
- **Format**: This field shows if the partition being created will be formatted.

3.12.4 Recommended Partitioning Scheme

Unless you have a reason for doing otherwise, we recommend that you create the following partitions:

- A swap partition (at least 32 MB) — swap partitions are used to support virtual memory. In other words, data is written to a swap partition when there is not enough RAM to store the data your system is processing. The size of your swap partition should be equal to twice your computer's RAM, or 32 MB, whichever amount is larger, but no more than 2048 MB (or 2 GB). In Disk Druid, the partition field for swap should look similar to the following:

  ```
  <Swap>  hda6  64M   64M  Linux swap
  ```

  For example, if you have 1 GB of RAM or less, your swap partition should be at least equal to the amount of RAM on your system, up to two times the RAM. For more than 1 GB of RAM, 2 GB of swap is recommended. Creating a large swap space partition will be especially helpful if you plan to upgrade your RAM at a later time.

- A `/boot` partition (50 MB) — the partition mounted on `/boot` contains the operating system kernel (which allows your system to boot Red Hat Linux Advanced Server), along with files used during the bootstrap process. Due to the limitations of most PC BIOSes, creating a small partition to hold these files is a good idea. For most users, a 50 MB boot partition is sufficient. In Disk Druid, the partition field for `/boot` should look similar to:

  ```
  /boot  hda1  50M   50M  Linux native
  ```

CAUTION

If your hard drive is more than 1024 cylinders (and your system was manufactured more than two years ago), you may need to create a `/boot` partition if you want the `/` (root) partition to use all of the remaining space on your hard drive.

CAUTION

If you have a RAID card, be aware that some BIOSes do not support booting from the RAID card. In cases such as these, the `/boot` partition must be created on a partition outside of the RAID array, such as on a separate hard drive.

- A `root` partition (350 MB - 3.2 GB) — this is where "`/`" (the root directory) will be located. In this setup, all files (except those stored in `/boot`) are on the root partition.

  A 350 MB partition will allow you to install a minimal Custom installation (without the Advanced Server packages), a 1.0 GB root partition will permit the equivalent of a default Advanced Server (with GNOME or KDE) installation (with *very* little free space), while a 3.2 GB root partition will let you perform a full Custom installation, choosing everything.

  In Disk Druid, the partition field for `/` should look similar to:

```
/  hda5  3734M   3734M   Linux native
```

3.12.5 Adding Partitions

To add a new partition, select the **New** button. A dialog box appears (see Figure 3–8, *Creating a New Partition*).

Note

You must dedicate at least one partition to Red Hat Linux Advanced Server, and optionally more. For more information, see Appendix F, *An Introduction to Disk Partitions*.

Figure 3–8 Creating a New Partition

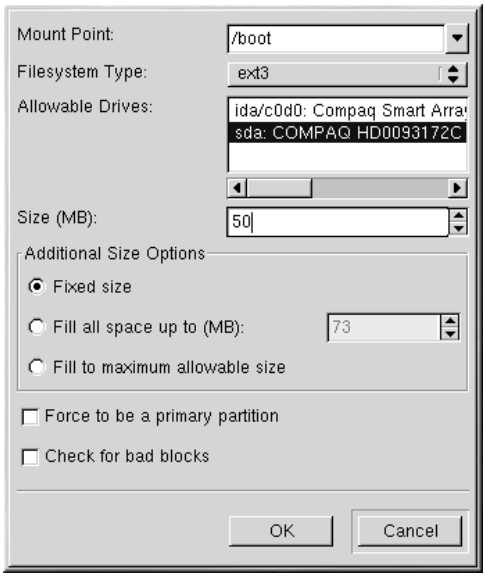

- **Mount Point**: Enter the partition's mount point. For example, if this partition should be the root partition, enter **/**; enter **/boot** for the **/boot** partition, and so on. You can also use the pulldown menu to choose the correct mount point for your partition.
- **Filesystem Type**: Using the pulldown menu, select the appropriate filesystem type for this partition. For more information on filesystem types, see *Filesystem Types* in Section 3.12.5.
- **Allowable Drives**: This field contains a list of the hard disks installed on your system. If a hard disk's box is highlighted, then a desired partition can be created on that hard disk. If the box is *not* checked, then the partition will *never* be created on that hard disk. By using different checkbox settings, you can have Disk Druid place partitions as you see fit, or let Disk Druid decide where partitions should go.
- **Size (Megs)**: Enter the size (in megabytes) of the partition. Note, this field starts with a "1" (one); unless changed, only a 1 MB partition will be created.
- **Additional Size Options**: Choose whether to keep this partition at a fixed size, to allow it to "grow" (fill up the available hard drive space) to a certain point, or to allow it to grow to fill any remaining hard drive space available.

If you choose **Fill all space up to (MB)**, you must give size constraints in the field to the right of this option. This allows you to keep a certain amount of space free on your hard drive for future use.

- **Force to be a primary partition**: Select whether the partition you are creating should be one of the first four partitions on the hard drive. If unselected, the partition created will be a logical partition. See Section F.1.3, *Partitions within Partitions — An Overview of Extended Partitions* , for more information.
- **Check for bad blocks**: Checking for bad blocks can help prevent data loss by locating the bad blocks on a drive and making a list of them to prevent using them in the future. If you wish to check for bad blocks while formatting each filesystem, please make sure to select this option.

  Selecting **Check for bad blocks** may dramatically increase your total installation time. Since most newer hard drives are quite large in size, checking for bad blocks may take a long time; the length of time depends on the size of your hard drive. If you choose to check for bad blocks, you can monitor your progress on virtual console #6.
- **Ok**: Select **Ok** once you are satisfied with the settings and wish to create the partition.
- **Cancel**: Select **Cancel** if you do not want to create the partition.

Filesystem Types

Red Hat Linux Advanced Server allows you to create different partition types, based on the filesystem they will use. The following is a brief description of the different filesystems available, and how they can be utilized.

- **ext2** — An ext2 filesystem supports standard Unix file types (regular files, directories, symbolic links, etc). It provides the ability to assign long file names, up to 255 characters. Versions prior to Red Hat Linux Advanced Server 2.1 used ext2 filesystems by default.
- **ext3** — The ext3 filesystem is based on the ext2 filesystem and has one main advantage — journaling. Using a journaling filesystem reduces time spent recovering a filesystem after a crash as there is no need to fsck[1] the filesystem.
- **software RAID** — Creating two or more software RAID partitions allows you to create a RAID device. For more information regarding RAID, refer to the chapter *RAID (Redundant Array of Independent Disks)* in the *Official Red Hat Linux Customization Guide*.
- **swap** — Swap partitions are used to support virtual memory. In other words, data is written to a swap partition when there is not enough RAM to store the data your system is processing.

1 The fsck application is used to check the filesystem for metadata consistency and optionally repair one or more Linux filesystems.

- **vfat** — The VFAT filesystem is a Linux filesystem that is compatible with Windows 95/NT long filenames on the FAT filesystem.

3.12.6 Editing Partitions

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

Note

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

3.12.7 Deleting a Partition

To delete a partition, highlight it in the **Partitions** section and click the **Delete** button. You will be asked to confirm the deletion.

Skip to Section 3.14, *Boot Loader Installation* for further installation instructions.

3.13 Partitioning with fdisk

This section applies only if you chose to use fdisk to partition your system.

To partition your system without using fdisk, please skip to Section 3.11, *Automatic Partitioning* for automatic partitioning or Section 3.12, *Partitioning Your System* for partitioning with Disk Druid.

If you have already completed disk partitioning, skip to Section 3.14, *Boot Loader Installation* for further installation instructions.

Unless you have previously used fdisk and understand how it works, we do not recommend that you use it. It is much easier for new users to accidentally corrupt or lose data using fdisk.

Disk Druid is easier to understand than fdisk. To exit fdisk, click **Back** to return to the previous screen, deselect fdisk, and then click **Next**.

If you have chosen to use fdisk, the next screen will prompt you to select a drive to partition using fdisk. Once you have chosen which drive to partition, you will be presented with the fdisk command screen. If you do not know what command to use, type [m] at the prompt for help.

When you are finished making partitions, type [w] to save your changes and quit. You will be taken back to the original fdisk screen where you can partition another drive or continue the installation.

Note

None of the changes you make take effect until you save them and exit fdisk using the w command. You can quit fdisk at any time without saving changes using the q command.

After you have partitioned your drive(s), click **Next**. You will need to use Disk Druid to assign mount points to the partitions you just created with fdisk.

You will not be able to add new partitions using Disk Druid, but you can edit mount points for the partitions you have already created. For each partition created with fdisk, click on the **Edit** button, choose the appropriate mount point for that partition from the pulldown menu, and click on **OK**.

3.14 Boot Loader Installation

In order to boot your Red Hat Linux Advanced Server system without a boot disk, you usually need to install a boot loader. You can choose to install either GRUB (selected by default) or LILO.

GRUB is a software boot loader that can be used to start Red Hat Linux Advanced Server on your computer. It can also start other operating systems, such as Windows 9*x*. Here, you will be asked how (or whether) you want to configure a boot loader and which one (GRUB or LILO).

Choose which boot loader you want to install. If you would rather use LILO, make sure it is selected instead of GRUB.

If you choose not to install GRUB or LILO for any reason, you will not be able to boot your Red Hat Linux Advanced Server system directly, and you will need to use another boot method (such as a boot diskette). Use this option only if you are sure you have another way of booting your Red Hat Linux Advanced Server system!

Figure 3–9 Boot Loader Installation

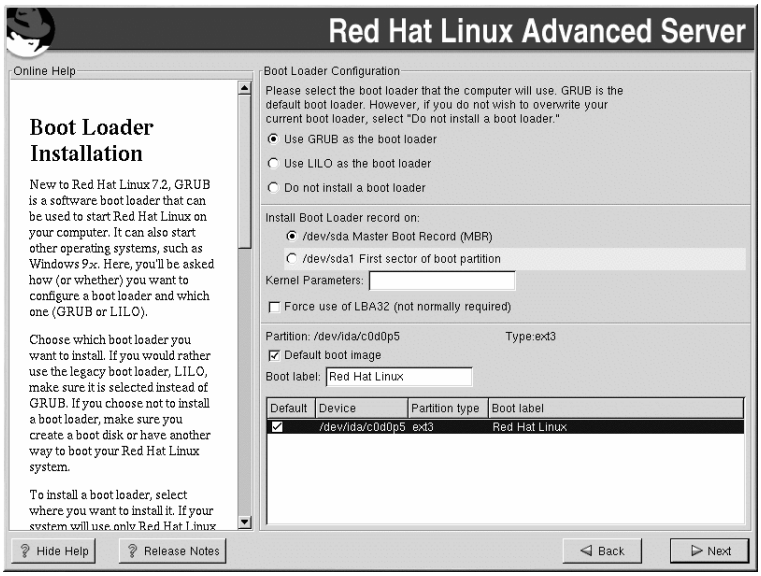

If you chose to install a boot loader (GRUB or LILO), you must determine where it will be installed. You may install your boot loader in one of two places:

The master boot record (MBR)

The recommended place to install a boot loader, unless the MBR already starts another operating system loader, such as System Commander or OS/2's Boot Manager. The MBR is a special area on your hard drive that is automatically loaded by your computer's BIOS, and is the earliest point at which the boot loader can take control of the boot process. If you install it in the MBR, when your machine boots, GRUB (or LILO) will present a boot prompt. You can then boot Red Hat Linux Advanced Server or any other operating system that you have configured the boot loader to boot.

The first sector of your root partition

Recommended if you are already using another boot loader on your system (such as OS/2's Boot Manager). In this case, your other boot loader will take control first. You can then configure that boot loader to start GRUB (or LILO), which will then boot Red Hat Linux Advanced Server.

If you have a RAID card, be aware that some BIOSes do not support booting from the RAID card. In cases such as these, the boot loader *should not* be installed on the MBR of the RAID array. Rather, the boot loader should be installed on the MBR of the same drive as the `/boot` partition was created.

Select where you would like GRUB (or LILO) to be installed on your system. If your system will use only Red Hat Linux Advanced Server, you should choose the MBR.

If you wish to add default options to GRUB or LILO's boot command, enter them into the **Kernel parameters** field. Any options you enter will be passed to the Linux kernel every time it boots.

The **Force use of LBA32 (not normally required)** option allows you to exceed the 1024 cylinder limit for the `/boot` partition. If you have a system which supports the LBA32 extension for booting operating systems above the 1024 cylinder limit, you shouldplace your `/boot` partition above this limit. If the installation program has not already detected this extension from your BIOS, you should select the **Force use of LBA32 (not normally required)** option.

Every bootable partition is listed, including partitions used by other operating systems. The partition holding your Red Hat Linux Advanced Server system's root filesystem will have a **Boot label** of `Red Hat Linux`. Other partitions may also have boot labels. If you would like to add boot labels for other partitions (or change an existing boot label), click once on the partition to select it. Once selected, you can change the boot label by editing the name in the **Boot label** text field.

Note

The **Boot label** column lists what you must enter at the boot prompt in order to boot the desired operating system. However, if you forget the boot labels defined on your system, you can always press [Tab] at the prompt to display a list of defined boot labels.

3.14.1 Rescue Mode

If you need to use rescue mode, there are several options available to you.

- Using the CD-ROM to boot, type **`linux rescue`** at the `boot:` prompt.
- Using the network boot disk you have created, type **`linux rescue`** at the `boot:` prompt. You will then be prompted to retrieve the rescue image from the network.

- Using the boot disk you have created, type **`linux rescue`** at the `boot:` prompt. You then pick an installation method and choose a valid installation tree to load from.

For more information regarding rescue mode, refer to the *Official Red Hat Linux Customization Guide*.

3.14.2 Alternative Boot Loaders

If you do not wish to use GRUB or LILO to boot your Red Hat Linux Advanced Server system, you have several alternatives:

Boot disk

You can use the boot disk created by the installation program (if you elected to create one).

LOADLIN

You can load Linux from MS-DOS. Unfortunately, it requires a copy of the Linux kernel (and an initial RAM disk, if you have a SCSI adapter) to be available on an MS-DOS partition. The only way to accomplish this is to boot your Red Hat Linux Advanced Server system using some other method (for example, from LILO on a diskette) and then copy the kernel to an MS-DOS partition. LOADLIN is available from ftp://metalab.unc.edu/pub/Linux/system/boot/dualboot/ and associated mirror sites.

SYSLINUX

SYSLINUX is an MS-DOS program very similar to LOADLIN. It is also available from ftp://metalab.unc.edu/pub/Linux/system/boot/loaders/ and associated mirror sites.

Some commercial boot loaders

You can load Linux using commercial boot loaders. For example, System Commander and Partition Magic are able to boot Linux (but still require GRUB or LILO to be installed in your Linux root partition).

3.14.3 SMP Motherboards, GRUB, and LILO

This section is specific to SMP motherboards only. If the installation program detects an SMP motherboard on your system, it will create two `/boot/grub/grub.conf` or `/etc/lilo.conf` entries (depending on the boot loader you installed), rather than the usual single entry.

The two entries in `grub.conf` will be `Red Hat Linux (`*`kernel version`*`)` and `Red Hat Linux (`*`kernel versions`*`-smp)`. The `Red Hat Linux (`*`kernel version`*`-smp)` will boot by default. However, if you have trouble with the SMP kernel, you can elect to boot the `Red Hat Linux (`*`kernel version`*`)` entry instead. You will retain all the functionality as before, but you will only be operating with a single processor.

The two entries in `lilo.conf` will be `linux` and `linux-up`. The `linux` entry will boot by default. However, if you have trouble with the SMP kernel, you can elect to boot the `linux-up` entry instead. You will retain all the functionality as before, but you will only be operating with a single processor.

3.15 GRUB Password

If you did not select GRUB as your boot loader, or did not install a boot loader, skip to Section 3.16, *Network Configuration*.

GRUB passwords provide a security mechanism in an environment where physical access to your server is available.

If you are installing GRUB as your boot loader, you should create a password to protect your system. Otherwise, users may be able to pass options to the kernel which can compromise your system security.

Figure 3–10 GRUB Password

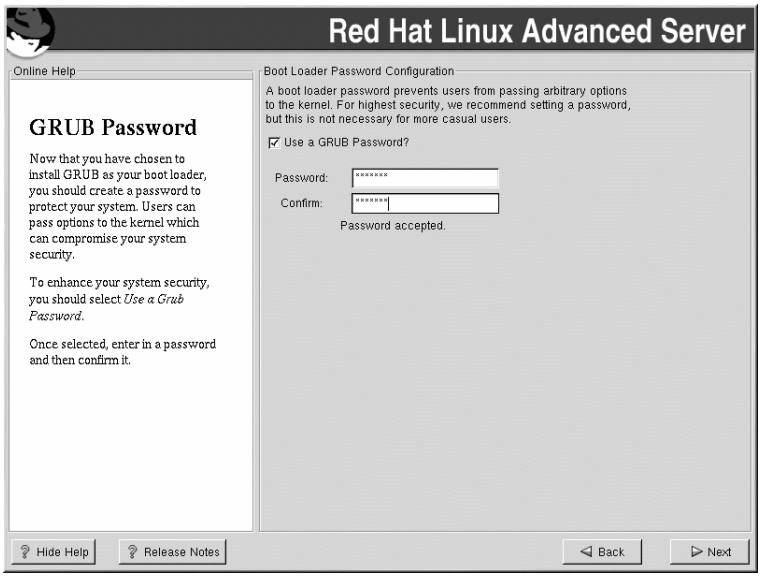

If you choose to use a GRUB password to enhance your system security, be sure to select the checkbox labeled **Use a GRUB Password**.

Once selected, enter a password and confirm it.

3.16 Network Configuration

If you do not have a network device, you will not see this screen. Skip ahead to Section 3.17, *Firewall Configuration*.

If you have a network device and you have not already configured your networking (such as booting from a network boot disk you created and entering in your network information as prompted), you now have the opportunity (as shown in Figure 3–11, *Network Configuration*) to do so.

If you have multiple devices, you will see a tab for each device. You may switch between devices (for example, between eth0 and eth1) and the information you provide on each tab will be specific to each device.

Indicate if you would like to configure your IP address using DHCP (Dynamic Host Configuration Protocol). If you select **Activate on boot**, your network interface will be started when you boot. If you do not have DHCP client access or you are unsure what to provide here, please contact your network administrator.

Next enter, where applicable, the **IP Address**, **Netmask**, **Network**, and **Broadcast** addresses. If you are unsure about any of these, please contact your network administrator.

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

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

Figure 3–11 Network Configuration

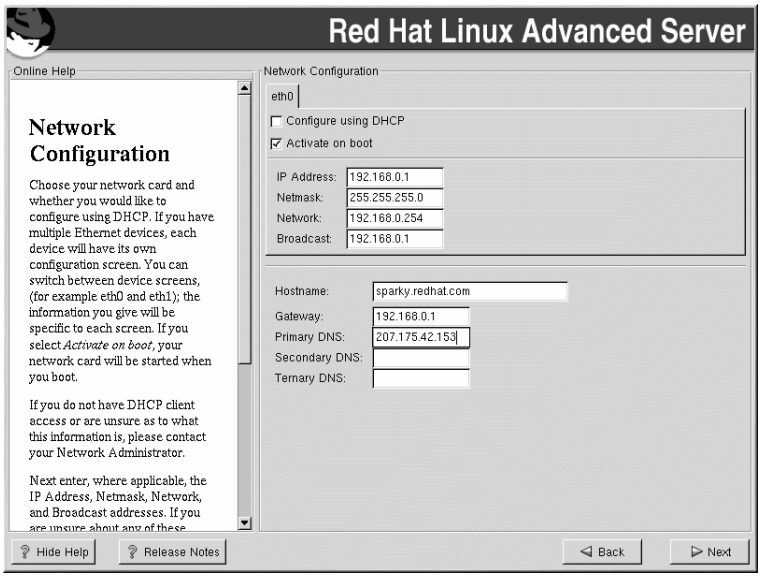

Note

Do not use the numbers as seen in this sample configuration. These values will not work for your own network configuration. If you are not sure what values to enter, contact your network administrator for assistance.

Tip

Even if your computer is not part of a network, you can enter a hostname for your system. If you do not take this opportunity to enter a name, your system will be known as `localhost`.

3.17 Firewall Configuration

Red Hat Linux Advanced Server offers firewall protection for enhanced system security. A firewall exists between your computer and the network, and determines which resources on your computer remote users on the network can access. A properly configured firewall can greatly increase the security of your system.

Figure 3–12 Firewall Configuration

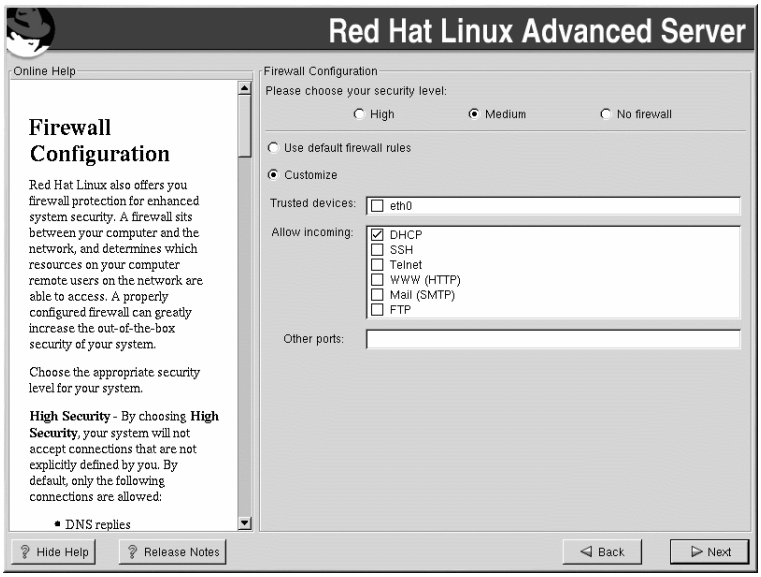

Choose the appropriate security level for your system.

High

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

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

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

- Active mode FTP (passive mode FTP, used by default in most clients, should still work)
- IRC DCC file transfers

- RealAudio™
- Remote X Window System clients

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

Medium

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

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

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

No Firewall

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

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

Trusted Devices

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

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

Allow Incoming

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

DHCP

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

SSH

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

Telnet

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

WWW (HTTP)

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

Mail (SMTP)

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

FTP

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

Other ports

You can allow access to ports which are not listed here, by listing them in the **Other ports** field. Use the following format: **`port:protocol`**. For example, if you want to allow

IMAP access through your firewall, you can specify **imap:tcp**. You can also explicitly specify numeric ports; to allow UDP packets on port 1234 through the firewall, enter **1234:udp**. To specify multiple ports, separate them with commas.

3.18 Language Support Selection

Red Hat Linux Advanced Server can install and support multiple languages for use on your system.

You must select a language to use as the default language. The default language will be used on your Red Hat Linux Advanced Server system once installation is complete. If you choose to install other languages, you can change your default language after the installation.

If you are only going to use one language on your system, selecting only that language will save significant disk space. The default language is the language you selected to use during the installation. However, if you select only one language, you will only be able to use that specified language after the Red Hat Linux Advanced Server installation is complete.

Figure 3–13 Language Support Selection

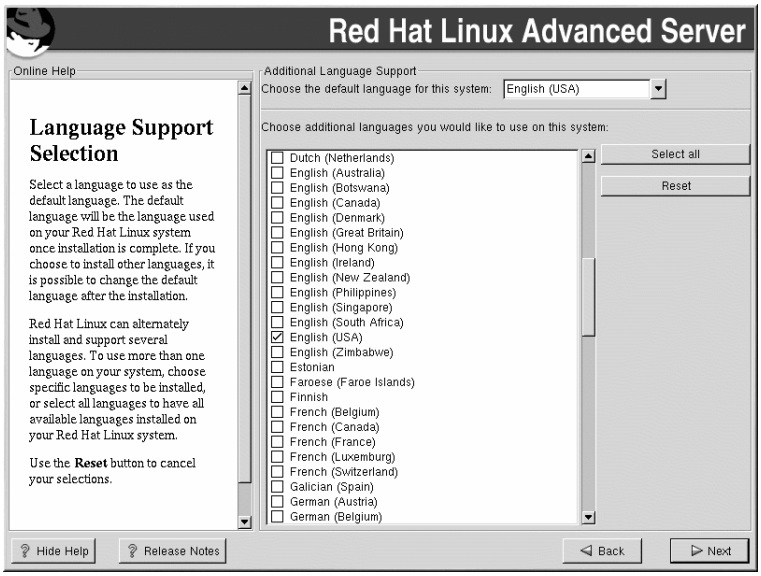

To use more than one language on your system, choose specific languages to be installed or select all languages to have all available languages installed on your Red Hat Linux Advanced Server system.

Use the **Reset** button to cancel your selections. Resetting will revert to the default; only the language you selected for use during the installation will be installed.

3.19 Time Zone Configuration

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

Notice the two tabs at the top of the screen (see Figure 3–14, *Configuring the Time Zone*). The first tab allows you to configure your time zone by your location. You can specify different areas to view: **World**, **North America**, **South America**, **Pacific Rim**, **Europe**, **Africa**, and **Asia**.

Figure 3–14 Configuring the Time Zone

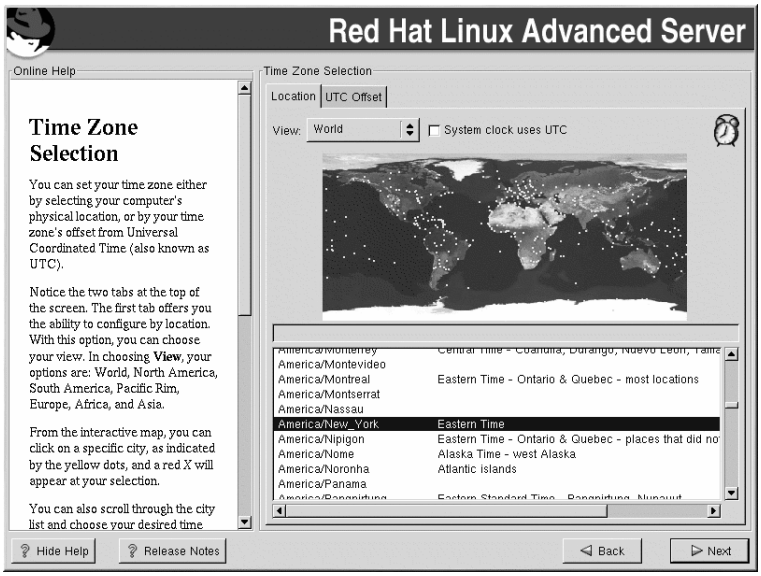

On the interactive map, you can also click on a specific city, which is marked by a yellow dot; a red **X** will appear indicating your selection. You can also scroll through a list and choose a time zone.

The second tab allows you to specify a UTC offset. The tab displays a list of offsets to choose from, as well as an option to set daylight saving time.

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

Tip

If you wish to change your time zone configuration after you have booted your Red Hat Linux Advanced Server system, become root and use the `/usr/sbin/timeconfig` command.

3.20 Account Configuration

The **Account Configuration** screen allows you to set your root password. Additionally, you can set up user accounts for you to log in to once the installation is complete (see Figure 3–15, *Account Creation*).

