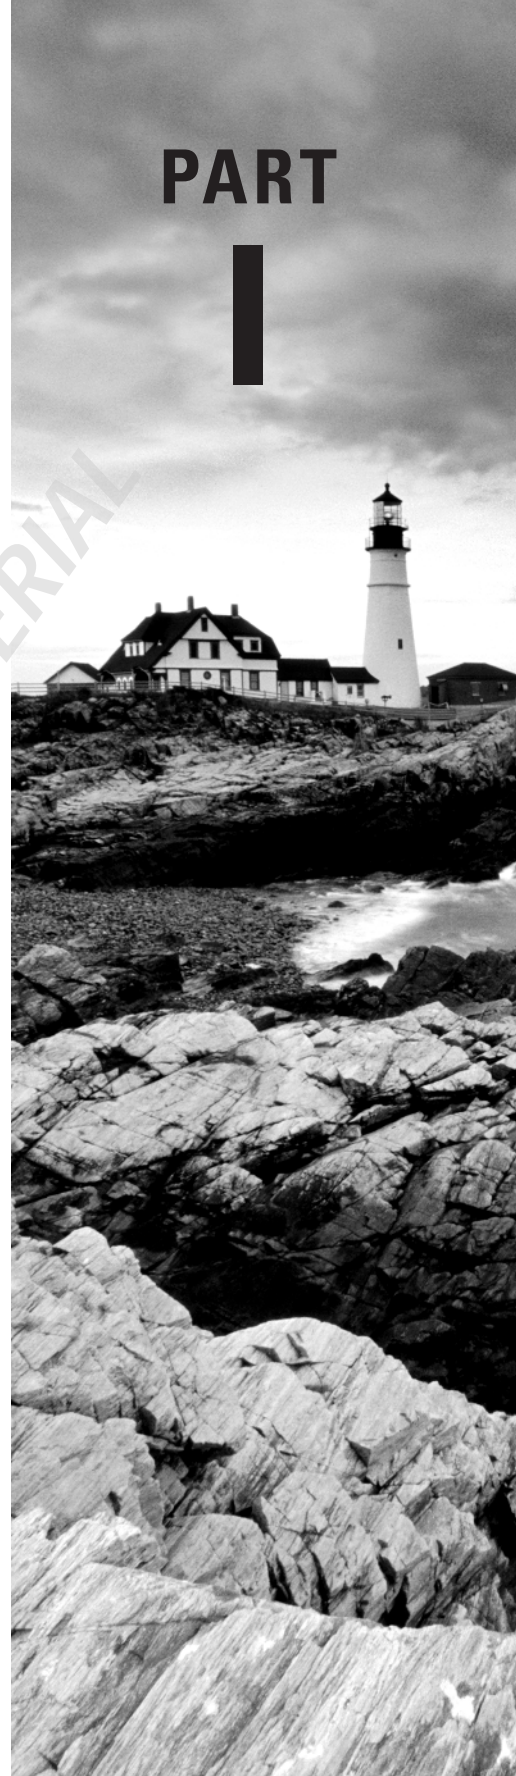


Getting Familiar with Red Hat Enterprise Linux

PART

I

COPYRIGHTED MATERIAL



Chapter

1

Getting Started with Red Hat Enterprise Linux

TOPICS COVERED IN THIS CHAPTER:

- ✓ Linux, Open Source, and Red Hat
- ✓ Red Hat Enterprise Linux and Related Products
- ✓ Installing Red Hat Enterprise Linux Server
- ✓ Exploring the GNOME User Interface





Red Hat Enterprise Linux is in use at most Fortune 500 companies, and it takes care of mission-critical tasks in many of them. This chapter introduces Red Hat Enterprise Linux. It begins with a brief history, where you'll learn about Linux in general and the role of Red Hat in the Linux story. Following that, it provides an overview of Red Hat Enterprise Linux (RHEL) and its related products. Finally, you'll learn how to install RHEL so that you can start building your RHEL skills.

Linux, Open Source, and Red Hat

If you want to work with Red Hat, it helps to understand a little bit about its background. In this introduction, you'll learn about the rise of UNIX, the Linux kernel and open source, and the founding of Red Hat.

Origins of Linux

The late 1960s and early 1970s were the dawn of the modern computing era. It was the period of proprietary stacks, where a vendor would build a “closed” computer system and create the operating software to run on it. Computers were extremely expensive and rare among businesses. In that period, scientists were still looking for the best way to operate a computer, and that included developing the best programming language. It was normal for computer programmers to address the hardware directly, using very complex assembly programming languages.

An important step forward was the development of the general-purpose programming language C by Dennis Richie at Bell Telephone Laboratories in 1969. This language was developed for use with the UNIX operating system.

The UNIX operating system was the first operating system where people from different companies tried to work together to build instead of competing with each other, keeping their efforts secret. This spirit brought UNIX to scientific, government, and higher-education institutions. There it also became the basis for the rise of another phenomenon, the Internet Protocol (IP) and the Internet. One of the huge contributors to the success of UNIX was the spirit of openness of the operating system. Everyone could contribute to it, and the specifications were freely available to anyone.

