

Red Hat Enterprise Linux 4

Installation Guide for the IBM® POWER Architecture



Red Hat Enterprise Linux 4: Installation Guide for the IBM® POWER Architecture

Copyright © 2005 Red Hat, Inc.



Red Hat, Inc.

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

rhel-ig-ppc(EN)-4-Print-RHI (2004-09-24T13:10)

Copyright © 2005 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

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

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

Red Hat and the Red Hat "Shadow Man" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

All other trademarks referenced herein are the property of their respective owners.

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

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

Table of Contents

Introduction	i
1. Document Conventions	i
2. How to Use This Manual	iv
2.1. We Need Feedback!	iv
3. Accessibility Solutions	v
1. Steps to Get You Started	1
1.1. Where to Find Other Manuals.....	1
1.2. Preparation for eServer OpenPower, p5, and i5	1
1.3. Preparation for iSeries.....	2
1.4. Preparation for pSeries.....	3
1.5. Do You Have Enough Disk Space?.....	4
1.6. Can You Install Using the CD-ROM?.....	4
1.7. Preparing for a Network Installation.....	4
1.7.1. Using ISO Images for NFS Installs	6
1.8. Preparing for a Hard Drive Installation.....	7
2. Installing Red Hat Enterprise Linux	9
2.1. The Graphical Installation Program User Interface	9
2.1.1. Booting the OpenPower, p5, i5, and pSeries Installation Program.....	10
2.1.2. Booting the iSeries Installation Program	11
2.1.3. A Note about Virtual Consoles on OpenPower, p5, and pSeries	12
2.1.4. Using the HMC vterm.....	12
2.1.5. Using the iSeries Virtual Console	13
2.2. The Text Mode Installation Program User Interface	13
2.2.1. Using the Keyboard to Navigate	15
2.3. Installing from CD-ROM.....	16
2.3.1. What If the IDE CD-ROM Was Not Found?.....	17
2.4. Installing from a Hard Drive	17
2.5. Performing a Network Installation.....	18
2.6. Installing via NFS	18
2.7. Installing via FTP.....	19
2.8. Installing via HTTP.....	20
2.9. Welcome to Red Hat Enterprise Linux	21
2.10. Language Selection	21
2.11. Keyboard Configuration	22
2.12. Disk Partitioning Setup	23
2.13. Automatic Partitioning	25
2.14. Partitioning Your System	27
2.14.1. Graphical Display of Hard Drive(s).....	29
2.14.2. Disk Druid 's Buttons	29
2.14.3. Partition Fields	30
2.14.4. Recommended Partitioning Scheme	31
2.14.5. Adding Partitions	32
2.14.6. Editing Partitions	34

2.14.7. Deleting a Partition	35
2.15. Network Configuration	35
2.16. Firewall Configuration	38
2.17. Language Support Selection	41
2.18. Time Zone Configuration	42
2.19. Set Root Password	44
2.20. Package Group Selection	46
2.21. Preparing to Install	48
2.22. Installing Packages	49
2.23. Installation Complete	49
2.24. Activate Your Subscription	50
2.24.1. Provide a Red Hat Login.....	51
2.24.2. Provide Your Subscription Number	51
2.24.3. Connect Your System.....	51
2.25. Rescue Mode on POWER Systems	52
2.25.1. Special Considerations for Accessing the SCSI Utilities from Rescue Mode	52
A. Upgrading Your Current System.....	55
A.1. Determining Whether to Upgrade or Re-Install.....	55
A.2. Upgrading Your System.....	56
A.3. Upgrading Packages.....	57
A.4. Upgrade Complete	57
B. Troubleshooting Your Installation of Red Hat Enterprise Linux	59
B.1. You are Unable to Boot Red Hat Enterprise Linux	59
B.1.1. Is Your System Displaying Signal 11 Errors?.....	59
B.2. Trouble Beginning the Installation	60
B.2.1. Is Your Mouse Not Detected?	60
B.3. Trouble During the Installation	60
B.3.1. No devices found to install Red Hat Enterprise Linux	61
B.3.2. Saving Traceback Messages Without a Diskette Drive	61
B.3.3. Trouble with Partition Tables	61
B.3.4. Other Partitioning Problems for IBM™ POWER System Users	62
B.3.5. Are You Seeing Python Errors?	62
B.4. Problems After Installation	63
B.4.1. Unable to IPL from *NWSSTG	63
B.4.2. Booting into a Graphical Environment	63
B.4.3. Problems with the X Window System (GUI).....	65
B.4.4. Problems with the X Server Crashing and Non-Root Users	65
B.4.5. Problems When You Try to Log In	65
B.4.6. Your Printer Does Not Work	66
B.4.7. Apache-based httpd service/Sendmail Hangs During Startup.....	66

Error

C. An Introduction to Disk Partitions	67
C.1. Hard Disk Basic Concepts.....	67
C.1.1. It is Not What You Write, it is How You Write It	68
C.1.2. Partitions: Turning One Drive Into Many	69
C.1.3. Partitions within Partitions — An Overview of Extended Partitions	
72	
C.1.4. Making Room For Red Hat Enterprise Linux	72
C.1.5. Partition Naming Scheme.....	74
C.1.6. Disk Partitions and Other Operating Systems	75
C.1.7. Disk Partitions and Mount Points.....	75
C.1.8. How Many Partitions?	76
D. Driver Media	77
D.1. Why Do I Need Driver Media?	77
D.1.1. So What Is Driver Media Anyway?	77
D.1.2. How Do I Obtain Driver Media?.....	77
D.1.3. Using a Driver Image During Installation	78
E. Additional Boot Options	79
Index.....	83
Colophon.....	87

Introduction

Welcome to the *Red Hat Enterprise Linux Installation Guide*. This guide contains useful information to assist you during the installation of Red Hat Enterprise Linux 4. 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 Enterprise Linux.

1. Document Conventions

In this manual, 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 file names). In these cases, they are considered to be part of the command, so the entire phrase is 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.

file name

File names, directory names, paths, and RPM package names are represented this way. This style indicates that a particular file or directory exists with that name on your 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 file systems.

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

application

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

Use **Mozilla** to browse the Web.

[key]

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

To use [Tab] completion, type in a character and then press the [Tab] key. Your terminal displays 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 exits your graphical session and returns you to the graphical login screen or the console.

text found on a GUI interface

A title, word, or phrase found on a GUI interface screen or window is shown in this style. Text shown in this style indicates that 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

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

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

Instructions to type in a sequence of commands from a GUI menu look like the following example:

Go to **Applications** (the main menu on the panel) => **Programming** => **Emacs Text Editor** to start the **Emacs** text editor.

button on a GUI screen or window

This style indicates that the text can 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

Text in this style indicates text displayed to a shell prompt such as error messages and responses to commands. For example:

The `ls` command displays the contents of a directory. For example:

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


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

prompt

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

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

user input

Text that the user types, 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 must type in the **text** command at the `boot:` prompt.

<replaceable>

Text used in examples that is meant to be replaced with data provided by the user is displayed in this style. In the following example, *<version-number>* is displayed in this style:

The directory for the kernel source is `/usr/src/kernels/<version-number>/`, where *<version-number>* is the version and type of kernel installed on this system.

Additionally, we use several different strategies to draw your attention to certain pieces of information. In order of urgency, these items are marked as a note, tip, important, caution, or 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 do not take effect until you restart the DHCP daemon.

**Caution**

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

**Warning**

Be careful to remove only the necessary partitions. Removing other partitions could result in data loss or a corrupted system environment.

2. How to Use This Manual

This manual focuses on a CD-ROM-based installation and is ideal for users (both new and old) who want a quick and simple installation solution. It helps you prepare your system and walk you through the installation of Red Hat Enterprise Linux.

Red Hat Enterprise Linux includes multiple installation CD-ROMs. Note that only the first CD-ROM (CD #1) is bootable. The additional installation CD-ROMs are required, even for a minimal installation. Red Hat also provides supplementary CD-ROMs containing source RPMs and documentation for all the packages, as well as a Linux Applications CD (LACD).

In general, this document distinguishes between i5/p5 which use the Power 5 processor, and iSeries/pSeries, which use earlier versions of the POWER processors. Any references to iSeries do not include i5 systems, nor do references to pSeries include p5 systems.

If you are an experienced user and you do not need a review of the basics, you can skip ahead to Chapter 2 *Installing Red Hat Enterprise Linux* to begin the installation process.

2.1. We Need Feedback!

If you discover a typo in the *Red Hat Enterprise Linux Installation Guide* or have thought of a way to make this manual better, we would love to hear from you. Submit a bug report against the component `rhel-ig-ppc` in Bugzilla at:

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

When submitting a bug report, be sure to mention the manual's identifier:

`rhel-ig-ppc(EN)-4-Print-RHI (2004-09-24T13:10)`

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]), use the online support system by registering your subscriptions at:

<http://www.redhat.com/apps/activate/>

3. Accessibility Solutions

While the graphic user interface (GUI) is convenient for sighted users, it is often inhibiting to those with visual impairments because of the difficulty speech synthesizers have interpreting graphics. Red Hat Enterprise Linux is an ideal operating system for users with visual limitations because the GUI is not required by the kernel. Most modern tools including email, news, Web browsers, calendars, calculators, and much more can run on Linux without a graphical environment. The working environment can also be customized to meet the hardware or software needs of the user.

Red Hat, Inc. is the distribution of choice for people with special needs because of the outstanding support that is offered with the purchase of any boxed set. Many Linux distributions provide limited or nonexistent support to customers. Red Hat's installation support is deliverable via email or via the telephone and special circumstances will be considered and addressed for users with physical limitations. Customers should inform the support technician if they require specialized support.

For more information, refer to:

- <http://www.tldp.org/HOWTO/Accessibility-HOWTO/>
- <http://www.tldp.org/HOWTO/Keyboard-and-Console-HOWTO.html>

- The `README-Accessibility` file provides more information and instructions for using some of the tools available for users with special needs. This file can be found in `/usr/share/doc/redhat-release-X/`, where `X` is the release number of your installed subscription.

Chapter 1.

Steps to Get You Started

1.1. Where to Find Other Manuals

Red Hat Enterprise Linux manuals are available online or on the Red Hat Enterprise Linux Documentation CD included with your Red Hat Enterprise Linux subscription.

If you are a new user and need information on getting started with Red Hat Enterprise Linux, you may find the *Red Hat Enterprise Linux Step By Step Guide* to be helpful.

For an overview of system administration, refer to the *Red Hat Enterprise Linux Introduction to System Administration*. If you are looking for information on system configuration, you may find the *Red Hat Enterprise Linux System Administration Guide* to be helpful.

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

If you need information on securing your system, refer to the *Red Hat Enterprise Linux Security Guide*.

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



Note

Although this manual reflects the most current information possible, read the *Red Hat Enterprise Linux Release Notes* for information that may not have been available prior to our documentation being finalized. The Release Notes can be found on the Red Hat Enterprise Linux CD #1, online at <http://www.redhat.com/docs/>, or in the `/usr/share/doc/redhat-release-4<product>/` directory after installation, where `<product>` is AS, ES, WS, or Desktop.

1.2. Preparation for eServer OpenPower, p5, and i5

The eServer OpenPower, p5, and i5 systems introduce many options for partitioning, virtual or native devices, and consoles. All three versions of the system use the same kernel and have many of the same options available, depending on the system configuration.

If you are using a non-partitioned OpenPower or p5 system, you do not need any pre-installation setup. For systems using the HVSI serial console, hook up your console to the T2 serial port.

If using a partitioned system, whether OpenPower, p5, or i5, the steps to create the partition and start the installation are largely the same. You should create the partition at the HMC and assign some CPU and memory resources, as well as SCSI and Ethernet resources, which can be either virtual or native. The HMC create partition wizard steps you through the creation.

For more information on creating the partition, refer to IBM's Infocenter article on *Configuring Linux logical partitions* available online at:

```
http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/  
index.htm?info/iphbi/iphbiconfig.htm
```

Note that due to the length of this URL, it has been broken into two lines.

If you are using virtual SCSI resources, rather than native SCSI, you must configure a 'link' to the virtual SCSI serving partition, and then configure the virtual SCSI serving partition itself. You create a 'link' between the virtual SCSI client and server slots using the HMC. You can configure a virtual SCSI server on either AIX or i5/OS, depending on which model and options you have.

For more information on these topics, refer to IBM's Infocenter article, *Creating a Linux logical partition using i5/OS virtual I/O resources* available online at:

```
http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/  
index.htm?info/iphbl/iphblioserverpartition.htm
```

Note that due to the length of this URL, it has been broken into two lines.

**Note**

Additional Information on virtualization eServer i5 and p5 can be found in the Redbooks SG24-6388-00 and SG24-7940-00, *Linux on IBM eServer i5 Implementation Guide* and *Introduction to Advanced POWER Virtualization on IBM eServer p5 Servers*. These can be accessed from <http://www.redbooks.ibm.com/> and searching on the book number.

Once you have your system configured, you need to Activate from the HMC or power it on. Depending on what type of install you are doing, you may need to configure SMS to correctly boot the system into the installation program.

1.3. Preparation for iSeries

The installation process assumes basic familiarity with the IBM iSeries platform and makes reference to IBM's *Configuring Linux in a Guest Partition*. In general, this section assumes that you have read the *Configuring Linux in a Guest Partition* document (<http://publib.boulder.ibm.com/series/v5r2/ic2924/info/rzalm/rzalmlinuxkickoff.htm>) in its entirety.



