Installation of Red Hat Linux

Red Hat Linux is the most popular Linux distribution that is used today. The software is available both commercially as well as free on Red Hat's Web site, *http://www.redhat.com*. In this document, we discuss the installation procedure of Red Hat Linux 6.2. It's possible that you are using a different version, but the techniques will be broadly similar. The italicized references like *(23.1)* point to the sections in *Your UNIX: The Ultimate Guide*, where the concept is explained.

The installation steps have been grouped into the following categories:

A. Pre-installation Preparations
B. Preliminary Settings
C. Preparing the Disk
D. Host and Network Configuration
E. User Configuration
F. Package Installation
G. Post-install Configuration

The simplest method of installing Linux is to acquire a CD-ROM containing the software. The default installation of Red Hat occupies around 450 MB of disk space, but since you are likely to load other programs as and when you find them, have at least 1 GB of free space on your hard disk for accommodating Linux. If you can afford to provide more space, then by all means do so. This space may be available on your hard disk containing

1. no partitions, which means Linux can be installed without fear of disturbing anything as there are no partitions and data to protect.
2. another operating system (e.g., Windows) on one partition and unpartitioned space that can be used for installing Linux.
3. another operating system on a partition that uses the *entire* disk, but has free disk space available. Using a special utility, **`FIPS.EXE`** that comes with every Linux distribution, this space can be converted into a partition on which to install Linux.

Every operating system requires a separate partition, and normally, you can't use the free space of one partition to install another operating system. If your machine satisfies conditions 1 or 2, then you are ready to install Linux. Skip section A1 and move to A2. If you are in a situation that requires `FIPS`, then read the next section.

Note: If you are satisfying condition 1, but would nonetheless like to have both Windows and Linux on your machine, then install Windows first. Before you do that, create a separate partition for Windows using the **`fdisk`** command of DOS.

A1. `FIPS.EXE`: Extracting a Partition Out of an Existing One

In this scenario, your computer may already be running Windows, whose partition uses the complete disk (condition 3), but there is sufficient free space on your `C:` drive that you could use for installing Linux. You need to use the program **`FIPS.EXE`**, which you'll find in the `\DOSUTILS` directory of your Linux CD-ROM. This is a DOS program which converts the free space of your existing DOS partition to another partition.

Before you use **`FIPS.EXE`**, you'll have to defragment your partition containing Windows. You may know that already if you have been using Windows for some time; use *Start>Programs>Accessories>System Tools>Disk Defragmenter*. Once that is done, boot from a DOS disk and then run the **`FIPS.EXE`** from the CD-ROM (the settings in your `AUTOEXEC.BAT` and

CONFIG.SYS must be able to detect the CD-ROM as a separate drive (probably D:). The program displays the partition table and prompts you for the partition number to split. Eventually, FIPS will repartition the drive and make a partition available for you. You are now a user satisfying condition 2.

A2. RAWWRITE.EXE: Preparing the Boot Disk

Your CD-ROM containing the Linux distribution may or may not be bootable, and even if it is, your computer may lack the capability to boot from it. If you are in such a situation, then you'll need to use a boot disk. If this disk has not been supplied with your distribution, then you can create it by using the **RAWRITE.EXE** utility that's available in the \DOSUTILS directory of your CD-ROM. Use this DOS program to create a binary image on the boot diskette from the file available as \IMAGES\BOOT.IMG in the CD-ROM. This is a file having 1,474,560 bytes—the exact capacity of a 3.5", 1.44 MB floppy diskette.

If your CD-ROM is represented by the drive D:, then move to it and then use this command:

```
[D:\>] \DOSUTILS\RAWRITE[Enter]
Enter disk image source file name: \IMAGES\BOOT.IMG
Enter target diskette drive: A:
Please insert a formatted diskette into drive A: and press -ENTER- : [Enter]
```

Make sure that you have inserted a *formatted* diskette. After the light on the floppy drive goes out, take out the diskette and label it your boot diskette, showing the version number of the distribution. If your system boots from a floppy before the CD-ROM (if at all), then let the CD-ROM remain in the drive. You are now ready to start the installation.

A3. Booting the Machine

Now, reboot the machine with or without the diskette in the drive, depending on whether you are booting from it or the CD-ROM. You'll first be greeted with a screen full of options, followed by the boot: prompt at the bottom of the screen. Don't bother about these options except to note that at this stage you really have two choices:

- Press *[Enter]* to install Linux in the graphics mode.
- Enter the string text to use the text mode (recommended).