Figure 3–15 Account Creation

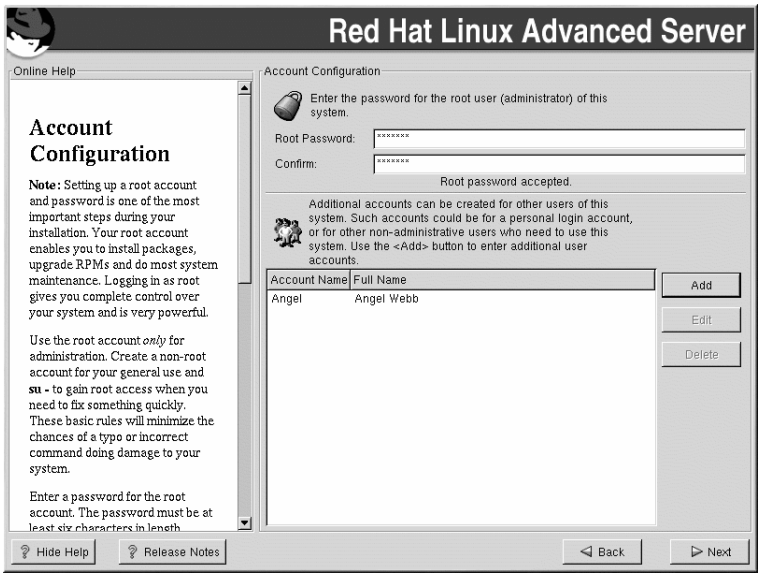

3.20.1 Setting the Root Password

Setting up a root account and password is one of the most important steps during your installation. Your root account is similar to the administrator account used on Windows NT machines. The root account is used to install packages, upgrade RPMs, and perform most system maintenance. Logging in as root gives you complete control over your system.

Use the root account only for system administration. Create a non-root account for your general use and `su -` to root when you need to fix something quickly. These basic rules will minimize the chances of a typo or an incorrect command doing damage to your system.

Tip

To become root, type `su -` at the shell prompt in a terminal window and then press [Enter]. Then, enter the root password and press [Enter].

The installation program will prompt you to set a root password[2] for your system. You must enter a root password. The installation program will not let you proceed to the next section without entering a root password.

The root password must be at least six characters long; the password you type is not echoed to the screen. You must enter the password twice; if the two passwords do not match, the installation program will ask you to enter them again.

You should make the root password something you can remember, but not something that is easy for someone else to guess. Your name, your phone number, **qwerty**, **password, root**, **123456**, and **anteater** are all examples of bad passwords. Good passwords mix numerals with upper and lower case letters and do not contain dictionary words: **Aard387vark** or **420BMttNT**, for example. Remember that the password is case-sensitive. If you write down your password, keep it in a secure place. However, it is recommended that you do not write down this or any password you create.

Note

Do not use one of the example passwords offered in this manual. Using one of these passwords could be considered a security risk.

Note

The root user (also known as the superuser) has complete access to the entire system; for this reason, logging in as the root user is best done *only* to perform system maintenance or administration.

---

[2] A root password is the administrative password for your Red Hat Linux Advanced Server system. You should only log in as root when needed for system maintenance. The root account does not operate within the restrictions placed on normal user accounts, so changes made as root can have implications for your entire system.

3.20.2 Setting Up User Accounts

If you choose to create a user account now, you will have an account to log in to once the installation has completed. This allows you to safely and easily log into your computer without having to be root to create your user account.

Enter an account name. Then enter and confirm a password for that user account. Enter the full name of the account user and press **Add**. Your account information will be added to the account list, and the user account fields will be cleared so that you can add another user.

Figure 3–16 Creating a User Account

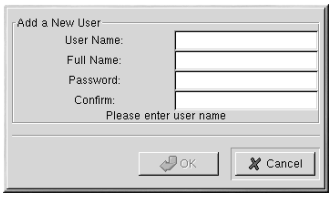

Choose **New** to add a new, non-root, user. Enter the user's information and use the **Add** button to add the user to the account list.

You can also **Edit** or **Delete** the user accounts you have created and no longer want.

3.21 Package Group Selection

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

Note

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

GNOME and KDE are both graphical desktop environments[3] that handle the overall look and feel of your system. You must choose one of these to have a default graphical setup, but you can also install both to determine for yourself which you prefer.

[3] A desktop environment in Linux is similar to the environment you might see in other operating systems. However, environments differ in their look and feel and are easily customized for your individual needs.

You can select components, which group packages together according to function (for example, **GNOME**, **KDE**, or **Software Development**), individual packages, or a combination of the two.

To select a component, click on the checkbox beside it (see Figure 3–17, *Package Group Selection*).

Figure 3–17 Package Group Selection

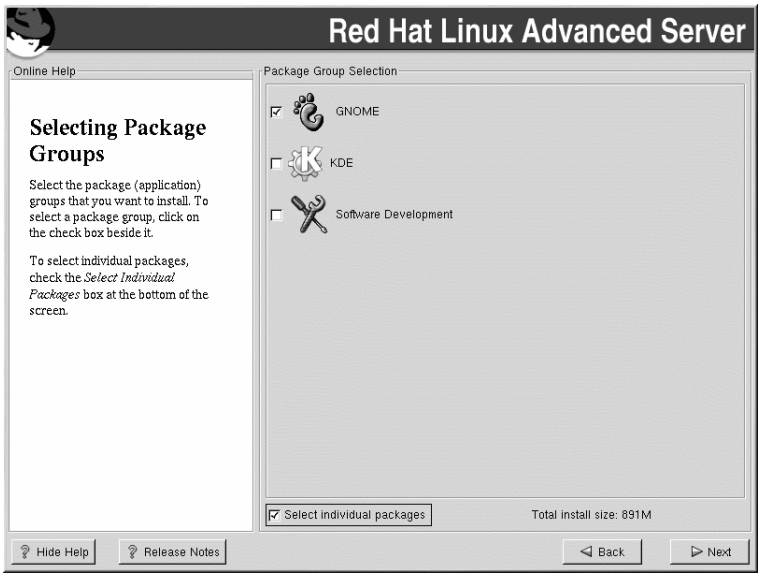

Select each component you wish to install. Selecting **Everything** (at the end of the component list) during a custom installation installs all packages included with Red Hat Linux Advanced Server. If you select every package, you will need approximately 1.7 GB of free disk space.

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

3.21.1 Selecting Individual Packages

After selecting the components you wish to install, you can select or deselect individual packages using your mouse (see Figure 3–18, *Selecting Individual Packages*).

Figure 3–18 Selecting Individual Packages

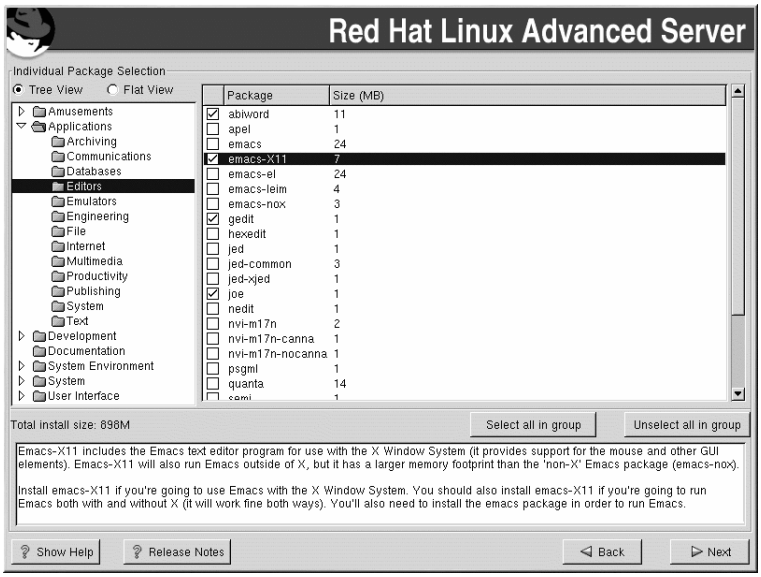

You can choose to view the individual packages in **tree view** or **flat view**. **Tree view** allows you to see the packages grouped by application type. When you expand this list and pick one group, the list of packages in that group appears in the panel on the right. **Flat view** allows you to see all of the packages in an alphabetical listing on the right of the screen.

To sort alphabetically, click on the **Package** tab. To sort packages by size, click on the **Size (MB)** tab.

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

For more information about a specific package, click on the individual package name. The package information will appear at the bottom of the screen.

You can also select or deselect all packages listed within a particular group, by clicking on the **Select all in group** or **Unselect all in group** buttons.

Note

Some packages (such as the kernel and certain libraries) are required for every Red Hat Linux Advanced Server system and are not available to select or deselect. These base packages are selected by default.

3.21.2 A Brief Introduction to GNOME

GNOME is a powerful graphics-driven desktop environment. GNOME includes a panel (for starting applications and displaying status), a desktop (where data and applications can be placed), multiple window managers (which control the look and feel of your desktop), and a standard set of desktop tools and applications.

GNOME allows you to setup your desktop the way you want it to look and "feel." GNOME's session manager remembers settings and currently running programs. So, once you have set things the way you like, they will stay that way.

Figure 3–19 Sample GNOME User Screen

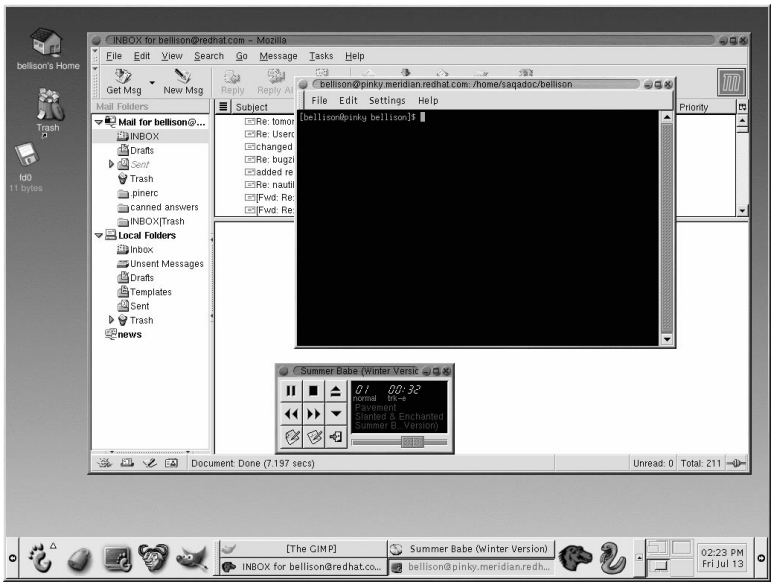

Figure 3–19, *Sample GNOME User Screen*, shows a typical graphical environment using the GNOME desktop environment. A typical graphical environment for KDE would be similar.

Refer to the *Official Red Hat Linux Getting Started Guide* to learn more about GNOME.

3.21.3 A Brief Introduction to KDE

KDE provides a complete desktop environment, including a file manager, a window manager, an integrated help system, a configuration system, numerous tools and utilities, and an ever-increasing number of applications.

KDE offers a contemporary desktop, a searchable help system with convenient access to help on the use of the KDE desktop and its applications, standardized menu and toolbars, keybindings, color schemes, and more.

Refer to the *Official Red Hat Linux Getting Started Guide* to learn more about KDE.

3.21.4 Unresolved Dependencies

Many software packages, in order to work correctly, depend on other software packages that must be installed on your system. For example, many of the graphical Red Hat system administration tools require the `python` and `pythonlib` packages. To make sure your system has all the packages it needs in order to be fully functional, Red Hat Linux Advanced Server checks these package **dependencies** each time you install or remove software packages.

If any package requires another package which you have not selected to install, the program presents a list of these unresolved dependencies and gives you the opportunity to resolve them (see Figure 3–20, *Unresolved Dependencies*).

The **Unresolved Dependencies** screen appears only if you are missing packages that are needed by the packages you have selected. At the bottom of the screen, under the list of missing packages, an **Install packages to satisfy dependencies** checkbox is selected by default. If you leave this checked, the installation program will resolve dependencies automatically by adding all required packages to the list of selected packages.

Figure 3–20 Unresolved Dependencies

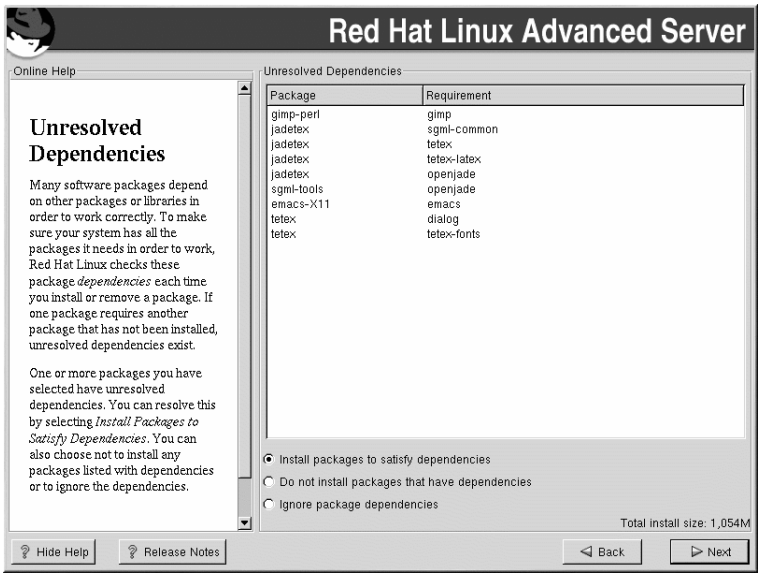

3.22 X Configuration — Video Card

The first part of X configuration deals with video card configuration.

3.22.1 Video Card Configuration

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

If you decided to install the X Window System packages, you now have the opportunity to configure an X server for your system. If you did not choose to install the X Window System packages, skip ahead to Section 3.23, *Preparing to Install*.

If your video card does not appear on the list (see Figure 3–21, *Video Card Setup*), X may not support it. However, if you have technical knowledge about your card, you may choose **Unlisted Card** and attempt to configure it by matching your card's video chipset with one of the available X servers.

Figure 3–21 Video Card Setup

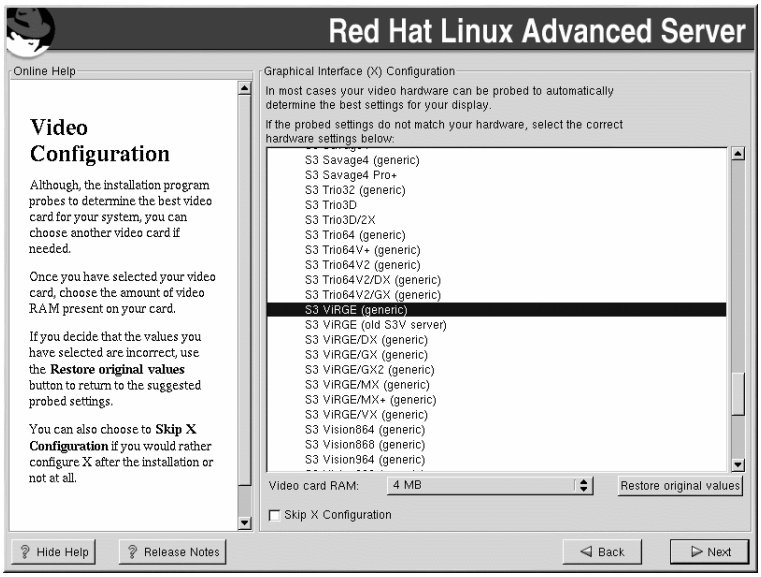

Next, enter the amount of video memory installed on your video card. If you are not sure, please consult the documentation accompanying your video card. You will not damage your video card by choosing more memory than is available, but the X server may not start correctly if you do.

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

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

3.23 Preparing to Install

You should now see a screen preparing you for the installation of Red Hat Linux Advanced Server.

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

WARNING

If, for some reason, you would rather not continue with the installation process, this is your last opportunity to safely cancel the process and reboot your machine. Once you press the Next button, partitions will be written and packages will be installed. If you wish to abort the installation, you should reboot now before any existing information on any hard drive is rewritten.

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

3.24 Installing Packages

At this point there is nothing left for you to do until all the packages have been installed (see Figure 3–22, *Installing Packages*). How quickly this happens depends on the number of packages you have selected and your computer's speed.

Figure 3–22 Installing Packages

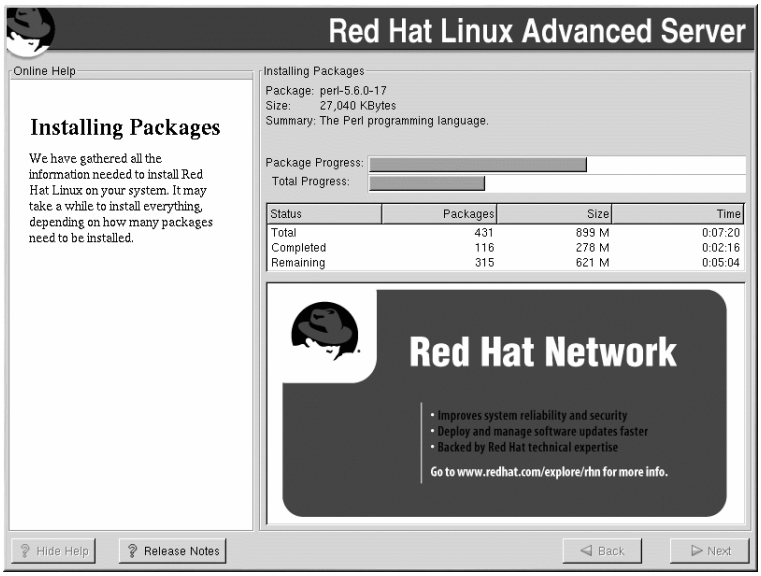

3.25 Boot Disk Creation

If you chose to create a boot disk, you should now insert a blank, formatted diskette into your diskette drive (see Figure 3–23, *Creating Your Boot Disk*).

It is highly recommended that you create a boot disk. If your system were not able to boot properly using GRUB or LILO, a boot disk would enable you to properly boot your Red Hat Linux Advanced Server system.

After a short delay, your boot disk will be created; remove it from your diskette drive and label it clearly. Note that if you would like to create a boot disk after the installation, you will be able to do so. For more information, please see the `mkbootdisk` man page, by typing `man mkbootdisk` at the shell prompt.

If you boot your system with the boot disk (instead of GRUB or LILO) , make sure you create a new boot disk if you make any changes to your kernel.

Figure 3–23 Creating Your Boot Disk

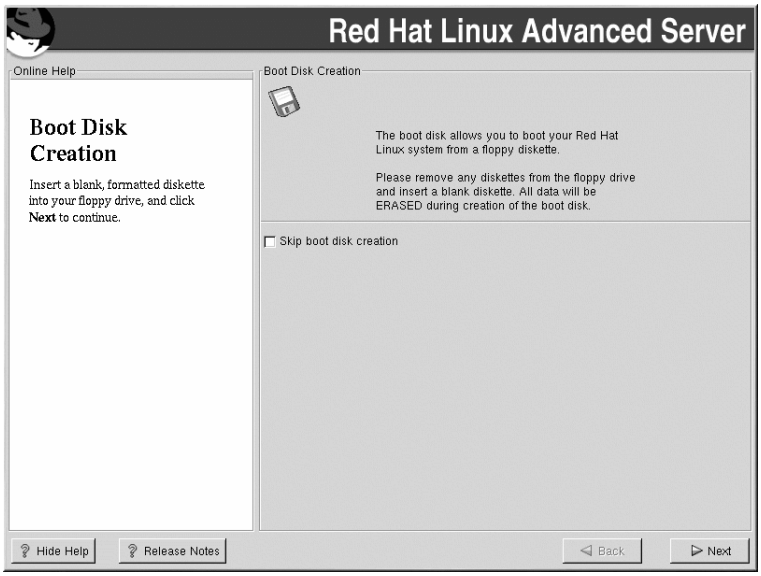

3.26 X Configuration — Monitor and Customization

In order to complete X configuration, you must configure your monitor and customize your X settings.

3.26.1 Configuring Your Monitor

If you chose to skip X configuration, go to Section 3.27, *Installation Complete*.

Xconfigurator, the X Window System configuration tool, presents a list of monitors for you to choose from. In the list, you can either use the monitor that is autodetected for you, or choose another monitor.

Figure 3–24 Monitor Selection

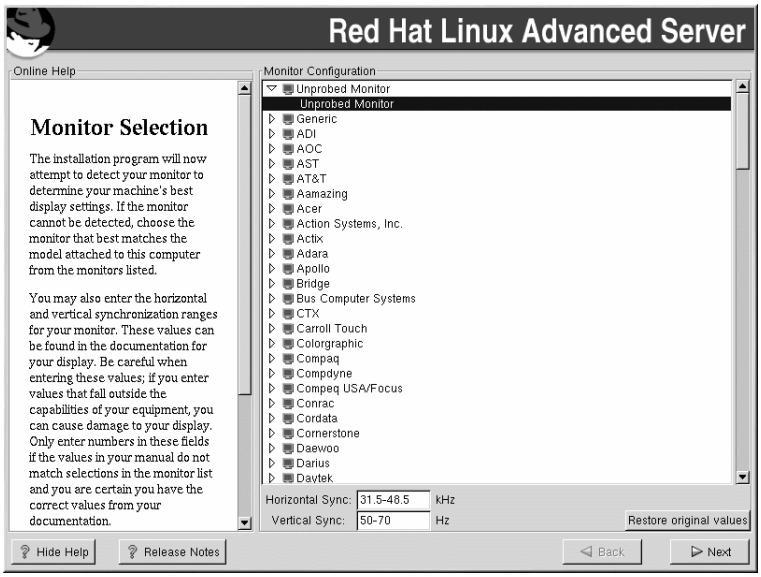

If your monitor does not appear on the list, select the most appropriate **Generic** model available. If you do select a **Generic** monitor, Xconfigurator will suggest horizontal and vertical sync ranges. These values are generally available in the documentation which accompanies your monitor, or from your monitor's vendor or manufacturer; please check your documentation to make sure these values are set correctly.

CAUTION

Do not select a monitor *similar* to your monitor unless you are certain that the monitor you are selecting does not exceed the capabilities of your monitor. Doing so may overclock your monitor and damage or destroy it.

The horizontal and vertical ranges that Xconfigurator suggests for your monitor are also displayed on this screen.

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

Click **Next** when you have finished configuring your monitor.

3.26.2 Custom Configuration

Choose the correct color depth and resolution for your X configuration. Click **Test Setting** to try out this configuration. If you do not like what you see during the test, click **No** to choose another resolution.

Note

If you need to exit out of the X test, use the [Ctrl]-[Alt]-[Backspace] key combination. Also note that this will not work in some test cases.

We recommend that you test your configuration, to make sure the resolution and color settings are usable.

If you installed both GNOME and KDE, you can choose which one to use as your default desktop environment. If you installed one or the other, it will only show GNOME or KDE as the desktop default.

You can also choose whether you want to boot your system into a text or graphical environment once Red Hat Linux Advanced Server is installed. Unless you have special needs, booting into a graphical environment (similar to a Windows environment) is recommended. If you choose to boot into a text environment, you will be presented with a command prompt (similar to a DOS environment).

Figure 3–25 X Customization

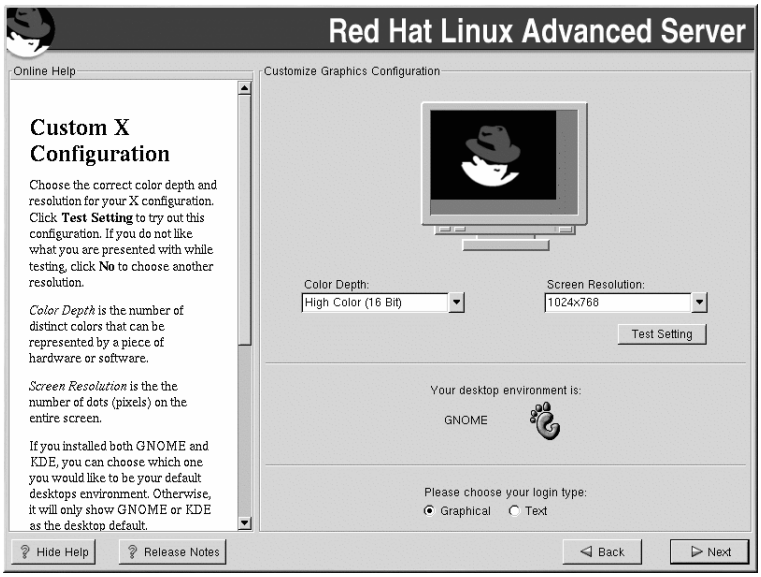

3.27 Installation Complete

Congratulations! Your Red Hat Linux Advanced Server 2.1 installation is now complete!

The installation program will prompt you to prepare your system for reboot. Do not forget to remove any diskette in the diskette drive or CD in the CD-ROM drive. If you did not install a boot loader, you will need to use your boot disk now.

After your computer's normal power-up sequence has completed, you should see the graphical boot loader prompt, at which you can do any of the following things:

- Press [Enter] — causes the default boot entry to be booted.
- Select a boot label, followed by [Enter] — causes the boot loader to boot the operating system corresponding to the boot label. (Press [?] at the text mode boot loader prompt for a list of valid boot labels.)
- Do nothing — after the boot loader's timeout period, (by default, five seconds) the boot loader will automatically boot the default boot entry.

Do whatever is appropriate to boot. If your system was not able to boot properly using GRUB or LILO, or a third-party boot loader, a boot disk would enable you to properly boot your Red Hat Linux

system. You should see one or more screens of messages scroll by. Eventually, you should see a `login:` prompt or a GUI login screen (if you installed the X Window System and chose to start X automatically).

Tip

If you are not sure what to do next, we suggest you begin with the *Official Red Hat Linux Getting Started Guide* (available online at http://www.redhat.com/docs if not included as part of your product), which covers topics relating to the basics of your system and is an introduction to using Red Hat Linux Advanced Server.

If you are a more experienced user looking for information on administration topics, you may find the *Official Red Hat Linux Reference Guide* to be more helpful.

If you are looking for information on system configuration, you may find the *Official Red Hat Linux Customization Guide* to be helpful.

---

4 Installing Red Hat Linux Advanced Server via Text Mode

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

- If you are new to Linux installations, read Chapter 3, *Installing Red Hat Linux Advanced Server*, first. The main focus of that chapter is the graphical installation process, but most of the concepts apply to the text mode installation as well. After reading that chapter, Section 4.1, *Things You Should Know*, will give you more information regarding the aspects of installing Red Hat Linux Advanced Server that do not apply to the graphical installation process.

  Additionally, Appendix F, *An Introduction to Disk Partitions* may be helpful to you, since it discusses disk partition resizing. If you plan to install Red Hat Linux Advanced Server on a disk where another operating system is currently installed, this knowledge will be crucial.

- If you plan to install over a network (via NFS, FTP, or HTTP), you must make a network boot disk. Chapter 1, *Steps to Get You Started*, explains how to do this.

- If you have never used the text mode installation program, or need a refresher on its user interface, read this chapter.

4.1 Things You Should Know

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

Chapter 2, *System Requirements Table* provides a table for you to fill out with your specific system requirements, which helps you keep up with any information needed during your installation. Please review the hardware table at Table 2–1, *System Requirements Table*.

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

Note

You can perform a text mode installation of Red Hat Linux Advanced Server 2.1 by following the instructions in this chapter. However, if you are installing from a CD-ROM, you might prefer to use the graphical installation mode, which is easy to use and provides a flexible, custom installation mode. For more information on graphical installations, turn to Chapter 3, *Installing Red Hat Linux Advanced Server*.

4.1.1 Basic Hardware Configuration

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

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

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

4.1.2 Video Configuration

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

- Your video card — the card's make and model number (or the video chipset it uses) and the amount of video RAM it has. (Most PCI-based cards are auto-detected by the installation program.)