Tip

Additional information on the iSeries can be found online at:

<http://www-1.ibm.com/servers/eserver/iseries/linux/>

To prepare your iSeries system for installation, follow the instructions in the *Planning to run Linux on a logical partition* and *Creating a logical partition to run Linux* sections of the *Configuring Linux in a Guest Partition* document. Consult your OS/400 command documentation for information on configuring the OS/400 facilities (<http://publib.boulder.ibm.com/series/v5r2/ic2924/info/rbam6/rbam6clmain.htm>). The basic steps for installation preparation on your iSeries are as follows:

First, create a Logical Partition in SST using the screens in **Work with system partitions**. The partition must be of type Guest, have some amount of processor and memory, and either a native LAN adapter or virtual Ethernet assigned. After you create your Linux partition, assign a hosting partition to it. Select option **13** on the new Linux partition in the **Work with Partition Configuration** screen. Next, IPL your primary partition after creating the new partition.

If you are using shared processors with V5R2 OS/400, you can dynamically change the amount of processor units assigned to your partition, as long as it is within the minimum and maximum processor units you have set for your partition. You cannot dynamically go below a minimum of 0.10 times the number of processors. Changing the processor or memory limits requires a Primary Partition IPL.

Next, create an NWSD by executing the command `crtnwsd` from the OS/400 command line. The Partition field has the name of the Logical Partition you created in SST; it should be of type Guest and Codepage 437. For a basic NWSD, use `*None` for the Resource Name and keep the defaults for all the other fields — the boot fields are setup in Section 2.1.2 *Booting the iSeries Installation Program*.

Assuming you plan on using a virtual disk rather than native SCSI using the `ipr` driver, create an NWSSTG (virtual disk) by executing the command `wrknwsstg` from the OS/400 command line. Once you create it, you must add a link to the NWSD you created using the name of the NWSD.

1.4. Preparation for pSeries

If you are using a partitioned pSeries system, you must configure the system using the HMC. For more information on this HMC, refer to the *Hardware Management Console for pSeries Installation and Operations Guide* available online at:

```
http://publib16.boulder.ibm.com/pseries/en_US/  
infocenter/base/hardware_docs/pdf/380590.pdf
```

Note that due to the length of this URL, it has been broken into two lines.

1.5. Do You Have Enough Disk Space?

Nearly every modern-day operating system (OS) uses *disk partitions*, and Red Hat Enterprise Linux is no exception. When you install Red Hat Enterprise Linux, 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), refer to Appendix C *An Introduction to Disk Partitions* before proceeding.

The disk space used by Red Hat Enterprise Linux must be separate from the disk space used by other OSes you may have installed on your system.

Before you start the installation process, you must

- have enough *unpartitioned*¹ disk space for the installation of Red Hat Enterprise Linux, or
- have one or more partitions that may be deleted, thereby freeing up enough disk space to install Red Hat Enterprise Linux.

To gain a better sense of how much space you really need, refer to the recommended partitioning sizes discussed in Section 2.14.4 *Recommended Partitioning Scheme*.

1.6. Can You Install Using the CD-ROM?

Installing from a CD-ROM requires that you have purchased a Red Hat Enterprise Linux 4 product, or you have a Red Hat Enterprise Linux CD-ROM, and you have a CD-ROM drive. If your system supports booting from the CD-ROM, it is an easy way to begin a local CD-ROM installation.

1. Unpartitioned disk space means that available disk space on the hard drive(s) you are installing to has not been divided into sections for data. When you partition a disk, each partition behaves like a separate disk drive.

1.7. Preparing for a Network Installation



Note

Make sure an installation CD (or any other type of CD) is not in your hosting partition's drive if you are performing a network-based installation. Having a CD in the drive may cause unexpected errors.

The Red Hat Enterprise Linux installation media must be available for either a network installation (via NFS, FTP, or HTTP) or installation via local storage. Use the following steps if you are performing an NFS, FTP, or HTTP installation.

The NFS, FTP, or HTTP server to be used for installation over the network must be a separate machine which can provide the complete `RedHat/` directory. Both the `RedHat/base/` and `RedHat/RPMS/` directories must be available and populated with all files from all installation CD-ROMs.



Note

The directory specified in the following refers to `/location/of/disk/space/`. This means it is the directory up to, but *not* including, the `RedHat/` distribution directory. For example, if you have Red Hat Enterprise Linux 4 installation files located in `/export/rhel/` on the installation staging server, `/location/of/disk/space/` would be `/export/rhel/`.

To copy the `RedHat/` directory from the installation CD-ROMs to a Linux machine which acts as an installation staging server, perform the following steps:

- For each binary CD-ROM, 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/rhel/`
 - `umount /mnt/cdrom/`
- Note that the Release Notes are not included in the `RedHat` directory. Unless they are specifically copied over, the Release Notes will not be available during your installation

of Red Hat Enterprise Linux. The Release Notes are formatted in HTML files located at the root of the disc. Copy the files to your installation directory. For example:

```
cp /mnt/cdrom/RELEASE-NOTES*.html /location/of/disk/space/
```

The Release Notes are also available online from <http://www.redhat.com/docs/>.

- Next, make `/location/of/disk/space/` available for network installation via NFS, FTP, or HTTP and verify access from a client system.
- For NFS, export the directory by adding an entry to `/etc/exports` to export to a specific system:

```
/location/of/disk/space client.ip.address(ro,no_root_squash)
```

To export to all machines (not appropriate for all NFS systems), add:

```
/location/of/disk/space *(ro,no_root_squash)
```

Start the NFS daemon (on a Red Hat Enterprise Linux machine, use `/sbin/service nfs start`). If NFS is already running, reload the configuration file (on a Red Hat Enterprise Linux system, use `/sbin/service nfs reload`).

Be sure to test the NFS share following the directions in the *Red Hat Enterprise Linux System Administration Guide*.

If the `RedHat/` directory does not appear in the NFS shared directory, the wrong path may have been mounted and/or exported.

- FTP and HTTP installations also support a second type of tree structure. To make it easier to access the contents of the installation CD-ROMs, mount each CD-ROM or ISO image with the following mount point on the FTP or HTTP server (where *X* is the number of the CD-ROM or ISO image):

```
/location/of/disk/space/discX/
```

For example:

```
mount -o loop CD1.iso /location/of/disk/space/disc1/
```

1.7.1. Using ISO Images for NFS Installs

NFS installations can use ISO (or CD-ROM) images rather than copying an entire installation tree. After placing the required ISO images (the binary Red Hat Enterprise Linux CD-ROMs) in a directory, choose to install via NFS. 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 an NFS 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.

**Note**

The Red Hat Enterprise Linux installation program has the ability to test the integrity of the installation media. It works with the CD, DVD, hard drive ISO, and NFS ISO installation methods. Red Hat recommends that you test all installation media before starting the installation process, and before reporting any installation-related bugs (many of the bugs reported are actually due to improperly-burned CDs). To use this test, type the following command at the `boot:` prompt:

```
linux mediacheck
```

Additionally, if a file called `updates.img` exists in the directory from which you install, then it will be used for installation program updates. Refer to the file `install-methods.txt` in the `anaconda` RPM package for detailed information on the various ways to install Red Hat Enterprise Linux, as well as how to apply the installation program updates.

**Note**

You can only have the ISO images for one release and one variant of Red Hat Enterprise Linux in the directory.

1.8. Preparing for a Hard Drive Installation

**Note**

Hard drive installations only work from `ext2`, `ext3`, or `FAT` file systems. If you have a file system 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. An ISO image is a file containing an exact copy of a CD-ROM disk image. Because Red Hat Enterprise Linux has so many packages included with its distribution, there are several ISO images available. After placing the required ISO images (the binary Red Hat Enterprise Linux CD-ROMs) in a directory, choose to install from the hard drive. You can 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.

To prepare your system for a hard drive installation, you must set the system up in one of the following ways:

- Using a set of CD-ROMs — Create CD-ROM ISO image files from each installation CD-ROM. For each CD-ROM, execute the following command on a Linux system:

```
dd if=/dev/cdrom of=/tmp/file-name.iso
```

- Using ISO images — transfer these images to the system to be installed.

Verifying that ISO images are intact before you attempt an installation, helps to avoid problems. 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 Linux machine as the ISO images.



Note

The Red Hat Enterprise Linux installation program has the ability to test the integrity of the installation media. It works with the CD, DVD, hard drive ISO, and NFS ISO installation methods. Red Hat recommends that you test all installation media before starting the installation process, and before reporting any installation-related bugs (many of the bugs reported are actually due to improperly-burned CDs). To use this test, type the following command at the `boot:` prompt:

```
linux mediacheck
```

iSeries system users must remember to add boot options, as are listed in this manual for other operating systems, by first logging into an OS/400 system as an iSeries user and then setting any IPL parameters within their partition's NWSD.

Additionally, if a file called `RedHat/base/updates.img` exists in the directory from which you install, it is used for installation program updates. Refer to the file `install-methods.txt` in the `anaconda` RPM package for detailed information on the various ways to install Red Hat Enterprise Linux, as well as how to apply the installation program updates.

Chapter 2.

Installing Red Hat Enterprise Linux

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

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

2.1. The Graphical Installation Program User Interface

If you have used a *graphical user interface (GUI)* before, you are already familiar with this process; use your mouse to navigate the screens, click buttons, or enter text fields.

You can also navigate through the installation using the keyboard. The [Tab] key allows you to move around the screen, the Up and Down arrow keys to scroll through lists, [+] and [-] keys expand and collapse lists, while [Space] and [Enter] selects or removes from selection a highlighted item. You can also use the [Alt]-[X] key command combination as a way of clicking on buttons or making other screen selections, where [X] is replaced with any underlined letter appearing within that screen.

If you would like to use a graphical installation with a system that does not have that capability, such as a partitioned system, you can use VNC or display forwarding. Both the VNC and display forwarding installation options require an active network during the installation and the use of boot time arguments. For more information on available boot time options, refer to Appendix E *Additional Boot Options*.



Note

If you do not wish to use the GUI installation program, the text mode installation program is also available. To start the text mode installation program, use the following boot command:

```
linux text
```

Please refer to Section 2.2 *The Text Mode Installation Program User Interface* for a brief overview of text mode installation instructions.

It is highly recommended that installs be performed using the GUI installation program. The GUI installation program offers the full functionality of the Red Hat Enterprise Linux installation program, including LVM configuration which is not available during a text mode installation.

Users who must use the text mode installation program can easily follow the GUI installation instructions and obtain all needed information.

iSeries system users must remember to add boot options, as are listed in this manual (Appendix E *Additional Boot Options*) for other operating systems, by first logging into an OS/400 system as an iSeries user and then setting any IPL parameters within their OS/400 session.

2.1.1. Booting the OpenPower, p5, i5, and pSeries Installation Program

To boot an OpenPower, p5, i5, or pSeries system from a CD-ROM, you must specify the install boot device in the **System Management Services** (SMS) menu.

To enter the **System Management Services** GUI, press the [1] key during the boot process when you hear the chime sound. This brings up a graphical interface similar to the one described in this section.

On a text console, press [1] when the self test is displaying the banner along with the tested components:

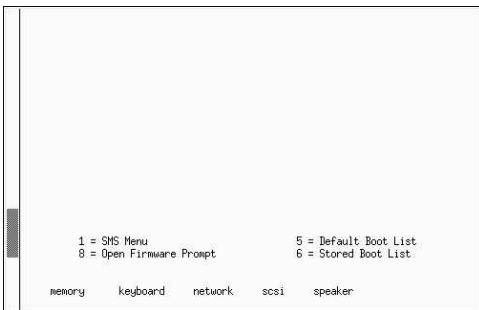


Figure 2-1. SMS console

Once in the SMS menu, select the option for **Select Boot Options**. In that menu, specify **Select Install or Boot a Device**. There, select **CD/DVD**, and then the bus type (in most

cases SCSI). If you are uncertain, you can select to view all devices. This scans all available buses for boot devices, including network adapters and hard drives.

Finally, select the device containing the installation CD. YABOOT is loaded from this device and you are presented with a `boot:` prompt. Press [Enter] or wait for the timeout to expire for the installation to begin.

If you are booting a non-iSeries system (or booting a *STMF on i5) via the network, use the `images/netboot.img` file on CD #1.

2.1.2. Booting the iSeries Installation Program

Insert the installation CD into the CD-ROM drive. Determine the volume name by which OS/400 refers to the CD. This can be found by executing the command `wrklnk '/qopt/*'` at the OS/400 command line. `WRKOPTVOL` gives more information on the optical media. Unless there is another CD with the same label in another drive, the volume name is **Red_Hat**. Note, if the volume name includes a timestamp (as it does when there are name conflicts), ejecting and re-inserting the CD causes the volume name to change and you must determine what it is again.

Ensure that your Network Server (NWS) is varied off using the OS/400 `WRKCFGSTS (*nws)` command.

Edit the Network Server Description (NWS). Set the IPL source field to `*STMF`.

Set the IPL stream file field to `/QOPT/RED_HAT/PPC/ISERIES/BOOT.IMG` (case is not significant), where **Red_Hat** is the name of the CD volume as it appears under `/QOPT`.

Set the IPL parameters option to `'dd'` (including quotes) if you are using a driver. Otherwise, set the option to `*NONE`.



Note

Red Hat recommends using the `vnc` or `display=<workstation>` boot commands to begin a graphical installation. Otherwise, set the field to `*NONE` to perform a text mode installation. On your workstation, be sure to set the `xhost+` command to the hostname of the system being installed. For more information on available boot commands, refer to Appendix E *Additional Boot Options*.

Use a Telnet client to connect to the virtual console of the OS/400 partition hosting your Linux partition using a command like `telnet <hosting partition> 2301`. Refer to Section 2.1.5 *Using the iSeries Virtual Console* for more details on using the virtual console.

Vary on the NWS. The boot sequence begins on the virtual console in a few seconds. The installation program begins automatically. If you are performing a VNC or X-Host based installation, network information can be configured via a virtual console.

2.1.3. A Note about Virtual Consoles on OpenPower, p5, and pSeries

This information only applies to users of non-partitioned OpenPower, p5, and pSeries systems using a video card as their console. Users of partitioned OpenPower, p5, i5, and pSeries systems should skip to Section 2.1.4 *Using the HMC vterm*, while users of iSeries should skip to Section 2.1.5 *Using the iSeries Virtual Console*.