The graphics mode enables you to use the mouse both for navigation and selection, but in the text mode you'll have to use the keyboard for doing the same work. UNIX is essentially a command-driven system which makes extensive use of the keyboard, so we'll use the text mode. Just enter text at the prompt and then press *[Enter]*:

```
boot: text[Enter]
```

The system now goes through the booting sequence that loads the kernel and identifies the hardware devices. The system should be able to identify your hard disk and CD-ROM drive, and you should see something like this in the boot messages:

```
hda: ST36421A, ATA DISK drive
hdc: Creative CD1220E, ATAPI CDROM drive
```

These are the device names on a machine with an IDE controller. If these two essential devices have been identified by the system, the rest of the installation procedure should be smooth sailing.

If the installation messages fail to display the existence of the hard disk, which often is the case for SCSI disks, then you'll have to pass some options to the kernel at the boot: prompt. For instance, if you know your disk geometry, you can pass this information at the prompt:

```
boot: linux hd=784,255,63
```

This hard disk has 784 cylinders, 255 heads and 63 sectors/track.

A4. The Basic Navigation Rules

During the installation process, you'll encounter a number of windows with buttons on them, and

you'll also encounter dialog boxes. The following guidelines should help you interact with the machine during the installation process:

- To move from one button to another, or even among the various window elements, you'll have to use the *[Tab]* key. If you overshoot, use *[Alt][Tab]* to move in the reverse direction.
- If the cursor is on a list, then you can navigate across the list elements using the cursor Up and Down keys. Sometimes, you can do that even faster using *[PageUp]* and *[PageDown]*.
- If you have made a mistake in a previous window, you can go back to the window using the *Back* button. However, there are a few things that can't be undone in this way.
- Generally, the *OK* button takes you to the next window.

We'll now interact with Red Hat Linux and learn to handle its windows and dialog boxes. The headings in the following sections often reflect the ones you see at the top of the window. For instance, the next one entitled *Language Selection* is, in fact, the header of the first window of the installation process. In any case, keep an eye on the last line of the screen; you'll often find helpful information there.

B1. Language Selection

The first window (Exhibit A) prompts you to choose your language. By default, the highlight is on English. As explained before, you can use the cursor motion keys (Up and Down) to choose other languages. If the list is large, then you can use the *[PageUp]* and *[PageDown]* keys for faster access. After you have made your selection, press the *[Tab]* key to highlight the *OK* button. In case you find that you have to change your selection yet again, you can move the cursor to the language list by pressing *[Tab]* again. Now press *[Enter]* on the *OK* button to confirm your selection.

**********Exhibit A: Selecting the Language

B2. Keyboard Selection

The keyboard functions vary with the language you choose (Exhibit B). The default keyboard type is `us`, which you need not disturb if you have chosen English as your language. This time, you'll see the *Back* button on the window. This button lets you go back to previous screens and change your selections. Now, follow the usual procedure of hitting the *[Tab]* key to highlight the *OK* button before finally pressing *[Enter]*.

**********Exhibit B: Selecting the Keyboard

B3. Greetings from Red Hat Linux

Red Hat Software now welcomes you and advises you to look up the installation guide and register your software (in case you have purchased it) at Red Hat's site *http://www.redhat.com*.

B4. Installation Type

This screen offers you five options (Exhibit C) which determines the components of the system to be installed. GNOME and KDE govern the look and feel of your X Window desktop. You can use your Linux machine as a workstation, or even as a server. It's best, however, to install the components that are also discussed in *Your UNIX*, so choose the fourth option (*Install Custom System*) that will let you select each and every package individually.

**********Exhibit C: Selecting the Type of Installation

C1. Disk Setup

You now have to partition your disk. This is a very important part of the setup, and if you are not careful about what you do, your machine may not even boot. This window offers you to choose the method of creating or modifying partitions using two tools (Exhibit D):

1. Disk Druid
2. **`fdisk`**

All Linux flavors use the **`fdisk`** command to create a partition, but Red Hat also offers the Disk

Druid utility which is a lot easier to use. We'll discuss both these tools, but we'll start with Disk Druid (option 1). **`fdisk`**, however, gives you much greater control over your partitions, and if you want to start using it straightaway, then skip to C3A.

**********Exhibit D: Choose between Disk Druid and **`fdisk`**

C2A. Current Disk Partitions—Disk Druid (Option 1)

The screen is divided into two parts (Exhibit E). The upper portion shows the existing partitions of all disks. Here, there's just one DOS partition (`hda1`, running Windows). The lower portion (headed *Drive Summaries*) displays the attributes of the entire disk in five columns:

1. Drive—`hda` here represents the device name of the first IDE hard disk. If you had another hard disk here, then you would have seen a second entry which could show `hdb` (or even `hdc` or `hdd`) in this column. The device names have this significance:

   `/dev/hda`—Primary master (Bootable drive)

   `/dev/hdb`—Primary slave

   `/dev/hdc`—Secondary master

   `/dev/hdd`—Secondary slave

These represent the ports that hard disks and CD-ROMs connect to, so `/dev/hdc` could well be a CD-ROM rather than a hard disk. However, the primary master is always connected to the bootable hard disk, so `/dev/hda` must be a hard disk. SCSI hard disks using the prefix `sd` rather `hd`.