- Your monitor — the unit's make and model number, along with allowable ranges for horizontal and vertical refresh rates. (Newer models may be auto-detected by the installation program.)

4.1.3 Network-related Information

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

- IP address — usually represented as a set of four numbers separated by dots, such as `10.0.2.15`.
- Netmask — another set of four numbers separated by dots; an example netmask would be `255.255.248.0`.
- Gateway IP address — another set of four dot-separated numbers; for example, `10.0.2.254`.
- One or more name server IP addresses — one or more sets of dot-separated numbers; for example, `10.0.2.1` might be the address of a name server.
- Domain name — the name your organization uses; for example, Red Hat has a domain name of `redhat.com`.
- Hostname — the name assigned to your individual system; for example, a computer might be named `pooh`.

Note

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

4.2 The Installation Program User Interface

The Red Hat Linux Advanced Server text mode installation program uses a screen-based interface that includes most of the on-screen "widgets" commonly found on graphical user interfaces. Figure 4–1, *Installation Program Widgets as seen in* **Configure TCP/IP**, and Figure 4–2, *Installation Program Widgets as seen in Disk Druid* , illustrate the screens you will see.

Figure 4–1 Installation Program Widgets as seen in Configure TCP/IP

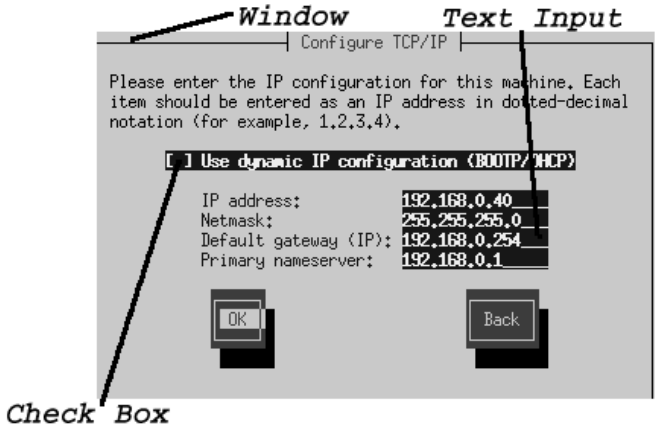

Figure 4–2 Installation Program Widgets as seen in Disk Druid

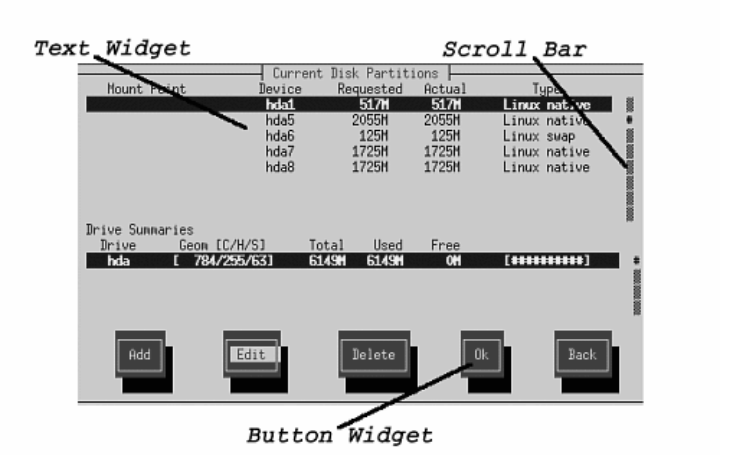

Here is a list of the most important widgets shown in Figure 4–1, *Installation Program Widgets as seen in* **Configure TCP/IP**, and Figure 4–2, *Installation Program Widgets as seen in Disk Druid* :

- Window — windows (usually referred to as **dialogs** in this manual) will appear on your screen throughout the installation process. At times, one window may overlay another; in these cases,

you can only interact with the window on top. When you are finished in that window, it will disappear, allowing you to continue working in the window underneath.

- Text Input — text input lines are regions where you can enter information required by the installation program. When the cursor rests on a text input line, you may enter and/or edit information on that line.
- Checkbox — checkboxes allow you to select or deselect a feature. The box displays either an asterisk (selected) or a space (unselected). When the cursor is within a checkbox, press [Space] to select an unselected feature or to deselect a selected feature.
- Text widget — text widgets are regions of the screen for the display of text. At times, text widgets may also contain other widgets, such as checkboxes. If a text widget contains more information than can be displayed in the space reserved for it, a scroll bar appears; if you position the cursor within the text widget, you can then use the [Up] and [Down] arrow keys to scroll through all the information available. Your current position is shown on the scroll bar by a **#** character, which moves up and down the scroll bar as you scroll.
- Button widget — button widgets are the primary method of interacting with the installation program. You progress through the windows of the installation program by navigating these buttons, using the [Tab] and [Enter] keys. Buttons can be selected when they are highlighted.
- Cursor — although not a widget, the cursor is used to select (and interact) with a particular widget. As the cursor is moved from widget to widget, it may cause the widget to change color, or you may only see the cursor itself positioned in or next to the widget. In Figure 4–1, *Installation Program Widgets as seen in* **Configure TCP/IP**, the cursor is positioned on the **OK** button. Figure 4–2, *Installation Program Widgets as seen in Disk Druid* , shows the cursor on the **Edit** button.

4.2.1 Using the Keyboard to Navigate

Navigation through the installation dialogs is performed through a simple set of keystrokes. To move the cursor, use [Left], [Right], [Up], and [Down] arrow keys. Use [Tab], and [Alt]-[Tab] to cycle forward or backward through each widget on the screen. Along the bottom, most screens display a summary of available cursor positioning keys.

To "press" a button, position the cursor over the button (using [Tab], for example) and press [Space] or [Enter]. To select an item from a list of items, move the cursor to the item you wish to select and press [Enter]. To select an item with a checkbox, move the cursor to the checkbox and press [Space] to select an item. To deselect, press [Space] a second time.

Pressing [F12] accepts the current values and proceeds to the next dialog; it is equivalent to pressing the **OK** button.

Unless a dialog box is waiting for your input, do not press any keys during the installation process (doing so may result in unpredictable behavior).

4.3 Starting the Installation Program

The following methods can be used to start the installation:

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

While the installation program loads, messages will scroll on your screen. When the installation program has loaded, this prompt appears:

```
boot:
```

4.3.1 Text Mode Boot Options

If you press [Enter] at the `boot:` prompt, or if you take no action within the first minute after the `boot:` prompt appears, the graphical installation program will start. Pressing one of the help screen function keys as described in Section 4.3.2, *Displaying Online Help* disables this autostart feature.

To start the text mode installation program, before pressing [Enter], type:

```
boot: text
```

If the installation program does not properly detect your hardware, you may need to restart the installation in "expert" mode. To start an expert mode installation, type:

```
boot: text expert
```

Expert mode disables most hardware probing, and gives you the option of entering options for the drivers loaded during the installation.

Note

The initial boot messages will not contain any references to SCSI or network cards. This is normal; these devices are supported by modules that are loaded during the installation process.

Note that the command to start a serial installation has changed. If you must perform the installation in serial mode, use the following command:

```
boot: linux text console=<device>
```

In this command, <*device*> should be the device you are using (such as ttyS0 or ttyS1).

You can also pass options to the kernel as you are booting the installation program. For example, to instruct the kernel to use all the RAM in a 128 MB system, enter:

```
boot: linux text mem=128M
```

4.3.2 Displaying Online Help

Once the installation program is loaded into memory, you can obtain information about the installation process and options by pressing [F1] through [F6]. For example, press [F2] to see general information about the online help screens.

4.4 Installation Cross-Reference Table

Note

This installation cross-reference table only documents the screens seen when performing an Advanced Server installation.

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

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

Table 4–1 Installation Cross-reference Table

Text Mode Screen	Text Mode Reference Point
Language Screen	Section 3.5, *Language Selection*
Keyboard Screen	Section 3.6, *Keyboard Configuration* — Layout
Installation Method	Section 3.3, *Selecting an Installation Method*
Installing from a Hard Disk	Section 4.5, *Installing from a Hard Drive*
Installing over a Network	Section 4.6, *Installing over a Network*
Mouse Selection	Section 3.7, *Mouse Configuration*
Welcome Screen	Section 3.8, *Welcome to Red Hat Linux*
Installation Type	Section 3.9, *Install Options*
Disk Partitioning Setup	Section 3.10, *Disk Partitioning Setup*
Automatic Partitioning	Section 3.11, *Automatic Partitioning*
Disk Druid	Section 3.12, *Partitioning Your System*
fdisk	Section 3.13, *Partitioning with fdisk*
Boot Loader Installation	Section 3.14, *Boot Loader Installation*
GRUB Password	Section 3.15, *GRUB Password*
Boot Loader Configuration — Where to install boot loader, kernel options, boot labels	Section 3.14, *Boot Loader Installation*
Hostname Configuration	Section 3.16, *Network Configuration*
Firewall Configuration	Section 3.17, *Firewall Configuration*
Network Configuration	Section 3.16, *Network Configuration*
Language Support and Default Language	Section 3.18, *Language Support Selection*
Time Zone Selection	Section 3.19, *Time Zone Configuration*
Root Password	Section 3.20, *Account Configuration*
Add User/User Account Setup	Section 3.20, *Account Configuration*

Text Mode Screen	Text Mode Reference Point
Package Installation (Group, Individual, Dependencies)	Section 3.21, *Package Group Selection*
Video Card Configuration	Section 3.26, *X Configuration — Monitor and Customization*
Package Installation	Section 3.23, *Preparing to Install* and Section 3.24, *Installing Packages*
Boot Disk Creation	Section 3.25, *Boot Disk Creation*
Monitor Configuration	Section 3.26.1, *Configuring Your Monitor*
Custom X Configuration	Section 3.26.2, *Custom Configuration*
Installation Complete	Section 3.27, *Installation Complete*

4.5 Installing from a Hard Drive

Note

Hard drive installations only work from ext2, ext3, or FAT filesystems. If you have a filesystem other than those listed here, such as reiserfs, you will not be able to perform a hard drive installation.

Hard drive installations require the use of the ISO (or CD-ROM) images rather than copying an entire installation tree. After placing the required ISO images (the binary Red Hat Linux Advanced Server CD-ROMs) in a directory, choose to install from the hard drive. You will then point the installation program at that directory to perform the installation.

Verifying that the ISO images are intact before you attempt an installation will help to avoid problems that are often encountered during a hard drive installation. To verify the ISO images are intact prior to performing an installation, use an md5sum program (many md5sum programs are available for various operating systems). An md5sum program should be available on the same server as the ISO images.

The **Select Partition** screen (Figure 4–3, *Selecting Partition Dialog for Hard Drive Installation*) applies only if you are installing from a disk partition (that is, if you selected **Hard Drive** in the **Installation Method** dialog). This dialog allows you to name the disk partition and directory from which you are installing Red Hat Linux Advanced Server.

Enter the device name of the partition containing the Red Hat ISO images. There is also a field labeled **Directory holding images**. If the ISO images are not in the root directory of that partition, enter the path to the ISO images (for example, if the ISO images are in `/test/new/RedHat`, you would enter `/test/new`).

After you have identified the disk partition, you will next see the **Welcome** dialog. See Table 4–1, *Installation Cross-reference Table*, for more information.

Figure 4–3 Selecting Partition Dialog for Hard Drive Installation

4.6 Installing over a Network

If you are performing a network installation, the **Configure TCP/IP** dialog appears; for an explanation of this dialog, go to Section 3.16, *Network Configuration*, and then return here.

4.6.1 Setting Up the Server

Because the Red Hat Linux Advanced Server 2.1 installation program is capable of installing Red Hat Linux Advanced Server from multiple CD-ROMs, if you intend to support NFS, FTP, or HTTP installations you must copy the `RedHat` directory from each CD-ROM comprising Red Hat Linux Advanced Server 2.1 onto a disk drive:

- Insert CD-ROM 1 and execute the following commands:
  - `mount /mnt/cdrom`
  - `cp -var /mnt/cdrom/RedHat /location/of/disk/space`

Where */location/of/disk/space* is a directory you create such as */export/2.1/*.

– `umount /mnt/cdrom`

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

    Where */location/of/disk/space* is a directory you create such as */export/2.1/*.
  – `umount /mnt/cdrom`
- Next, make */location/of/disk/space* accessible to the installation program (for example, exporting it for NFS installations).

```
Export /location/of/disk/space
```

If you are not sure how to do this, refer to the *Official Red Hat Linux Customization Guide* and the *Official Red Hat Linux Reference Guide* for more information.

4.6.2 NFS Setup

The NFS dialog (Figure 4–4, *NFS Setup Dialog*) applies only if you are installing from an NFS server (if you booted from a network or PCMCIA boot disks and selected **NFS Image** in the **Installation Method** dialog).

Figure 4–4 NFS Setup Dialog

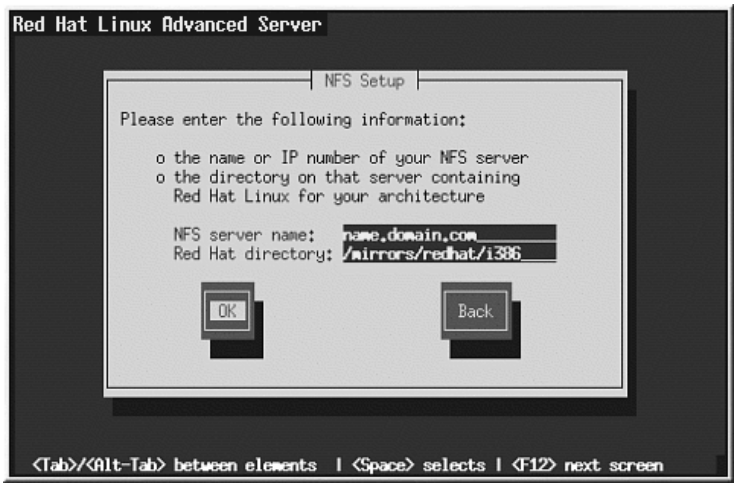

Enter the fully-qualified domain name or IP address of your NFS server. For example, if you are installing from a host named `eastcoast` in the domain `redhat.com`, enter `eastcoast.redhat.com` in the **NFS Server** field.

Next, enter the name of the exported directory. If you followed the setup described in Section 4.6, *Installing over a Network*, you would enter the directory */location/of/disk/space/* which contains the `RedHat` directory.

If the NFS server is exporting a mirror of the Red Hat Linux Advanced Server installation tree, enter the directory which contains the `RedHat` directory. (If you do not know this directory path, ask your system administrator.) For example, if your NFS server contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386`.

Next you will see the **Welcome** dialog. See Table 4–1, *Installation Cross-reference Table*, for more information.

4.6.3 FTP Setup

The FTP dialog (Figure 4–5, *FTP Setup Dialog*) applies only if you are installing from an FTP server (if you selected **FTP** in the **Installation Method** dialog). This dialog allows you to identify the FTP server from which you are installing Red Hat Linux Advanced Server.

Figure 4–5 FTP Setup Dialog

Enter the name or IP address of the FTP site you are installing from, and the name of the directory containing the `RedHat` installation files for your architecture. For example, if the FTP site contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386`.

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

Next you will see the **Welcome** dialog. See Table 4–1, *Installation Cross-reference Table*, for more information.

4.6.4 HTTP Setup

The HTTP dialog (Figure 4–6, *HTTP Setup Dialog*) applies only if you are installing from an HTTP server (if you selected **HTTP** in the **Installation Method** dialog). This dialog prompts you for information about the HTTP server from which you are installing Red Hat Linux Advanced Server.

Figure 4–6 HTTP Setup Dialog

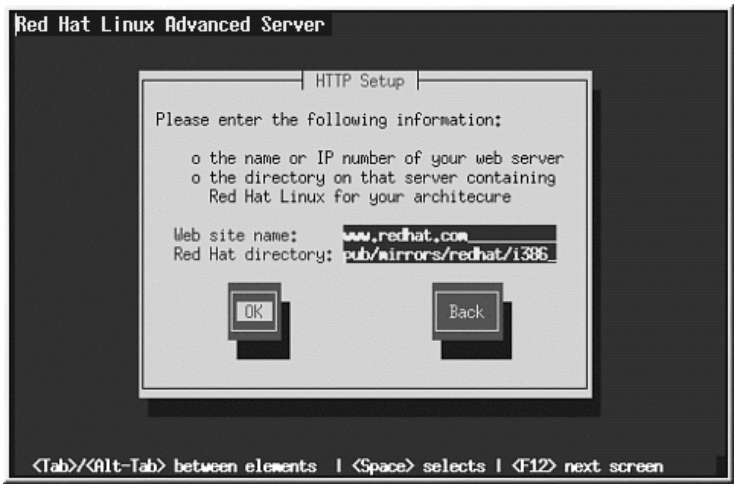

Enter the name or IP address of the HTTP site you are installing from, and the name of the directory there containing the `RedHat` installation files for your architecture. For example, if the HTTP site contains the directory `/mirrors/redhat/i386/RedHat`, enter `/mirrors/redhat/i386`.

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

Next you will see the **Welcome** dialog. See Table 4–1, *Installation Cross-reference Table*, for more information.

Part II Configuring Red Hat Linux Advanced Server

5 Introduction

Using Red Hat Linux Advanced Server 2.1, it is possible to create highly available server clustering solutions able to withstand many common hardware and software failures with little or no interruption of critical services. In addition, Red Hat Linux Advanced Server, by allowing multiple computers to work together in offering these critical services, can help system administrators plan and execute system maintenance and upgrades without any service interruption.

This manual guides you through the following steps in understanding and deploying a clustering solution based on Red Hat Linux Advanced Server's **Linux Virtual Server** (LVS) technology:

- Explains the Linux Virtual Server technology used by Red Hat Linux Advanced Server to create a load-balancing cluster.
- Explains how to configure a Red Hat Linux Advanced Server LVS cluster.
- Walks you through the Piranha Configuration Tool, a graphical interface used for configuring and monitoring an LVS cluster.

An accompanying manual, the *Red Hat Cluster Manager Installation and Administration Guide*, explains how to deploy a clustering solution based on **Red Hat Cluster Manager**:

- Learn about the Red Hat Cluster Manager's high-availability services technology.
- Explain how to set up the appropriate hardware and software to use Red Hat Cluster Manager.

5.1 Technology Overview

Red Hat Linux Advanced Server implements highly available server solutions via clustering. It is important to note that **cluster** computing consists of three distinct branches:

- **Compute clustering** (such as Beowulf) uses multiple machines to provide greater computing power for computationally intensive tasks. This type of clustering is not addressed by Red Hat Linux Advanced Server.
- **High-availability (HA) clustering** uses multiple machines to add an extra level of reliability for a service or group of services.
- **Load-balance clustering** uses a router to dispatch traffic to a pool of servers.

Red Hat Linux Advanced Server addresses the latter two types of clustering technology. Using a collection of programs to monitor the health of the systems and services in the cluster.

Note

Red Hat Linux Advanced Server's clustering technology is not synonymous with fault tolerance. Fault tolerant systems use highly specialized and often very expensive hardware to implement a fully redundant environment in which services can run uninterrupted by hardware failures.

However, fault tolerant systems do not account for operator and software errors which Red Hat Linux Advanced Server can address through service redundancy. Also, since Red Hat Linux Advanced Server is designed to run on readily available hardware, it creates an environment with a high level of system availability at a fraction of the cost of fault tolerant hardware.

5.2 Basic Configurations

While Red Hat Linux Advanced Server can be configured in a variety of different ways, the configurations can be broken into two major categories:

- High-availability clusters using Red Hat Cluster Manager.
- Load-balancing clusters using Linux Virtual Servers.

5.2.1 High-Availability Clusters Using Red Hat Cluster Manager

High-availability clusters based on Red Hat Cluster Manager utilize two Linux servers or **nodes** and a shared storage device to enhance the availability of key services on the network. Each of these key services in the cluster is assigned its own virtual server IP address (VIP). The VIP address, or floating IP, is an IP address that is distinct from the either node's normal IP address and is associated with the service rather than any particular machine in the cluster. If a monitored service on one of the nodes fails, then that node is removed from the cluster and the remaining server starts the appropriate services — maintaining their floating IP addresses and causing minimal disruption to the end user. This procedure is called **failover**.

Each node in a Red Hat Cluster Manager high-availability cluster must have access to a shared storage device for two reasons:

- Either node in the cluster must be able to reach the service data of a failed node.
- Either node in the cluster must check the health of the other node via mutually accessible **quorum partitions** [1]located on the shared storage device.

[1] Quorum partitions are small raw devices used by each node in Red Hat Cluster Manager to check the health of the other node. See *Red Hat Cluster Manager Installation and Administration Guide* for more details.

Having access to the same data source helps Red Hat Cluster Manager more effectively handle failover situations because after a failure occurs the functional node's newly activated services have access to the exact same data used by the failed node. However, to protect the integrity of data on shared devices, services within a high-availability cluster are only allowed to run on one node at any given time.

Red Hat Cluster Manager's use of shared storage also gives administrators flexibility in how they use each node in the cluster. For example, one can either run different services on each server — a configuration known as **active-active** — or run all services on one node while the other sits idle — a configuration known as **hot-standby**.

The shared storage device in a Red Hat Cluster Manager cluster also enables each node to verify the health of the other by regularly updated status information on mutually accessible quorum disk partitions. If the quorum partition is not updated properly by a member of the cluster, the other node can verify the integrity of that member by pinging it through a **heartbeat channel**. Heartbeat channels can be configured on one or more Ethernet interfaces or a serial connection or on both interfaces concurrently.

For more information about configuring Red Hat Cluster Manager clusters, please see the accompanying manual titled *Red Hat Cluster Manager Installation and Administration Guide*.

5.2.2 Load-Balancing Clusters Using Linux Virtual Servers

To the outside world, an LVS cluster appears as one server, but in reality, a user from the World Wide Web who is accessing a group of servers behind a pair of redundant LVS routers.

An LVS cluster consists of at least two layers. The first layer is composed of a pair of similarly configured Linux machines or nodes. One of these nodes acts as an **LVS router**, directing requests from the Internet to the second layer — a pool of servers called **real servers**. The real servers provide the critical services to the end-user while the LVS router balances the load to these servers.

For a detailed overview of LVS clustering, see Chapter 6, *Linux Virtual Server Overview*.

6 Linux Virtual Server Overview

As mentioned in the introduction, Red Hat Linux Advanced Server LVS clustering uses a Linux machine called the active router to send requests from the Internet to a pool of servers. To accomplish this, LVS clusters consist of two basic machine classifications — the LVS routers (one active and one backup) and a pool of real servers which provide the critical services.

The active router serves two roles in the cluster:

- To balance the load on the real servers.
- To check the integrity of the services on each of the real servers.

The backup router's job is to monitor the active router and assume its role in the event of failure.

To explain how LVS clusters perform these functions, let us look at a basic configuration.

6.1 A Basic LVS Configuration

Figure 6–1, *A Basic LVS Configuration* shows a simple LVS cluster consisting of two layers. On the first layer are two LVS routers — one active and one backup. Each of the LVS routers have two network interfaces per machine, one interface on the Internet and one on the private network, enabling them to regulate traffic between the two networks. For this example the active router is using **Network Address Translation** or **NAT** to direct traffic from the Internet to a variable number of real servers on the second layer, which in turn provide the necessary services. Therefore, the real servers in this example are connected to a dedicated private network segment and pass all public traffic back and forth through the active LVS router. To the outside world, the server farm appears as one entity.

Figure 6–1 A Basic LVS Configuration

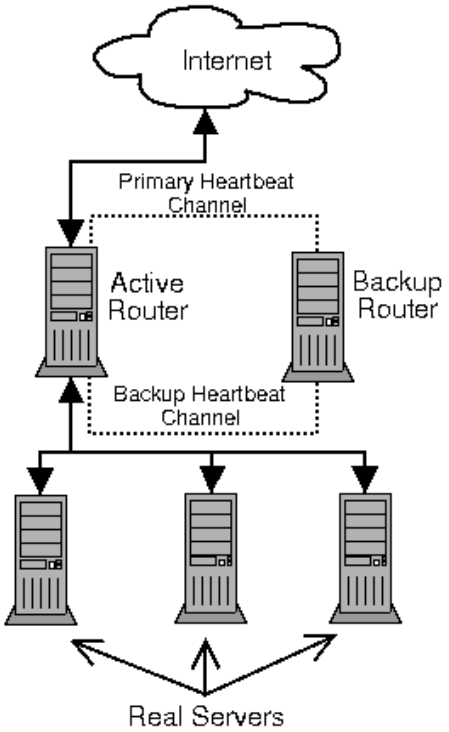

Service requests arriving at the LVS cluster are addressed to a *virtual IP address* or VIP. This is a publicly-advertised address the administrator of the site associates with a fully-qualified domain name, such as www.example.com, and which is assigned to one or more **virtual server** [1]. Figure 6–1, *A Basic LVS Configuration* shows only one virtual IP address, but there may be more than one. The important thing to remember is that a VIP address will migrate from one LVS router to the other during a failover, thus maintaining a presence at that IP address. As such, they can be considered *floating IP addresses*.

VIP addresses may be aliased to the same device which connects the LVS router to the Internet. For instance, if eth0 is connected to the Internet, than multiple virtual servers can be aliased to `eth0:1`.

[1] A virtual server is a service configured to listen on a specific virtual IP. See Section 9.6, **VIRTUAL SERVERS** for more on configuring a virtual server using the Piranha Configuration Tool.

Alternatively, each virtual server can be associated with a separate device per service. For example, HTTP traffic can be handled on `eth0:1`, and FTP traffic can be handled on `eth0:2`.

Only one LVS router is active at a time. The role of the active router is to redirect service requests from virtual IP addresses to the real servers. The redirection is based on one of eight supported load-balancing algorithms described further in Section 6.3, *LVS Scheduling Overview*.

4 The active router also dynamically monitors the overall health of the specific services on the real servers through simple **send/expect scripts**. To aid in detecting the health of services that require dynamic data, such as HTTPS or SSL, the administrator can also call external executables. If a service on a real server malfunctions, the active router stops sending jobs to that server until it returns to normal operation.

The backup router performs the role of a hot-standby system. Periodically, the LVS routers exchange "I'm alive" heartbeat messages through the primary external public interface and, in a failover situation, the private interface. Should the backup node fail to receive a heartbeat message within an expected interval, it initiates a failover and assumes the role of the active router. During failover, the backup router takes over the VIP addresses serviced by the failed router using a technique known as **ARP spoofing** — where the backup LVS router announces itself as the destination for IP packets addressed to the failed node. When the failed node returns to active service, the backup node assumes its hot-backup role again.

The simple, two-layered configuration used in Figure 6–1, *A Basic LVS Configuration* is best for clusters serving data which does not change very frequently — such as static Web pages — because the individual real servers do not automatically sync data between each node. The next section discusses approaches to replicating data across the server pool.

6.1.1 Data Replication and Data Sharing Between Real Servers

Since there is no built-in component in LVS clustering to share the same data between the real servers, the administrator has two basic options:

- Synchronize the data across the real server pool.
- Add a third layer to the topology for shared data access.

The first option is better for servers which do not allow large numbers of users to upload or change data on the real servers. If the cluster allows large numbers of users to modify data, such as an e-commerce website, the latter option is preferable.

Configuring Real Servers to Synchronize Data

There are many ways an administrator can choose to synchronize data across the pool of real servers. For instance, shell scripts can be employed so that if a Web engineer updates a page, the page is posted

to all of the servers simultaneously. Also, the cluster administrator can use programs such as `rsync` to replicate changed data across all nodes at a set interval.