The Red Hat Enterprise Linux installation program offers more than the dialog boxes of the installation process. Several kinds of diagnostic messages are available to you, as well as 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.

A virtual console is a shell prompt in a non-graphical environment, accessed from the physical machine, not remotely. Multiple virtual consoles can be accessed simultaneously.

These virtual consoles can be helpful if you encounter a problem while installing Red Hat Enterprise Linux. Messages displayed on the installation or system consoles can help pinpoint a problem. Refer to Table 2-1 for a listing of the virtual consoles, keystrokes used to switch to them, and their contents.

Generally, there is no reason to leave the default console (virtual console #7 for graphical installations or virtual console #1 for text-based installations) unless you are attempting to diagnose installation problems.

Console	Keystrokes	Contents
1	[Ctrl]-[Alt]-[F1]	installation dialog
2	[Ctrl]-[Alt]-[F2]	shell prompt
3	[Ctrl]-[Alt]-[F3]	install log (messages from installation program)
4	[Ctrl]-[Alt]-[F4]	system-related messages
5	[Ctrl]-[Alt]-[F5]	other messages
7	[Ctrl]-[Alt]-[F7]	X graphical display

Table 2-1. Console, Keystrokes, and Contents

2.1.4. Using the HMC vterm

The HMC vterm is the console for any partitioned OpenPower, p5, i5, or pSeries system. This is opened by right clicking on the partition on the HMC, and then selecting **Open Terminal Window**. Only a single vterm can be connected to the console at one time and there is no console access for partitioned system besides the vterm. This often is referred to as a 'virtual console', but is different from the virtual consoles in Section 2.1.3 *A Note about Virtual Consoles on OpenPower, p5, and pSeries* or the iSeries virtual console in Section 2.1.5 *Using the iSeries Virtual Console*.

2.1.5. Using the iSeries Virtual Console

Before you can use the iSeries virtual console for installation, you must create a DST userid and give it authority to access the virtual console. In V5R2 you simply start SST (**STRSST** from the command line of the hosting OS/400 partition) and select option 8, **Work with service tools user IDs**. You cannot use the QSECOFR userid to access the virtual console, you must create a new userid. Once you have created your userid, give it **Partition remote panel key** access for your Linux partition.

To access the virtual console, use the command `telnet <systemname> 2301` where `<systemname>` is the hostname of the OS/400 partition. When it connects, the virtual console lists the Linux partitions hosted by that OS/400 partition. Select the partition number and enter your DST userid and password as prompted.

The virtual console is usable from a variety of Telnet clients, but it works best when used either from Linux-based X terminals (or GNOME terminals) using the standard `telnet` client or the MS Windows **PuTTY** client. Other Telnet clients, such as the standard Windows Telnet client, have problems with the virtual console's rendering of the text based installation program or the control sequences in programs such as `vi`, so their use is not recommended.

2.2. The Text Mode Installation Program User Interface

The Red Hat Enterprise Linux text mode installation program uses a screen-based interface that includes most of the on-screen *widgets* commonly found on graphical user interfaces. Figure 2-2, and Figure 2-3, illustrate the screens that appear during the installation process.



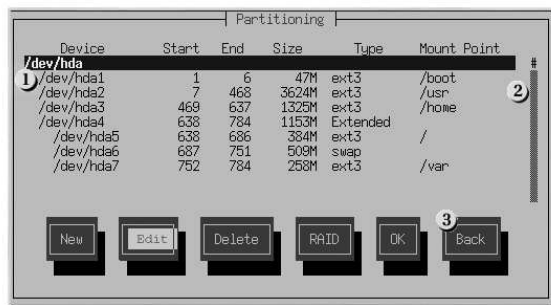
Note

While text mode installations are not explicitly documented, those using the text mode installation program can easily follow the GUI installation instructions.



- 1 Window 2 Check Box 3 Text Input

Figure 2-2. Installation Program Widgets as seen in Boot Loader Configuration



- 1 Text Widget 2 Scroll Bar 3 Button Widget

Figure 2-3. Installation Program Widgets as seen in Disk Druid

Here is a list of the most important widgets shown in Figure 2-2 and Figure 2-3:

- Window — Windows (usually referred to as *dialogs* in this manual) 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 disappears, allowing you to continue working in the window underneath.

- **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 or deselect a feature.
- **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.
- **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.
- **Scroll Bar** — Scroll bars appear on the side or bottom of a window to control which part of a list or document is currently in the window's frame. The scroll bar makes it easy to move to any part of a file.
- **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 the cursor itself may only appear positioned in or next to the widget. In Figure 2-2, the cursor is positioned on the **OK** button. Figure 2-3, shows the cursor on the **Edit** button.

2.2.1. Using the Keyboard to Navigate

Navigation through the installation dialogs is performed through a simple set of keystrokes. To move the cursor, use the [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.

**Caution**

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).

2.3. Installing from CD-ROM

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

**Note**

The first CD-ROM is required when booting the installation program, and again after subsequent CD-ROMs have been processed.

The installation program then probes your system and attempts to identify your CD-ROM drive. It starts by looking for an IDE (also known as an ATAPI) CD-ROM drive. If found, continue to the next stage of the installation process (refer to Section 2.10 *Language Selection*).

**Note**

To cancel the installation process at this time, iSeries users should vary off the NWS and non-iSeries users should press the computer's Reset button or use the [Control]-[Alt]-[Delete] key combination to restart your machine. You can safely cancel the installation at any point before the **About to Install** screen. Refer to Section 2.21 *Preparing to Install* for more information.

The following applies to non-iSeries systems only:

If your CD-ROM drive is not detected, and it is a SCSI CD-ROM, the installation program prompts 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 detect your SCSI adapter automatically.

**Tip**

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

2.3.1. 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 **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.

2.4. Installing from a Hard Drive

The **Select Partition** screen 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 Enterprise Linux.



Figure 2-4. Selecting Partition Dialog for Hard Drive Installation

Enter the device name of the partition containing the Red Hat Enterprise Linux ISO images. There is also a field labeled **Directory holding images**.

If the ISO images are in the root (top-level) directory of a partition, enter a `/`. If the ISO images are located in a subdirectory of a mounted partition, enter the name of the directory holding the ISO images within that partition. For example, if the partition on which the ISO images is normally mounted as `/home/`, and the images are in `/home/new/`, you would enter `/new/`.

After you have identified the disk partition, the **Welcome** dialog appears.

2.5. Performing a Network Installation

If you are performing a network installation, the **Configure TCP/IP** dialog appears. This dialog asks for your IP and other network addresses. You can choose to configure the IP address and Netmask of the device via DHCP or manually. If manually, enter the IP address you are using during installation and press [Enter]. The installation program attempts to guess your Netmask based on your IP address; you may change the netmask if it is incorrect. Press [Enter]. The installation program guesses the Default gateway and Primary nameserver addresses from your IP address and netmask; you may change them if they are incorrect.

2.6. Installing via NFS

The NFS dialog applies only if you are installing from an NFS server (if you selected **NFS Image** in the **Installation Method** dialog).

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

Next, enter the name of the exported directory. If you followed the setup described in Section 1.7 *Preparing for a Network Installation*, 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 Enterprise Linux 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 the NFS site contains the directory `/mirrors/redhat/arch/RedHat/`, enter `/mirrors/redhat/arch/` (where `arch` is replaced with the architecture type of your system, such as `i386`, `ia64`, `ppc`, or `s390`). If everything was specified properly, a message appears indicating that the installation program for Red Hat Enterprise Linux is running.




Figure 2-5. NFS Setup Dialog

If the NFS server is exporting the ISO images of the Red Hat Enterprise Linux CD-ROMs, enter the directory which contains the ISO images.

Next, the **Welcome** dialog appears.

2.7. Installing via FTP

The FTP 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 Enterprise Linux.

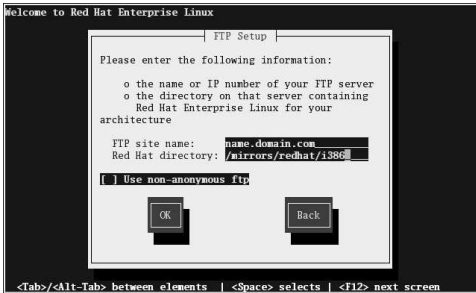


Figure 2-6. FTP Setup Dialog

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/arch/RedHat/`, enter `/mirrors/redhat/arch/` (where *arch* is replaced with the architecture type of your system, such as *i386*, *ia64*, *ppc*, or *s390*). If everything was specified properly, a message box appears indicating that `base/hdlist` is being retrieved.

Next, the **Welcome** dialog appears.



Tip

You can save disk space by using the ISO images you have already copied to the server. To accomplish this, install Red Hat Enterprise Linux using ISO images without copying them into a single tree by loopback mounting them. For each ISO image:

```
mkdir discX
mount -o loop example-1.iso discX
```

2.8. Installing via HTTP

The HTTP 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 Enterprise Linux.

Enter the name or IP address of the HTTP site you are installing from, and the name of the directory containing the RedHat/ installation files for your architecture. For example, if the HTTP site contains the directory `/mirrors/redhat/arch/RedHat/`, enter `/mirrors/redhat/arch/` (where *arch* is replaced with the architecture type of your

system, such as i386, ia64, ppc, or s390). If everything was specified properly, a message box appears indicating that `base/hdlist` is being retrieved.

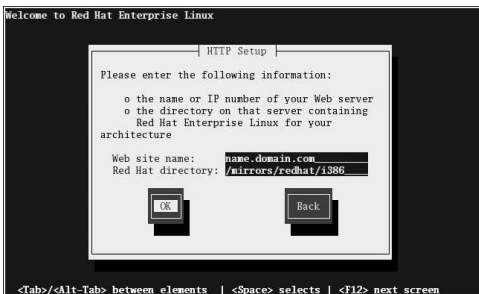


Figure 2-7. HTTP Setup Dialog

Next, the **Welcome** dialog appears.



Tip

You can save disk space by using the ISO images you have already copied to the server. To accomplish this, install Red Hat Enterprise Linux using ISO images without copying them into a single tree by loopback mounting them. For each ISO image:

```
mkdir discX  
  
mount -o loop example-1.iso discX
```

2.9. Welcome to Red Hat Enterprise Linux

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

Notice the **Hide Help** button at the bottom left corner of the screen. The help screen is open by default. To minimize the help text, click on **Hide Help**.

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

2.10. Language Selection

Using your mouse, select a language to use for the installation (refer to Figure 2-8).

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




Figure 2-8. Language Selection

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

2.11. Keyboard Configuration

Using your mouse, select the correct layout type (for example, U.S. English) for the keyboard you would prefer to use for the installation and as the system default (refer to Figure 2-9).

Once you have made your selection, click **Next** to continue.

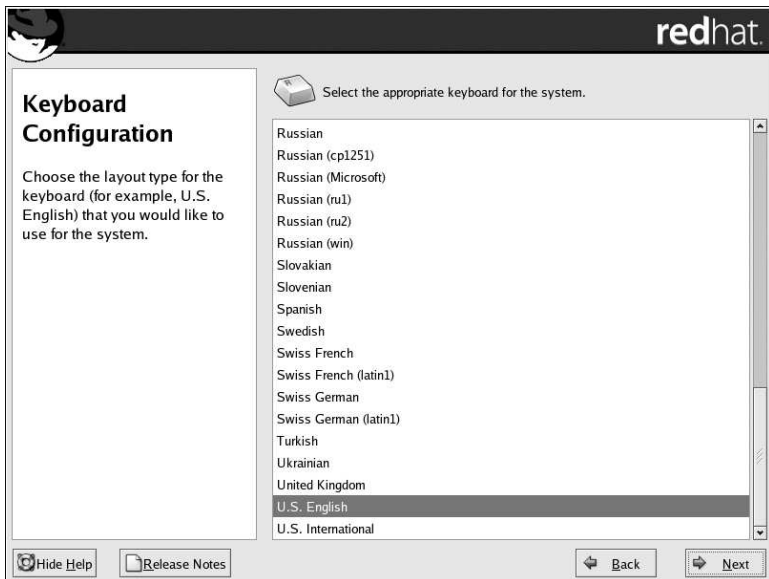


Figure 2-9. Keyboard Configuration



Tip

To change your keyboard layout type after you have completed the installation, use the **Keyboard Configuration Tool**.

Type the `system-config-keyboard` command in a shell prompt to launch the **Keyboard Configuration Tool**. If you are not root, it prompts you for the root password to continue.

2.12. 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 multiple operating systems. If you are not sure how you want your system to be partitioned, read *Appendix C An Introduction to Disk Partitions* for more information.

On this screen, you can choose to perform automatic partitioning, or manual partitioning using **Disk Druid**.

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 the **Disk Druid** partitioning tool.

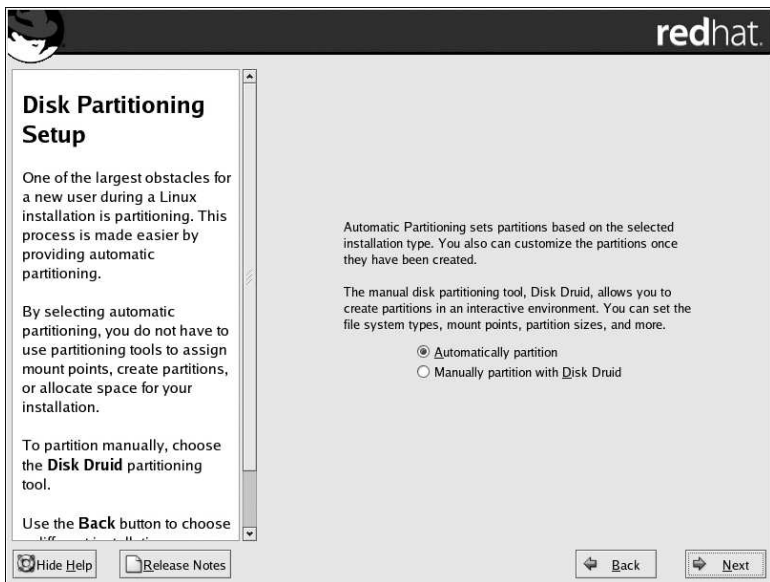


Figure 2-10. Disk Partitioning Setup

If you chose to manually partition using **Disk Druid**, refer to Section 2.14 *Partitioning Your System*.