2. Disk Geometry—You can know the geometry of your hard drive from here. The entry `556/240/63` shows 556 cylinders, 240 heads and 63 sectors per track.

3. Total space available in the disk—This disk has a capacity of 4104 megabytes.

4. Used space—Here, it represents the space allocated to the partitions. You can see that 2045 MB has been allocated to the DOS partition.

5. Free space—This shows the space available for *creating* a partition, and doesn't in any way reflect the free space actually available in the partitions. We'll use this space of 2059 MB to create two partitions to hold our Linux system.

Using the Up and Down keys, you can move both highlights so that once a disk has been selected in the lower window, you can move the upper highlight to select the respective partition.

**********Exhibit E: Using Disk Druid

C2B Creating the Swap and root Partitions

Since we are going to create a partition and not modify or delete one, use *[Tab]* to select the *Add* button to add a partition. There are two types of partitions that we need to create:

- A *swap* partition. Every Linux system needs to have a swap partition which holds active processes when they are not running in the CPU. They are loaded back into memory from this partition when they are ready to run once again *(21.7)*.
- The *root* partition containing the entire Linux software *(21.7)*.

First, we'll create the swap partition of 64 MB (Exhibit F). To do that, you have to choose `Linux swap` as the partition type. This partition doesn't have a "mount point", i.e., it can't be accessed by users *(21.10)*.

**********Exhibit F: Creating the Swap Partition

After you hit *OK*, you'll be back to the previous window (Exhibit G), which now shows `hda2` as the second partition in the upper portion of the window. The free space now available in the disk to create the root partition reduces to 1993 MB.

**********Exhibit G: Disk Druid Window after Creating the Swap Partition

To create the root partition, use *Add* once again to see a window similar to the one used in creating the swap partition (Exhibit H). This partition has to be mounted on the / directory, so we just enter a / (frontslash, not backslash) to specify the mount point. We also have to specify 1993 as the partition size (in MB) and select its type as `Linux native`.

**********Exhibit H: Creating the root Partition

Use *[Tab]* to move between these fields, and after you have arrived at *Allowable Drives*, use the cursor motion keys to navigate between the disk drives (in case you see more than one disk drive here). Use the *[Spacebar]* as a toggle switch; every time you press it, an asterisk appears or disappears in the field representing the hard drive (`hda`). Make sure the asterisk is there in the field. Finally highlight the *OK* button and press *[Enter]* to return to the previous window.

**********Exhibit I: Disk Druid Window After Creating the root Partition

This window now shows yet another entry at the top (Exhibit I). A new Linux native partition has been created that uses / (the root directory) as the mount point. The *Drive Summaries* section now shows that we have no more space to create any partitions, so the figures we worked with were correct. Now, just press *OK* when you'll be asked to confirm to save the changes in the partition table. You can't install Linux unless the new partition table is written to disk. Now move to C4.

Note: The Linux native partition now shows the device name as `hda2`, the same name shown by the swap partition earlier. The swap partition is now shown as `hda5`. You need not bother about the significance of this device name right now *(21.4.3—Linux)*. Section C3C of this document also explains that `hda5` is a *logical* partition.

C3A. Current Disk Partitions—`fdisk` (Option 2)

`fdisk` is the standard Linux command that you can use to handle partitions. If you opt for using **`fdisk`** rather than Disk Druid, then a separate window (*Disk Setup*) appears to let you select the hard disk on which you would like to create the partitions. This will show you one or more hard drives on your disk, and you have to select the one. Under normal circumstances, it will be the first disk on your system (`/dev/hda`). This time, use the *Edit* button to move to **`fdisk`**.