Because of the huge success of UNIX, companies started claiming parts of this operating system in the 1970s. They succeeded fairly well, and that was the beginning of the development of different flavors of UNIX, such as BSD, Sun Solaris, and HP AIX. Instead of working together, these UNIX flavors worked beside one another, with each sponsoring organization trying to develop the best version for a specific solution.

As a reaction to the closing of UNIX, Richard Stallman of MIT announced in 1984 the GNU operating system project. The goal of this project was to develop “a sufficient body of free software [...] to get along without any software that is not free.”

During the 1980s, many common Unix commands, tools, and applications were developed until, in 1991, the last gap was filled in with the launch of the Linux kernel by a student at the University of Helsinki in Finland, Linus Torvalds. The interesting fact about the Linux kernel is that it was never developed to be part of the GNU project. Rather, it was an independent initiative. Torvalds just needed a license to ensure that the Linux kernel would be free software forever, and he chose to use the GNU General Public License (GPL) for this purpose. The GPL is a *copyleft license*, which means that derived works can be distributed only under the same license terms. Using GPL made it possible to publish open source software where others could freely add to or modify lines of code.

Torvalds also made an announcement on Usenet, a very popular news network that was used to communicate information about certain projects in the early 1990s. In his Usenet message, Torvalds asked others to join him working on the Linux kernel, a challenge that was very soon taken up by many programmers around the world.

Distributions

With the adoption of the Linux kernel, finally everything that was needed to create a complete operating system was in place. There were many GNU utilities to choose from, and those tools, together with a kernel, made a complete operating system. The only thing enthusiastic users still needed to do was to gather this software, compile it from source code, and install the working parts on a computer. Because this was a rather complicated task, some initiatives started soon to provide ready-to-install Linux distributions. Among the first was MCC Interim Linux, a distribution made available for public download in February 1992, shortly after the release of the Linux kernel itself. In 1993, Patrick Volkerding released a distribution called Slackware, a distribution that could be downloaded to floppy disk images in the early days. It is still available and actively being developed today.

In 1993, Marc Ewing and Bob Young founded Red Hat, the first Linux distributor operating as a business. Since then, Red Hat has acquired other companies to integrate specific Linux-related technologies.

Red Hat went public in 1999, thus becoming the first Linux-based company on Wall Street. Because of the publicity stemming from its IPO, Red Hat and Linux received great exposure, and many companies started using it for their enterprise IT environments. It was

initially used for applications, such as intranet web servers running Apache software. Soon Linux was also used for core financial applications.

Today Linux in general and Red Hat Linux in particular is at the heart of the IT organization in many companies. Large parts of the Internet operate on Linux, using popular applications such as the Apache web server or the Squid proxy server. Stock exchanges use Linux in their real-time calculation systems, and large Linux servers are running essential business applications on top of Oracle and SAP. Linux has largely replaced UNIX, and Red Hat is a leading force in Linux.

One reason why Red Hat has been so successful since the beginning is the level of support the company provides. Red Hat offers three types of support, and this gives companies the confidence they need to run vital business applications on Linux.

The three types of Linux support provided by Red Hat are as follows:

Hardware Support Red Hat has agreements with every major server hardware vendor to make sure that whatever server a customer buys, the hardware vendor will assist them in fixing hardware issues, when Red Hat is installed on it.

Software Support Red Hat has agreements with every major enterprise software vendor to make sure that their software runs properly on top of the Red Hat Linux operating system and that the enterprise software is also guaranteed to run on Red Hat Linux by the vendor of the operating system.

Hands-on Support This means that if a customer is experiencing problems accomplishing tasks with Red Hat software, the Red Hat Global Support organization is there to help them by fixing bugs and providing technical assistance.

It is also important to realize that Red Hat is doing much more than just gathering the software pieces and putting them together on the installation media. Red Hat employs hundreds of developers who work on developing new solutions that will run on Red Hat Enterprise Linux in the near future.

Fedora

Even as Red Hat is actively developing software to be part of Red Hat Linux, it still is largely involved in the open source community. The most important approach to do this is by sponsoring the Fedora project. Fedora is a freely available Linux distribution that is completely comprised of open source software, and Red Hat is providing the funds and people to tackle this project. Both Red Hat and Fedora are free of charge; with Red Hat you pay only for updates and support.

Fedora is used as a development platform for the latest and greatest version of Linux, which is provided free of charge for users who are interested. As such, Fedora can be used as a test platform for features that will eventually be included in Red Hat Enterprise Linux. If you want to know what will be included in future versions of Red Hat Linux, Fedora is the best place to look. Also, Fedora makes an excellent choice to install on your personal computer, because it offers all the functions you would expect from a modern operating system—even some functions that are of interest only to home users.

Red Hat Enterprise Linux and Related Products

Red Hat offers several products, of which Red Hat Enterprise Linux and JBoss are the most important solutions. There are other offerings in the product catalog as well. In the following sections, you can read about these products and their typical application.

Red Hat Enterprise Linux Server Edition

The core of the Red Hat offering is Red Hat Enterprise Linux. This is the basis for two editions: a server edition and a workstation edition. The RHEL Server edition is the highly successful Red Hat product that is used in companies around the globe.