 **Warning**

If you receive an error after the **Disk Partitioning Setup** phase of the installation saying something similar to

The partition table on device hda was unreadable. To create new partitions it must be initialized, causing the loss of ALL DATA on this drive.

you may not have a partition table on that drive or the partition table on the drive may not be recognizable by the partitioning software used in the installation program.

No matter what type of installation you are performing, backups of the existing data on your systems should always be made.

2.13. 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 does 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).



Caution

If you select this option, all data on the selected hard drive(s) is 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 Enterprise Linux.

- **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).

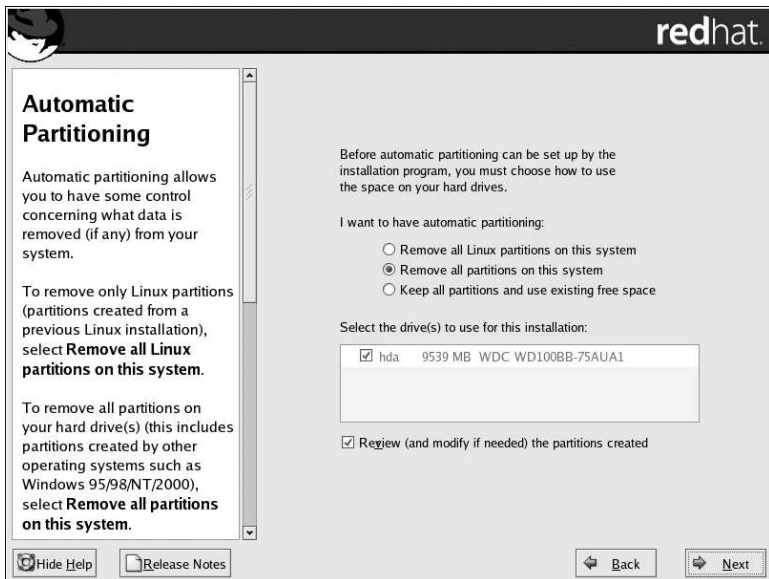


Figure 2-11. Automatic Partitioning

Using your mouse, choose the hard drive(s) on which you want Red Hat Enterprise Linux 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, are not touched.



Caution

It is always a good idea to back up any data that you have on your systems. For example, if you are upgrading or creating a dual-boot system, you should back up any data you wish to keep on your hard drive(s). Mistakes do happen and can result in the loss of all your data.



Tip

This text does not apply to iSeries system users.

If you have a RAID card, be aware that some firmware 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.

A `/boot/` partition is also necessary for software RAID setups.

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, the partitions created for you in **Disk Druid** appear. You can make modifications to these partitions if they do not meet your needs.

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

2.14. Partitioning Your System

If you chose automatic partitioning and did not select **Review**, skip ahead to Section 2.15 *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 to partition manually, you must tell the installation program where to install Red Hat Enterprise Linux. This is done by defining mount points for one or more disk partitions in which Red Hat Enterprise Linux is installed.



Note

If you have not yet planned how to set up your partitions, refer to Appendix C *An Introduction to Disk Partitions* and Section 2.14.4 *Recommended Partitioning Scheme*. At a bare minimum, you need an appropriately-sized root (`/`) partition, a `/boot/` partition for non-iSeries systems (iSeries systems do not require a `/boot/` partition), PPC PReP boot partition, and a swap partition equal to twice the amount of RAM you have on the system.

For iSeries users:

It is recommended to setup a PReP boot partition, either as your primary or backup boot method. The size range of a PReP boot partition should be between 16 and 32 MB and it needs to be a primary partition marked active.

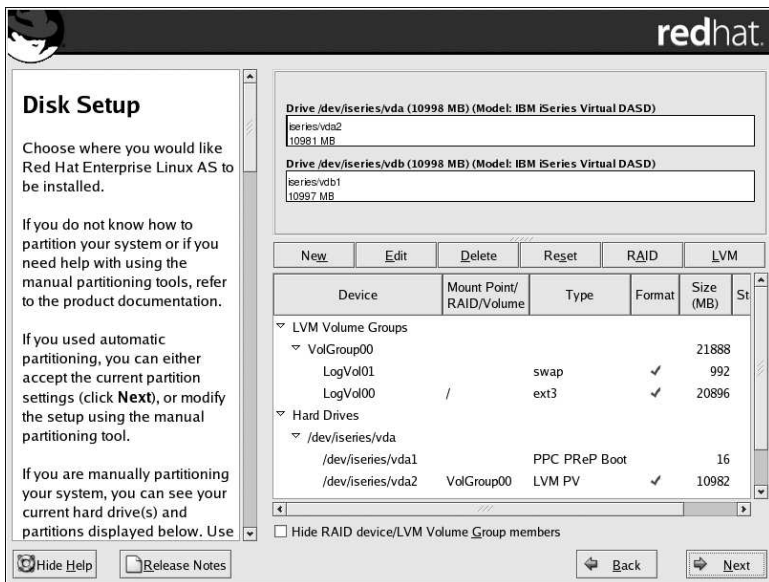


Figure 2-12. Partitioning with Disk Druid on iSeries systems

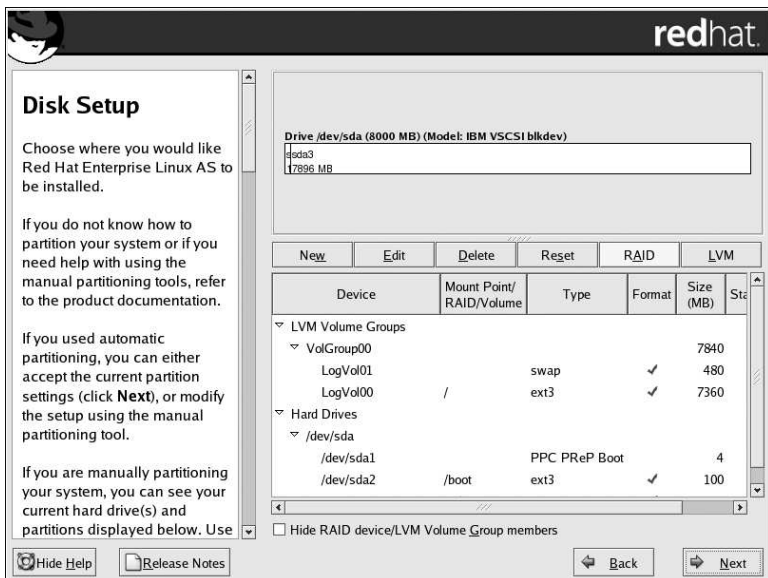


Figure 2-13. Partitioning with Disk Druid on pSeries systems

The partitioning tool used by the installation program is **Disk Druid**. With the exception of certain esoteric situations, **Disk Druid** can handle the partitioning requirements for a typical installation.

2.14.1. 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 can review 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.

2.14.2. Disk Druid's Buttons

These buttons control **Disk Druid's** actions. They are used to change the attributes of a partition (for example the file system 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 the mount point and size fields) 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.

- To make a RAID device, you must first create (or reuse existing) 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.
- **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.
- **RAID:** Used to provide redundancy to any or all disk partitions. *It should only be used if you have experience using RAID.* To read more about RAID, refer to the *Red Hat Enterprise Linux System Administration Guide*.

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

- **LVM:** Allows you to create an LVM logical volume. The role of LVM (Logical Volume Manager) is to present a simple logical view of underlying physical storage space, such as a hard drive(s). LVM manages individual physical disks — or to be more precise, the individual partitions present on them. *It should only be used if you have experience using LVM.* To read more about LVM, refer to the *Red Hat Enterprise Linux System Administration Guide*. Note, LVM is only available in the graphical installation program.

To create an LVM logical volume, you must first create partitions of type physical volume (LVM). Once you have created one or more physical volume (LVM) partitions, select **LVM** to create an LVM logical volume.

2.14.3. 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.
- **Mount Point/RAID/Volume:** 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 is 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.
- **Type:** This field shows the partition's file system type (for example, ext2 or ext3).
- **Format:** This field shows if the partition being created will be formatted.
- **Size (MB):** This field shows the partition's size (in MB).
- **Start:** This field shows the cylinder on your hard drive where the partition begins.
- **End:** This field shows the cylinder on your hard drive where the partition ends.

Hide RAID device/LVM Volume Group members: Select this option if you do not want to view any RAID device or LVM Volume Group members that have been created.

2.14.4. Recommended Partitioning Scheme

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

- A swap partition (at least 256 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.

If you are unsure about what size swap partition to create, make it twice the amount of RAM on your machine. It must be of type swap.

Creation of the proper amount of swap space varies depending on a number of factors including the following (in descending order of importance):

- The applications running on the machine.
- The amount of physical RAM installed on the machine.
- The version of the OS.

Swap should equal 2x physical RAM for up to 2 GB of physical RAM, and then 1x physical RAM for any amount above 2 GB, but never less than 32 MB.

Using this formula, a system with 2 GB of physical RAM would have 4 GB of swap, while one with 3 GB of physical RAM would have 5 GB of swap. Creating a large swap space partition can be especially helpful if you plan to upgrade your RAM at a later time.

For systems with really large amounts of RAM (more than 32 GB) you can likely get away with a smaller swap partition (around 1x, or less, of physical RAM).

- A PPC PReP boot partition on the first partition of the hard drive — the PPC PReP boot partition contains the operating system kernel (on iSeries systems) or the YABOOT boot loader (which allows other POWER systems to boot Red Hat Enterprise Linux). *For POWER systems excluding iSeries systems:* Unless you plan to boot from a floppy or network source, you must have a PPC PReP boot partition to boot Red Hat Enterprise Linux.

For OpenPower, p5, i5, and pSeries users: The PPC PReP boot partition should be between 4-8 MB, not to exceed 10 MB.

For iSeries users: It is recommended to setup a PPC PReP boot partition, either as your primary or backup boot method. The size range of a PPC PReP boot partition should be between 16 and 32 MB (but not to exceed 32 MB). The PPC PReP boot partition must be a primary partition, marked active.

- *For non-iSeries users:* A `/boot/` partition (100 MB) (`/boot/` is not needed on iSeries systems) — the partition mounted on `/boot/` contains the operating system kernel (which allows your system to boot Red Hat Enterprise Linux), along with files used during the bootstrap process. Due to the limitations of most PC firmware, creating a small partition to hold these is a good idea. For most users, a 100 MB boot partition is sufficient.



Caution

If you have a RAID card, be aware that some firmware 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. This is the case PReP on iSeries systems.

- A `root` partition (500 MB - 5.0 GB) — this is where `/` (the root directory) is located. In this setup, all files (except those stored in `/boot`) are on the root partition.

A 500 MB partition allows you to install a minimal installation, while a 5.0 GB root partition lets you perform a full installation, choosing all package groups.


2.14.5. Adding Partitions

To add a new partition, select the **New** button. A dialog box appears (refer to Figure 2-14).



Note

You must dedicate at least two partitions for this installation, and optionally more. For more information, refer to Appendix C *An Introduction to Disk Partitions*.



The screenshot shows the 'Add Partition' dialog box. It has the following fields and options:

- Mount Point:** A text field containing '/boot' and a pull-down arrow.
- File System Type:** A pull-down menu showing 'ext3'.
- Allowable Drives:** A list box containing one entry: 'hda 9539 MB WDC WD100BB-75AUA1' with a checked checkbox.
- Size (MB):** A text field containing '100' and a pull-down arrow.
- Additional Size Options:** A group box containing three radio buttons:
 - Fixed size
 - Fill all space up to (MB): A text field containing '1'.
 - Fill to maximum allowable size
- Force to be a primary partition
- Buttons: 'Cancel' and 'OK' at the bottom right.

Figure 2-14. 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 pull-down menu to choose the correct mount point for your partition.
- **File System Type:** Using the pull-down menu, select the appropriate file system type for this partition. For more information on file system types, refer to Section 2.14.5.1 *File System Types*.
- **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 where you need them, or let **Disk Druid** decide where partitions should go.
- **Size (MB):** Enter the size (in megabytes) of the partition. Note, this field starts with 100 MB; unless changed, only a 100 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 is created as a logical partition. Refer to

Section C.1.3 *Partitions within Partitions — An Overview of Extended Partitions*, for more information.

- **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.

2.14.5.1. File System Types

Red Hat Enterprise Linux allows you to create different partition types, based on the file system they will use. The following is a brief description of the different file systems available, and how they can be utilized.

- **ext2** — An ext2 file system supports standard Unix file types (regular files, directories, symbolic links, etc). It provides the ability to assign long file names, up to 255 characters.
- **ext3** — The ext3 file system is based on the ext2 file system and has one main advantage — journaling. Using a journaling file system reduces time spent recovering a file system after a crash as there is no need to `fsck`¹ the file system. The ext3 file system is selected by default and is highly recommended.
- **physical volume (LVM)** — Creating one or more physical volume (LVM) partitions allows you to create an LVM logical volume. LVM can improve performance when using physical disks. For more information regarding LVM, refer to the *Red Hat Enterprise Linux System Administration Guide*.
- **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 *Red Hat Enterprise Linux System Administration 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. Refer to the *Red Hat Enterprise Linux System Administration Guide* for additional information.
- **vfat** — The VFAT file system is a Linux file system that is compatible with Microsoft Windows long filenames on the FAT file system.
- **PPC PReP Boot** — The PPC PReP boot partition contains the operating system kernel (for booting from the *NWSSSTG IPL Source on iSeries) or the YABOOT boot loader (which allows your pSeries system to boot Red Hat Enterprise Linux).