However, this type of data synchronization does not work well if the cluster is extremely busy with users uploading files or issuing database transactions. For a cluster with a high load, a **three-tiered topology** is the best solution.

6.2 A Three Tiered LVS Configuration

Figure 6–2, *A Three Tiered LVS Configuration* shows a typical three tiered LVS cluster topology. In this example, the active LVS router routes the requests from the Internet to the pool of real servers. Each of the real servers then accesses a shared data source over the network.

Figure 6–2 A Three Tiered LVS Configuration

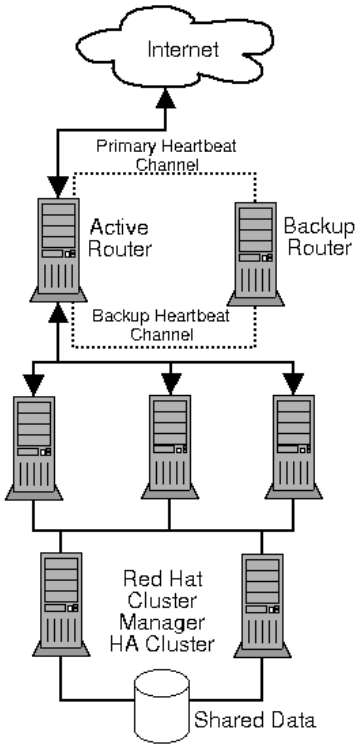

This configuration is ideal for busy FTP servers, where accessible data is stored on a central, highly available server and accessed by each real server via an NFS exported directory or Samba share. This topography is also great for websites that access a central, highly available database for transactions. And using an active-active configuration with Red Hat Cluster Manager, the administrator can configure one high-availability cluster to serve both of these roles simultaneously.

The third tier in the above example does not have to use Red Hat Cluster Manager, but failing to use a highly available solution would introduce a critical single point of failure. For more information on setting up Red Hat Cluster Manager high availability cluster refer to the *Red Hat Cluster Manager Installation and Administration Guide*.

6.3 LVS Scheduling Overview

One of the advantages of using an LVS cluster is its ability to perform flexible, IP-level load balancing on the real server pool. This flexibility is due to the variety of scheduling algorithms an administrator can choose from when configuring a cluster. LVS load balancing is superior to less flexible methods, such as **Round-Robin DNS** where the hierarchical nature of DNS and the caching by client machines can lead to load imbalances. Also, the low level filtering employed by the LVS router has advantages over application-level request forwarding because balancing loads at the network packet level causes minimal computational overhead and allows for greater scalability.

Using scheduling, the active router can take into account the real servers' activity and, optionally, an administrator-assigned **weight** factor when routing service requests. Thus it is possible to create a group of real servers using a variety of hardware and software combinations and the active router can load each real server evenly.

The scheduling mechanism for an LVS cluster is provided by a collection of kernel patches called **IP Virtual Server** or **IPVS** modules. These modules enable L4 switching, which is designed to work well with multiple servers on a single IP address.

To track and route packets to the real servers efficiently, IPVS builds an **IPVS table** in the kernel. This table is used by the active LVS router to redirect requests from a virtual server address to and returning from real servers in the pool. The IPVS table is constantly updated by a daemon called **ipvsadm** — adding and removing cluster members depending on their availability.

6.3.1 Scheduling Algorithms

The form the IPVS table takes depends on the scheduling algorithm the administrator chooses for any given virtual server. To allow for maximum flexibility in both the types of services you can cluster and how these services are scheduled, Red Hat Linux Advanced Server 2.1 provides the eight scheduling algorithms listed below. For instructions on how to assign scheduling algorithms see Section 9.6.1, *The* **VIRTUAL SERVER** *Subsection*.

Round-Robin Scheduling

Distributes each request sequentially around the pool of real servers. Using this algorithm, all the real servers are treated as equals without regard to capacity or load. This scheduling model resembles round-robin DNS but is more granular due to the fact that it is network-connection based and not host-based. LVS round-robin scheduling also does not suffer the imbalances caused by cached DNS queries.

Weighted Round-Robin Scheduling

Distributes each request sequentially around the pool of real servers but gives more jobs to servers with greater capacity. Capacity is indicated by a user-assigned weight factor, which is then adjusted upward or downward by dynamic load information. See Section 6.3.2, *Server Weight and Scheduling* for more on weighting real servers.

Weighted round-robin scheduling is a better choice if there are significant differences in the capacity of certain real servers in the pool. However, if the request load varies dramatically, the more heavily weighted server may answer more than its share of requests.

Least-Connection

Distributes more requests to real servers with fewer active connections. Because it keeps track of live connections to the real servers through the IPVS table, least-connection is a type of dynamic scheduling algorithm, making it a better choice if there is a high degree of variation in the request load. It is best suited for a real server pool where each member node has roughly the same capacity. If a group of servers have different capabilities, weighted least-connection scheduling is a better choice.

Weighted Least-Connections (default)

Distributes more requests to servers with fewer active connections relative to their capacities. Capacity is indicated by a user-assigned weight, which is then adjusted upward or downward by dynamic load information. The addition of weighting makes this algorithm ideal when the real server pool contains hardware of varying capacity. See Section 6.3.2, *Server Weight and Scheduling* for more on weighting real servers.

Locality-Based Least-Connection Scheduling

Distributes more requests to servers with fewer active connections relative to their destination IPs. This algorithm is designed for use in a proxy-cache server cluster. It routes the packets for an IP address to the server for that address unless that server is above its capacity and has a server in its half load, in which case it will assign the least loaded real server to the IP address.

Locality-Based Least-Connection Scheduling with Replication Scheduling (R)

Distributes more requests to servers with fewer active connections relative to their destination IPs. This algorithm is also designed for use in a proxy-cache server cluster. It differs from Locality-Based Least-Connection Scheduling by mapping the target IP address to a subset of real server nodes. Requests are then routed to the server in this subset with the lowest number

of connections. If all the nodes for the destination IP are above capacity, it replicates a new server for that destination IP address by adding the real server with the least connections from the overall pool of real servers to the subset of real servers for that destination IP. The most loaded node is then dropped from the real server subset to prevent over-replication.

Destination Hash Scheduling

Distributes requests to the pool of real servers by looking up the destination IP in a static hash table. This algorithm is designed for use in a proxy-cache server cluster.

Source Hash Scheduling

Distributes requests to the pool of real servers by looking up the source IP in a static hash table. This algorithm is designed for LVS routers with multiple firewalls.

6.3.2 Server Weight and Scheduling

The administrator of an LVS cluster can assign a **weight** to each node in the real server pool. This weight is an integer value which is factored into any *weight-aware* scheduling algorithms (such as weighted least-connections) and helps the LVS router load hardware with different capabilities more evenly.

Weights work as a ratio relative to one another. For instance, if one real server has a weight of 1 and the other server has a weight of 5, then the server with a weight of 5 will get 5 connections for every 1 connection the other server gets. The default value for a real server weight is 1.

Although adding weight to varying hardware configurations in a real server pool can help load-balance the cluster more efficiently, it can cause temporary imbalances when a real server is introduced to the real server pool and the virtual server is scheduled using weighted least-connections. To illustrate, let us say there are three servers in the real server pool. Servers A and B are weighted at 1 and the third, server C, is weighted at 2. If server C goes down for some reason, servers A and B will take up the slack. But once server C comes back online, the LVS router will see it has zero connections and flood the server with all incoming requests until it is on par with servers A and B.

To prevent this phenomenon, the administrator can make the virtual server a **quiesce** server — anytime a new real server node comes online, the least-connections table is reset to zero so the LVS router routes requests as if all the real servers were freshly added to the cluster.

6.4 Routing Methods

Red Hat Linux Advanced Server 2.1 uses **Network Address Translation** or **NAT routing** for LVS clustering — affording the administrator tremendous flexibility when utilizing available hardware and integrating the cluster into an existing network.

6.4.1 NAT Routing

Figure 6–3, *An LVS Cluster Implemented with NAT Routing*, illustrates an LVS cluster utilizing NAT routing to move requests between the Internet and a private network.

Figure 6–3 An LVS Cluster Implemented with NAT Routing

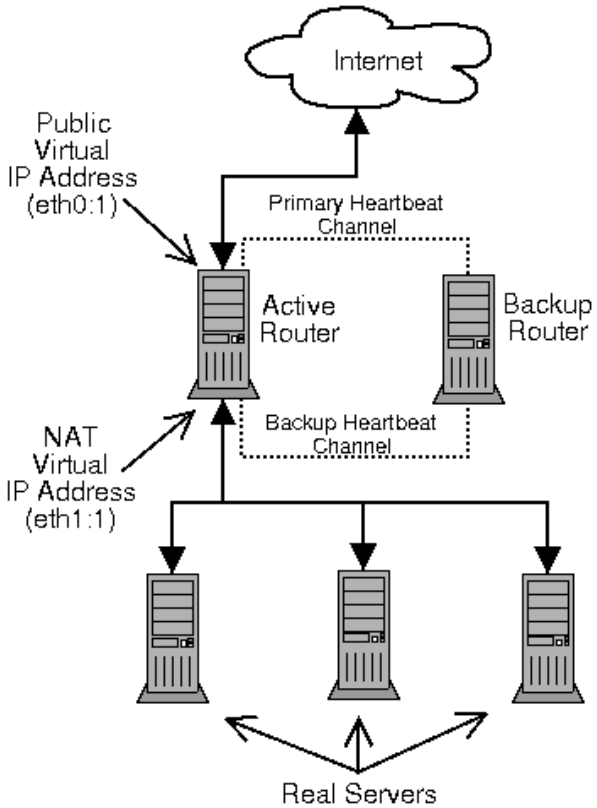

In the example, there are two NICs in the LVS router. The NIC for the Internet has a **real IP address** on eth0 and has a floating IP address aliased to eth0:1. The NIC for the private network interface has a real IP address on eth1 and has a floating IP address aliased to eth1:1. In the event of failover, the virtual interface facing the Internet and the private facing virtual interface are taken-over by the

backup LVS router simultaneously. All of the cluster's real servers located on the private network use the floating IP for the NAT router as their default route to communicate with the active LVS router so that their ability to respond to requests from the Internet never goes down.

In this example, the LVS router's public LVS floating IP address and private NAT floating IP address are aliased to two physical NICs. While it is possible to associate each floating IP address to its own physical device on the LVS router nodes, having more than two NICs is not a requirement.

Under this topography, the active LVS router receives the request and routes it to the appropriate server. The real server then processes the request and returns the packets to the LVS router which uses network address translation to replace the address of the real server in the packets with the LVS routers public VIP address. This process is called **IP masquerading** because the actual IP addresses of the real servers is hidden from the requesting clients.

NAT routing is easy to set up and quite flexible. Using this topography, the real servers may be any kind of machine running any operating system. The chief disadvantage is that, once the real server pool contains more than twenty members, the LVS router may become a bottleneck because it must process outgoing as well as incoming requests.

6.5 Persistence and Firewall Marks

In certain situations, it may be desirable for a client to reconnect repeatedly to the same real server, rather than have an LVS load balancing algorithm send that request to the best available server. Examples of such situations include multi-screen web forms, cookies, SSL, and FTP connections. In these cases, a client may not work properly unless the transactions are being handled by the same server in order to retain context. LVS provides two different features to handle this: **persistence** and **firewall marks**.

6.5.1 Persistence

When enabled, persistence acts like a timer. When a client connects to a service, LVS remembers the last connection for a specified period of time. If that same client IP address connects again within that period, it will be sent to the same server it connected to previously — bypassing the load-balancing mechanisms. When a connection occurs outside the time window, it is handled according to the scheduling rules in place.

Persistence also allows the administrator to specify a subnet mask to apply to the client IP address test as a tool for controlling what addresses have a higher level of persistence, thereby grouping connections to that subnet.

Grouping connections destined for different ports can be important for protocols which use more than one port to communicate, such as FTP. However, persistence is not the most efficient way to deal with the problem of grouping together connections destined for different ports. For these situations, it is best to use **firewall marks**.

6.5.2 Firewall Marks

Firewall marks are an easy and efficient way to a group ports used for a protocol or group of related protocols. For instance, if one is running an e-commerce site, firewall marks can be used to bundle HTTP connections on port 80 and secure, HTTPS connections on port 443. By assigning the same firewall mark to the virtual server for each protocol, state information for the transaction can be preserved because the LVS router will route all requests to the same real server after a connection is opened.

Because of its efficiency and ease-of-use, administrators of LVS clusters should use firewall marks instead of persistence whenever possible for grouping connections. However, administrators should still add persistence to the virtual servers in conjunction with firewall marks to ensure the clients are reconnected to the same server for an adequate period of time.

6.6 LVS Cluster — A Block Diagram

LVS routers use a collection of programs to monitor cluster members and cluster services. Figure 6–4, *Components of a Running LVS Cluster*, illustrates how these various programs on both the active and backup LVS routers work together to manage the cluster.

Figure 6–4 Components of a Running LVS Cluster

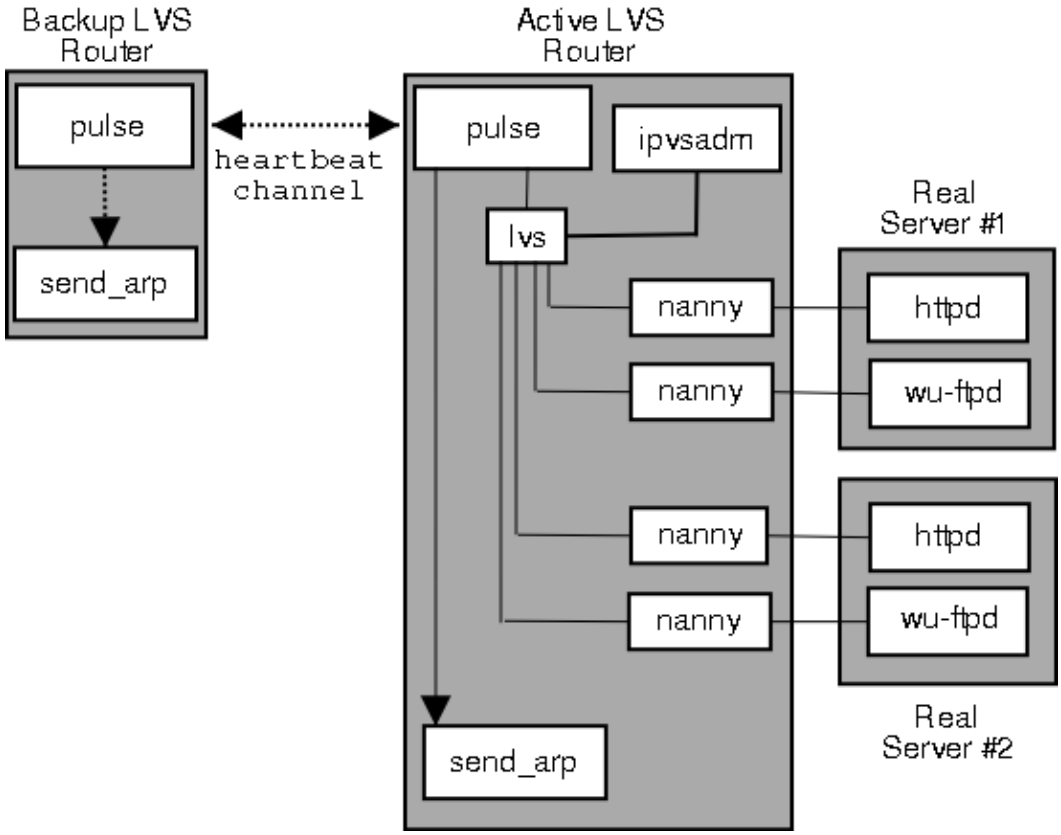

The `pulse` daemon is runs on both the active and passive LVS routers. On the backup router, `pulse` sends a **heartbeat** to the public interface of the active router to make sure the active router is still functioning correctly. On the active router, `pulse` starts the `lvs` daemon and responds to **heartbeat** queries from the backup LVS router.

Once started, the `lvs` daemon calls the `ipvsadm` service to configure and maintain the IPVS routing table in the kernel and starts a `nanny` process for each configured virtual server on each real server. Each `nanny` process checks the state of one configured service on one real server, and tells the `lvs` daemon if the service on that real server is malfunctioning. If a malfunction is detected, the `lvs` daemon instructs `ipvsadm` to remove that real server from the kernel's IPVS routing table.

If the backup router does not receive a response from the active router, it will initiate failover by calling `send_arp` to reassign all virtual IP addresses to the NIC hardware addresses (**MAC** address) of the backup node, sends a command to the active router via both the public and private network interfaces to shut down the `lvs` daemon on the active router, and starts the `lvs` daemon on the backup node to accept requests for the configured virtual servers.

6.6.1 Components of an LVS Cluster

Below is a list and a description for each software component of an LVS router.

`pulse`

This is the controlling process which starts all other daemons related to LVS routers. At boot time, the daemon is started by the `/etc/rc.d/init.d/pulse` script. It then reads the configuration file `/etc/sysconfig/ha/lvs.cf`. On the active router, `pulse` starts the LVS daemon. On the backup router, `pulse` determines the health of the active router by executing a simple heartbeat at a user-configurable interval. If the active router fails to respond after a user-configurable interval, it initiates failover. During failover, `pulse` on the backup router instructs the `pulse` daemon on the active router to shut down all LVS services, starts the `send_arp` program to reassign the floating IP addresses to the backup router's MAC address, and starts the `lvs` daemon.

`lvs`

The `lvs` daemon runs on the active LVS router once called by pulse. It reads the configuration file `/etc/sysconfig/ha/lvs.cf`, calls the `ipvsadm` service to build and maintain the IPVS routing table, and assigns a `nanny` process for each configured LVS service. If `nanny` reports a real server is down, `lvs` instructs the `ipvsadm` service to remove the real server from the kernel's IPVS routing table.

`ipvsadm`

This service updates the IPVS routing table in the kernel. The `lvs` daemon sets up and administers an LVS cluster by calling `ipvsadm` to add, change, or delete entries in the IPVS routing table.

`nanny`

The `nanny` monitoring daemon runs on the active LVS router. Through this daemon, the active router determines the health of each real server and, optionally, monitors its workload. A separate process runs for each service defined on each real server.

`/etc/sysconfig/ha/lvs.cf`

This is the LVS cluster configuration file. Directly or indirectly, all daemons get their configuration information from this file.

Piranha Configuration Tool

The Web-based tool for monitoring, configuring, and administering an LVS cluster. Normally this is the tool used to maintain `/etc/sysconfig/ha/lvs.cf`.

`send_arp`

This program sends out ARP broadcasts when the floating IP address changes from one node to another during failover.

The next chapter reviews important post-installation configuration steps you should take before configuring Red Hat Linux Advanced Server to be an LVS router.

7 Initial Configuration

After installing Red Hat Linux Advanced Server, you must take some basic steps to set up both the LVS routers and the real servers in the LVS cluster. This chapter reviews these basic steps.

Note

The LVS router node that *will* become the active node once the cluster is started is also referred to as the **primary node**. When configuring an LVS cluster, use the Piranha Configuration Tool on the primary node.

7.1 Configuring Services on the LVS Routers

The Red Hat Linux Advanced Server installation program installs all of the components needed to set up an LVS cluster, but the appropriate services must be activated before configuring the cluster. For both of the LVS routers, set the appropriate services to start at boot time. There are three primary tools for setting services to activate at boot time under Red Hat Linux: the command line program `chkconfig`, the ncurses-based program `ntsysv`, or the graphical application serviceconf. All of these tools require root access.

Tip

To attain root access, open a shell prompt and type the following command followed by the machine's root password:

```
su -
```

On the LVS routers, there are at least three services which need to be set to activate at boot time. They are:

- The `piranha-gui` service (primary node only).
- The `pulse` service.
- The `sshd` service.

If you are clustering multi-port services or using firewall marks, you must also enable *either* the `ipchains` or `iptables` service, but not both.

Important

If both `ipchains` and `iptables` are selected to activate on a particular runlevel, `ipchains` will take precedence. Although `iptables` is the default network packet filtering mechanism under the 2.4 kernel, it cannot be used concurrently with `ipchains`. At boot time, the kernel activates `ipchains` then attempts to activate `iptables`. If `ipchains` rules are present, the kernel will issue an error and fail to start `iptables`.

It is best to set these services to activate in both runlevel 3 and runlevel 5. To accomplish this using `chkconfig`, type the following command for each service:

```
/sbin/chkconfig --level 35 daemon on
```

In the above command, replace *daemon* with the name of the service you are activating. To get a list of services on the system as well as what runlevel they are set to activate on, issue the following command:

```
/sbin/chkconfig --list
```

WARNING

Turning any of the above services on using `chkconfig` does not actually start the daemon. To do this use the `/sbin/service`. See Section 7.3, *Starting the Piranha Configuration Tool Service* for an example of how to use the `service` command.

For more information on runlevels and configuring services with `ntsysv` and serviceconf, refer to the chapter titled *Controlling Access to Services* in the *Official Red Hat Linux Customization Guide*.

7.2 Setting a Password for the Piranha Configuration Tool

Before using the Piranha Configuration Tool for the first time on the primary LVS router, you must restrict access to it by creating a password. To do this, login as root and issue the following command:

```
/usr/sbin/piranha-passwd
```

After entering this command, you will be prompted to create the administrative password — taking care to select a password that is secure.

WARNING

For a password to be more secure, it should not contain proper nouns, commonly used acronyms, or words in a dictionary from any language. Also, do not leave the password in clear text anywhere on the system.

If the password is changed during an active Piranha Configuration Tool session, the administrator will be prompted to provide the new password.

7.3 Starting the Piranha Configuration Tool Service

After you have set the password for the Piranha Configuration Tool, start or restart the `piranha-gui` service located in `/etc/rc.d/init.d/piranha-gui`. To do this, type the following command as root:

```
/sbin/service piranha-gui start
```

Issuing this command starts a private session of Apache (`httpd`) by calling the symbolic link `/usr/sbin/piranha_gui -> /usr/sbin/httpd`. For security reasons, the `piranha-gui` version of `httpd` runs as the piranha user in a separate process. The fact the `piranha-gui` leverages the `httpd` service means two things:

1. Apache must be installed on the system.
2. Stopping or restarting Apache via the `service` command will stop the `piranha-gui` service.

WARNING

If the command `service httpd stop` or `service httpd restart` is issued on an LVS router, you will need to start the `piranha-gui` service by issuing the following command:

```
/sbin/service piranha-gui start
```

The `piranha-gui` service is all that is necessary to begin configuring an LVS cluster. However, if you are configuring the cluster remotely, the `sshd` service is also required. You do *not* need to start the `pulse` service until configuration using the Piranha Configuration Tool is complete. See Section 9.8, *Starting the Cluster* for information on starting the `pulse` service.

7.3.1 Configuring the Piranha Configuration Tool Web Server Port

The Piranha Configuration Tool runs on port 3636 by default. To change this port number, change the line `Port 3636` in Section 2 of the `piranha-gui` Web server configuration file `/etc/sysconfig/ha/conf/httpd.conf`.

To use the Piranha Configuration Tool you will need at minimum a text-only Web browser. If you start a Web browser on the primary LVS router, open the location **`http://localhost:3636`**. You can reach the Piranha Configuration Tool from anywhere on through a Web browser by replacing *localhost* with the hostname or IP address of the primary LVS router.

When your browser connects to the Piranha Configuration Tool, you will notice that you must login to access the cluster configuration services. Enter **`piranha`** in the **Username** field and the password set with `piranha-passwd` in the **Password** field.

Now that the Piranha Configuration Tool is running, you may wish to consider limiting who has access to the tool over the network. The next section reviews ways to accomplish this task.

7.4 Limiting Access To the Piranha Configuration Tool

The Piranha Configuration Tool prompts for a valid username and password combination. However, because all of the data passed to the Piranha Configuration Tool is in clear-text, it is recommended that you restrict access to trusted networks or only to the local machine.

The easiest way to restrict access is to use Apache's built in access control mechanisms by editing `/etc/sysconfig/ha/web/secure/.htaccess`. After altering the file you do not have to restart the `piranha-gui` service because the server checks the `.htaccess` file each time it accesses the directory.

By default, the access controls for this directory allow anyone to view the contents of the directory. Here is what the default access looks like:

```
Order deny,allow
Allow from all
```

To limit access to the Piranha Configuration Tool to only the localhost change the `.htaccess` file to allow access from only the loopback device (127.0.0.1). For more information on the loopback device, see the chapter titled *Network Scripts* in the *Official Red Hat Linux Reference Guide*.

```
Order deny,allow
Deny from all
Allow from 127.0.0.1
```

You can also allow specific hosts or subnets as in this example:

```
Order deny,allow
Deny from all
Allow from 192.168.1.100
Allow from 172.16.57
```

In the above example, only Web browsers from the machine with the IP address of 192.168.1.100 and machines on the 172.16.57/24 network can access the Piranha Configuration Tool.

Editing this file limits access to the configuration pages in `/etc/sysconfig/ha/web/secure/` directory but not to the login and the help pages in `/etc/sysconfig/ha/web/`. To limit access to this directory, create a `.htaccess` file in the `/etc/sysconfig/ha/web/` with order, allow, and deny lines identical to `/etc/sysconfig/ha/web/secure/.htaccess`.

7.5 Turning on Packet Forwarding

In order for the LVS router to forward network packets properly to the real servers, each LVS router node must have IP forwarding turned on in the kernel. Log in as root and change the line which reads `net.ipv4.ip_forward = 0` in `/etc/sysctl.conf` to the following:

```
net.ipv4.ip_forward = 1
```

The changes will take effect when you reboot the system.

Tip

If this is the first time booting into Red Hat Linux Advanced Server, you will also need to manually turn on forwarding by issuing the following command:

```
echo 1 > /proc/sys/net/ipv4/ip_forward
```

To turn on IP forwarding manually, issue the following command as root:

```
cat /proc/sys/net/ipv4/ip_forward
```

If the above command returns a `1`, IP forwarding is enabled. If it returns a `0`, then you must issue the `echo` command listed above.

7.6 Configuring Services on the Real Servers

If the real servers in the cluster are Red Hat Linux systems, set the appropriate the server daemons to activate at boot time. These daemons can include `httpd` for Web services or `xinetd` for FTP or telnet services.

It may also be useful to access the real servers remotely, so the `sshd` daemon should also be running.

8 Setting Up a Red Hat Linux Advanced Server LVS Cluster

A Red Hat Linux Advanced Server LVS cluster consists of two basic groups: the LVS routers and the real servers. To prevent a single point of failure, each groups should contain at least two member systems.

The LVS router group should consist of two identical or very similar systems running Red Hat Linux Advanced Server. One will act as the active LVS router while the other stays in hot standby mode, so they need to have as close to the same capabilities as possible.

Before choosing and configuring the hardware for the real server group, you most decide what which of the three types of LVS topographies to use.

8.1 The NAT LVS Cluster

The NAT topography allows for great latitude in utilizing existing hardware, but it is limited in its ability to handle large loads due to the fact that all packets going into and coming out of the cluster pass through the LVS router.

Network Layout

The topography for an LVS cluster utilizing NAT routing is the easiest to configure from a network layout perspective because the cluster needs only one access point to the public network. The real servers pass all requests back through the LVS router so they are on their own private network.