You'll now see **`fdisk`**'s prompt on the screen. The cursor here is represented by the underscore (_) character:

```
Command (m for help): _
```

You'll now have to work with the internal commands of **`fdisk`**. First have a look at the list of its commands with the **`m`** command. For the purpose of installation, you need to remember these commands only:

`d` — Delete a partition
`l` — List known partition types
`n` — Add a new partition
`p` — Print the partition table
`q` — Quit without saving changes
`t` — Change a partition's system id
`w` — Write table to disk and exit

Before we create partitions, use the **`p`** command to display the existing partition table:

```
Command (m for help): p

Disk /tmp/hda: 240 heads, 63 sectors, 556 cylinders
Units = cylinders of 15128 * 512 bytes

   Device Boot    Start       End    Blocks   Id  System
/tmp/hda1   *         1       277   2094088+   6  FAT16

Command (m for help): _
```

This shows the existence of the DOS partition (FAT16). Note that one cylinder here comprises 15128 * 512 = 7.74 MB (approx.). Disk Druid doesn't let you know the details of the partitions, but **fdisk** lets you create either a *primary* or an *extended* partition *(21.4.3)*. While you can have up to four primary partitions on an IDE disk, you can have just one extended partition (the rest would be primary).

C3B. Creating a Primary Partition for the root File System

We'll first create one *primary* partition to hold the root file system. We'll next create an *extended* partition, "inside" which we can have one or more *logical* partitions. To create a new partition in the vacant disk space, use the **n** command:

```
Command (m for help): n
Command action
   e   extended
   p   primary (1-4)
p                                                    Creating a primary partition
Partition number (1-4): 2                            The second partition
First cylinder (278-556, default 278): 278
Last cylinder or +size or +sizeM or +sizeK (278-556, default 556): 547
```

fdisk lets you specify either the cylinder values or the size in kilo- or megabytes. Note that it also provides default values which you can make use of by simply pressing *[Enter]*. We have just created a Linux native primary partition.

C3C. Creating an Extended and Logical Partition for the Swap File System

We now have to create an extended partition. We have just adequate space (556 - 547 + 1 = 10 cylinders) to create a small partition. This time we'll work with default values:

```
Command (m for help): n
Command action
   e   extended
   p   primary (1-4)
e                                                    Creating an extended partition
Partition number (1-4): 3                            The third partition
First cylinder (548-556, default 548): [Enter]
Using default value 548
Last cylinder or +size or +sizeM or +sizeK (548-556, default 556): [Enter]
Using default value 556
```

The extended partition we created is of no use unless we create one or more logical partitions in it. We have adequate space here for just one swap partition:

```
Command (m for help): n
Command action
   l   logical (5 or over)                                  Note this change
   p   primary partition (1-4)
l                                                    fdisk selects partition 5
First cylinder (548-556, default 548): [Enter]
Using default value 548
Last cylinder or +size or +sizeM or +sizeK (548-556, default 556): [Enter]
Using default value 556
```

Note: The extended partition is required to hold *multiple* logical partitions, though here we created just a single and a small one. You can have some important directory structures like /var, /home and /tmp in their own separate partitions. The advantage with this approach is that when one partition becomes full, it can't encroach into another partition and occupy space there *(21.4.1)*.

C3D. The Final Partition Table

Now that we have created a logical partition, let's have a look at the final form of the partition table:

```
Command (m for help): p

Disk /tmp/hda: 240 heads, 63 sectors, 556 cylinders
Units = cylinders of 15128 * 512 bytes

   Device Boot    Start       End    Blocks   Id  System
/tmp/hda1   *         1       277   2094088+   6  FAT16
/tmp/hda2           278       547   2041200   83  Linux
/tmp/hda3           548       556     68040    5  Extended
/tmp/hda5           548       556     68008+  83  Linux
```

We see here the swap partition as `Linux` type. We need to change this to reflect a Linux swap partition. Use **`fdisk`**'s **`t`** command to change the type of the `hda5` partition:

```
Command (m for help): t
Partition number (1-5): 5
Hex code (type L to list codes): l                    Case insignificant
```

The swap partition has the code 82 and the Linux native partition has 83. In case, you want to see an exhaustive list, use the list code (**`l`**) and then make the selection at the prompt that appears at the bottom of the screen:

```
Hex code (type L to list codes): 82
Changed system type of partition 5 to 82 (Linux swap)
```

Have a look at the partition table once more; the last line should now look like this:

```
/tmp/hda5           548       556     68008+  82  Linux swap
```

And after you are satisfied that you have got things right, write the partition table to disk:

```
Command (m for help): w

The partition table has been altered!
Calling ioctl to re-read partition table
.......
```

You are now asked to select the hard drive where you'd like to install Linux. The fact that you used **`fdisk`** on the first hard disk is no indication that you'd like to load Linux there. (An expert can install multiple flavors of Linux on different hard drives.) So just select the proper drive (in case there's more than one) and press *Done*.

Note: Now that your partition table has been written to disk, you can't undo this action by pressing the *[Back]* button. It will take you back to the previous screens all right, but will always reflect the new partition table.