1. The `fsck` application is used to check the file system for metadata consistency and optionally repair one or more Linux file systems.

2.14.6. 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 can only change the partition's mount point. To make any other changes, you must delete the partition and recreate it.

2.14.7. Deleting a Partition

To delete a partition, highlight it in the **Partitions** section and click the **Delete** button. Confirm the deletion when prompted.

2.15. Network Configuration

If you do not have a network device, physical LAN card, or virtual LAN, this screen does not appear during your installation and you should advance to Section 2.16 *Firewall Configuration*.




Figure 2-15. Network Configuration

The installation program automatically detects any network devices you have and displays them in the **Network Devices** list.

Once you have selected a network device, click **Edit**. From the **Edit Interface** pop-up screen, you can choose to configure the IP address and Netmask of the device via DHCP (or manually if DHCP is not selected) and you can choose to activate the device at boot time. If you select **Activate on boot**, your network interface is 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.

Figure 2-16 shows a dialog box titled "Edit Interface eth0". It contains the following elements:

- Section: **Configure eth0**
- Checkbox: Configure using DHCP
- Checkbox: Activate on boot
- IP Address: 192 . 0 . 2 . 0
- Netmask: 255 . 255 . 255 . 0
- Buttons: Cancel, OK

Figure 2-16. Editing a Network Device



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.

If you have a hostname (fully qualified domain name) for the network device, you can choose to have DHCP (Dynamic Host Configuration Protocol) automatically detect it or you can manually enter the hostname in the field provided.

Finally, if you entered the IP and Netmask information manually, you may also enter the Gateway address and the Primary, Secondary, and Tertiary DNS addresses.



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`.

**Tip**

To change your network configuration after you have completed the installation, use the **Network Administration Tool**.

Type the `system-config-network` command in a shell prompt to launch the **Network Administration Tool**. If you are not root, it prompts you for the root password to continue.

2.16. Firewall Configuration

Red Hat Enterprise Linux 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 2-17. Firewall Configuration

Next, you can decide whether to enable a firewall for your Red Hat Enterprise Linux system.

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.

Enable firewall

If you choose **Enable firewall**, connections are not accepted by your system (other than the default settings) that are not explicitly defined by you. By default, only connections in response to outbound requests, such as DNS replies or DHCP requests, are allowed. If access to services running on this machine is needed, you can choose to allow specific services through the firewall.

If you are connecting your system to the Internet, this is the safest option to choose.

Next, select which services, if any, should be allowed to pass through the firewall.

Enabling these options allow the specified services to pass through the firewall. Note, these services may *not* be installed on the system by default. Make sure you choose to enable any options that you may need.

Remote Login (SSH)

Secure *Shell* (SSH) is a suite of tools for logging in to 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.

Web Server (HTTP, HTTPS)

The HTTP and HTTPS protocols are 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 must install the `httpd` package if you want to serve webpages.

File Transfer (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 must install the `vsftpd` package in order to publicly serve files.

Mail Server (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 Internet Service Provider'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.

**Note**

By default, the Sendmail mail transport agent (MTA) does not accept network connections from any host other than the local computer. To configure Sendmail as a server for other clients, you must edit `/etc/mail/sendmail.mc` and change the `DAEMON_OPTIONS` line to also listen on network devices (or comment out this option entirely using the `dnl` comment delimiter). You must then regenerate `/etc/mail/sendmail.cf` by running the following command (as root):

```
make -C /etc/mail
```

You must have the `sendmail-cf` package installed for this to work.

Additionally, you can now setup *SELinux* (Security Enhanced Linux) during your installation of Red Hat Enterprise Linux.

SELinux allows you to provide granular permissions for all subjects (users, programs, and processes) and objects (files and devices). You can safely grant an application only the permissions it needs to do its function.

The SELinux implementation in Red Hat Enterprise Linux is designed to improve the security of various server daemons while minimizing the impact on the day-to-day operations of your system.

Three states are available for you to choose from during the installation process:

- **Disable** — Select **Disable** if you do not want SELinux security controls enabled on this system. The **Disabled** setting turns enforcing off and does not set up the machine for the use of a security policy.
- **Warn** — Select **Warn** to be notified of any denials. The **Warn** state assigns labels to data and programs, and logs them, but does not enforce any policies. The **Warn** state is a good starting place for users who eventually want a fully active SELinux policy, but who first want to see what effects the policy would have on their general system operation. Note that users selecting the **Warn** state may notice some false positive and negative notifications.
- **Active** — Select **Active** if you want SELinux to act in a fully active state. The **Active** state enforces all policies, such as denying access to unauthorized users for certain files and programs, for additional system protection. Choose this state only if you are sure that your system can still properly function with SELinux fully enabled.

For additional information about SELinux, refer to the following URLs:

- <http://www.redhat.com/docs/>
- <http://www.nsa.gov/selinux/>

**Tip**

To change your security configuration after you have completed the installation, use the **Security Level Configuration Tool**.

Type the `system-config-securitylevel` command in a shell prompt to launch the **Security Level Configuration Tool**. If you are not root, it prompts you for the root password to continue.

2.17. Language Support Selection

You 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 is the language used on the system once the installation is complete. Typically, the default language is the language you selected to use during the installation.

If you choose to install other languages during this installation, 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 saves significant disk space.

**Caution**

If you select only one language, you can only use that specified language after the installation is complete.

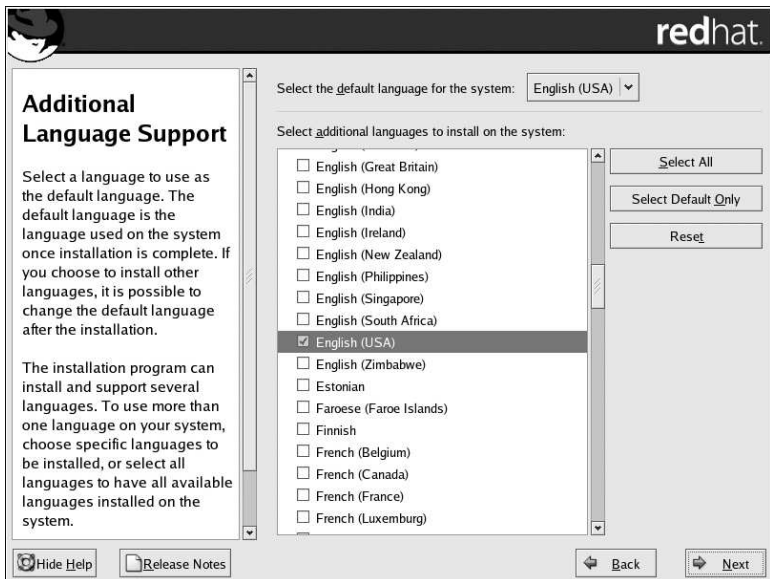


Figure 2-18. Language Support Selection

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 Enterprise Linux system.

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



Tip

To change the language configuration after you have completed the installation, use the **Language Configuration Tool**.

Type the `system-config-language` command in a shell prompt to launch the **Language Configuration Tool**. If you are not root, it prompts you for the root password to continue.

2.18. Time Zone Configuration

Set your time zone by selecting the city closest to your computer's physical location.

There are two ways for you to select your time zone:

- Using your mouse, click on the interactive map to select a specific city (represented by a yellow dot). A red X appears indicating your selection.
- You can also scroll through the list at the bottom of the screen to select your time zone. Using your mouse, click on a location to highlight your selection.

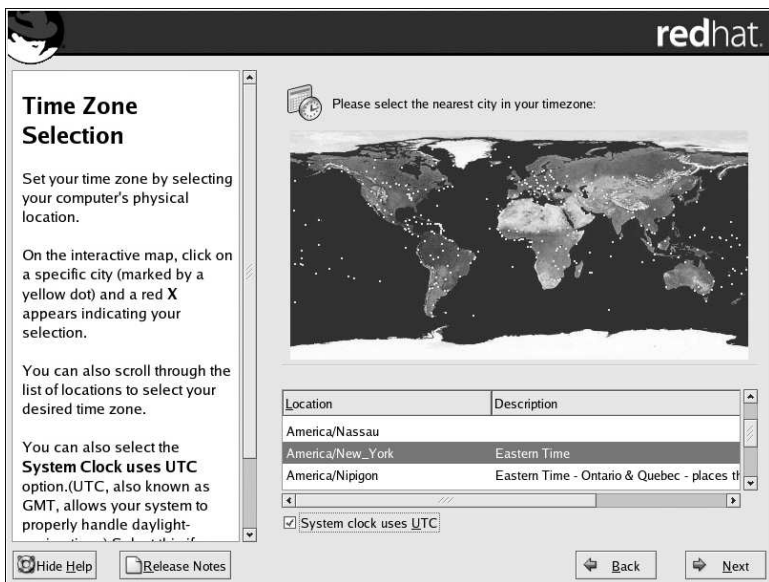


Figure 2-19. Configuring the Time Zone

Select **System Clock uses UTC** if you know that your system is set to UTC.



Note

iSeries users must select **System Clock uses UTC**.

**Tip**

To change your time zone configuration after you have completed the installation, use the **Time and Date Properties Tool**.

Type the `system-config-date` command in a shell prompt to launch the **Time and Date Properties Tool**. If you are not root, it prompts you for the root password to continue.

To run the **Time and Date Properties Tool** as a text-based application, use the command `timeconfig`.

2.19. Set 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.

**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.




Figure 2-20. Root Password

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 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 prompts you to set a root password² for your system. *You cannot proceed to the next stage of the installation process without entering a root password.*

2. A root password is the administrative password for your Red Hat Enterprise Linux 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.

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 asks 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 *420BMTNT*, 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.



Tip

To change your root password after you have completed the installation, use the **Root Password Tool**.

Type the `system-config-rootpassword` command in a shell prompt to launch the **Root Password Tool**. If you are not root, it prompts you for the root password to continue.

2.20. Package Group Selection

Now that you have made most of the choices for your installation, you are ready to confirm the default package selection or customize packages for your system.

The **Package Installation Defaults** screen appears and details the default package set for your Red Hat Enterprise Linux installation. This screen varies depending on the version of Red Hat Enterprise Linux you are installing.

If you choose to accept the current package list, skip ahead to Section 2.21 *Preparing to Install*.

To customize your package set further, select **Customize the set of packages to be installed** option on the screen. Clicking **Next** takes you to the **Package Group Selection** screen.

You can select package groups, which group components together according to function (for example, **X Window System** and **Editors**), individual packages, or a combination of the two.



Note

Users who want support for developing or running 64-bit applications are encouraged to select the **Compatibility Arch Support** and **Compatibility Arch Development Support** packages to install architecture specific support for their systems.

To select a component, click on the checkbox beside it (refer to Figure 2-21).

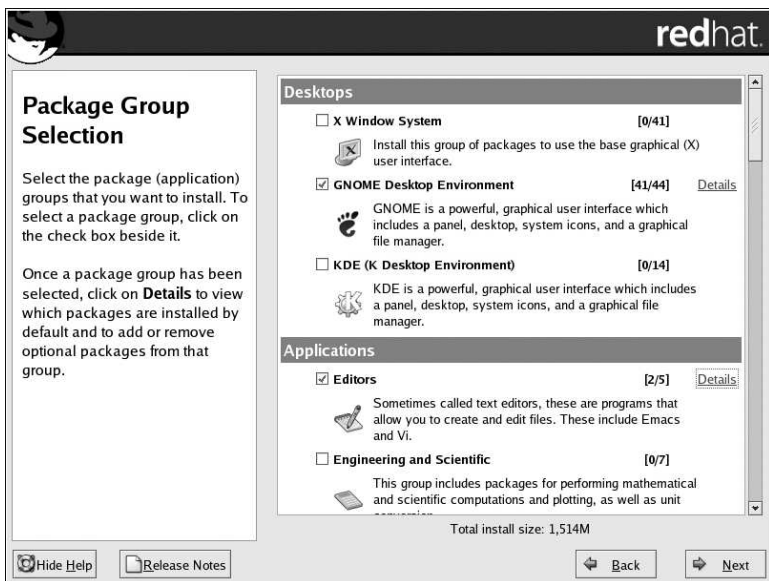


Figure 2-21. Package Group Selection

Select each component you wish to install. Selecting **Everything** (at the end of the component list) installs all packages included with Red Hat Enterprise Linux.

Once a package group has been selected, click on **Details** to view which packages are installed by default, and to add or remove optional packages from that group.

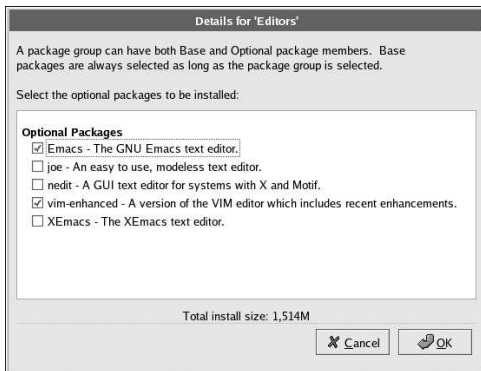


Figure 2-22. Package Group Details

2.21. Preparing to Install

A screen preparing you for the installation of Red Hat Enterprise Linux now appears.

For your reference, a complete log of your installation can be found in `/root/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 are written and packages are installed. If you wish to abort the installation, you should reboot now before any existing information on any hard drive is rewritten.

To cancel the installation process at this time, use the following instruction methods:

- iSeries users should vary off the NWS at the green screen

- Non-iSeries users should press the computer's Reset button or use the [Control]-[Alt]-[Delete] key combination to restart your machine

2.22. Installing Packages

At this point there is nothing left for you to do until all the packages have been installed. How quickly this happens depends on the number of packages you have selected and your computer's speed.

2.23. Installation Complete

Congratulations! Your Red Hat Enterprise Linux installation is now complete!

The installation program prompts you to prepare your system for reboot.