Hardware

The NAT topography is the most flexible in regards to cluster hardware because the real servers do not need to be Linux machines to function correctly in the cluster. In a NAT cluster, each real server only needs one NIC since it will only be responding to the LVS router. The LVS routers, on the other hand, need two NICs each to route traffic between the two networks. Because this topography creates a network bottleneck at the LVS router, gigabit Ethernet NICs can be employed on each LVS router to increase the bandwidth the LVS routers can handle. If gigabit Ethernet is employed on the LVS routers, any switch connecting the real servers to the LVS routers must have at least two gigabit Ethernet ports to handle the load efficiently.

Software

Because the NAT topography requires the use of `iptables` or `ipchains` for some configurations, there can be a fair amount of software configuration outside of Piranha Configuration

Tool. In particular, FTP services and the use of firewall marks requires extra manual configuration of the LVS routers to route requests properly.

8.1.1 Putting the Cluster Together

After determining which of the above routing methods to use, the hardware for the LVS cluster should be linked together on the network.

Important

The adapter devices on the LVS routers must be configured to access the same networks. For instance if `eth0` connects to public network and `eth1` connects to the private network, then these same devices on the backup LVS router must connect to the same networks.

Also the gateway listed in the first interface to come up at boot time is added to the routing table and subsequent gateways listed in other interfaces are ignored. This is especially important to consider when configuring the real servers.

After physically connecting together the cluster hardware, configure the network interfaces on the primary and backup LVS routers. This can be done using a graphical application such as redhat-config-network or by editing the network scripts manually. For more information about adding devices using redhat-config-network, see the chapter titled *Network Configuration* in the *Official Red Hat Linux Customization Guide*. For more information on editing network scripts by hand, see the chapter titled *Network Scripts* in the *Official Red Hat Linux Reference Guide*. For the remainder of the chapter, example alterations to network interfaces are made either manually or through the Piranha Configuration Tool.

General LVS Networking Tips

Configure the real IP addresses for both the public and private networks on the LVS routers before attempting to configure the cluster using the Piranha Configuration Tool. The sections on each topography give example network addresses, but the actual network addresses are needed. Below are some useful commands for bringing up network interfaces or checking their status.

Bringing Up Real Network Interfaces

The best way to bring up any real network interface is to use the following commands as root replacing *N* with the number corresponding to the interface (`eth0` and `eth1`):

```
/sbin/ifup ethN
```

WARNING

Do *not* use the `ifup` scripts to bring up any floating IP addresses you may configure using Piranha Configuration Tool (`eth0:1` or `eth1:1`). Use the `service` command to start `pulse` instead (see Section 9.8, *Starting the Cluster* for details).

To bring a network interface down, simply type:

```
/sbin/ifdown ethN
```

Again, replace *N* in the above command with the number corresponding to the interface you wish to bring down.

Checking the Status of Network Interfaces

If you need to check which network interfaces are up at any given time, type:

```
/sbin/ifconfig
```

To view the routing table for a machine, issue the following command:

```
/sbin/route
```

8.2 Configuring Network Interfaces for a NAT LVS Cluster

To set up a NAT LVS cluster, the administrator must first configure the network interfaces for the public network and the private network on the LVS routers. In this example, the LVS routers' public interfaces (`eth0`) will be on the 192.168.26/24 network (I know, I know, this is not a routable IP, but let us pretend there is a firewall in front of the LVS router for good measure) and the private interfaces which link to the real servers (`eth1`) will be on the 10.11.12/24 network.

So on the active or **primary** LVS router node, the public interface's network script, `/etc/sysconfig/network-scripts/ifcfg-eth0`, could look something like this:

```
DEVICE=eth0
BOOTPROTO=static
ONBOOT=yes
IPADDR=192.168.26.9
NETMASK=255.255.255.0
GATEWAY=192.168.26.254
```

The `/etc/sysconfig/network-scripts/ifcfg-eth1` for the private NAT interface on the LVS router could look something like this:

```
DEVICE=eth1
BOOTPROTO=static
ONBOOT=yes
IPADDR=10.11.12.9
NETMASK=255.255.255.0
```

In this example, the VIP for the LVS router's public interface will be 192.168.26.10 and the VIP for the NAT or private interface will be 10.11.12.10. So, it is essential that the real servers route requests back to the VIP for the NAT interface.

Important

The sample Ethernet interface configuration settings in this section are for the real IP addresses of an LVS router and *not* the floating IP addresses. To configure the public and private floating IP addresses the administrator should use the Piranha Configuration Tool, as shown in Section 9.4, **GLOBAL SETTINGS** and Section 9.6.1, *The* **VIRTUAL SERVER** *Subsection*.

After configuring the primary LVS router node's network interfaces, configure the backup LVS router's real network interfaces — taking care that none of the IP address conflict with any other IP addresses on the network.

Important

Be sure each interface on the backup node services the same network as the interface on primary node. For instance, if eth0 connects to the public network on the primary node, it must also connect to the public network on the backup node as well.

8.2.1 Routing on the Real Servers

The most important thing to remember when configuring the real servers network interfaces in a NAT cluster is to set the gateway for the NAT floating IP address of the LVS router. In this example, that address will be 10.11.12.10.

Note

Once the network interfaces are up on the real servers, the machines will be unable to ping or connect in other ways to the public network. This is normal. You will, however, be able to ping the real IP for the LVS router's private interface, in this case 10.11.12.8.

So the real server's `/etc/sysconfig/network-scripts/ifcfg-eth0` file could look similar to this:

```
DEVICE=eth0
ONBOOT=yes
BOOTPROTO=static
IPADDR=10.11.12.1
NETMASK=255.255.255.0
GATEWAY=10.11.12.10
```

WARNING

If a real server has more than one network interface configured with a `GATEWAY=` line, the first one to come up will get the gateway. Therefore if both `eth0` and `eth1` are configured and `eth1` is used for LVS clustering, the real servers may not route requests properly.

It is best to turn off extraneous network interfaces by setting `ONBOOT=no` in their network scripts within the `/etc/sysconfig/network-scripts/` directory or by making sure the gateway is correctly set in the interface which comes up first.

8.2.2 Enabling NAT Routing on the LVS Routers

In a simple NAT LVS cluster where each clustered service uses only one port, like HTTP on port 80, the administrator needs only to enable packet forwarding on the LVS routers for the requests to be properly routed between the outside world and the real servers. See Section 7.5, *Turning on Packet Forwarding* for instructions on turning on packet forwarding. However, more configuration is necessary when the clustered services require more than one port to go to the same real server during a user session. For information on creating multi-port services using firewall marks, see Section 8.3, *Multi-port Services and LVS Clustering*.

Once forwarding is enabled on the LVS routers and the real servers are set up and have the clustered services running, use the Piranha Configuration Tool to configure the cluster as shown in Chapter 9, *Configuring the LVS Routers with Piranha Configuration Tool.*

WARNING

Do not configure the floating IP for `eth0:1` or `eth1:1` by manually editing network scripts or using a network configuration tool. Instead, use the Piranha Configuration Tool as shown in Section 9.4, GLOBAL SETTINGS and Section 9.6.1, *The* VIRTUAL SERVER *Subsection* to configure any cluster-related virtual interfaces.

When finished, start the `pulse` service as shown in Section 9.8, *Starting the Cluster*. Once `pulse` is up and running, the active LVS router will begin routing requests to the pool of real servers.

8.3 Multi-port Services and LVS Clustering

LVS routers under any topology require extra configuration when creating multi-port LVS services. Multi-port services can be created artificially by using firewall marks to bundle together different, but related protocols, such as HTTP (port 80) and HTTPS (port 443), or when LVS is used to cluster true multi-port protocols, such as FTP. In either case, the LVS router uses firewall marks to recognize that packets destined for different ports, but bearing the same firewall mark, should be handled identically. Also, when combined with persistence, firewall marks ensure connections from the client machine are routed to the same host, as long as the connections occur within the length of time specified by the persistence parameter. For more on assigning persistence to a virtual server, see Section 9.6.1, *The* **VIRTUAL SERVER** *Subsection*.

Unfortunately, the mechanism used to balance the loads on the real servers — IPVS — can recognize the firewall marks assigned to a packet, but cannot itself assign firewall marks. The job of *assigning* firewall marks must be performed by a network packet filter, such as `iptables` or `ipchains` outside of Piranha Configuration Tool.

8.3.1 Assigning Firewall Marks

To assign firewall marks to a packet destined for a particular port, the administrator must use `iptables` or `ipchains`.

Important

In the explanation below, both `iptables` and `ipchains` commands are used to illustrate marking packets. However, administrators should be aware that `iptables` is the preferred method for filtering network packets, because it is a more secure, robust, and flexible technology. Also, in future versions of Red Hat Linux `ipchains` may no longer be supported.

This section illustrates how to bundle HTTP and HTTPS as an example, however FTP is another commonly clustered multi-port protocol. If an LVS cluster is used for FTP services, see Section 8.4, *FTP In an LVS Cluster* for details on how to best configure the cluster.

The basic rule to remember when using firewall marks is that for every protocol using a firewall mark in Piranha Configuration Tool there must be a commensurate `iptables` or `ipchains` rule to assign marks to the network packets.

Before creating network packet filter rules, make sure there are no rules already in place. To do this, open a shell prompt, login as root, and type:

```
/sbin/service iptables status
/sbin/service ipchains status
```

If `iptables` is not running, the prompt will instantly reappear. If `ipchains` is not running, the following message will appear:

```
ipchains: Incompatible with this kernel
```

If either one of these network packet filters are active, it displays a set of rules. If rules are present, type whichever of the following commands correlates to the active network packet filter:

```
/sbin/service iptables stop
/sbin/service ipchains stop
```

If the rules already in place are important, check the contents of `/etc/sysconfig/iptables` or `/etc/sysconfig/ipchains` and copy any rules worth keeping to a safe place before proceeding.

Important

Although `iptables` is the default network packet filtering mechanism under the 2.4 kernel, it cannot be used concurrently with `ipchains`. If you attempt to `/sbin/modprobe ip_tables` and the `ipchains` module is already loaded, the kernel will issue an error and fail to start `iptables`. The inverse of this statement is also true. To see which modules are loaded, type `/sbin/lsmod`.

You are free to choose either `iptables` or `ipchains` as the network packet filter for the LVS routers, but be aware that `iptables` is preferred and that support for `ipchains` is not guaranteed in future releases.

Below are rules which assign the same firewall mark, 80, to incoming traffic destined for the floating IP address, *n.n.n.n*, on ports 80 and 443. For instructions on assigning the VIP to the public network interface, see Section 9.6.1, *The* **VIRTUAL SERVER** *Subsection*. Also note that you must log in as root and load the module for either `iptables` or `ipchains` before issuing rules for the first time.

`iptables`

```
/sbin/modprobe ip_tables
/sbin/iptables -t mangle -A PREROUTING -p tcp \
               -d n.n.n.n/32 \
               --dport 80 -j MARK --set-mark 80
/sbin/iptables -t mangle-A PREROUTING -p tcp \
               -d n.n.n.n/32 \
               --dport 443 -j MARK --set-mark 80
```

`ipchains`

```
/sbin/modprobe ipchains
/sbin/ipchains -A input -p tcp -d n.n.n.n/32 80 -m 80
/sbin/ipchains -A input -p tcp -d n.n.n.n/32 443 -m 80
```

In the above network packet filter commands, *n.n.n.n* should be replaced with the floating IP for your HTTP and HTTPS virtual servers. These commands have the net effect of assigning any traffic addressed to the VIP on the appropriate ports a firewall mark of 80, which in turn is recognized by IPVS and forwarded appropriately.

WARNING

The commands above will take effect immediately, but do not persist through a reboot of the system. To ensure network packet filter settings are restored upon reboot, refer to Section 8.5, *Saving Network Packet Filter Settings*

8.4 FTP In an LVS Cluster

File Transport Protocol (FTP) is an old and complex multi-port protocol that presents a distinct set of challenges to a clustered environment. To understand the nature of these challenges, you must first understand some key things about how FTP works.

8.4.1 How FTP Works

With most other server client relationships, the client machine opens up a connection to the server on a particular port and the server then responds to the client on that port. When an FTP client connects to an FTP server it opens a connection to the FTP control port 21. Then the *client* tells the FTP *server* whether to establish an **active** or **passive** connection. The type of connection chosen by the client determines how the server responds and on what ports transactions will occur.

The two types of data connections are:

Active Connections

When an active connection is established, the *server* opens a data connection to the client from port 20 to a high range port on the client machine. All data from the server is then passed over this connection.

Passive Connections

When an passive connection is established, the *client* asks the FTP server to establish a passive connection port, which can be on any port higher than 10,000. The server then binds to this high-numbered port for this particular session and relays that port number back to the client. The client then opens the newly bound port for the data connection. Each data request the client makes results in a seperate data connection. Most modern FTP clients attempt to establish a passive connection to FTP

The two important things to note about all of this in regards to clustering is:

1. The *client* determines the type of connection, not the server. This means, to effectively cluster FTP, you must configure the LVS routers to handle both active and passive connections.

2. The FTP client/server relationship can potentially open a large number of ports that the Piranha Configuration Tool and IPVS do not know about.

8.4.2 How This Affects LVS Routing

IPVS packet forwarding only allows connections in and out of the cluster based on it recognizing its port number or its firewall mark. If a client from outside the cluster attempts to open a port IPVS is not configured to handle, it drops the connection. Similarly, if the real server attempts to open a connection back out to the Internet on a port IPVS does not know about, it drops the connection. This means *all* connections from FTP clients on the Internet *must* have the same firewall mark assigned to them and all connections from the FTP server *must* be properly forwarded to the Internet using network packet filtering rules.

8.4.3 Creating Network Packet Filter Rules

Before assigning any `iptables` or `ipchains` rules for FTP service, review the information in Section 8.3.1, *Assigning Firewall Marks* concerning multi-port services and techniques for checking the existing network packet filtering rules.

Note

You must log in as root and load the module for either `iptables` or `ipchains` before issuing rules for the first time.

Below are rules which assign the same firewall mark, 21, to FTP traffic. For these rules to work properly, you must also use the **VIRTUAL SERVER** subsection of Piranha Configuration Tool to configure a virtual server for port 21 with a value of **21** in the **Firewall Mark** field. See Section 9.6.1, *The* **VIRTUAL SERVER** *Subsection* for details.

Rules for Active Connections

The rules for active connections tell the kernel to accept and forward connections coming to the *internal* floating IP address on port 20 — the FTP data port.

iptables

```
/sbin/iptables -t nat -A POSTROUTING -p tcp \
               -s n.n.n.0/24 --sport 20 -j MASQUERADE
```

ipchains

```
/sbin/ipchains -A forward  -p tcp -s n.n.n.0/24 20 -j MASQ
```

In the above network packet filter commands, *n.n.n* should be replaced with the first three values for the floating IP for the NAT interface's internal network interface defined in the **GLOBAL SETTINGS** panel of Piranha Configuration Tool. These commands allow the LVS router to accept outgoing connections from the real servers that IPVS does not know about.

Rules for Passive Connections

The rules for passive connections assign the appropriate firewall mark to connections coming in from the Internet to the floating IP for the service on a wide range of ports — 10,000 to 20,000.

WARNING

If you are limiting the port range for passive connections, you must also configure the FTP server to use a matching port range. This can be accomplished under Red Hat Linux by adding the following line to the end of `/etc/ftpaccess`:

```
passive ports 0.0.0.0/0 10000 20000
```

For other FTP servers, consult the documetation.

This range should be a wide enough for most situations; however, you can increase this number to include all available non-secured ports by changing `10000:20000` in the commands below to `1024:65535`.

`iptables`

```
/sbin/iptables -t mangle -A PREROUTING -p tcp \
               -d n.n.n.n/32  \
               --dport 21 -j MARK --set-mark 21
/sbin/iptables -t mangle -A PREROUTING -p tcp \
               -d n.n.n.n/32 \
               --dport 10000:20000 -j MARK --set-mark 21
```

`ipchains`

```
/sbin/ipchains -A input -p tcp -d n.n.n.n/32 21 -m 21
/sbin/ipchains -A input -p tcp -d n.n.n.n/32 \
               10000:20000 -m 21
```

In the above network packet filter commands, *n.n.n.n* should be replaced with the floating IP for the FTP virtual server defined in the **VIRTUAL SERVER** subsection of Piranha Configuration Tool. These commands have the net effect of assigning any traffic addressed to the floating IP on the appropriate ports a firewall mark of 21, which is in turn recognized by IPVS and forwarded appropriately.

WARNING

The commands above take effect immediately, but do not persist through a reboot of the system. To ensure network packet filter settings are restored after a reboot, see Section 8.5, *Saving Network Packet Filter Settings*

Finally, you will need to be sure the appropriate service is set to activate on the proper runlevels. For more on this, see Section 7.1, *Configuring Services on the LVS Routers*.

8.5 Saving Network Packet Filter Settings

After configuring the appropriate network packet filters for your situation, save the settings so they get restored after a reboot. For `iptables`, type the following command:

```
/sbin/iptables-save
```

With `ipchains`, type the following command:

```
/sbin/ipchains-save
```

This will save the settings in `/etc/sysconfig/iptables` or `/etc/sysconfig/ipchains` so they can be recalled at boot time.

Once this file is written, you will also be able to use the `/sbin/service` command to start, stop, and check the status (using the status switch) of `iptables` or `ipchains`. The `/sbin/service` will automatically load the appropriate module for you. For an example of how to use the `/sbin/service` command, see Section 7.3, *Starting the Piranha Configuration Tool Service*.

Finally, you will need to be sure the appropriate service is set to activate on the proper runlevels. For more on this, see Section 7.1, *Configuring Services on the LVS Routers*.

The next chapter will explain how to use the Piranha Configuration Tool to configure the LVS router and describe the steps necessary to active an LVS cluster.

9 Configuring the LVS Routers with Piranha Configuration Tool

The Piranha Configuration Tool provides a structured approach to creating the necessary configuration file for a Piranha cluster — `/etc/sysconfig/ha/lvs.cf`. This chapter describes the basic operation of the Piranha Configuration Tool and how to activate the cluster once configuration is complete.

Important

The configuration file for the LVS cluster follows strict formatting rules. Using the Piranha Configuration Tool is the best way to prevent syntax errors in the `lvs.cf` and therefore prevent software failures.

9.1 Necessary Software

The `piranha-gui` service must be running on the primary LVS router to use the Piranha Configuration Tool. To configure the cluster, you will minimally need a text-only Web browser, such as `links`. If you are accessing the LVS router from another machine, you will also need an `ssh` connection to the primary LVS router as the root user.

While configuring the primary LVS router it is a good idea to keep a concurrent `ssh` connection in a terminal window. This connection provides a secure way to restart `pulse` and other services, configure network packet filters, and monitor `/var/log/messages` during trouble shooting.

The next four sections walk through each of the configuration pages of the Piranha Configuration Tool and give instructions on using it to set up the LVS cluster.

9.2 Logging Into the Piranha Configuration Tool

When configuring an LVS cluster, you should always begin by configuring the primary router with the Piranha Configuration Tool. To do this,verify that the `piranha-gui` service is running and an administrative password has been set, as described in Section 7.2, *Setting a Password for the Piranha Configuration Tool.*

If you are accessing the machine locally, you can open **`http://localhost:3636`** in a Web browser to access the Piranha Configuration Tool. Otherwise, type in the hostname or real IP address for the server followed by **`:3636`**. Once the browser connects, you will see the screen shown in Figure 9–1, *The Welcome Panel*.

Figure 9–1 The Welcome Panel

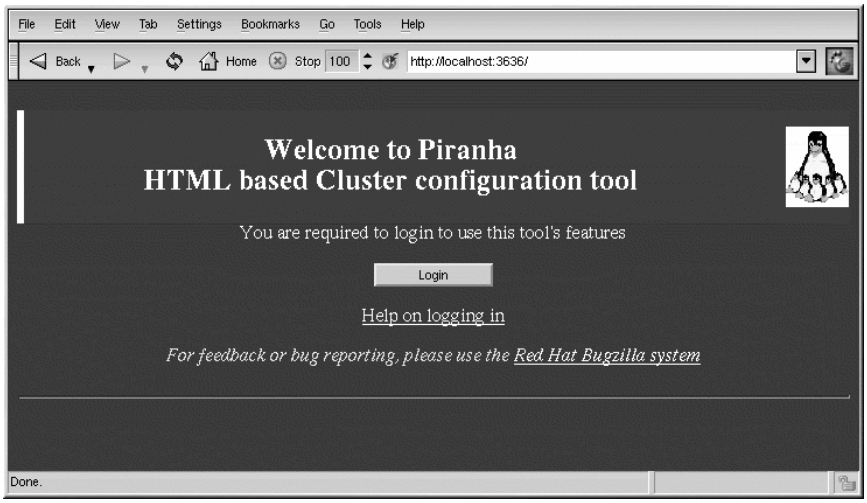

Click on the **Login** button and enter **piranha** for the **Username** and the administrative password you created in the **Password** field.

The Piranha Configuration Tool is made of four main screens or **panels**. In addition, the **Virtual Servers** panel contains four **subsections**. The **CONTROL/MONITORING** panel is the first panel after the login screen.

9.3 CONTROL/MONITORING

The **CONTROL/MONITORING** Panel presents the cluster administrator with a limited runtime status of the cluster. It displays the status of the `pulse` daemon, the LVS routing table, and the LVS-spawned `nanny` processes.

Note

The fields for **CURRENT LVS ROUTING TABLE** and **CURRENT LVS PROCESSES** remain blank until you actually start the cluster, as shown in Section 9.8, *Starting the Cluster*.

Figure 9–2 The CONTROL/MONITORING Panel

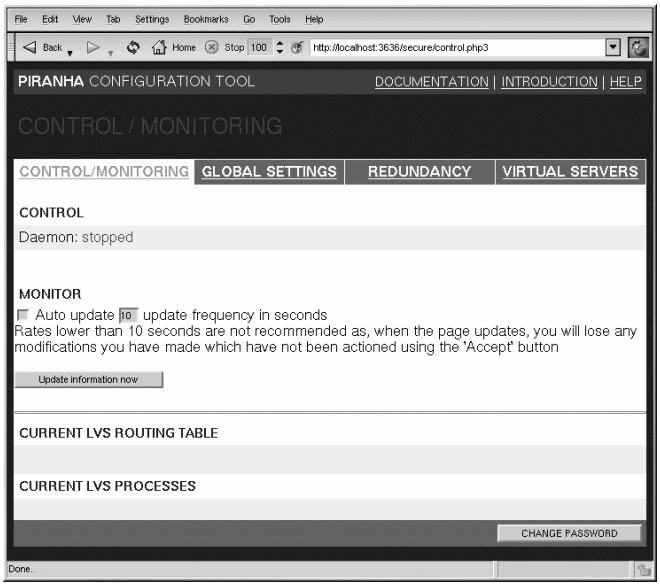

Auto update

The status display on this page can be updated automatically at a user configurable interval. To enable this feature, click on the **Auto update** checkbox and set the desired update frequency in the **Update frequency in seconds** text box (the default value is 10 seconds).

It is not recommended that you set the automatic update to an interval less than 10 seconds. Doing so may make it difficult to reconfigure the **Auto update** interval because the page will update too frequently. If you encounter this issue, simply click on another panel and then back on **CONTROL/MONITORING**.

The **Auto update** feature does not work with all browsers, such as Mozilla.

Update information now

You can manually update the status information manually by clicking this button.

CHANGE PASSWORD

Clicking this button takes you to a help screen with information on how to change the administrative password for the Piranha Configuration Tool.

9.4 GLOBAL SETTINGS

The **GLOBAL SETTINGS** panel is where the cluster administrator defines the networking details for the primary LVS router's public and private network interfaces.

Figure 9–3 The GLOBAL SETTINGS Panel

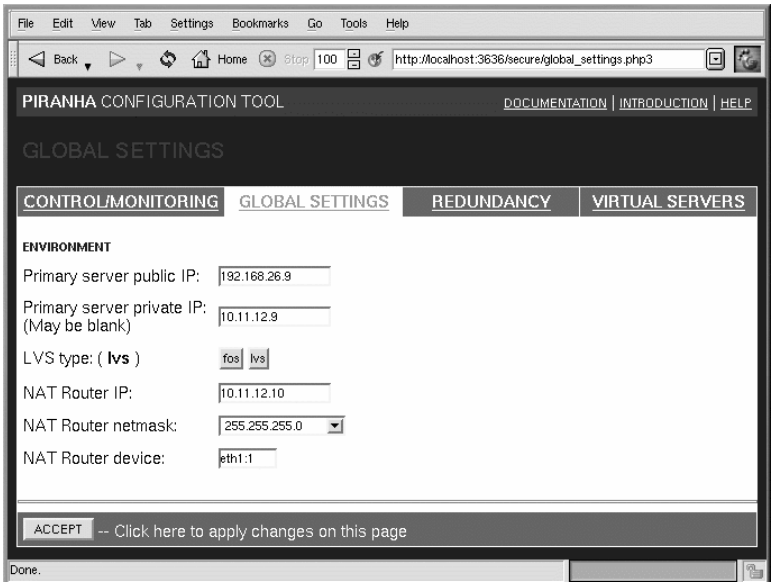

The top half of this panel sets up the primary LVS router's public and private network interfaces. These are the interfaces already configured in Section 8.2, *Configuring Network Interfaces for a NAT LVS Cluster*.

Primary server public IP

In this field, enter the publicly routable real IP address for the primary LVS node.

Primary server private IP

Enter the real IP address for an alternative network interface on the primary LVS node. This address is used solely as an alternative heartbeat channel for the backup router and does not have to correlate to the real private IP address assigned in Section 8.2, *Configuring Network Interfaces for a NAT LVS Cluster*. You may leave this field blank, but doing so will mean there is no alternate heartbeat channel for the backup LVS router to use and therefore will create a single point of failure.

Tip

The primary LVS router's private IP can be configured on any interface that accepts TCP/IP, whether it be an Ethernet adapter or a serial port.

LVS type

Click the **lvs** button to select LVS clustering.

WARNING

Only lvs clustering is supported by Red Hat. If you need a high-availability clustering solution, use Red Hat Cluster Manager. For more informations, see *Red Hat Cluster Manager Installation and Administration Guide*.

The next three fields deal specifically with the NAT router's virtual network interface connected the private network with the real servers.

NAT Router IP

Enter the private floating IP in this text field.

NAT Router Netmask

If the NAT router's floating IP needs a particular netmask, select it from drop-down list.

NAT Router Device

Use this text field to define the device name of the network interface for the floating IP address, such as **`eth1:1`**.

Tip

You should alias the NAT floating IP address to the Ethernet interface connected to the private network. In this example, the private network is on the `eth1` interface, so **`eth1:1`** is the floating IP address.

WARNING

After completing this page, click the ACCEPT button to make sure you do not lose any changes when selecting a new panel.

9.5 REDUNDANCY

The **REDUNDANCY** panel allows you to configure of the backup LVS router node and set various heartbeat monitoring options.

Tip

The first time you visit this screen, it displays an "inactive" **Backup** status and an **ENABLE** button. To configure the backup LVS router, click on the **ENABLE** button so that the screen matches Figure 9–4, *The REDUNDANCY Panel*.