C3E. Current Disk Partitions

We are now in a similar window shown by Disk Druid (Exhibit E), except that the options here are rather limited. All you have to do now is to highlight the Linux native partition and then choose *Edit*. The *Edit Partition* window appears as in Disk Druid—again with limited options. Here, you select / as the mount point and press *OK*. You are back to the previous window. Quit this window now with *OK*.

C4. Choose Partitions to Format

You have to format the Linux partition that you created just now (Exhibit J). By default, you'll find

the * against the device name (`/dev/hda2`), so make sure it's there. If you are not unduly worried about the state of your hard disk, then leave the other box unchecked. This formatting operation will not check for bad blocks during the formatting process.

**********Exhibit J: Choosing Partitions to Format

C5. LILO Configuration

The first window having this name will prompt for any special kernel options that you may like to use at boot time. If you are using an IDE disk, then remove the * from the checkbox to disable linear mode and press *OK*.

The next window asks for the place you would like to have the bootloader (LILO) installed (Exhibit K). This is a small program that is first loaded into memory during the system boot, and is responsible for the eventual loading of the Linux system. There are two places from where you have to choose:

- Master Boot Record (MBR)
- First sector of root partition

If you have a single Linux system on your hard disk with or without a DOS partition, then it's most safe to choose Master Boot Record. LILO, the Linux Loader, can boot other operating systems including Windows, so your existing Windows system is safe if you install LILO.

**********Exhibit K: LILO Configuration

Note: Under certain circumstances, however, you would like to have LILO in the same partition that holds the Linux root partition (second option). If you have another operating system on your disk which is loaded by its own bootloader, then you'll have to make a decision on whether you'd like to replace that system's bootloader with the new one.If you are booting Linux from the boot diskette, then you may not like to install LILO at all.

The next window (Exhibit L) lets you configure the system for booting other operating systems. The last column contains the boot label that you have to enter at the boot prompt. Since we are having a DOS partition on our system, we must be able to boot Windows as well. The default configuration shows `dos` and `linux` as the labels, with an * against `linux`. This means that pressing *[Enter]* or `linux` at the `boot:` prompt boots Linux, and entering `dos` starts Windows. You can change these labels in which case you have to use the *Edit* button.

**********Exhibit L: Selecting Other Operating Systems to Boot

D1. Hostname Configuration

Every computer needs a name without which it can't function in a network. If you are connected in a network, then enter a name (say, `parrots`) which is unique in the network. You need to enter a fully qualified domain name (like *parrots.bird.org*)*(11.9.3)* if your network is connected to the Internet or is using the FQDN even otherwise.

D2. Network Configuration

If your machine has an Ethernet network card installed, and which has been identified by Linux, then you'll find this window appear. It will ask you for four parameters:

- IP address
- Subnet mask
- IP address of the default gateway
- IP address of the primary name server

You'll first have to enter an IP address for your machine. If a DHCP server is not allotting this address dynamically from its pool of addresses, then uncheck the box at the top. This IP address is a set of four dot-delimited numbers, where the maximum value of each octet is 255.

You can't just choose any IP address you like. If the machines in the network have IP addresses of

the form `192.168.5.x` where *x* is the variable component, then you must also follow the same form. Just make sure that you have chosen a value of *x* (1–254) which is not used by any machine in the network.

The subnet mask appears automatically in the next line, and you'll not need to disturb the default value if your network is not subnetted *(23.1)*. But, if it is, then ask the network administrator and enter the value. Specification of the IP address and the netmask generates the **`ifconfig`** statement that activates the network interface during system boot *(23.3)*. This statement is placed in one of the system's startup scripts.

If your machine is to be allowed access to other networks, or even the Internet, then one of the machines in the network will have to act as a gateway. Find out the IP address of this gateway machine from the administrator and place the entry there. This generates the **`route`** statement that creates the *routing table* to define the default gateway *(23.5.1)*. So if you use the URL *www.xml.com* on your browser, TCP/IP packets traveling between your machine and the XML site will be routed through the gateway machine.

Finally, you'll have to enter the name of the primary name server. This is a machine that converts FQDNs to IP addresses, and vice versa *(24.2)*. There may be a name server running in your network, or if you are going to connect to the Internet through a dialup line, you may have to use the ISP's name server. You may also decide to set up your own name server on this host (after reading Chapter 24), so enter this IP address accordingly. Your basic network configuration is complete.

Note: If you don't find this window appearing after entering the hostname, then either your machine doesn't have a network interface card or has one that couldn't be identified by the Linux system. In that case, you have to install the card manually and set its parameters using the tools available in Red Hat Linux.