OpenPower, p5, i5, and pSeries — Completing the Installation

Do not forget to remove any boot media.

After rebooting, you must set the open firmware boot device to the disk containing your Red Hat Enterprise Linux PReP and / partitions. To accomplish this, wait until the LED indicator or HMC SRC says `E1F1`, then press [1] to enter the System Management Services GUI. Click on **Select Boot Options**. Select **Select Boot Devices**. Select **Configure 1st Boot Device**. Select the disk containing Red Hat Enterprise Linux. Set the other devices as you wish. Then exit the SMS menus to boot your new system.



Note

Steps in the SMS menu may be different depending on the machine model.

After your computer's normal power-up sequence has completed, YABOOT's prompt appears, at which you can do any of the following things:

- Press [Enter] — causes YABOOT's default boot entry to be booted.
- Select a boot label, followed by [Enter] — causes YABOOT to boot the operating system corresponding to the boot label. (Press [Tab] for non-iSeries systems at the `boot:` prompt for a list of valid boot labels.)
- Do nothing — after YABOOT's timeout period, (by default, five seconds) YABOOT automatically boots the default boot entry.

Once Red Hat Enterprise Linux has booted, one or more screens of messages should scroll by. Eventually, a `login:` prompt or a GUI login screen (if you installed the X Window System and chose to start X automatically) appears.

iSeries — Completing the Installation

When the installation program completes, it prompts you to press the [Enter] key. The LPAR then reboots with the installed Red Hat Enterprise Linux OS.

The installation program has completed the following tasks:

- Installed a kernel to the B slot and kernel parameters **ro root=LABEL=**
- Installed a kernel to the PPC PReP boot partition, if one exists
- The A slot has been left alone for the users discretion



Tip

iSeries users may also find the following information about kernel slots useful:

http://www-1.ibm.com/servers/eserver/iseries/linux/tech_faq.html

You should now vary off your partition and configure the NWSDD for running Red Hat Enterprise Linux, rather than installing it. You should change your IPL Source to either B or *NWSSTG. If you change it to *NWSSTG (booting off of your PPC PReP boot partition) you also need to add IPL Parameters. Now when you vary on your NWSDD it boots up to your installed system.

Note that users installing exclusively to native SCSI disks cannot boot from *NWSSTG.

Remember that the only place you can log in as root is at the console or other devices as listed in `/etc/securetty`. If you want to log in as root from remote systems, use `ssh`.

The first time you start your Red Hat Enterprise Linux system in run level 5 (the graphical run level), the **Setup Agent** is presented, which guides you through the Red Hat Enterprise Linux configuration. Using this tool, you can set your system time and date, install software, register your machine with Red Hat Network, and more. The **Setup Agent** lets you configure your environment at the beginning, so that you can get started using your Red Hat Enterprise Linux system quickly.

For more information on using the **Setup Agent**, refer to the chapter titled *Getting Started in the Red Hat Enterprise Linux Step By Step Guide*.

For information on registering your Red Hat Enterprise Linux subscription, refer to Section 2.24 *Activate Your Subscription*.

2.24. Activate Your Subscription

Before you can access service and software maintenance information, and the support documentation included in your subscription, you must activate your subscription by registering with Red Hat. Registration includes these simple steps:

- Provide a Red Hat login
- Provide a subscription number
- Connect your system

The first time you boot your installation of Red Hat Enterprise Linux, you are prompted to register with Red Hat using the **Setup Agent**. If you follow the prompts during the **Setup Agent**, you can complete the registration steps and activate your subscription.

If you can not complete registration during the **Setup Agent** (which requires network access), you can alternatively complete the Red Hat registration process online at <http://www.redhat.com/register/>.

2.24.1. Provide a Red Hat Login

If you do not have an existing Red Hat login, you can create one when prompted during the **Setup Agent** or online at:

<https://www.redhat.com/apps/activate/newlogin.html>

A Red Hat login enables your access to:

- Software updates, errata and maintenance via Red Hat Network
- Red Hat technical support resources, documentation, and Knowledgebase

If you have forgotten your Red Hat login, you can search for your Red Hat login online at:

https://rhn.redhat.com/help/forgot_password.pxt

2.24.2. Provide Your Subscription Number

Your subscription number is located in the package that came with your order. If your package did not include a subscription number, your subscription was activated for you and you can skip this step.

You can provide your subscription number when prompted during the **Setup Agent** or by visiting <http://www.redhat.com/register/>.

2.24.3. Connect Your System

The Red Hat Network Registration Client helps you connect your system so that you can begin to get updates and perform systems management. There are three ways to connect:

1. During the **Setup Agent** — Check the **Send hardware information** and **Send system package list** options when prompted.
2. After the **Setup Agent** has been completed — From **Applications** (the main menu on the panel), go to **System Tools**, then select **Red Hat Network**.
3. After the **Setup Agent** has been completed — Enter the following command from the command line as the root user:

- `/usr/bin/up2date --register`

2.25. Rescue Mode on POWER Systems

You can use the installation disks in rescue mode, in case your system does not boot. Rescue mode gives you access to the disk partitions on your system so you can make any changes necessary to rescue your installation.

To boot in rescue mode on iSeries or i5 with power control NWSD, follow the instructions for booting the installation program, with this change: In the NWSD, set the IPL parameters field to `'rescue'` (including the quotes), or to `'dd rescue'` if you need to load the SCSI driver. On other systems, specify the `rescue` or `dd rescue` (without quotes) after the default kernel name at the YABOOT prompt.

After the Language Selection screen (Section 2.10 *Language Selection*), the installation program attempts to mount the disk partitions on your system. It then presents you with a shell prompt where you can make the changes you need. These changes may include storing the kernel and command line into the IPL source, as described in the Installation Complete section (Section 2.23 *Installation Complete*).

When your changes are complete, you can exit the shell using `exit 0`. This causes a reboot from the C side. To reboot from the A or B side or from `*NWSSTG`, you should vary off the system instead of exiting the shell.

2.25.1. Special Considerations for Accessing the SCSI Utilities from Rescue Mode

If your system uses Native DASD disks, you may need access to the SCSI utilities from rescue mode. These utilities are located on the driver disc CD. The driver disc CD cannot be mounted from rescue mode unless special steps are taken. These steps are described below.

If you have a second CD-ROM drive assigned to your Linux system, you can mount the driver disc CD in the second drive.

If you have only one CD-ROM drive, you must set up an NFS boot, using the following steps:

1. *For OpenPower, p5, i5, and pSeries:*

Boot from the CD-ROM with the `linux rescue askmethod` command. This allows you to manually select NFS as the source of your rescue media instead of defaulting to the CD-ROM drive.

2. *For iSeries:*

Copy `/qopt/red_hat/ppc/iSeries/boot.img` to the IFS of the iSeries system containing your Linux system. You can do this with a command like:

```
COPY OBJ ('/qopt/red_hat/ppc/iSeries/boot.img') TODIR('/tmp')
```

Note that `/tmp` can be replaced with any location on your IFS.

3. Copy the first installation disc onto a file system of another Linux system.
4. Make this copy of the installation disc available through NFS or FTP.
5. Vary off or power down the system you need to rescue. Set its IPL parameters as instructed for booting the Installation discs in rescue mode, except that the IPL source should point to the copy of `boot.img` on your IFS (from step 1, above).
6. *Make sure the installation disc is not in your CD-ROM drive.*
7. IPL the Linux system.
8. Follow the prompts as described in Section 2.25 *Rescue Mode on POWER Systems*. An additional prompt for the installation source appears. Select NFS or FTP (as appropriate) and complete the following network configuration screen.
9. When the Linux system has booted into rescue mode, the CD-ROM drive is available for use and you can mount the driver media to access the SCSI utilities.

Appendix A.

Upgrading Your Current System

This appendix explains the various methods available for upgrading your Red Hat Enterprise Linux system.

A.1. Determining Whether to Upgrade or Re-Install

Although upgrades are supported from Red Hat Enterprise Linux 3, you are more likely to have a consistent experience by backing up your data and then installing this release of Red Hat Enterprise Linux 4 over your previous Red Hat Enterprise Linux installation.

This recommended reinstallation method helps to ensure the best system stability possible.

For more information about re-installing your Red Hat Enterprise Linux system, refer to the Whitepapers available online at <http://www.redhat.com/solutions/info/whitepapers/>.

If you currently use Red Hat Enterprise Linux 3, you can perform a traditional, installation program-based upgrade.

However, before you chose to upgrade your system, there are a few things you should keep in mind:

- Individual package configuration files may or may not work after performing an upgrade due to changes in various configuration file formats or layouts.
- If you have one of Red Hat's layered products (such as the Cluster Suite) installed, it may need to be manually upgraded after the Red Hat Enterprise Linux upgrade has been completed.
- Third party or ISV applications may not work correctly following the upgrade.

Upgrading your system installs the modular 2.6.x kernel as well as updated versions of the packages which are currently installed on your system.

The upgrade process preserves existing configuration files by renaming them with an `.rpmsave` extension (for example, `sendmail.cf.rpmsave`). The upgrade process also creates a log of its actions in `/root/upgrade.log`.



Caution

As software evolves, configuration file formats can change. It is very important to carefully compare your original configuration files to the new files before integrating your changes.

**Note**

It is always a good idea to back up any data that you have on your systems. For example, if you are upgrading or creating a dual-boot system, you should back up any data you wish to keep on your hard drive(s). Mistakes do happen and can result in the loss of all of your data.

Some upgraded packages may require the installation of other packages for proper operation. If you choose to customize your packages to upgrade, you may be required to resolve dependency problems. Otherwise, the upgrade procedure takes care of these dependencies, but it may need to install additional packages which are not on your system.

Depending on how you have partitioned your system, the upgrade program may prompt you to add an additional swap file. If the upgrade program does not detect a swap file that equals twice your RAM, it asks you if you would like to add a new swap file. If your system does not have a lot of RAM (less than 128 MB), it is recommended that you add this swap file.

A.2. Upgrading Your System

The **Upgrade Examine** screen appears if you have instructed the installation program to perform an upgrade.

**Note**

If the contents of your `/etc/redhat-release` file have been changed from the default, your Red Hat Enterprise Linux installation may not be found when attempting an upgrade to Red Hat Enterprise Linux 4.

You can relax some of the checks against this file by booting with the following boot command:

```
linux upgradeany
```

Use the `linux upgradeany` command if your Red Hat Enterprise Linux installation was not given as an option to upgrade.

To perform an upgrade, select **Perform an upgrade of an existing installation**. Click **Next** when you are ready to begin your upgrade.

To re-install your system, select **Perform a new Red Hat Enterprise Linux installation** and refer to <http://www.redhat.com/docs/wp/> as well as Chapter 2 *Installing Red Hat Enterprise Linux* for further instructions.

To perform a new installation of Red Hat Enterprise Linux on your system, select **Perform a new Red Hat Enterprise Linux installation** and refer to Chapter 2 *Installing Red Hat Enterprise Linux* for further instructions.

A.3. Upgrading Packages

At this point, there is nothing left for you to do until all the packages have been upgraded or installed.

A.4. Upgrade Complete

Congratulations! Your Red Hat Enterprise Linux 4 upgrade is now complete!

The installation program prompts you to prepare your system for reboot.

For information on registering your Red Hat Enterprise Linux subscription, refer to Section 2.24 *Activate Your Subscription*.

Appendix B.

Troubleshooting Your Installation of Red Hat Enterprise Linux

This appendix discusses some common installation problems and their solutions.

B.1. You are Unable to Boot Red Hat Enterprise Linux

B.1.1. Is Your System Displaying Signal 11 Errors?

A signal 11 error, commonly known as a *segmentation fault*, means that the program accessed a memory location that was not assigned.

If you receive a fatal signal 11 error 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. On non-iSeries systems, running the hardware diagnostics that were supplied with your system should help pinpoint any hardware problems (such as CPU cache or memory errors).

Ensure that you have the latest installation updates and images 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 firmware. You could also try to swap your memory around in the motherboard slots to check if the problem is either slot or memory related.

You can also try running the installation with only 256 MB of memory. This can be done by booting the installation program with the `mem=256M` boot option. To try this option, at the installation boot prompt, type:

```
mem=xxxM
```

where `xxx` should be replaced with the amount of memory in megabytes.

Note, iSeries system users must remember to add boot options, as are listed in this manual for other operating systems, by first logging into the hosting OS/400 partition as an iSeries user and then setting any IPL Parameters within the NWSD for their partition.

Another option is to perform a media check on your installation CD-ROMs. The Red Hat Enterprise Linux installation program has the ability to test the integrity of the installation media. It works with the CD, DVD, hard drive ISO, and NFS ISO installation methods. Red Hat recommends that you test all installation media before starting the installation process, and before reporting any installation-related bugs (many of the bugs reported are

actually due to improperly-burned CDs). To use this test, type the following command at the `boot:` prompt:

```
linux mediacheck
```

For more information concerning signal 11 errors, refer to:

<http://www.bitwizard.nl/sig11/>

B.2. Trouble Beginning the Installation

B.2.1. Is Your Mouse Not Detected?

This information does not apply to users of iSeries systems or users of LPAR-able pSeries systems.

If the **Mouse Not Detected** screen (refer to Figure B-1) 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 continue with the GUI installation, you are prompted to provide the installation program with your mouse configuration information.



Figure B-1. Mouse Not Detected

B.3. Trouble During the Installation

B.3.1. No devices found to install Red Hat Enterprise Linux Error Message

If you receive an error message stating `No devices found to install Red Hat Enterprise Linux`, there is probably a SCSI controller that is not being recognized by the installation program.

Check your hardware vendor's website to determine if a driver diskette image is available that fixes your problem. Second, add `dd` to the IPL parameters line in the Network Storage Description. For more general information on driver diskettes, refer to Appendix D *Driver Media*.

You can also refer to the *Red Hat Hardware Compatibility List*, available online at:

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

B.3.2. Saving Traceback Messages Without a Diskette Drive

If you receive a traceback error message during installation, you can usually save it to a diskette.