Figure 9–4 The REDUNDANCY Panel

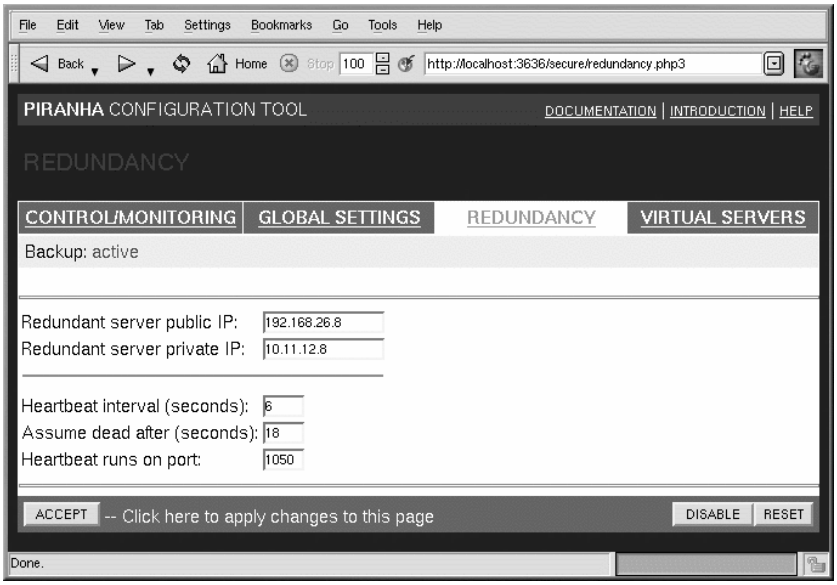

Redundant server public IP

Enter the public real IP address for the backup LVS router node.

Redundant server private IP

Enter the backup node's private real IP address in this text field.

If you do not see the field called **Redundant server private IP**, go back to the **GLOBAL SETTINGS** panel and enter a **Primary server private IP** address and click **ACCEPT**.

The rest of the panel is devoted to configuring the heartbeat channel, which is used by the backup node to monitor the primary node for failure.

Heartbeat Interval (seconds)

This field sets the number of seconds between heartbeats — the interval that the backup node will check the functional status of the primary LVS node.

Assume dead after (seconds)

If the primary LVS node does not respond after this number of seconds, then the backup LVS router node will initiate failover.

Heartbeat runs on port

This field sets the port at which the heartbeat communicates with the primary LVS node. The default is set to 539 if this field is left blank.

WARNING

Remember to click the ACCEPT button after making any changes in this panel to make sure you do not lose any changes when selecting a new panel.

9.6 VIRTUAL SERVERS

The **VIRTUAL SERVERS** panel displays information for each currently defined virtual server. Each table entry shows the status of the virtual server, the server name, the virtual IP assigned to the server, the netmask of the virtual IP, the port number to which the service communicates, the protocol used, and the virtual device interface.

Figure 9–5 The VIRTUAL SERVERS Panel

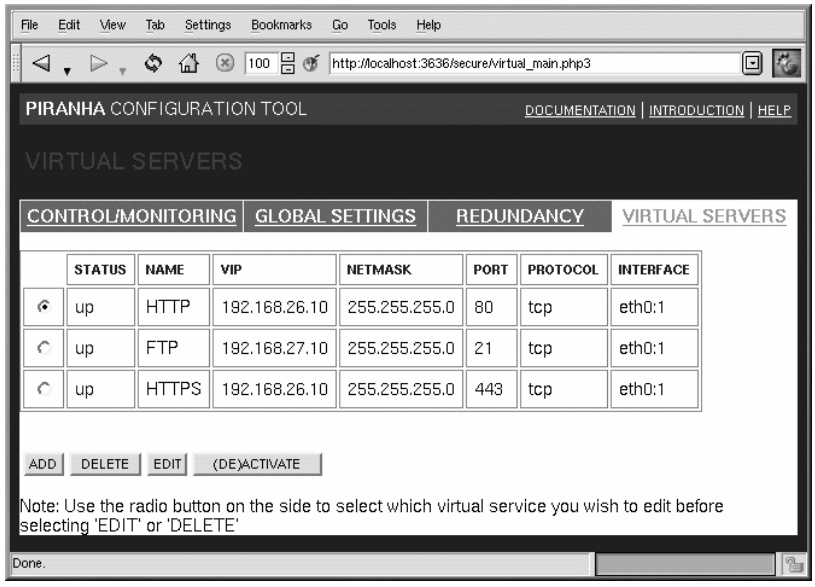

Each server displayed in the **VIRTUAL SERVERS** panel can be configured on subsequent screens or **subsections**.

To add a service, click the **ADD** button. To remove a service, select it by clicking the radio button next to the virtual server and click the **DELETE** button.

To enable or disable a virtual server in the table click its radio button and click the **(DE)ACTIVATE** button.

After adding a virtual server, you can configure it by clicking the radio button to its left and clicking the **EDIT** button to display the **VIRTUAL SERVER** subsection.

9.6.1 The VIRTUAL SERVER Subsection

The **VIRTUAL SERVER** subsection panel shown in Figure 9–6, *The VIRTUAL SERVERS Subsection* allows you to configure an individual virtual server. Links to subsections related specifically to this virtual server are located along the top of the page. But before configuring any of the subsections related to this virtual server, complete this page and click on the **ACCEPT** button.

Figure 9–6 The VIRTUAL SERVERS Subsection

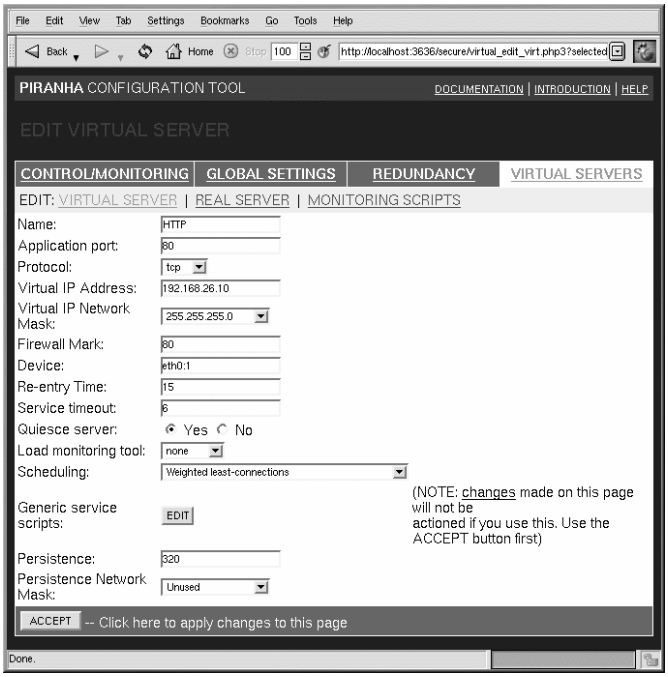

Name

Enter a descriptive name to identify the virtual server. This name is *not* the hostname for the machine, so make it descriptive and easily identifiable. You can even reference the protocol used by the virtual server, such as HTTP.

Application port

Enter the port number through which the service application will listen. Since this example is for HTTP services, port 80 is used.

Protocol

Choose between UDP and TCP in the drop-down menu. Web servers typically communicate via the TCP protocol, so this is selected in the example above.

Virtual IP Address

Enter the virtual server's floating IP address in this text field.

Virtual IP Network Mask

Set the netmask for this virtual server with the drop-down menu.

Firewall Mark

Do *not* enter a firewall mark integer value in this field unless you are bundling multi-port protocols or creating a multi-port virtual server for separate, but related protocols. In this example, the above virtual server has a **Firewall Mark** of 80 because we are bundling connections to HTTP on port 80 and to HTTPS on port 443 using the firewall mark value of 80. When combined with persistence, this technique will ensure users accessing both insecure and secure webpages are routed to the same real server, preserving state.

WARNING

Entering a firewall mark in this field allows IPVS to recognize that packets bearing this firewall mark are treated the same, but you must perform further configuration outside of the Piranha Configuration Tool to actually assign the firewall marks. See Section 8.3, *Multi-port Services and LVS Clustering* for instructions on creating multi-port services and Section 8.4, *FTP In an LVS Cluster* for creating a highly available FTP virtual server.

Device

Enter the name of the network device to which you want the floating IP address defined the **Virtual IP Address** field to bind.

You should alias the public floating IP address to the Ethernet interface connected to the public network. In this example, the public network is on the `eth0` interface, so **`eth0:1`** should be entered as the device name.

Re-entry Time

Enter an integer value which defines the length of time, in seconds, before the active LVS router attempts to bring a real server back into the cluster after a failure.

Service Timeout

Enter an integer value which defines the length of time, in seconds, before a real server is considered dead and removed from the cluster.

Quiesce server

When the **Quiesce server** radio button is selected, anytime a new real server node comes online, the least-connections table is reset to zero so the active LVS router routes requests as if all the real servers were freshly added to the cluster. This option prevents the a new server from becoming bogged down with a high number of connections upon entering the cluster.

Load monitoring tool

The LVS router can monitor the load on the various real servers by using either `rup` or `ruptime`. If you select `rup` from the drop-down menu, each real server must run the `rstatd` service. If you select `ruptime`, each real server must run the `rwhod` service.

Load monitoring is *not* the same as load balancing and can result in hard to predict scheduling behavior when combined with weighted scheduling algorithms. Also, if you use load monitoring, the real servers in the cluster must be Linux machines.

Scheduling

Select your preferred scheduling algorithm from the drop-down menu. The default is **`Weighted least-connection`**. For more information on scheduling algorithms, see Section 6.3.1, *Scheduling Algorithms*.

Persistence

If an administrator needs persistent connections to the virtual server during client transactions, enter the number of seconds of inactivity allowed to lapse before a connection times out in this text field.

Important

If you entered a value in the **Firewall Mark** field above, you should enter a value for persistence as well. Also, be sure that if you use firewall marks and persistence together, that the amount of persistence is the same for each virtual server with the firewall mark. For more on persistence and firewall marks, see Section 6.5, *Persistence and Firewall Marks*.

Persistence Network Mask

To limit persistence to particular subnet, select the appropriate network mask from the drop-down menu.

Note

Before the advent of firewall marks, persistence limited by subnet was a crude way of bundling connections. Now, it is best to use persistence in relation to firewall marks to achieve the same result.

WARNING

Remember to click the ACCEPT button after making any changes in this panel. To make sure you do not lose changes when selecting a new panel.

9.6.2 REAL SERVER Subsection

Clicking on the **REAL SERVER** subsection link at the top of the panel displays the **EDIT REAL SERVER** subsection. It displays the status of the physical server hosts for a particular virtual service.

Figure 9–7 The REAL SERVER Subsection

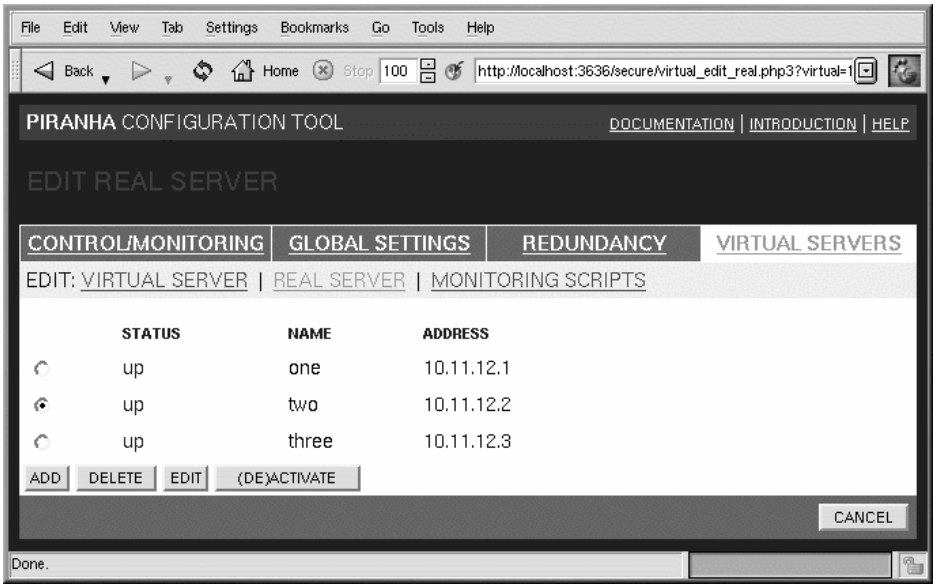

Click the **ADD** button to add a new server. To delete an existing server, select the radio button beside it and click the **DELETE** button. Click the **EDIT** button to load the **EDIT REAL SERVER** panel, as seen in Figure 9–8, *The REAL SERVER Configuration Panel*.

Figure 9–8 The REAL SERVER Configuration Panel

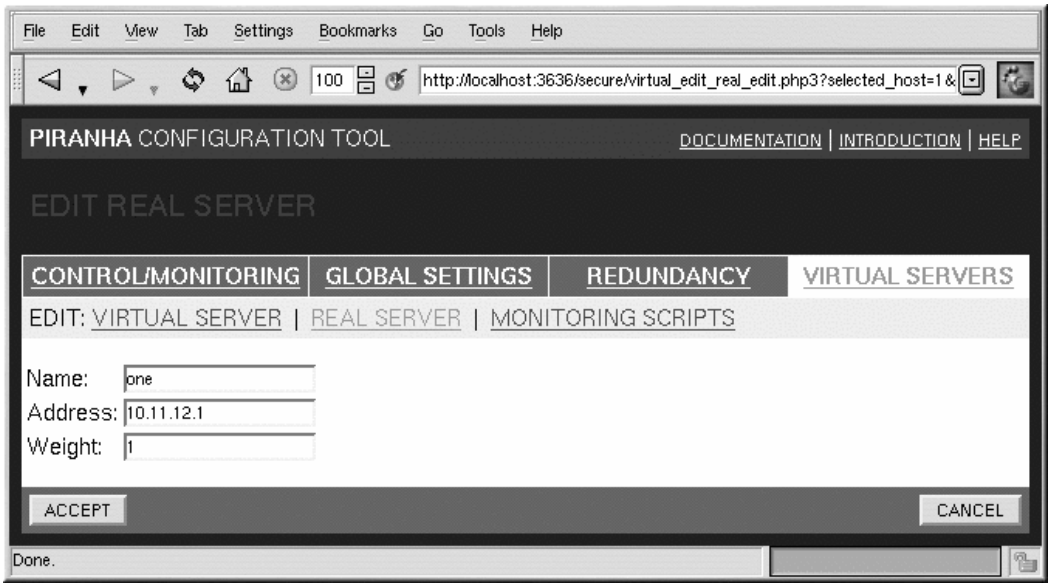

This panel consists of three entry fields:

Name

A descriptive name for the real server.

Tip

This name is *not* the hostname for the machine, so make it descriptive and easily identifiable.

Address

The real server's IP address. Since the listening port is already specified for the associated virtual server, do not add a port number.

Weight

An integer value indicating this host's capacity relative to that of other hosts in the pool. The value can be arbitrary, but treat it as a ratio in relation to other real servers in the cluster. For more on server weight, see Section 6.3.2, *Server Weight and Scheduling*.

WARNING

Remember to click the ACCEPT button after making any changes in this panel. To make sure you do not lose any changes when selecting a new panel.

9.6.3 EDIT MONITORING SCRIPTS Subsection

There are two ways to access the **EDIT MONITORING SCRIPTS** subsection: by clicking the **EDIT** button in the **Generic Server Scripts** field or by clicking on the **MONITORING SCRIPTS** link at the top of the page. If you choose the latter method, be sure to click the **ACCEPT** button before doing so or you will loss your changes.

The **EDIT MONITORING SCRIPTS** subsection allows the administrator to specify a send/expect string sequence to verify that the service for the virtual server is functional on each real server. It is also the place where the administrator can specify customized scripts to check services requiring dynamically changing data.

Figure 9–9 The EDIT MONITORING SCRIPTS Subsection

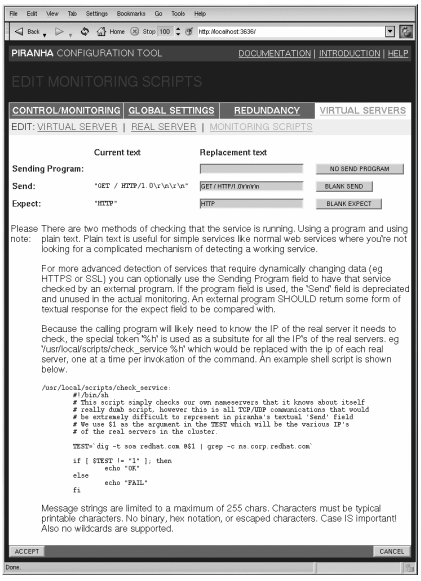

Send

Enter a string for the nanny daemon to send to each real server in this field. By default the send field is completed for HTTP. You can alter this value depending on your needs. If you leave this field blank, the nanny daemon attempts to open the port and assume the service is running if it succeeds.

Only one send sequence is allowed in this field, and it can only contain printable, ASCII characters as well as the following escape characters:

- \n for new line.
- \r for carriage return.
- \t for tab.
- \ to escape the next character which follows it.

Sending Program

For more advanced service verification, you can use this field to specify the path to a service-checking script. This functionality is especially helpful for services that require dynamically changing data, such as HTTPS or SSL.

To use this functionality, you must write a script that returns a textual response, set it to be executable, and type the path to it in the **Sending Program** field.

Tip

To ensure that each server in the real server pool is checked, use the special token **%h** after the path to the script in the **Sending Program** field. This token is replaced with each real server's IP address as the script is called by the nanny daemon.

Included in the **EDIT MONITORING SCRIPTS** subsection is a sample script you can use as a guide when composing an external service-checking script.

Note

If an external program is entered in the **Sending Program** field, then the **Send** field is ignored.

Expect

Enter a the textual response the server should return if it is functioning properly. If you wrote your own sending program, enter the response you told it to send if it was successful.

Tip

To determine what to send for a given service, you can open a `telnet` connection to the port on a real server and see what is returned. For instance, FTP reports 220 upon connecting, so could enter **`quit`** in the **Send** field and **`220`** in the **Expect** field.

WARNING

Remember to click the ACCEPT button after making any changes in this panel. To make sure you do not lose any changes when selecting a new panel.

Once you have configured virtual servers using the Piranha Configuration Tool, you must copy specific configuration files to the backup LVS router. See Section 9.7, *Synchronizing Configuration Files* for details.

9.7 Synchronizing Configuration Files

After configuring the primary LVS router, there are several configuration files that must be copied to the backup LVS router before you start the cluster.

These files include:

- `/etc/sysconfig/ha/lvs.cf` — the configuration file for the LVS routers.
- `/etc/sysctl` — the configuration file that, among other things, turns on packet forwarding in the kernel.
- `/etc/sysconfig/iptables` or `/etc/sysconfig/ipchains` — If you are using firewall marks, you should synchronize one of these files based on which network packet filter you are using.

Important

The `/etc/sysctl`, `/etc/sysconfig/iptables`, or `/etc/sysconfig/ipchains` files do *not* change when you configure the cluster using the Piranha Configuration Tool.

9.7.1 Synchronizing `lvs.cf`

Anytime the LVS configuration file, `/etc/sysconfig/ha/lvs.cf`, is created or updated, you must copy it to the backup LVS router node.

WARNING

Both the active and backup LVS router nodes must have identical `lvs.cf` files. Mismatched LVS configuration files between the LVS router nodes can prevent failover.

The best way to do this is to use the `scp` command.

Important

To use `scp` the `sshd` must be running on the backup router, see Section 7.1, *Configuring Services on the LVS Routers* for details on how to properly configure the necessary services on the LVS routers.

Issue the following command as the root user from the primary LVS router to sync the **`lvs.cf`** files between the router nodes:

```
scp /etc/sysconfig/ha/lvs.cf n.n.n.n:/etc/sysconfig/ha/lvs.cf
```

In the above command, replace *n.n.n.n* with the real IP address of the backup LVS router.

9.7.2 Synchronizing `sysctl`

The `sysctl` file is only modified once in most situations. This file is read at boot time and tells the kernel to turn on packet forwarding.

Important

If you are not sure whether or not packet forwarding is enabled in the kernel, see Section 7.5, *Turning on Packet Forwarding* for instructions on how to check and, if necessary, enable this key functionality.

9.7.3 Synchronizing Network Packet Filtering Rules

If you are using a network packet filter, such as `iptables` or `ipchains`, you will need to synchronize the appropriate configuration file on the backup LVS router.

If you alter the any network packet filter rules, enter the following command as root from the primary LVS router:

```
scp /etc/sysconfig/net_filter n.n.n.n:/etc/sysconfig/.
```

In the above command, replace *n.n.n.n* with the real IP address of the backup LVS router and *net_filter* with either `iptables` or `ipchains`.

Next either open an `ssh` session to the backup router or log into the machine as root and type the following command:

```
/sbin/service net_filter restart
```

In the above command, replace *net_filter* with either `iptables` or `ipchains`.

Once you have copied these files over to the backup router and started the appropriate services (see Section 7.1, *Configuring Services on the LVS Routers* for more on this topic) you are ready to start the cluster.

9.8 Starting the Cluster

To start the LVS cluster, it is best to have two root terminals open simultaneously or two simultaneous root open `ssh` sessions to the primary LVS router.

In one terminal, watch the kernel log messages with the command:

```
tail -f /var/log/messages
```

Then start the cluster by typing the following command into the other terminal:

```
/sbin/service pulse start
```

Follow the progress of the `pulse` service's startup in the terminal with the kernel log messages. When you see the following output, the pulse daemon has started properly:

```
gratuitous lvs arps finished
```

To stop watching `/var/log/messages`, type [Ctrl]-[c].

From this point on, the primary LVS router is also the active LVS router. While you can make requests to the cluster at this point, you should start the backup LVS router before putting the cluster into service. To do this, simply repeat the process described above on the backup LVS router node.

After completing this final step, the cluster will be up and running.

Part III Appendixes

A Additional Resources for LVS Clustering

The following are additional sources of information for LVS Clustering with Red Hat Linux Advanced Server.

Installed Documentation

- `/usr/share/doc/piranha-`*`version_number/`* — Installed with the Piranha package is a large amount of extra information pertaining to the care and feeding of an LVS cluster. The best place to start is to point a Web browser at the `index.html` file in this directory.

Helpful Websites

- http://www.redhat.com/ — The Red Hat website contains links to Red Hat Linux Advanced Server product information, documentation, and support.
- http://www.linuxvirtualserver.org/ — The community LVS project page and a good source for additional information and documentation. This site also links to many clustering-related pages for Linux.
- https://listman.redhat.com/ — To get help from the community and from developers, subscribe to the Piranha mailing list at this website.

Related Books

- *Red Hat Cluster Manager Installation and Administration Guide* — If you are creating a three tier cluster as mentioned in Section 6.2, *A Three Tiered LVS Configuration*, refer to this guide to configure a highly available shared data source.

B A Sample /etc/sysconfig/ha/lvs.cf File

The following is a sample `lvs.cf` file for a cluster using firewall mark 21 for FTP services and firewall mark 80 to bundle HTTP and HTTPS requests.

```
serial_no = 94
primary = 192.168.26.9
primary_private = 10.11.12.9
service = lvs
backup_active = 1
backup = 192.168.26.8
backup_private = 10.11.12.8
heartbeat = 1
heartbeat_port = 1050
keepalive = 6
deadtime = 18
network = nat
nat_router = 10.11.12.10 eth1:1
nat_nmask = 255.255.255.0
reservation_conflict_action = preempt
debug_level = NONE
virtual HTTP {
     active = 1
     address = 192.168.26.10 eth0:1
     vip_nmask = 255.255.255.0
     fwmark = 80
     port = 80
     persistent = 320
     send = "GET / HTTP/1.0\r\n\r\n"
     expect = "HTTP"
     load_monitor = none
     scheduler = wlc
     protocol = tcp
     timeout = 6
     reentry = 15
     quiesce_server = 1
     server one {
         address = 10.11.12.1
         active = 1
         weight = 1
     }
     server two {
```

```
          address = 10.11.12.2
          active = 1
          weight = 1
     }
     server three {
          address = 10.11.12.3
          active = 1
          weight = 1
     }
}
virtual FTP {
     active = 1
     address = 192.168.27.10 eth0:1
     vip_nmask = 255.255.255.0
     fwmark = 21
     port = 21
     persistent = 1
     send = "quit"
     expect = "220"
     load_monitor = none
     scheduler = rr
     protocol = tcp
     timeout = 6
     reentry = 15
     quiesce_server = 0
     server one {
          address = 10.11.12.1
          active = 1
          weight = 1
     }
     server two {
          address = 10.11.12.2
          active = 1
          weight = 1
     }
     server three {
          address = 10.11.12.3
          active = 1
          weight = 1
     }
}
virtual HTTPS {
     active = 1
     address = 192.168.26.10 eth0:1
     vip_nmask = 255.255.255.0
```

```
        fwmark = 80
        port = 443
        persistent = 320
        load_monitor = none
        scheduler = wlc
        protocol = tcp
        timeout = 6
        reentry = 15
        quiesce_server = 1
        server one {
            address = 10.11.12.1
            active = 1
            weight = 1
        }
        server two {
            address = 10.11.12.2
            active = 1
            weight = 1
        }
        server three {
            address = 10.11.12.3
            active = 1
            weight = 1
        }
}
```

C Removing Red Hat Linux

To uninstall Red Hat Linux Advanced Server from your system, you will need to remove the GRUB or LILO information from your master boot record (MBR).

In DOS, NT, and Windows 95 you can use fdisk to create a new MBR with the "undocumented" flag `/mbr`. This will ONLY rewrite the MBR to boot the primary DOS partition. The command should look like the following:

```
fdisk /mbr
```

If you need to remove Linux from a hard drive, and have attempted to do this with the default DOS fdisk, you will experience the "Partitions exist but they do not exist" problem. The best way to remove non-DOS partitions is with a tool that understands partitions other than DOS.

You can do this with the installation media by typing **`linux expert`** at the `boot:` prompt:

```
boot: linux expert
```

Select install (versus upgrade) and at the point when you should partition the drive, choose fdisk. In fdisk, type [p] to print out the partition numbers, and remove the Linux partitions with the [d] command. When you are satisfied with the changes you have made, you can quit with a [w] and the changes will be saved to disk. If you deleted too much, type [q] and no changes will be made.

Once you have removed the Linux partitions, you can reboot your computer using [Control]-[Alt]-[Delete] instead of continuing with the install.

D Getting Technical Support

D.1 Remember to Sign Up

If you have an official edition of Red Hat Linux Advanced Server 2.1 and/or an official Red Hat OEM partner kit, please remember to sign up for the benefits you are entitled to as a Red Hat customer.

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

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

To sign up, go to http://www.redhat.com/apps/activate/. You will find your **Product ID** on the Registration Information Card in your Official Red Hat Linux Advanced Server boxed set.

D.2 An Overview of Red Hat Support

Note

Refer to the service level agreement at http://www.redhat.com/support/sla/ for more information on how Red Hat's technical support staff can assist you.

Red Hat provides installation assistance for Official Red Hat Linux Advanced Server products and covers installation on a single computer. This assistance is intended to help customers successfully install Red Hat Linux Advanced Server. Assistance with installation is offered via telephone and the Web. Note, telephone support is only available with certain Red Hat Linux Advanced Server products. Please check your product to see what types of support are available to you.

Red Hat Support will attempt to answer any questions you may have before the installation process is initiated. Depending on the product purchased, it can include the following:

- Hardware compatibility questions
- Basic hard drive partitioning strategies

Red Hat Support can provide the following assistance during the installation process:

- Getting supported hardware recognized by the Red Hat Linux Advanced Server operating system

- Assistance with drive partitioning

We can also help you with basic post-installation tasks, such as:

- Successfully configuring the X Window System using Xconfigurator
- Configuring a local parallel port printer to print text
- Configuring a mouse

Our installation assistance service is designed to get Red Hat Linux Advanced Server running on your system as quickly and as easily as possible. However, there are many other things that you may want to do with your Red Hat Linux Advanced Server system, from compiling a custom kernel to setting up a Web server, which are not covered.