D3. Mouse Selection

You can do without a mouse on a Linux machine, but then you can't use its X Window system. You may have a mouse connected to the serial port or a dedicated PS2 port. The words `serial` and `PS/2` are shown against each mouse type (Exhibit M). If you are using a two-button mouse, then you have the option of emulating the behavior of a three-button one. If you are already using a three-button mouse, then uncheck the box that determines whether 3-button emulation is to be performed. If you have a mouse whose make you can't find in the list, then choose `Generic`—2 or 3 Button—PS/2 or serial.

**********Exhibit M: Selecting The Mouse Type

D4. Selecting the Mouse Device

If you have selected a serial mouse, you have to select the port that your mouse is connected to. Usually, it will be either COM1 or COM2. The installation process usually identifies the port correctly, but if you feel that the default selection is not correct, then change it. Just make sure that your mouse and modem are connected to different serial ports (unless the mouse is using the PS/2 port).

D5. Time Zone Selection

Now, select your time zone. The US has different time zones for the different geographical locations, so select the one that applies to your state. If you are located outside the US, then choose your country from the list. This list mostly contains a major group (like America, Africa, Europe, etc.), under which you'll find the city you live in or one that is close to you. If you are following Greenwich Mean Time, then simply check the box that shows up before the list.

E1. Creating a Password for root

All system and network administration work has to be done from the root user account. Since you are to be the administrator of your own machine, then enter an unmeaningful string here—but twice (Exhibit N). Don't choose names that can be easily guessed; rather have a combination of upper- and lowercase letters as well as numerals. However, you must remember the password yourself, so

note it in an inaccessible document. If you forget your root password, you can't perform any administrative activity and you have to reinstall Linux.

**********Exhibit N: Setting the root Password

E2. Add a User to the System

You also need an ordinary nonprivileged account where you would like to do your normal work—like editing and compiling C or Java programs, writing XML code or developing shell and **`perl`** scripts. Create this account in this window (Exhibit O). It is recommended that you enter the full name here because that name gets added at the front of your email address. For instance, if you create a user-id "romeo" with the full name "romeo agnelli", then romeo's email address would be *"romeo agnelli" <romeo@parrots.birds.org>*. Just as for root, the password has to be entered twice.

**********Exhibit O: Creating a User Account

E3. User Account Setup

The system now shows you one line of information for the user you just created. These details are stored in the file `/etc/passwd`, where there's an entry for every user of the machine *(22.3.3)*. If this machine is to be used by several users, then as administrator you would probably like to create user accounts for others at this initial stage itself. Using *Add*, you can create additional accounts, or delete or modify them using the buttons *Delete* and *Edit*.

These things are easily done later with the **`useradd`** or **`usermod`** or **`userdel`** commands *(22.3.2 and 22.3.4)*, so we'll be content with just a single non-root user account for the time being.

E4. Authentication Configuration

Previous UNIX systems stored the user's password in the file `/etc/passwd`, which is readable by all. Though the password was stored in an encrypted manner, it was not totally secure, as techniques are available that can guess the password—at least the ones that are derived from common names. Modern UNIX systems store the password in the file `/etc/shadow` which is readable only by the root user, so we'll opt for shadow passwords.

Regarding MD5 passwords, you may or may not check the box; it's needed for added security. We'll ignore NIS as we still have some way to go.

F1. Package Group Selection

We now move to the installation of software packages that are available in the Linux distribution. The package selection window now appears (Exhibit P), showing the broad groups under which the packages are classified. Some of these groups show an * against them, which means that there are one or more packages inside which also have a * against their names. These packages would be installed unless you deselect them.

We now have the option of making a quick default installation by jumping to *OK*. This would do for most purposes, but it won't install some important tools that are discussed in the book (the reason why we decided to go in for a customized setup in the first place).

**********Exhibit P: Selecting the Software Packages

This is the list of the groups that Red Hat offers in its distribution:

Printer Support
X Window System
GNOME
Mail/WWW/News Tools
DOS/Windows Connectivity
Graphics Manipulation
Multimedia Support
Networked Workstation
Dialup Workstation

News Server
NFS Server
SMB (Samba) Server
Anonymous FTP Server
Web Server
DNS Name Server
Authoring/Publishing
Emacs
Development
Kernel Development
Utilities

We'll leave the default selections undisturbed, but we'll also select some packages that are shown unchecked in the package list. To view the detailed list, check the box in the last line to opt for selecting individual packages before using *OK*.

**********Exhibit Q: Selecting Individual Packages in Group

Now we move on to the screen that lets you select individual packages in a group (Exhibit Q). Each group shows a <+> on its left, which means that the group has individual packages to offer. You can press either *[Enter]* or *[Spacebar]* to expand the group. On expanding it, you'll find asterisks against some of the packages. You can use the toggling behavior of the *[Spacebar]* and *[Enter]* to select or deselect a package. As you move down, you'll come across the next package group, which you'll have to expand using the same techniques.