If you do not have a diskette drive available in your system, you can `scp` the error message to a remote system.

This information does not apply to users of iSeries systems or headless pSeries systems.

When the traceback dialog appears, the traceback error message is automatically written to a file named `/tmp/anacdump.txt`. Once the dialog appears, switch over to a new tty (virtual console) by pressing the keys `[Ctrl]-[Alt]-[F2]` and `scp` the message written to `/tmp/anacdump.txt` to a known working remote system.

B.3.3. Trouble with Partition Tables

If you receive an error after the **Disk Partitioning Setup** (Section 2.12 *Disk Partitioning Setup*) phase of the installation saying something similar to

The partition table on device `hda` was unreadable. To create new partitions it must be initialized, causing the loss of ALL DATA on this drive.

you may not have a partition table on that drive or the partition table on the drive may not be recognizable by the partitioning software used in the installation program.

No matter what type of installation you are performing, backups of the existing data on your systems should always be made.

B.3.4. Other Partitioning Problems for IBM™ POWER System Users

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 / (root) partition
- A <swap> partition of type swap
- A PPC PReP Boot partition.
- For non-iSeries users only, a /boot/ partition.



Tip

When defining a partition's type as swap, you do not have to assign it a mount point. **Disk Druid** automatically assigns the mount point for you.

B.3.5. Are You Seeing Python Errors?

During some upgrades or installations of Red Hat Enterprise Linux, 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 the /tmp/directory. 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/>

The **anaconda** website may also be a useful reference and can be found online at:

<http://rhlinux.redhat.com/anaconda/>

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/>

B.4. Problems After Installation

B.4.1. Unable to IPL from *NWSSTG

If you are experiencing difficulties when trying to IPL from `*NWSSTG`, you may not have created a PReP Boot partition set as `active`.

B.4.2. Booting into a Graphical Environment

This information does not apply to users of iSeries systems.

If you have installed the X Window System but are not seeing a graphical desktop environment once you log into your Red Hat Enterprise Linux system, you can start the X Window System graphical interface using the command `startx`.

Once you enter this command and press [Enter], the graphical desktop environment is displayed.

Note, however, that this is just a one-time fix and does not change the log in process for future log ins.

To set up your system so that you can log in at a graphical login screen, you must edit one file, `/etc/inittab`, by changing just one number in the runlevel section. When you are finished, reboot the computer. The next time you log in, you are presented with a graphical login prompt.

Open a shell prompt. If you are in your user account, become root by typing the `su` command.

Now, type `gedit /etc/inittab` to edit the file with **gedit**. The file `/etc/inittab` opens. Within the first screen, a section of the file which looks like the following appears:

```
# Default runlevel. The runlevels used by RHS are:
# 0 - halt (Do NOT set initdefault to this)
# 1 - Single user mode
# 2 - Multiuser, without NFS (The same as 3, if you do not have networking)
# 3 - Full multiuser mode
# 4 - unused
# 5 - X11
# 6 - reboot (Do NOT set initdefault to this)
#
id:3:initdefault:
```

To change from a console to a graphical login, you should change the number in the line `id:3:initdefault:` from a 3 to a 5.



Warning

Change *only* the number of the default runlevel from 3 to 5.

Your changed line should look like the following:

```
id:5:initdefault:
```

When you are satisfied with your change, save and exit the file using the [Ctrl]-[Q] keys. A window appears and asks if you would like to save the changes. Click **Save**.

The next time you log in after rebooting your system, you are presented with a graphical login prompt.

B.4.3. Problems with the X Window System (GUI)

If you are having trouble getting X (the X Window System) to start, you may not have installed it during your installation.

If you want X, you can either install the packages from the Red Hat Enterprise Linux CD-ROMs or perform an upgrade.

If you elect to upgrade, select the X Window System packages, and choose GNOME, KDE, or both, during the upgrade package selection process.

B.4.4. Problems with the X Server Crashing and Non-Root Users

If you are having trouble with the X server crashing when anyone other than root logs in, you may have a full file system (or, a lack of available hard drive space).

To verify that this is the problem you are experiencing, run the following command:

```
df -h
```

The `df` command should help you diagnose which partition is full. For additional information about `df` and an explanation of the options available (such as the `-h` option used in this example), refer to the `df` man page by typing `man df` at a shell prompt.

A key indicator is 100% full or a percentage above 90% or 95% on a partition. The `/home/` and `/tmp/` partitions can sometimes fill up quickly with user files. You can make some room on that partition by removing old files. After you free up some disk space, try running X as the user that was unsuccessful before.

B.4.5. Problems When You Try to Log In

If you did not create a user account in the **Setup Agent**, log in as root and use the password you assigned to root.

If you cannot remember your root password, boot your system as `linux single`.

Once you have booted into single user mode and have access to the `#` prompt, you must type `passwd root`, which allows you to enter a new password for root. At this point you can type `shutdown -r now` to reboot the system with the new root 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 the graphical login screen does not appear, check your hardware for compatibility issues. The *Hardware Compatibility List* can be found at:

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

B.4.6. Your Printer Does Not Work

If you are not sure how to set up your printer or are having trouble getting it to work properly, try using the **Printer Configuration Tool**.

Type the `system-config-printer` command at a shell prompt to launch the **Printer Configuration Tool**. If you are not root, it prompts you for the root password to continue.

B.4.7. Apache-based httpd service/Sendmail Hangs During Startup

If you are having trouble with the Apache-based `httpd` service or Sendmail hanging at startup, make sure the following line is in the `/etc/hosts` file:

```
127.0.0.1 localhost.localdomain localhost
```

Appendix C.

An Introduction to Disk Partitions



Note

This appendix is not necessarily applicable to non-x86-based architectures. However, the general concepts mentioned here may apply.

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 pre-installed 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 Enterprise Linux installation is as simple and painless as possible.

If you are reasonably comfortable with disk partitions, you could skip ahead to Section C.1.4 *Making Room For Red Hat Enterprise Linux*, for more information on the process of freeing up disk space to prepare for a Red Hat Enterprise Linux installation. This section also discusses the partition naming scheme used by Linux systems, sharing disk space with other operating systems, and related topics.

C.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, this appendix uses a simplified diagram of a disk drive to help explain what is really happening when a disk drive is partitioned. Figure C-1, shows a brand-new, unused disk drive.



Figure C-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 is adequate. 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. . .

C.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 *file system*") writes information to the drive, creating order out of the empty space in an unformatted drive.




Figure C-2. Disk Drive with a File System

As Figure C-2, implies, the order imposed by a file system involves some trade-offs:

- A small percentage of the drive's available space is used to store file system-related data and can be considered as overhead.
- A file system splits the remaining space into small, consistently-sized segments. For Linux, these segments are known as *blocks*.¹

Given that file systems make things like directories and files possible, these trade-offs are usually seen as a small price to pay.

It is also worth noting that there is no single, universal file system. As Figure C-3, shows, a disk drive may have one of many different file systems written on it. As you might guess, different file systems tend to be incompatible; that is, an operating system that supports one file system (or a handful of related file system types) may not support another. This last statement is not a hard-and-fast rule, however. For example, Red Hat Enterprise Linux

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.

supports a wide variety of file systems (including many commonly used by other operating systems), making data interchange between different file systems easy.




Figure C-3. Disk Drive with a Different File System

Of course, writing a file system 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 C-4. Disk Drive with Data Written to It

As Figure C-4, shows, some of the previously-empty blocks are now holding data. However, by just looking at this picture, we cannot determine exactly how many files reside on this drive. There may only be one file or many, 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.

C.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 file systems were never designed to support anything above a certain capacity. Or the file systems *could* support larger drives with a greater capacity, but the overhead imposed by the file system 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 file system or user data. But for clarity, they are separate in our diagrams.




Figure C-5. Disk Drive with Partition Table

As Figure C-5 shows, the partition table is divided into four sections or four *primary* partitions. A primary partition is a partition on a hard drive that can contain only one logical drive (or section). 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" is 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 file system 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.

By this point, you might be wondering how all this additional complexity is normally used. Refer to Figure C-6, for an example.




Figure C-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 C-1, it is adequate for the purposes of this discussion.

Table C-1, contains a listing of some popular (and obscure) partition types, along with their hexadecimal numeric values.

Partition Type	Value	Partition Type	Value
Empty	00	Novell Netware 386	65
DOS 12-bit FAT	01	PIC/IX	75
XENIX root	02	Old MINIX	80
XENIX usr	03	Linux/MINUX	81
DOS 16-bit <=32M	04	Linux swap	82
Extended	05	Linux native	83
DOS 16-bit >=32	06	Linux extended	85
OS/2 HPFS	07	Amoeba	93
AIX	08	Amoeba BBT	94
AIX bootable	09	BSD/386	a5
OS/2 Boot Manager	0a	OpenBSD	a6
Win95 FAT32	0b	NEXTSTEP	a7
Win95 FAT32 (LBA)	0c	BSDI fs	b7
Win95 FAT16 (LBA)	0e	BSDI swap	b8
Win95 Extended (LBA)	0f	Syrinx	c7
Venix 80286	40	CP/M	db
Novell	51	DOS access	e1

Partition Type	Value	Partition Type	Value
PPC PReP Boot	41	DOS R/O	e3
GNU HURD	63	DOS secondary	f2
Novell Netware 286	64	BBT	ff

Table C-1. Partition Types

C.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 C-1, 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 C-7, shows a disk drive with one primary partition and one extended partition containing two logical partitions (along with some unpartitioned free space).




Figure C-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, due to the way in which partitions are accessed in Linux, you should avoid defining more than 12 logical partitions on a single disk drive.

Now that we have discussed partitions in general, let us review how to use this knowledge to install Red Hat Enterprise Linux.

C.1.4. Making Room For Red Hat Enterprise Linux

The following list presents some possible scenarios you may face when attempting to re-partition your hard disk:

- Unpartitioned free space is available
- An unused partition is available

Let us look at each scenario in order.



Note

Keep in mind that the following illustrations are simplified in the interest of clarity and do not reflect the exact partition layout that you encounter when actually installing Red Hat Enterprise Linux.

C.1.4.1. 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 C-8, shows what this might look like.




Figure C-8. Disk Drive with Unpartitioned Free Space

In Figure C-8, 1 represents an undefined partition with unallocated space and 2 represents a defined partition with allocated 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.

Next, we will discuss a slightly more common situation.

C.1.4.2. 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 C-9, illustrates such a situation.




Figure C-9. Disk Drive With an Unused Partition

In Figure C-9, 1 represents an unused partition and 2 represents reallocating an unused partition for Linux.

If you find yourself in this situation, you can use the space allocated to the unused partition. You first must delete the partition and then create the appropriate Linux partition(s) in its place. You can delete the unused partition and manually create new partitions during the installation process.

C.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 file system contained on that partition.

Red Hat Enterprise Linux uses a naming scheme that is more flexible and conveys more information than the approach used by other operating systems. The naming scheme is file-based, with file names in the form of `/dev/xxyn`.

Here is how to decipher the partition naming scheme:

/dev/

This 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, usually 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 Enterprise Linux. Of course, this does not mean that Red Hat Enterprise Linux can access data on every type of partition, but in many cases it is possible to access data on a partition dedicated to another operating system.

Keep this information in mind; it makes things easier to understand when you are setting up the partitions Red Hat Enterprise Linux requires.

C.1.6. Disk Partitions and Other Operating Systems

If your Red Hat Enterprise Linux partitions are 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.

C.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` is mounted on `/usr/`, that would mean that all files and directories under `/usr/` 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`.

C.1.8. How Many Partitions?

At this point in the process of preparing to install Red Hat Enterprise Linux, you must 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: `swap`, `/boot/` (a `/boot/` file is not necessary for iSeries systems, but is needed for all others), PPC PReP Boot, and `/` (root). Note that new virtual disks can be added dynamically on OS/400 release V5R2 and later.

For more information, refer to Section 2.14.4 *Recommended Partitioning Scheme*.

Appendix D.

Driver Media

D.1. Why Do I Need Driver Media?

While the Red Hat Enterprise Linux installation program is loading, a screen may appear asking you for driver media. The driver media screen is most often seen in the following scenarios:

- If you need to perform an installation from a network device
- If you need to perform an installation from a block device
- If you need to perform an installation from a PCMCIA device
- If you run the installation program by entering `linux dd` at the installation boot prompt or in the IPL Parameter field of the NWS D
- If you run the installation program on a computer which does not have any PCI devices

D.1.1. So What Is Driver Media Anyway?

Driver media can add support for hardware that may or may not be supported by the installation program. Driver media could include a driver diskette or image produced by Red Hat, it could be a diskette or CD-ROM you make yourself from driver images found on the Internet, or it could be a diskette or CD-ROM that a hardware vendor includes with a piece of hardware.

Driver media is used if you need access to a particular device to install Red Hat Enterprise Linux. Drivers can be used for network (NFS) installations, installations using a PCMCIA or block device, non-standard or very new CD-ROM drives, SCSI adapters, NICs, and other uncommon devices.



Note

If an unsupported device is not needed to install Red Hat Enterprise Linux on your system, continue with the installation and add support for the new piece of hardware once the installation is complete.

D.1.2. How Do I Obtain Driver Media?

Driver images can be obtained from several sources. They may be included with Red Hat Enterprise Linux, or they may be available from a hardware or software vendor's website. If you suspect that your system may require one of these drivers, you should create a driver diskette or CD-ROM before beginning your Red Hat Enterprise Linux installation.



Tip

It is also possible to use a driver image via a network file. Instead of using the `linux dd` boot command, use the `linux dd=url` command, where `url` is replaced by an HTTP, FTP, or NFS address of the driver image to be used.

Another option for finding specialized driver 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 Enterprise Linux that does not work with drivers already in the installation program or included on the driver images on the Red Hat Enterprise Linux CD #1. In such cases, the Red Hat website may contain a link to a driver image.