For assistance with these tasks, there is a wealth of online information available in the form of HOWTO documents, Linux-related websites, and commercial publications. The various Linux HOWTO documents are included with Red Hat Linux Advanced Server on the Documentation CD in the `/HOWTOS` directory. These HOWTOS are provided in text files that can easily be read from within Red Hat Linux Advanced Server and other operating systems.

A large number of Linux-related websites are available. The best starting point for finding information on Red Hat Linux Advanced Server is the Red Hat, Inc. website:

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

Many Linux-related books are available. If you are new to Linux, a book that covers Linux basics will be invaluable. We can recommend several titles: The *Official Red Hat Linux Getting Started Guide*; *Using Linux*, by Bill Ball; *Linux Clearly Explained*, by Bryan Pfaffenberger; *Linux for Dummies*, by Jon "maddog" Hall; and *Learning Red Hat Linux*, by Bill McCarty.

Certain Red Hat Linux products include additional support programs covering advanced configurations. Please see the Red Hat Support website for more information. The Red Hat technical support website is located at the following URL:

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

D.3 Scope of Red Hat Support

Red Hat, Inc. can only provide installation assistance to customers who have purchased an Official Red Hat Linux Advanced Server boxed set and/or an Official Red Hat OEM partner kit. If you have obtained Linux from any other company, you must contact that company for support. Other companies include:

- Macmillan
- Sams/Que

- Linux Systems Labs (LSL)
- Mandrake
- CheapBytes

Additionally, Red Hat Linux Advanced Server obtained via any of the following methods does not qualify for support from Red Hat:

- Red Hat Linux PowerTools Archive
- Downloaded via FTP on the Internet
- Included in a package such as Motif or Applixware
- Copied or installed from another user's CD
- A CD-ROM (or CD-ROM set) included in a Linux book or other publication.

D.4 How to Get Technical Support

In order to receive technical support for your Official Red Hat product, you must register your product on Red Hat's website.

Every Official Red Hat product comes with a Product Identification code: a 16-character alphanumeric string. The Product ID for Red Hat Linux Advanced Server 2.1 is located on the Registration Information Card that can be found inside the box. Your Product ID is included in your boxed set, and you should keep it in a safe place. You need this code, so do not lose the card!

Note

Do not throw away the card with your Product ID. You need the Product ID to get technical support. If you lose the certificate, you may not be able to receive support.

The Product ID is the code that will enable your technical support and any other benefits or services that you purchased from Red Hat, depending upon which Red Hat product you purchased.

D.4.1 Signing up for Technical Support

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

1. Create a customer profile at http://www.redhat.com/apps/activate/. You may have already completed this step; if you have, continue to the next step.

2. Using the login name and password you created during the customer profile, please log in at the Red Hat Support website at http://www.redhat.com/support.

   If you created a new customer profile, once you activate your product you will see a webpage that shows your registered products. There is also a button, **Access Web Support**, on this page that will take you to the support website.

3. Update your contact information if necessary.

   Note

   If your email address is not correct, communications regarding your technical support requests CANNOT be delivered to you, and you will not be able to retrieve your login and password by email. Be sure that you give us your correct email address.

   If you are worried about your privacy, please see Red Hat's privacy statement at http://www.redhat.com/legal/privacy_statement.html.

4. Add a product to your profile. Please enter the following information:
   - The Product ID for your boxed set product
   - The Support Certificate Number or Entitlement Number if the product is a contract
5. Set your customer preferences.
6. Answer the optional customer questionnaire.
7. Submit the form.

If the previous steps were completed successfully, you can now login at http://www.redhat.com/support and open a new technical service request. However, you must still use your Product ID in order to obtain technical support via telephone (if the product you purchased came with phone support). You will also be asked for your login name when contacting the technical support team via telephone.

D.5 Questions for Technical Support

Technical support is both a science and a mystical art form. In most cases, support technicians must rely on customer observations and communications with the customer in order to diagnose and solve the problem. Therefore, it is extremely important that you are as detailed and clear as possible when you state your questions and report your problems. Examples of what you should include are:

- Symptoms of the problem (for example: "Linux is not able to access my CD-ROM drive. When it tries, I get timeout errors.")

- When the problem began (for example: "My system was working fine until yesterday, when a lightning storm hit my area.")
- Any changes you made to your system (for example: "I added a new hard drive and used Partition Wizzo to add Linux partitions.")
- Other information that may be relevant to your situation, such as the installation method (CD-ROM, NFS, HTTP)
- Specific hardware devices that may be relevant to your problem (for example: If you cannot setup networking, what kind of network card do you have?)

Note

Refer to the service level agreement at http://www.redhat.com/support/sla/ for more information on how Red Hat's technical support staff can assist you.

D.5.1 How to Send Support Questions

Please login at http://www.redhat.com/support and open a new service request, or call the phone number for support. If your product came with phone support, or you have purchased a phone support contract, the phone number you will need to call will be provided to you during the sign up process.

For more information on using Red Hat's online support system go to http://www.redhat.com/support/services/access.html.

E Troubleshooting Your Installation of Red Hat Linux Advanced Server

This appendix discusses some common installation problems and their solutions.

E.1 You are Unable to Boot Red Hat Linux Advanced Server

E.1.1 Are You Unable to Boot from the CD-ROM?

If you cannot boot from your Red Hat Linux Advanced Server CD-ROM, you have two options:

1. You can change your BIOS so that booting from the CD-ROM is recognized first in the boot order, or
2. Boot using a boot disk you have created.

Note

There are a few cases where the system BIOS will not allow the Red Hat Linux Advanced Server CD-ROM to boot because of the size of the boot image on the CD-ROM itself. In cases such as these, a boot disk should be made to boot Red Hat Linux Advanced Server. Once booted, the CD-ROMs will work properly for the installation.

To change your BIOS, refer to your system manual for the correct keyboard combination that allows you to access your BIOS, or you can read the key sequence needed while the system begins to boot.

To create a boot disk, follow the instructions in Section 1.4.2, *Making Installation Diskettes*.

To boot Red Hat Linux Advanced Server using a boot disk, insert the diskette you have created into your floppy drive and then boot/reboot your computer. Make sure that your BIOS is set to use the floppy or removable disk (A:) to boot.

E.1.2 Are You Unable to Boot from the Local Boot Disk?

If you are experiencing difficulties in getting a local boot disk to boot your system correctly, you may need an updated boot disk.

Check the online errata (http://www.redhat.com/support/errata) for updated diskette images (if available) and follow the instructions in Section 1.4.2, *Making Installation Diskettes*, to make an updated boot disk for your system.

E.1.3 Is Your System Displaying Signal 11 Errors?

If you receive a fatal signal 11 during your installation, it is probably due to a hardware error in memory on your system's bus. A hardware error in memory can be caused by problems in executables or with the system's hardware. Like other operating systems, Red Hat Linux Advanced Server places its own demands on your system's hardware. Some of this hardware may not be able to meet those demands, even if they work properly under another OS.

Check to see if you have the latest installation and supplemental boot diskettes from Red Hat. Review the online errata to see if newer versions are available. If the latest images still fail, it may be due to a problem with your hardware. Commonly, these errors are in your memory or CPU-cache. A possible solution for this error is turning off the CPU-cache in the BIOS. You could also try to swap your memory around in the motherboard slots to see if the problem is either slot or memory related.

For more information concerning signal 11 errors, refer to http://www.bitwizard.nl/sig11/.

E.1.4 Are You Unable to Boot from a Network Boot Disk?

If you are experiencing difficulties in getting the network boot disk you made to boot your system correctly, you may need an updated boot disk.

Check the online errata for updated diskette images (if available) and follow the instructions provided to make an updated boot disk for your system.

E.1.5 Are You Unable to Boot With Your RAID Card?

If you have performed an installation and cannot boot your system properly, you may need to reinstall and create your partitions differently.

Some BIOSes do not support booting from RAID cards. At the end of an installation, a text-based screen showing the boot loader prompt (for example, `GRUB:`) and a flashing cursor may be all that appears. If this is the case, you will need to repartition your system.

Whether you choose automatic or manual partitioning, you will need to install your `/boot` partition outside of the RAID array, such as on a separate hard drive. An internal hard drive is necessary to use for partition creation with problematic RAID cards.

You must also install your perferred boot loader (GRUB or LILO) outside of the RAID array — not on the MBR. The boot loader should be installed on the MBR of the same drive as the `/boot` partition was created.

Once these changes have been made, you should be able to finish your installation and boot the system properly.

E.2 Trouble Beginning the Installation

E.2.1 Is Your Mouse Not Detected?

If the **Mouse Not Detected** screen (see Figure E–1, *Mouse Not Detected*) appears, then the installation program was not able to identify your mouse correctly.

You can choose to continue with the GUI installation or use the text mode installation, which does not require using a mouse. If you choose to continue with the GUI installation, you will need to provide the installation program with your mouse configuration information (see Figure 3–3, *Mouse Configuration*).

Figure E–1 Mouse Not Detected

For an overview of text mode installation instructions, please refer to Chapter 4, *Installing Red Hat Linux Advanced Server via Text Mode*.

E.2.2 Problems with Booting into the Graphical Installation

The Red Hat Linux Advanced Server installation program uses frame buffers by default. However, there are some video cards that will not work with this setting. The end result will be a problem booting into the graphical installation program.

The installation program will first try to run in frame buffer mode. If that fails, it will try to run in a lower resolution mode. If that still fails, the installation program will run in text mode.

Users who have video cards that will not run at 800 x 600 resolution should type **lowres** at the boot: prompt to run the installation program in 640 x 480 resolution.

If this still does not work, you can run the installation program without frame buffers by typing **nofb** at the boot: prompt.

E.3 Trouble During the Installation

E.3.1 Partition Creation Problems

If you are having trouble creating a partition (for example, a root (/) partition), make sure you are setting its partition type to Linux Native.

Unless your BIOS supports otherwise, make sure /boot does not exceed the 1023 cylinder head. If you do not, the installation program will not allow you to create a /boot or / partition. Some new systems allow you to exceed the 1023 limit (with GRUB and the newer LILO versions that are available), but most machines with older BIOS will not.

E.3.2 Using Remaining Space

You have a swap and a / (root) partition created, and you have selected the root partition to use the remaining space, but it does not fill the hard drive.

If your hard drive is more than 1024 cylinders, you must create a /boot partition if you want the / (root) partition to use all of the remaining space on your hard drive.

E.3.3 Other Partitioning Problems

If you are using Disk Druid to create partitions, but cannot move to the next screen, you probably have not created all the partitions necessary for Disk Druid's dependencies to be satisfied.

You must have the following partitions as a bare minimum:

- A /boot partition of type Linux native
- A / (root) partition of type Linux native
- A <swap> partition of type Linux swap

Tip

When defining a partition's type as Linux swap, you do not have to assign it a mount point. Disk Druid automatically assigns the mount point for you.

E.3.4 Are You Seeing Python Errors?

During some installations of Red Hat Linux Advanced Server, the installation program (also known as Anaconda) may fail with a Python or traceback error. This error may occur after the selection of individual packages or while trying to save the upgrade log in `/tmp`. The error may look similar to:

```
Traceback (innermost last):
   File "/var/tmp/anaconda-7.1//usr/lib/anaconda/iw/progress_gui.py",
line 20, in run
     rc = self.todo.doInstall ()
   File "/var/tmp/anaconda-7.1//usr/lib/anaconda/todo.py", line 1468, in
doInstall
     self.fstab.savePartitions ()
   File "fstab.py", line 221, in savePartitions
     sys.exit(0)
 SystemExit: 0

 Local variables in innermost frame:
 self: <fstab.GuiFstab instance at 8446fe0>
 sys: <module 'sys' (built-in)>

 ToDo object:
 (itodo
 ToDo
 p1
 (dp2
 S'method'
 p3
 (iimage
 CdromInstallMethod
 p4
 (dp5
 S'progressWindow'
 p6

 <failed>
```

This error occurs in some systems where links to `/tmp` are symbolic to other locations or have been changed since creation. These symbolic or changed links are invalid during the installation process, so the installation program cannot write information and fails.

If you experience such an error, first try to download any available errata for Anaconda. Errata can be found at http://www.redhat.com/support/errata.

You can also search for bug reports related to this problem. To search Red Hat's bug tracking system, go to http://bugzilla.redhat.com/bugzilla.

Finally, if you are still facing problems related to this error, register your product and contact our support team. To register your product, go to http://www.redhat.com/apps/activate.

E.4 Problems After Installation

E.4.1 Trouble With the Graphical GRUB Screen?

If, for some reason, you need to disable the graphical boot screen, you can do so, as root, by editing the `/boot/grub/grub.conf` file and then rebooting your system.

To do this, comment out the line which begins with `splashimage` in the `grub.conf` file. To comment out a line, insert the `;` character at the beginning of the line.

Once you reboot, the `grub.conf` file will be reread and your changes will take place.

You may re-enable the graphical boot screen by uncommenting (or adding) the above line back into the `grub.conf` file and rebooting.

E.4.2 Trouble With the Graphical LILO Screen?

If, for some reason, you need to disable the graphical boot screen, you can do so, as root, by editing the `/etc/lilo.conf` file and then rerunning LILO.

First, as root, comment out (or delete) the line which reads `message=/boot/message` in the `/etc/lilo.conf` file. To comment out a line, insert the `;` character at the beginning of the line. Next, rerun LILO by typing `/sbin/lilo -v`. The next time you boot, you will see the text `LILO:` prompt, as used in previous Red Hat Linux Advanced Server releases.

You may re-enable the graphical boot screen by adding the above line back into the `lilo.conf` file and rerunning LILO.

E.4.3 Problems with Server Installations and X

If you performed a server installation and you are having trouble getting X to start, you may not have installed the X Window System during your installation.

If you want the X Window System, you can perform an upgrade to install X. During the upgrade, select the X Window System packages, and choose GNOME, KDE, or both.

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

E.4.4 Problems When You Try to Log In

If you did not create a user account during the installation you will need to log in as root and use the password you assigned to root.

If you cannot remember your root password, you will need to boot your system as `linux single` at the LILO `boot:` or GRUB prompt. Then at the # prompt, you will need to type `passwd root`, which will allow you to enter a new password for root. At this point you can type `shutdown -r now` and the system will reboot with your new password.

If you cannot remember your user account password, you must become root. To become root, type `su -` and enter your root password when prompted. Then, type `passwd <username>`. This allows you to enter a new password for the specified user account.

If you selected either the custom or workstation installation and do not see the graphical login screen, check your hardware for compatibility issues. The *Hardware Compatibility List* can be found at http://hardware.redhat.com/hcl/.

E.4.5 Does Netscape Navigator Crash on JavaScript Pages?

If Netscape Navigator continuously crashes on pages that contain JavaScripts, you may need to edit your `~/.mailcap` file.

Edit the file using pico by typing `pico ~/.mailcap` at the prompt in a terminal window. (You may use any text editor.)

Remove the following lines from the file:

```
application/x-javascript;;\
 x-mozilla-flags=save
```

You can also turn off JavaScript within Netscape Navigator itself. Click on **Edit**=>**Preferences**=>**Advanced** and make sure the **Enable JavaScript** checkbox is not selected.

If these do not help, you can also try to use a newer version of Netscape Navigator if available. Check the Red Hat Linux errata website under security advisories for more information.

E.4.6 Your Printer Will Not Work Under X

If you are not sure how to set up your printer or are having trouble getting it to work properly, try using the graphical printconf program. Log in as root, open a terminal window, and type `printconf-gui`.

E.4.7 Is Your RAM Not Being Recognized?

Sometimes, the kernel does not recognize all of your memory (RAM). You can check this with the following command:

```
cat /proc/meminfo
```

Find out if the displayed quantity is the same as the known amount of RAM in your system. If they are not equal, add the following line to the `/boot/grub/grub.conf` or `/etc/lilo.conf` file, depending on the boot loader you installed:

`append="mem=`*`xx`*`M"`

Replace *xx* with the amount of RAM you have in megabytes. Remember that per-image append lines completely overwrite the global append line. It might be worth adding this to the per-image descriptions, as shown in this example:

`append="mem=128M"`

In `/boot/grub/grub.conf`, the above example would look similar to the following:

```
#NOTICE: You have a /boot partition. This means that
#        all kernel paths are relative to /boot/
default=0
timeout=30
splashimage=(hd0,0)/grub/splash.xpm.gz
title Red Hat Linux (2.4.6-2)
        root (hd0,0)
        kernel /vmlinuz-2.2.6-2 ro root=/dev/hda3
        append="mem=128M"
```

Once you reboot, the changes made to `grub.conf` will be reflected on your system.

In `/etc/lilo.conf`, the above example would look similar to the following:

```
boot=/dev/sda
      map=/boot/map
      install=/boot/boot.b
      prompt
      timeout=50

      image=/boot/vmlinuz-2.2.12-20
```

```
label=linux
root=/dev/sda1
initrd=/boot/initrd-2.2.12-20.img
read-only
append="mem=128M"
```

Remember to run `/sbin/lilo -v` after changing `/etc/lilo.conf`.

Note that you can also produce the same effect by actually passing this option when you are specifying the label/image to use in GRUB or LILO. For example if you had a label named *linux*, at the boot loader menu you could type:

`linux mem=`*`xx`*`M`

Remember to replace *xx* with the amount of RAM in your system.

E.4.8 Problems with Sound Configuration

If you do not have sound after your installation, you may need to run the sound configuration utility. As root, type `sndconfig` in a terminal window.

Note

`sndconfig` must be run in runlevel 3. More information about runlevels can be found in the *Official Red Hat Linux Reference Guide*, in chapter *Boot Process, Init, and Shutdown*.

If the sndconfig application does not help, you may need to select the `Enable sound server startup` option under the **Multimedia**=>**Sound** in the GNOME Control Center.

To do this, click on **Panel**=>**Programs**=>**Settings**=>**GNOME Control Center** to launch the GNOME Control Center. In the GNOME Control Center, select the **Sound** submenu of the **Multimedia** menu. On the right, a **General** sound menu will appear. Select `Enable sound server startup` and then click **OK**.

F An Introduction to Disk Partitions

Disk partitions are a standard part of the personal computer landscape and have been for quite some time. However, with many people purchasing computers featuring preinstalled operating systems, relatively few people understand how partitions work. This chapter attempts to explain the reasons for and use of disk partitions so your Red Hat Linux Advanced Server installation will be as simple and painless as possible.

If you are reasonably comfortable with disk partitions, you could skip ahead to Section F.1.4, *Making Room For Red Hat Linux Advanced Server*, for more information on the process of freeing up disk space to prepare for a Red Hat Linux Advanced Server installation. This section also discusses the partition naming scheme used by Linux systems, sharing disk space with other operating systems, and related topics.

F.1 Hard Disk Basic Concepts

Hard disks perform a very simple function — they store data and reliably retrieve it on command.

When discussing issues such as disk partitioning, it is important to know a bit about the underlying hardware. Unfortunately, it is easy to become bogged down in details. Therefore, we will use a simplified diagram of a disk drive to help explain what is really happening when a disk drive is partitioned. Figure F–1, *An Unused Disk Drive*, shows a brand-new, unused disk drive.

Figure F–1 An Unused Disk Drive

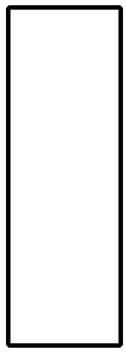

Not much to look at, is it? But if we are talking about disk drives on a basic level, it will do. Say that we would like to store some data on this drive. As things stand now, it will not work. There is something we need to do first…

F.1.1 It is Not What You Write, it is How You Write It

Experienced computer users probably got this one on the first try. We need to **format** the drive. Formatting (usually known as "making a **filesystem**") writes information to the drive, creating order out of the empty space in an unformatted drive.

Figure F–2 Disk Drive with a Filesystem

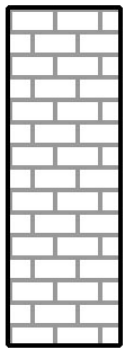

As Figure F–2, *Disk Drive with a Filesystem*, implies, the order imposed by a filesystem involves some trade-offs:

- A small percentage of the drive's available space is used to store filesystem-related data and can be considered as overhead.
- A filesystem splits the remaining space into small, consistently-sized segments. For Linux, these segments are known as **blocks**. [1]

Given that filesystems make things like directories and files possible, these tradeoffs are usually seen as a small price to pay.

It is also worth noting that there is no single, universal filesystem. As Figure F–3, *Disk Drive with a Different Filesystem*, shows, a disk drive may have one of many different filesystems written on it. As you might guess, different filesystems tend to be incompatible; that is, an operating system that supports one filesystem (or a handful of related filesystem types) may not support another. This last statement is not a hard-and-fast rule, however. For example, Red Hat Linux Advanced Server supports

[1] Blocks really *are* consistently sized, unlike our illustrations. Keep in mind, also, that an average disk drive contains thousands of blocks. But for the purposes of this discussion, please ignore these minor discrepancies.

a wide variety of filesystems (including many commonly used by other operating systems), making data interchange between different filesystems easy.

Figure F–3 Disk Drive with a Different Filesystem

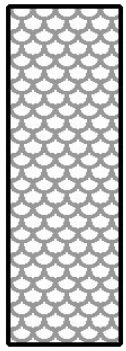

Of course, writing a filesystem to disk is only the beginning. The goal of this process is to actually *store* and *retrieve* data. Let us take a look at our drive after some files have been written to it.

Figure F–4 Disk Drive with Data Written to It

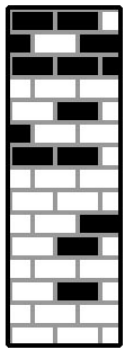

As Figure F–4, *Disk Drive with Data Written to It*, shows, 14 of the previously-empty blocks are now holding data. However, by simply looking at this picture, we cannot determine exactly how many files reside on this drive. There may be as few as one or as many as 14 files, as all files use at least one block and some files use multiple blocks. Another important point to note is that the used blocks do not have to form a contiguous region; used and unused blocks may be interspersed. This is known as **fragmentation**. Fragmentation can play a part when attempting to resize an existing partition.

As with most computer-related technologies, disk drives changed over time after their introduction. In particular, they got bigger. Not larger in physical size, but bigger in their capacity to store information. And, this additional capacity drove a fundamental change in the way disk drives were used.

F.1.2 Partitions: Turning One Drive Into Many

As disk drive capacities soared, some people began to wonder if having all of that formatted space in one big chunk was such a great idea. This line of thinking was driven by several issues, some philosophical, some technical. On the philosophical side, above a certain size, it seemed that the additional space provided by a larger drive created more clutter. On the technical side, some filesystems were never designed to support anything above a certain capacity. Or the filesystems *could* support larger

drives with a greater capacity, but the overhead imposed by the filesystem to track files became excessive.

The solution to this problem was to divide disks into **partitions**. Each partition can be accessed as if it was a separate disk. This is done through the addition of a **partition table**.

Note

While the diagrams in this chapter show the partition table as being separate from the actual disk drive, this is not entirely accurate. In reality, the partition table is stored at the very start of the disk, before any filesystem or user data. But for clarity, we will keep it separate in our diagrams.

Figure F–5 Disk Drive with Partition Table

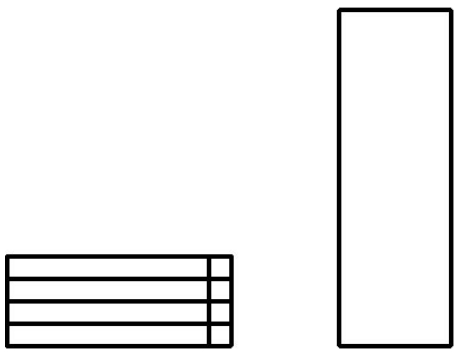

As Figure F–5, *Disk Drive with Partition Table*, shows, the partition table is divided into four sections. Each section can hold the information necessary to define a single partition, meaning that the partition table can define no more than four partitions.

Each partition table entry contains several important characteristics of the partition:

- The points on the disk where the partition starts and ends

- Whether the partition is "active"
- The partition's type

Let us take a closer look at each of these characteristics. The starting and ending points actually define the partition's size and location on the disk. The "active" flag is used by some operating systems' boot loaders. In other words, the operating system in the partition that is marked "active" will be booted.

The partition's type can be a bit confusing. The type is a number that identifies the partition's anticipated usage. If that statement sounds a bit vague, that is because the meaning of the partition type is a bit vague. Some operating systems use the partition type to denote a specific filesystem type, to flag the partition as being associated with a particular operating system, to indicate that the partition contains a bootable operating system, or some combination of the three.

Table F–1, *Partition Types*, contains a listing of some popular (and obscure) partition types, along with their numeric values.

Table F–1 Partition Types

Partition Type	Value	Partition Type	Value
Empty	00	Novell Netware 386	65
DOS 12-bit FAT	01	PIC/IX	75
XENIX root	02	Old MINIX	80
XENIX usr	03	Linux/MINUX	81
DOS 16-bit <=32M	04	Linux swap	82
Extended	05	Linux native	83
DOS 16-bit >=32	06	Linux extended	85
OS/2 HPFS	07	Amoeba	93
AIX	08	Amoeba BBT	94
AIX bootable	09	BSD/386	a5
OS/2 Boot Manager	0a	OpenBSD	a6
Win95 FAT32	0b	NEXTSTEP	a7
Win95 FAT32 (LBA)	0c	BSDI fs	b7
Win95 FAT16 (LBA)	0e	BSDI swap	b8
Win95 Extended (LBA)	0f	Syrinx	c7

Partition Type	Value	Partition Type	Value
Venix 80286	40	CP/M	db
Novell	51	DOS access	e1
Microport	52	DOS R/O	e3
GNU HURD	63	DOS secondary	f2
Novell Netware 286	64	BBT	ff

By this point, you might be wondering how all this additional complexity is normally used. See Figure F–6, *Disk Drive With Single Partition*, for an example.

Figure F–6 Disk Drive With Single Partition

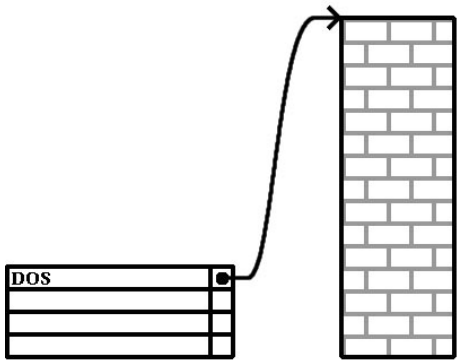

In many cases, there is only a single partition spanning the entire disk, essentially duplicating the method used before partitions. The partition table has only one entry used, and it points to the start of the partition.

We have labeled this partition as being of the "DOS" type. Although it is only one of several possible partition types listed in Table F–1, *Partition Types*, it is adequate for the purposes of this discussion.

This is a typical partition layout for most newly purchased computers with a consumer version of Microsoft Windows™ preinstalled.

F.1.3 Partitions within Partitions — An Overview of Extended Partitions

Of course, over time it became obvious that four partitions would not be enough. As disk drives continued to grow, it became more and more likely that a person could configure four reasonably-sized partitions and still have disk space left over. There needed to be some way of creating more partitions.

Enter the extended partition. As you may have noticed in Table F–1, *Partition Types*, there is an "Extended" partition type. It is this partition type that is at the heart of extended partitions.

When a partition is created and its type is set to "Extended," an extended partition table is created. In essence, the extended partition is like a disk drive in its own right — it has a partition table that points to one or more partitions (now called **logical partitions**, as opposed to the four **primary partitions**) contained entirely within the extended partition itself. Figure F–7, *Disk Drive With Extended Partition*, shows a disk drive with one primary partition and one extended partition containing two logical partitions (along with some unpartitioned free space).