These are some of the packages that you would like to have on your system; some of them would be selected already:

Group	*Suggested Packages*
Amusements/Games	`fortune-mod`
Applications/Archiving	`dump`, `zip` and `unzip`
Applications/Communications	`dip`, `mgetty`, `minicom`
Applications/Cryptography	`gpgp`
Applications/Editors	`emacs`, `emacs-X11`, `gedit`, `gnotepad+`
Applications/Internet	`elm`, `fetchmail`, `finger`, `ftp`, `ircii`, `lynx`, `metamail`, `mutt`, `netscape-communicator`, `pine`, `rsh`, `talk`, `telnet`, `tin`, `traceroute`, `urlview`
Applications/Publishing	`enscript`, `ghostscript`, `ghostscript-fonts`, `gv`, `printtool`, `rhs-printfilters`
Applications/System	`bind-utils`, `control-panel`, `gnome-linuxconf`, `gnome-utils`, `gnorpm`, `kernelcfg`, `modemtool`, `mtools`, `netcfg`, `samba-client`, `samba-common`, `tksysv`
Applications/Text	`ispell`, `m4`
Development/Languages	`cpp`, `expect`, `perl`, `python`, `tcl`
Documentation	`apache-manual`, `nag`, `sag`, `sendmail-doc`
System Environment/Daemons	`anonftp`, `apache`, `bind`, `finger-server`, `inetd`, `iputils`, `lpr`, `nfs-utils`, `ppp`, `rsh-server`, `samba`, `sendmail-cf`, `squid`, `talk-server`, `tcp_wrappers`, `telnet-server`, `wu-ftpd`, `wvdial`
System Environment/Shells	`bash2`, `mc`, `pdksh`, `tcsh`
User Interface/Desktops	All KDE packages

F2. Package Dependencies

Many of the packages in the distribution require other packages to be installed. When you select packages in this way, this may not happen. In that case, the next screen (Exhibit R) alerts you to

these inadequacies and presents a list of the packages that additionally need to be installed. “To satisfy dependencies”, check that box that installs the unselected but required packages; these packages will be installed automatically.

**********Exhibit R: Resolving Package Dependencies

F3. X Probe Results

Linux now tries to identify your video card. In most cases, it recognizes the card correctly, but in case it can’t, it uses a generic card driver that should work for the time being. You can change this card definition after the system is installed, and when you configure X Window (with the **`xf86config`** command).

F4. Package Installation

The installation process will be logged in the file `/tmp/install.log`, which eventually will contain a list of all installed packages. The system now formats the file systems you created, and proceeds with the package installation (Exhibit S). The time required for the installation process depends on the number of packages you have selected, the speed of your CD-ROM drive and the age of your hardware. A typical installation that includes all the components listed previously takes about 15 minutes on average-quality hardware.

**********Exhibit S: Package Installation

It will be worthwhile to keep staring at this window because each package description is shown against the *Summary:* heading as it is being installed. This is something you couldn’t do at the time of installation (a strength of SuSE Linux, there), and you may even note some of them down so that you can try out the commands later. In many cases, the actual commands are similar to the package names, and it doesn’t take much time to identify them once you start using the command completion feature of the bash shell *(17.7.2)*.

G1. Creating a Boot Disk

After the packages are installed on the disk, the system performs a “post install configuration” and prompts for the creation of a boot disk (Exhibit T). A boot disk is an essential accessory that you should always have with your machine. Sometimes, the master boot record may get overwritten by another operating system, or you may not like to install LILO on the existing one which handles other operating systems. Sometimes, MBR or disk corruption may prevent the system from booting. You then have to boot the machine using this boot disk.

**********Exhibit T: Creating a Boot Disk

You can create a boot disk later too, but in case you want to do it right now, then insert a formatted disk and create your boot disk.

G2. Monitor Setup

You now have to select the type of your monitor from a list (Exhibit U). Monitors from Compaq, Dell, IBM and HP are well-represented, so it’s quite possible that you’ll find your monitor in the list. In case you don’t find it, choose *Custom*, which is the first item in the list. We’ll consider both options.

**********Exhibit U: Selecting the Monitor Type

G3. Selecting a Particular Monitor—Probing to Begin (Option 1)

Once the system knows the monitor type, it also knows its attributes like the horizontal and vertical sync rates, and you don’t have to provide them. Linux next needs to know more about the VGA card on your computer—like the amount of video memory, the color depth possible, the clockchip settings, etc. It probes the card with the **`Xconfigurator`** program (available only in Red Hat) and attempts to extract the details.