D.1.3. Using a Driver Image During Installation

If you need to use a driver image, such as during a PCMCIA device or NFS installation, the installation program prompts you to insert the driver (as a diskette, CD-ROM, or file name) when it is needed.

However, there are some cases where you must specifically tell the Red Hat Enterprise Linux installation program to load that driver diskette and use it during the installation process.

For example, to specifically load a driver diskette that you have created, begin the installation process by booting from the Red Hat Enterprise Linux CD #1 (or using boot media you have created). At the `boot:` prompt or in the IPL Parameter field of the Linux Partition's NWSD, enter `linux dd`. Refer to Chapter 2 *Installing Red Hat Enterprise Linux* for details on booting the installation program.

The Red Hat Enterprise Linux installation program asks you to insert the driver diskette. Once the driver diskette is read by the installation program, it can apply those drivers to hardware discovered on your system later in the installation process.

Appendix E.

Additional Boot Options

This appendix discusses additional boot and kernel boot options available for the Red Hat Enterprise Linux installation program.

To use any of the boot options presented here, type the command you wish to invoke at the installation `boot:` prompt.

iSeries system users must remember to add boot options, as are listed in this manual for other operating systems, by first logging into the hosting OS/400 partition as an iSeries user and then setting any IPL Parameters within the NWSD for their partition.

Boot Time Command Arguments

`askmethod`

This command asks you to select the installation method you would like to use when booting from the Red Hat Enterprise Linux CD-ROM.

`dd`

This argument causes the installation program to prompt you to use a driver diskette.

`dd=url`

This argument causes the installation program to prompt you to use a driver image from a specified HTTP, FTP, or NFS network address.

`display=IP:0`

This command allows remote display forwarding. In this command, *IP* should be replaced with the IP address of the system on which you want the display to appear.

On the system you want the display to appear on, you must execute the command `xhost +remotehostname`, where *remotehostname* is the name of the host from which you are running the original display. Using the command `xhost +remotehostname` limits access to the remote display terminal and does not allow access from anyone or any system not specifically authorized for remote access.

`driverdisk`

This command performs the same function as the `dd` command and also prompts you to use a driver diskette during the installation of Red Hat Enterprise Linux.

`ide=nodma`

This command disables DMA on all IDE devices and may be useful when having IDE-related problems.

`mediacheck`

This command gives you the option of testing the integrity of the install source (if an ISO-based method). This command works with the CD, DVD, hard drive ISO, and NFS ISO installation methods. Verifying that the ISO images are intact before you attempt an installation helps to avoid problems that are often encountered during an installation.

`mem=xxxM`

This command allows you to override the amount of memory the kernel detects for the machine. This may be needed for some older systems where only 16 MB is detected and for some new machines where the video card shares the video memory with the main memory. When executing this command, `xxx` should be replaced with the amount of memory in megabytes.

`nopass`

This command disables the passing of keyboard and mouse information to stage 2 of the installation program. It can be used to test keyboard and mouse configuration screens during stage 2 of the installation program when performing a network installation.

`nopcmcia`

This command ignores any PCMCIA controllers in system.

`noprobe`

This command disables hardware detection and instead prompts the user for hardware information.

`noshell`

This command disables shell access on virtual console 2 during an installation.

`nousb`

This command disables the loading of USB support during the installation. If the installation program tends to hang early in the process, this command may be helpful.

`nousbstorage`

This command disables the loading of the `usbstorage` module in the installation program's loader. It may help with device ordering on SCSI systems.

rescue

This command runs rescue mode. Refer to the *Red Hat Enterprise Linux System Administration Guide* for more information about rescue mode.

resolution=

Tells the installation program which video mode to run. It accepts any standard resolution, such as **640x480**, **800x600**, **1024x768**, and so on.

serial

This command turns on serial console support.

text

This command disables the graphical installation program and forces the installation program to run in text mode.

updates

This command prompts you to insert a floppy diskette containing updates (bug fixes). It is not needed if you are performing a network installation and have already placed the updates image contents in `RHupdates/` on the server.

vnc

This command allows you to install from a VNC server.

vncpassword=

This command sets the password used to connect to the VNC server.

Index

Symbols

- /boot/ partition
 - recommended partitioning, 31
- /root/install.log
 - install log file location, 48
- /var/ partition
 - recommended partitioning, 31

A

- accessibility, v
- activating your subscription, 51
- adding partitions, 32
 - file system type, 34
- ATAPI CD-ROM
 - unrecognized, problems with, 17
- automatic partitioning, 23, 25

B

- boot options
 - additional, 79
 - kernel, 79
 - linux mediacheck, 8
- booting the installation program
 - OpenPower, p5, i5, and pSeries, 10
- booting the installation program on an iSeries system, 11

C

- canceling the installation, 16
- CD-ROM
 - ATAPI, 16
 - unrecognized, problems with, 17
 - IDE, 16
 - unrecognized, problems with, 17
 - installation from, 16
 - SCSI, 16

- clock, 43
- configuration
 - clock, 43
 - network, 35
 - time, 43
 - time zone, 43
- console, virtual
 - iSeries, 13
- consoles, virtual, 12
- conventions
 - document, i

D

- DASD installation, 17
- Disk Druid
 - adding partitions, 32
 - file system type, 34
 - buttons, 30
 - deleting partitions, 35
 - editing partitions, 35
 - partitions, 27
- disk partitioning, 23
- disk space, 4
- diskette
 - media, 77
- documentation
 - other manuals, 1
- driver media, 77
 - produced by others, 78
 - produced by Red Hat, 78
 - using a driver image, 78

E

- extended partitions, 72

F

- feedback
 - contact information for this manual, v
- file system
 - formats, overview of, 68
- file system types, 34
- firewall configuration, 38
 - customize incoming services, 39
 - customize trusted services, 39
 - security levels
 - enable firewall, 39
 - no firewall, 39
- FTP
 - installation, 5, 19

H

- hard disk
 - basic concepts, 67
 - extended partitions, 72
 - file system formats, 68
 - partition introduction, 69
 - partition types, 71
 - partitioning of, 67
- hard drive installation, 17
 - preparing for, 7
- hardware preparation, eServer i5, 1
- hardware preparation, eServer OpenPower, 1
- hardware preparation, eServer p5, 1
- hardware preparation, iSeries, 3
- hardware preparation, pSeries, 4
- HMC vterm, 13
- hostname configuration, 37
- how to use this manual, iv
- HTTP
 - installation, 5, 20

I

- IDE CD-ROM
 - unrecognized, problems with, 17
- install log file
 - /root/install.log, 48
- installation

- aborting, 16
- can you install with a CD-ROM, 4
- CD-ROM, 16
- DASD, 17
- disk space, 4
- FTP, 5, 19
- GUI
 - CD-ROM, 9
 - hard drive, 7, 17
 - HTTP, 5, 20
 - keyboard navigation, 15
 - network, 5
 - NFS, 5, 18
 - server information, 19
 - partitioning, 27
 - problems
 - IDE CD-ROM related, 17
 - program
 - graphical user interface, 9
 - text mode user interface, 13
 - virtual consoles, 12
 - starting, 16
- installation media
 - testing, 8
- installing packages, 46
- introduction, i
- iSeries
 - booting the installation program, 11
 - iSeries virtual console, 13

K

- kernel
 - boot options, 79
- keyboard
 - configuration, 22
 - navigating the installation program using, 15
- keymap
 - selecting type of keyboard, 22

L

- language
 - selecting, 22
 - support for multiple languages, 41

M

- manuals, 1
- mount points
 - partitions and, 75
- mouse
 - not detected, 60

N

- network
 - configuration, 35
 - installations
 - FTP, 19
 - HTTP, 20
 - NFS, 18
 - network installation
 - performing, 18
 - preparing for, 5
- NFS
 - installation, 5, 18

O

- online help
 - hiding, 21

P

- packages
 - groups, 46
 - selecting, 46
 - installing, 46
 - selecting, 46
- partition
 - extended, 72
- partitioning, 27

- automatic, 23, 25
- basic concepts, 67
- creating new, 32
 - file system type, 34
- deleting, 35
- editing, 35
- extended partitions, 72
- how many partitions, 70, 76
- introduction to, 69
- making room for partitions, 73
- mount points and, 75
- naming partitions, 74
- numbering partitions, 74
- other operating systems, 75
- primary partitions, 70
- recommended, 31
- types of partitions, 71
- using free space, 73
- using unused partition, 74
- password
 - setting root, 44
- POWER systems rescue mode, 52
 - accessing SCSI utilities, 52

R

- re-installation, 55
- recursion
 - (see recursion)
- registering your subscription, 51
- rescue mode, POWER systems, 52
 - accessing SCSI utilities, 52
- root / partition
 - recommended partitioning, 31
- root password, 44

S

- security levels
 - firewall configuration, 38
 - SELinux, 40
- selecting
 - packages, 46
- SELinux
 - security levels, 40
- starting
 - installation, 16
- steps
 - disk space, 4
 - eServer i5 hardware preparation, 1
 - eServer OpenPower hardware preparation, 1
 - eServer p5 hardware preparation, 1
 - installing with CD-ROM, 4
 - iSeries hardware preparation, 3
 - pSeries hardware preparation, 4
- subscription registration, 51
- swap file
 - upgrade, 56
- swap partition
 - recommended partitioning, 31

T

- TCP/IP configuration, 18
- time zone
 - configuration, 43
- traceback messages
 - saving traceback messages without a floppy drive, 61
- troubleshooting, 59
 - after the installation
 - Apache-based httpd service hangs during startup, 66
 - after the installation, 63
 - booting into a graphical environment, 64
 - booting into GNOME or KDE, 64
 - booting into the X Window System, 64
 - logging in, 65
 - printers, 66
 - Sendmail hangs during startup, 66
 - unable to IPL from *NWSSTG, 63

- X (X Window System), 65
 - X server crashes, 65
- beginning the installation, 60
 - mouse not detected, 60
- booting, 59
 - signal 11 error, 59
- CD-ROM failure
 - CD-ROM verification, 8
- during the installation
 - No devices found to install Red Hat Enterprise Linux error message, 61
 - saving traceback messages without a diskette drive, 61
- during the installation, 61
 - completing partitions, 62
 - partition tables, 61
 - Python errors, 62

U

- upgrade, 55
 - adding a swap file, 56
- user interface, graphical
 - installation program, 9
- user interface, text mode
 - installation program, 13

V

- virtual consoles, 12

Colophon

The manuals are written in DocBook SGML v4.1 format. The HTML and PDF formats are produced using custom DSSSL stylesheets and custom jade wrapper scripts. The DocBook SGML files are written using **Emacs** with the help of PSGML mode.

Garrett LeSage created the admonition graphics (note, tip, important, caution, and warning). They may be freely redistributed with the Red Hat documentation.

The Red Hat Product Documentation Team consists of the following people:

Sandra A. Moore — Primary Writer/Maintainer of the *Red Hat Enterprise Linux Installation Guide for x86, Itanium™, AMD64, and Intel® Extended Memory 64 Technology (Intel® EM64T)*; Primary Writer/Maintainer of the *Red Hat Enterprise Linux Installation Guide for the IBM® POWER Architecture*; Primary Writer/Maintainer of the *Red Hat Enterprise Linux Installation Guide for the IBM® S/390® and IBM® eServer™ zSeries® Architectures*

John Ha — Primary Writer/Maintainer of the *Red Hat Cluster Suite Configuring and Managing a Cluster*; Co-writer/Co-maintainer of the *Red Hat Enterprise Linux Security Guide*; Maintainer of custom DocBook stylesheets and scripts

Edward C. Bailey — Primary Writer/Maintainer of the *Red Hat Enterprise Linux Introduction to System Administration*; Primary Writer/Maintainer of the *Release Notes*; Contributing Writer to the *Red Hat Enterprise Linux Installation Guide for x86, Itanium™, AMD64, and Intel® Extended Memory 64 Technology (Intel® EM64T)*

Karsten Wade — Primary Writer/Maintainer of the *Red Hat SELinux Guide*; Contributing Writer to the *Red Hat Enterprise Linux System Administration Guide*

Andrius T. Benokraitis — Primary Writer/Maintainer of the *Red Hat Enterprise Linux Reference Guide*; Co-writer/Co-maintainer of the *Red Hat Enterprise Linux Security Guide*; Contributing Writer to the *Red Hat Enterprise Linux System Administration Guide*

Paul Kennedy — Primary Writer/Maintainer of the *Red Hat GFS Administrator's Guide*; Contributing Writer to the *Red Hat Cluster Suite Configuring and Managing a Cluster*

Mark Johnson — Primary Writer/Maintainer of the *Red Hat Desktop Deployment Guide*; Contributing Writer of Red Hat Network documentation

Melissa Goldin — Primary Writer/Maintainer of the *Red Hat Enterprise Linux Step By Step Guide*; Contributing Writer of Red Hat Network Documentation

Lucy Ringland — Red Hat Enterprise Linux Documentation Editor.

The Red Hat Localization Team consists of the following people:

Amanpreet Singh Alam — Punjabi translations

Jean-Paul Aubry — French translations

David Barzilay — Brazilian Portuguese translations

- Runa Bhattacharjee — Bengali translations
Chester Cheng — Traditional Chinese translations
Verena Fuehrer — German translations
Kiyoto Hashida — Japanese translations
N. Jayaradha — Tamil translations
Michelle Jiyeen Kim — Korean translations
Yelitza Louze — Spanish translations
Noriko Mizumoto — Japanese translations
Ankitkumar Rameshchandra Patel — Gujarati translations
Rajesh Ranjan — Hindi translations
Nadine Richter — German translations
Audrey Simons — French translations
Francesco Valente — Italian translations
Sarah Wang — Simplified Chinese translations
Ben Hung-Pin Wu — Traditional Chinese translations
Tongjie Tony Fu — Simplified Chinese Translations
Manuel Ospina — Spanish Translations