Figure F–7 Disk Drive With Extended Partition

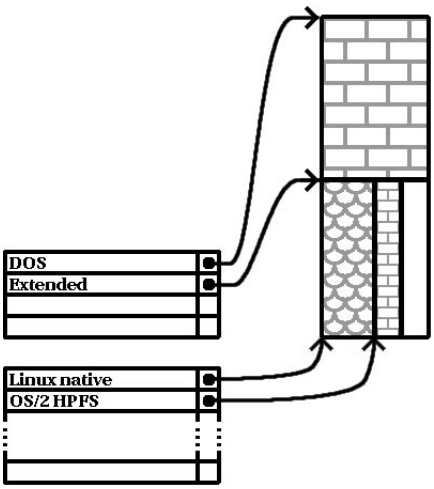

As this figure implies, there is a difference between primary and logical partitions — there can only be four primary partitions, but there is no fixed limit to the number of logical partitions that can exist. (However, in reality, it is probably not a good idea to try to define and use more than 12 logical partitions on a single disk drive.)

Now that we have discussed partitions in general, let us see how to use this knowledge to install Red Hat Linux Advanced Server.

F.1.4 Making Room For Red Hat Linux Advanced Server

There are three possible scenarios you may face when attempting to repartition your hard disk:

- Unpartitioned free space is available
- An unused partition is available
- Free space in an actively used partition is available

Let us look at each scenario in order.

Note

Please keep in mind that the following illustrations are simplified in the interest of clarity and do not reflect the exact partition layout that you will encounter when actually installing Red Hat Linux Advanced Server.

Using Unpartitioned Free Space

In this situation, the partitions already defined do not span the entire hard disk, leaving unallocated space that is not part of any defined partition. Figure F–8, *Disk Drive with Unpartitioned Free Space* shows what this might look like.

Figure F–8 Disk Drive with Unpartitioned Free Space

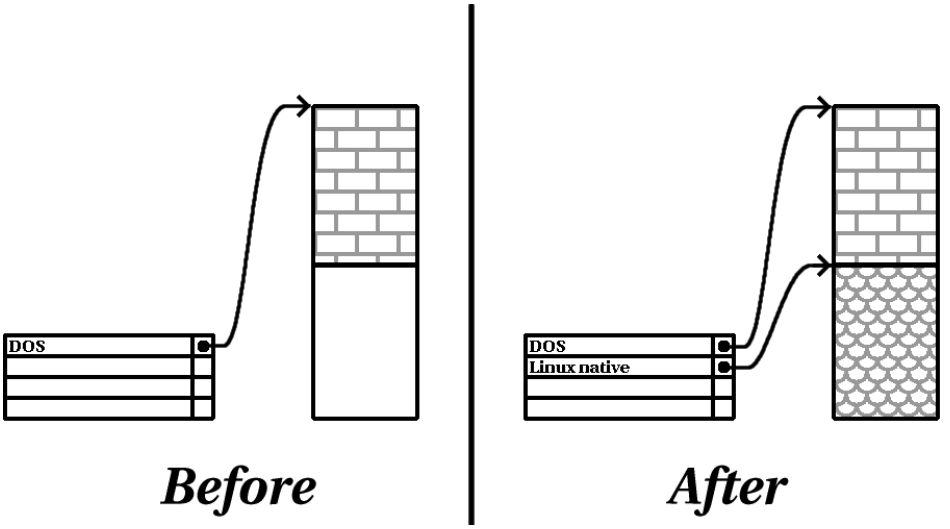

If you think about it, an unused hard disk also falls into this category. The only difference is that *all* the space is not part of any defined partition.

In any case, you can simply create the necessary partitions from the unused space. Unfortunately, this scenario, although very simple, is not very likely (unless you have just purchased a new disk just for Red Hat Linux Advanced Server). Most pre-installed operating systems are configured to take up all available space on a disk drive (see *Using Free Space from an Active Partition* in Section F.1.4).

Next, we will discuss a slightly more common situation.

Using Space from an Unused Partition

In this case, maybe you have one or more partitions that you do not use any longer. Perhaps you have dabbled with another operating system in the past, and the partition(s) you dedicated to it never seem to be used anymore. Figure F–9, *Disk Drive With an Unused Partition* , illustrates such a situation.

Figure F–9 Disk Drive With an Unused Partition

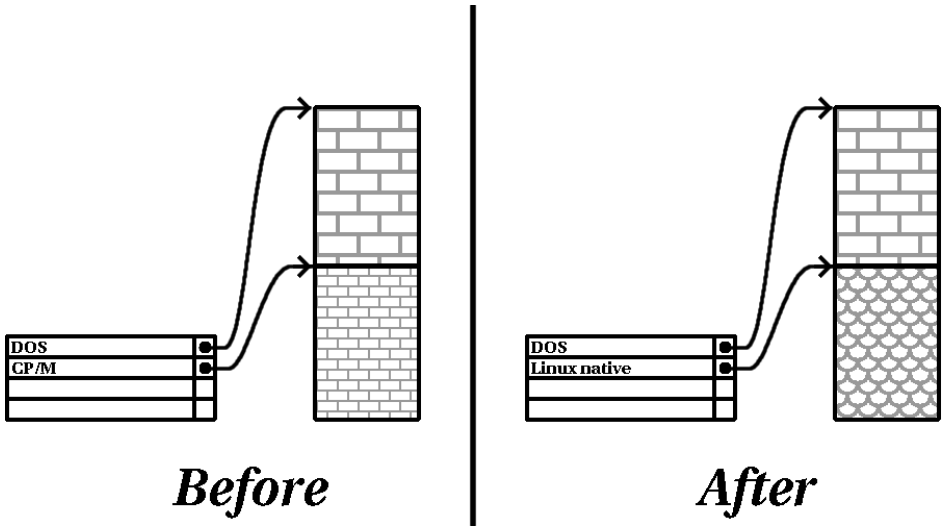

If you find yourself in this situation, you can use the space allocated to the unused partition. You will first need to delete the partition, and then create the appropriate Linux partition(s) in its place. You can either delete the partition using the DOS `fdisk` command, or you will be given the opportunity to do so during a custom installation.

Using Free Space from an Active Partition

This is the most common situation. It is also, unfortunately, the hardest to handle. The main problem is that, even if you have enough free space, it is presently allocated to a partition that is already in use. If you purchased a computer with pre-installed software, the hard disk most likely has one massive partition holding the operating system and data.

Aside from adding a new hard drive to your system, you have two choices:

Destructive Repartitioning

> Basically, you delete the single large partition and create several smaller ones. As you might imagine, any data you had in the original partition is destroyed. This means that making a complete backup is necessary. For your own sake, make two backups, use verification (if available in your backup software), and try to read data from your backup *before* you delete the partition.

CAUTION

If there was an operating system of some type installed on that partition, it will need to be reinstalled as well. Be aware that some computers sold with pre-installed operating systems may not include the CD-ROM media to reinstall the original operating system. The best time to notice if this applies to your system is *before* you destroy your original partition and its operating system installation.

After creating a smaller partition for your existing software, you can reinstall any software, restore your data, and continue your Red Hat Linux Advanced Server installation. Figure F–10, *Disk Drive Being Destructively Repartitioned* shows this being done.

Figure F–10 Disk Drive Being Destructively Repartitioned

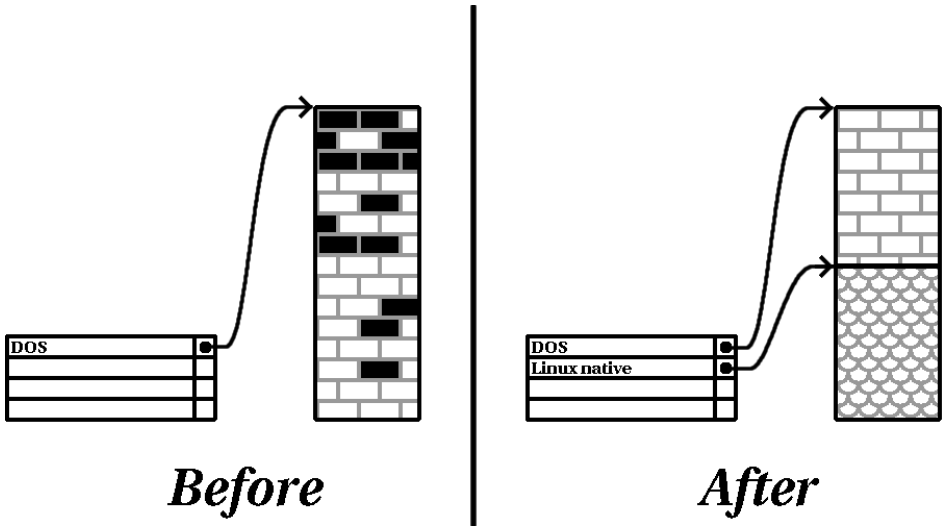

As Figure F–10, *Disk Drive Being Destructively Repartitioned*, shows, any data present in the original partition will be lost without proper backup!

Non-Destructive Repartitioning

Here, you run a program that does the seemingly impossible: it makes a big partition smaller without losing any of the files stored in that partition. Many people have found this method to be reliable and trouble-free. What software should you use to perform this feat? There are several disk management software products on the market. You will have to do some research to find the one that is best for your situation.

While the process of non-destructive repartitioning is rather straightforward, there are a number of steps involved:

- Compress existing data
- Resize the existing partition
- Create new partition(s)

Next we will look at each step in a bit more detail.

Compress existing data

As Figure F–11, *Disk Drive Being Compressed*, shows, the first step is to compress the data in your existing partition. The reason for doing this is to rearrange the data such that it maximizes the available free space at the "end" of the partition.

Figure F–11 Disk Drive Being Compressed

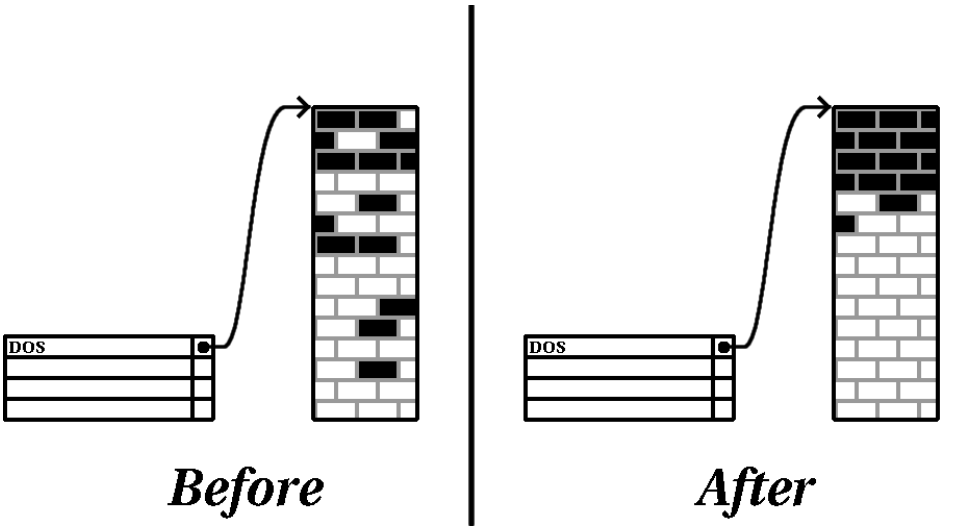

This step is crucial. Without it, the location of your data could prevent the partition from being resized to the extent desired. Note also that, for one reason or another, some data cannot be moved. If this is the case (and it severely restricts the size of your new partition(s)), you may be forced to destructively repartition your disk.

Resize the existing partition

Figure F–12, *Disk Drive with Partition Resized*, shows the actual resizing process. While the actual result of the resizing operation varies depending on the software used, in most cases the newly freed space is used to create an unformatted partition of the same type as the original partition.

Figure F–12 Disk Drive with Partition Resized

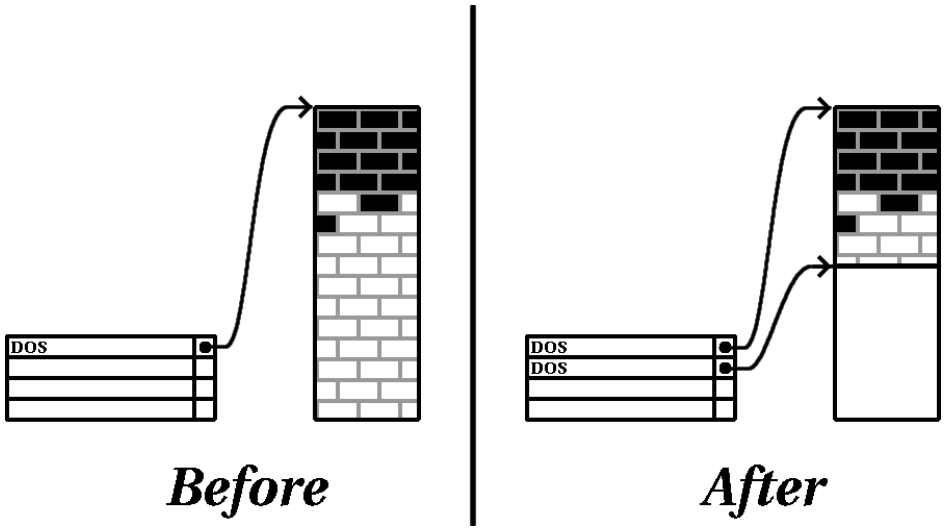

It is important to understand what the resizing software you use does with the newly freed space, so that you can take the appropriate steps. In the case we have illustrated, it would be best to simply delete the new DOS partition and create the appropriate Linux partition(s).

Create new partition(s)

As the previous step implied, it may or may not be necessary to create new partitions. However, unless your resizing software is Linux-aware, it is likely you will need to delete the partition that was created during the resizing process. Figure F–13, *Disk Drive with Final Partition Configuration*, shows this being done.

Figure F–13 Disk Drive with Final Partition Configuration

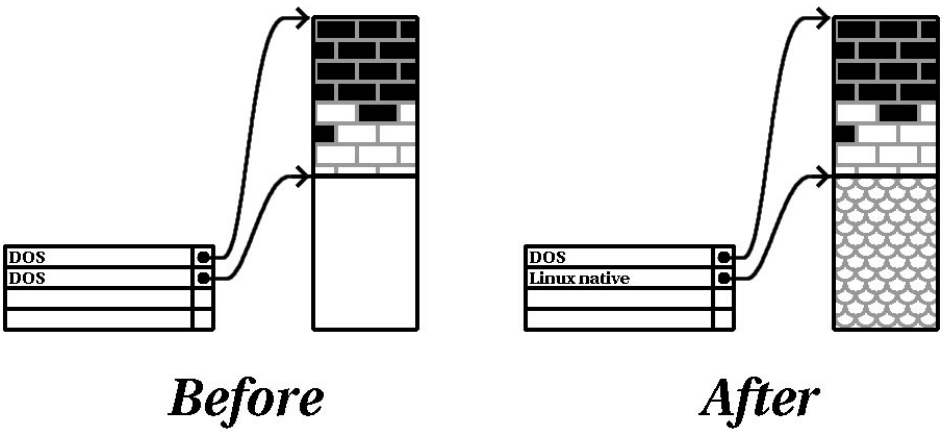

Note

The following information is specific to Intel-based computers only.

As a convenience to Red Hat Linux Advanced Server users, the DOS `fips` utility is included on the Red Hat Linux/x86 CD 1 in the `dosutils` directory. This is a freely available program that can resize FAT (File Allocation Table) partitions.

WARNING

Many people have successfully used `fips` to resize their hard drive partitions. However, because of the nature of the operations carried out by `fips` and the wide variety of hardware and software configurations under which it must run, Red Hat cannot guarantee that `fips` will work properly on your system. Therefore, no installation support is available for `fips`. *Use it at your own risk.*

That said, if you decide to repartition your hard drive with `fips`, it is *vital* that you do two things:

- *Perform a backup* — Make two copies of all the important data on your computer. These copies should be to removable media (such as tape or diskettes), and you should make sure they are readable before proceeding.
- *Read the documentation* — Completely read the `fips` documentation, located in the `dosutils/fipsdocs` subdirectory on Red Hat Linux/x86 CD 1.

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

F.1.5 Partition Naming Scheme

Linux refers to disk partitions using a combination of letters and numbers which may be confusing, particularly if you are used to the "C drive" way of referring to hard disks and their partitions. In the DOS/Windows world, partitions are named using the following method:

- Each partition's type is checked to determine if it can be read by DOS/Windows.
- If the partition's type is compatible, it is assigned a "drive letter." The drive letters start with a "C" and move on to the following letters, depending on the number of partitions to be labeled.
- The drive letter can then be used to refer to that partition as well as the filesystem contained on that partition.

Red Hat Linux Advanced Server uses a naming scheme that is more flexible and conveys more information than the approach used by other operating systems. The naming scheme is file-based, with filenames in the form:

`/dev/xxyN`

Here is how to decipher the partition naming scheme:

`/dev/`

This string is the name of the directory in which all device files reside. Since partitions reside on hard disks, and hard disks are devices, the files representing all possible partitions reside in `/dev/`.

`xx`

The first two letters of the partition name indicate the type of device on which the partition resides. You will normally see either `hd` (for IDE disks) or `sd` (for SCSI disks).

y

This letter indicates which device the partition is on. For example, `/dev/hda` (the first IDE hard disk) or `/dev/sdb` (the second SCSI disk).

N

The final number denotes the partition. The first four (primary or extended) partitions are numbered `1` through `4`. Logical partitions start at `5`. So, for example, `/dev/hda3` is the third primary or extended partition on the first IDE hard disk, and `/dev/sdb6` is the second logical partition on the second SCSI hard disk.

Note

There is no part of this naming convention that is based on partition type; unlike DOS/Windows, *all* partitions can be identified under Red Hat Linux Advanced Server. Of course, this does not mean that Red Hat Linux Advanced Server can access data on every type of partition, but in many cases it is possible to access data on a partition dedicated to another operating system.

Keep this information in mind; it will make things easier to understand when you are setting up the partitions Red Hat Linux Advanced Server requires.

F.1.6 Disk Partitions and Other Operating Systems

If your Red Hat Linux Advanced Server partitions will be sharing a hard disk with partitions used by other operating systems, most of the time you will have no problems. However, there are certain combinations of Linux and other operating systems that require extra care. Information on creating disk partitions compatible with other operating systems is available in several HOWTOs and Mini-HOWTOs, available on the Red Hat Linux Advanced Server Documentation CD in the `HOWTO` and `HOWTO/mini` directories. In particular, the Mini-HOWTOs whose names start with `Linux+` are quite helpful.

F.1.7 Disk Partitions and Mount Points

One area that many people new to Linux find confusing is the matter of how partitions are used and accessed by the Linux operating system. In DOS/Windows, it is relatively simple: Each partition gets a "drive letter." You then use the correct drive letter to refer to files and directories on its corresponding partition.

This is entirely different from how Linux deals with partitions and, for that matter, with disk storage in general. The main difference is that each partition is used to form part of the storage necessary to

support a single set of files and directories. This is done by associating a partition with a directory through a process known as **mounting**. Mounting a partition makes its storage available starting at the specified directory (known as a **mount point**).

For example, if partition `/dev/hda5` were mounted on `/usr`, that would mean that all files and directories under `/usr` would physically reside on `/dev/hda5`. So the file `/usr/share/doc/FAQ/txt/Linux-FAQ` would be stored on `/dev/hda5`, while the file `/etc/X11/gdm/Sessions/Gnome` would not.

Continuing our example, it is also possible that one or more directories below `/usr` would be mount points for other partitions. For instance, a partition (say, `/dev/hda7`) could be mounted on `/usr/local`, meaning that `/usr/local/man/whatis` would then reside on `/dev/hda7` rather than `/dev/hda5`.

F.1.8 How Many Partitions?

At this point in the process of preparing to install Red Hat Linux Advanced Server, you will need to give some consideration to the number and size of the partitions to be used by your new operating system. The question of "how many partitions" continues to spark debate within the Linux community and, without any end to the debate in sight, it is safe to say that there are probably as many partition layouts as there are people debating the issue.

Keeping this in mind, we recommend that, unless you have a reason for doing otherwise, you should at least create the following partitions:

- *A swap partition* — Swap partitions are used to support virtual memory. In other words, data is written to swap when there is not RAM to hold the data your system is processing. You *must* create a swap partition to correctly use Red Hat Linux Advanced Server. The minimum size of your swap partition should be equal to twice the amount of your computer's RAM or 32 MB, whichever is larger.
- *A `/boot` partition* — The partition mounted on `/boot` contains the operating system kernel (which allows your system to boot Red Hat Linux Advanced Server), along with a few other files used during the bootstrap process.

Make sure you read Section F.1.9, *One Last Wrinkle: Using GRUB or LILO* — the information there applies to the `/boot` partition!

Due to the limitations of most PC BIOSes, creating a small partition to hold these files is a good idea. For most users, a 32 MB boot partition is sufficient.

- *A root partition (/)* — The root partition is where / (the root directory) resides. In this partitioning layout, all files (except those stored in /boot) reside on the root partition. Because of this, it is in your best interest to maximize the size of your root partition. A 1.0 GB root partition will permit the equivalent of an Advanced Server installation (with *very* little free space), while a 3.2 GB root partition will let you install every package. Obviously, the more space you can give the root partition, the better.

Specific recommendations concerning the proper size for various Red Hat Linux Advanced Server partitions can be found in Section 1.5, *Which Installation Class is Best For You?*.

F.1.9 One Last Wrinkle: Using GRUB or LILO

GRUB and LILO are the most commonly used methods to boot Red Hat Linux Advanced Server on Intel-based systems. As operating system loaders, they operate "outside" of any operating system, using only the Basic I/O System (or BIOS) built into the computer hardware itself. This section describes GRUB and LILO's interactions with PC BIOSes and is specific to Intel-compatible computers.

BIOS-Related Limitations Impacting GRUB and LILO

GRUB and LILO are subject to some limitations imposed by the BIOS in most Intel-based computers. Specifically, most BIOSes cannot access more than two hard drives, and they cannot access any data stored beyond cylinder 1023 of any drive. Note that some recent BIOSes do not have these limitations, but this is by no means universal.

All the data GRUB and LILO need to access at boot time (including the Linux kernel) is located in the /boot directory. If you follow the partition layout recommended above or are performing an Advanced Server install, the /boot directory will be in a small, separate partition. Otherwise, it may reside in the root partition (/). In either case, the partition in which /boot resides must conform to the following guidelines if you are going to use GRUB or LILO to boot your Red Hat Linux Advanced Server system:

On First Two IDE Drives

If you have 2 IDE (or EIDE) drives, /boot must be located on one of them. Note that this two-drive limit also includes any IDE CD-ROM drives on your primary IDE controller. So, if you have one IDE hard drive, and one IDE CD-ROM on your primary controller, /boot must be located on the first hard drive *only*, even if you have other hard drives on your secondary IDE controller.

On First IDE or First SCSI Drive

If you have one IDE (or EIDE) drive and one or more SCSI drives, /boot must be located either on the IDE drive or the SCSI drive at ID 0. No other SCSI IDs will work.

On First Two SCSI Drives

If you have only SCSI hard drives, `/boot` must be located on a drive at ID 0 or ID 1. No other SCSI IDs will work.

Partition *Completely* Below Cylinder 1023

No matter which of the above configurations apply, the partition that holds `/boot` must be located entirely below cylinder 1023. If the partition holding `/boot` straddles cylinder 1023, you may face a situation where GRUB and LILO will work initially (because all the necessary information is below cylinder 1023) but will fail if a new kernel is to be loaded and that kernel resides above cylinder 1023.

As mentioned earlier, it is possible that some of the newer BIOSes may permit GRUB and LILO to work with configurations that do not meet these guidelines. Likewise, some of GRUB and LILO's more esoteric features may be used to get a Linux system started, even if the configuration does not meet our guidelines. However, due to the number of variables involved, Red Hat cannot support such efforts.

Note

Disk Druid, as well as the Advanced Server installation, takes these BIOS-related limitations into account.

G Driver Disks

G.1 Why Do I Need a Driver Disk?

While the Red Hat Linux Advanced Server installation program is loading, you may see a screen that asks you for a driver disk. The driver disk screen is most often seen in three scenarios:

- If you run the installation program in `expert mode`
- If you run the installation program by entering **`linux dd`** at the `boot:` prompt
- If you run the installation program on a computer which does not have any PCI devices

G.1.1 So What Is a Driver Disk Anyway?

A driver disk adds support for hardware that is not otherwise supported by the installation program. The driver disk could be produced by Red Hat, it could be a disk you make yourself from drivers found on the Internet, or it could be a disk that a hardware vendor includes with a piece of hardware.

There is no need to use a driver disk unless you need a particular device in order to install Red Hat Linux Advanced Server. Driver disks are most often used for non-standard or very new CD-ROM drives, SCSI adapters, or NICs. These are the only devices used during the installation that might require drivers not included on the Red Hat Linux Advanced Server CD-ROMs (or boot disk, if you created an installation boot disk to begin the install process).

Note

If an unsupported device is not needed to install Red Hat Linux Advanced Server on your system, continue with the installation and add support for the new piece of hardware once the installation is complete.

G.1.2 How Do I Obtain a Driver Disk?

The Red Hat Linux Advanced Server CD-ROM 1 includes driver disk images (`images/drvnet.img` — network card drivers and `images/drvblock.img` — drivers for SCSI controllers) containing many rarely used drivers. If you suspect that your system may require one of these drivers, you should create the driver disk before beginning your Red Hat Linux Advanced Server installation.

Another option for finding specialized driver disk information is on Red Hat's website at http://www.redhat.com/support/errata under the section called **Bug Fixes**. Occasionally, popular hardware may be made available after a release of Red Hat Linux Advanced Server that will not

work with drivers already in the installation program or included on the driver disk images on the Red Hat Linux Advanced Server CD-ROM 1. In such cases, the Red Hat website may contain a link to a driver disk image.

Creating a Driver Disk from an Image File

If you have a driver disk image that you need to write to a floppy disk, this can be done from within DOS or Red Hat Linux Advanced Server.

To create a driver disk from a driver disk image using Red Hat Linux Advanced Server:

1. Insert a blank, formatted floppy disk into the first floppy drive.
2. From the same directory containing the driver disk image, such as *dd.img*, type `dd if=`*`dd.img`* `of=/dev/fd0` as root.

To create a driver disk from a driver disk image using DOS:

1. Insert a blank, formatted floppy disk into the a: drive.
2. From the same directory containing the driver disk image, such as *dd.img*, type `rawrite` *`dd.img`* `a:` at the command line.

G.1.3 Using a Driver Disk During Installation

Having a driver disk is not enough; you must specifically tell the Red Hat Linux Advanced Server installation program to load that driver disk and use it during the installation process.

Note

A driver disk is different than a boot disk. If you require a boot disk to begin the Red Hat Linux Advanced Server installation, you will still need to create that floppy and boot from it before using the driver disk.

If you do not already have an installation boot disk and your system does not support booting from the CD-ROM, you should create an installation boot disk. For instructions on how make a boot disk, see Section 1.4.2, *Making Installation Diskettes*.

Once you have created your driver disk, begin the installation process by booting from the Red Hat Linux Advanced Server CD-ROM 1 (or the installation boot disk). At the `boot:` prompt, enter either **`linux expert`** or **`linux dd`**. Refer to Section 3.2.1, *Booting the Installation Program* for details on booting the installation program.

The Red Hat Linux Advanced Server installation program will ask you to insert the driver disk. Once the driver disk is read by the installer, it can apply those drivers to hardware discovered on your system later in the installation process.

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

X