This starts the X server, and you’ll find the screen blinking a number of times. Linux may either recommend a setting or report a problem. If Linux is not able to start the X server and display a test message, you may have to customize the monitor and VGA card settings, in which case go to G4A.

G4A. Custom Monitor Setup (Option 2)

In this setup (G4A to G4D), you do all configuration yourself. Linux wants to know:

1. The vertical refresh rate. This is the frequency of refreshing the screen.

2. The horizontal sync rate. This is the number of times the scan lines are drawn on the screen every second.

Both these parameters are available in the manual that accompanied your monitor. Have these figures in front of you before you continue with this setup. Both the horizontal and vertical rates have to be selected from lists, and it's important that you make the selections correctly; a serious mistake here could damage your monitor.

If you are not sure of what you are doing, then choose the `Standard VGA` or `Super VGA` for the horizontal sync range and `50-70` for the vertical sync rate. (You can also set these values later using the **`xf86config`** program.)

G4B. Screen Configuration

Linux now needs to know the default resolution and color depth of the monitor. You may ask Linux to probe the video card and suggest the values. This may not be possible on some (older) systems, and may cause it to hang. If you are not sure that this won't happen to you, then select *Don't Probe* and move to G4C. If you select *Probe*, then **`Xconfigurator`** starts the X server as in G3. If the probing is successful, you should now jump to G5.

G4C. Video Memory

You now have to specify the amount of video memory that's available in the graphics card. This is typically around 4 MB, but it's quite common to find around 1 MB on older systems. (Note that this window may not appear at all.)

G4D. Clockchip Configuration and Probe for Clocks

Don't specify any parameters here; just use use *No Clockchip Setting*. The next screen prompts for probing the clocks. Ignore this screen too by using *Skip*. (Note that these windows may not appear at all.) If you opt for *Probe* this time, the X server will start as in G3.

G5. Color Depth and Monitor Resolution Suggested by Linux

After **`Xconfigurator`** has finished probing the video card, it may ask you to confirm the parameters that it thinks should work properly on your machine (Exhibit V). There are two parameters to consider here—the color depth and the resolution in pixels.

The color resolution depends on the memory you have in your video card. 8-bit color lets you display 256 colors, whereas 16-bit and 24-bit enable the use of 65,536 and over 16 million colors, respectively. Here, **`Xconfigurator`** has decided that it would work well with 8-bit color.

Next, you have to decide whether the suggested value of 800 X 600 for the monitor resolution would also do for you. If you want to change either parameter, then select *Let Me Choose* and move to G6. Otherwise, simply press *Use Default* and move to G7.

**********Exhibit V: Suggested Values for the Video Card

G6. Select Video Modes and Color

You have decided to set the video modes of your display. The number of horizontal and vertical pixels in the monitor are shown against the (initially) blank checkboxes (Exhibit W). You can easily obtain a resolution of 1024 X 768 pixels on a 17" monitor, but may not be able to go beyond 800 X 600 on a 14" monitor. Don't make too many selections. 800 X 600 in the 8- or 16-bit column should suffice for most purposes.

**********Exhibit W: Specifying the Video Modes and Color

Note: In any case, don't select 640 X 480 because some of the non-resizable windows won't fit. For instance, you can't use Netscape's configuration screen *(13.7.8—Fig. 13.8)* because that won't fit in 640 X 480 space.

G7. Testing the Configuration—Starting X

Linux now starts the X server and confirms your seeing a small message on the screen in graphics mode. You have also to decide whether X should start automatically when the system boots. Since, we are going to use Linux as a UNIX machine with its command line tools, you should not let X start automatically.

Note: Even if you have opted for starting X during boot, either by design or by mistake, there's nothing to worry. You can start virtual terminal sessions *(22.13.3)* on your screen by pressing *[Ctrl][Alt][F2]* or *[Ctrl][Alt][F3]*. (You can go up to *[F5]* in the default configuration.) Each of these sessions presents a `login:` prompt on the screen, which you can use to work at the shell prompt. You can return to the graphical mode by pressing *[Ctrl][Alt][F7]*.

G8. Installation Complete

"Congratulations!", says Red Hat, which means that you have completed the installation (Exhibit X). The CD-ROM is ejected from the drive automatically. If there's any disk in the floppy drive, then you should remove it too before the rebooting takes place automatically. Watch for any error messages that the system may throw up, and in all likelihood, you'll see the `login:` prompt on the screen. You are now ready to begin your UNIX experience!

**********Exhibit X: Completion of Installation