At the time of this writing, the current RHEL release is version 6.2.

With the Red Hat Enterprise Linux Server edition, there is a major new release about every three to four years. In between the major updates, there are minor ones, represented by the number after the dot in the version number. Apart from these releases, Red Hat provides patches to fix bugs and to apply security updates. Typically, these patches are applied by using the Red Hat Network, a certified collection of repositories where Red Hat makes patches available after verifying them.

To download and install repositories from the Red Hat Network (RHN), a current subscription is required. Without a current subscription, you can still run RHEL, but no updates will be installed through RHN. As an alternative to connecting each server directly to RHN, Red Hat provides a solution called Satellite. Satellite works as a proxy to RHN, and just the Satellite server is configured to fetch updates from RHN, after which the Red Hat nodes in the network connect to Satellite to access their updates. Be aware that there is also a product called RHN Proxy, which is a real caching proxy, whereas Satellite is a versioning and deployment tool.

Red Hat Enterprise Linux for Free

If you want updates and support, you have to pay for Red Hat Enterprise Linux, so how come people have to buy licenses for GPL software that is supposed to be available for free? Well, the fact is that the sources of all the software in RHEL are indeed available for free. As with any other Linux vendor, Red Hat provides source code for the software in RHEL. What customers typically buy, however, is a subscription to the compiled version of the software that is in RHEL. In the compiled version, the Red Hat logo is included.

This is more than just a logo; it's the guarantee of quality that customers expect from the leader in Linux software.

Still, the fact is that the sources of the software contained in RHEL are available for free. Some Linux distributions have used these sources to create their own distributions. The two most important distributions are CentOS (short for Community Enterprise Operating System) and Scientific Linux. Because these distributions are built upon Red Hat Linux with the Red Hat logo removed, the software is basically the same. However, small binary differences do exist, such as the integration of the software with RHN. The most important difference, however, is that these distributions don't offer the same level of support as in RHEL. So, you're better off going for the real thing.

You can download a free version of RHEL with 30 days of access to RHN at www.redhat.com. Alternatively, you can download CentOS at www.centos.org or Scientific Linux at www.scientificlinux.org.

Red Hat Enterprise Linux Workstation Edition

The other product that falls under Red Hat Enterprise Linux is the Workstation edition. This solution is based on the same code as RHEL Server. Also, the same license conditions apply for RHEL Workstation as for RHEL Server, and you need a current subscription to access and install updates from RHN. To date, Red Hat Linux Workstation hasn't experienced the same level of success as Red Hat Linux Enterprise Server.

Red Hat Add-Ons

RHEL includes everything most people need to run a Linux server. Some components require an extra effort, though, and for that reason they are offered as add-ons in RHEL. The two most significant kinds of add-on are the Enterprise File System (XFS) and Red Hat Cluster Services.

Enterprise File System (XFS) The Enterprise File System offers full scalability for large environments where many files or very large files have to be handled on large file systems. Even though ext4, the default file system in Red Hat Enterprise Linux, has been optimized significantly over time, it still doesn't fit well in environments that have very specific storage needs, such as the need to stream multimedia files or to handle hundreds of thousands of files per day.

Red Hat Cluster Services (RHCS) RHCS offers high-availability clustering to vital services in the network. In an RHCS cluster, you run specialized cluster software on multiple nodes that are involved in the cluster, and that software monitors the availability of vital services. If anything goes down with such a service, Red Hat Cluster Services takes over and makes sure that the service is launched on another node.

Red Hat Directory Server

In a corporate environment where many user accounts have to be managed, it doesn't make sense to manage these accounts in stand-alone databases on individual servers. One solution is to have servers handle their authentication on external directory servers. An example of this approach is to connect RHEL to Microsoft Active Directory, an approach that is used frequently by many Red Hat customers. Another approach is to use Red Hat Directory Server, a dedicated LDAP directory service that can be used to store and manage corporate identities.

Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization (RHEV) provides a virtualization platform that can be compared with other solutions, such as VMware vSphere. In RHEV, several dedicated servers running the KVM hypervisor are managed through RHEV-M, the management server for the virtual environment. In the RHEV infrastructure, fully installed RHEL servers as well as dedicated on-iron hypervisors (the RHEV-H) can be used. A major reason why companies around the world are using RHEV is because it offers the same functionality as VMware vSphere, but for a fraction of the price.

JBoss Enterprise Middleware

JBoss Enterprise Middleware is an application layer that can be installed on top of any operating system, including RHEL. The platform is used to build custom-made applications which can offer their services to perform any tasks you can think of. JBoss is an open platform, and therefore its adoption level is high. Red Hat has had huge success selling JBoss solutions on top of Red Hat Enterprise Linux.

Red Hat Cloud

Red Hat Cloud is the solution where everything comes together. In the lower layers of the cloud infrastructure, Red Hat can offer Platform as a Service services that are based on RHEV or any other virtualization platform. At the PaaS layer, Red Hat Cloud helps deploy virtual machines on demand easily. In the higher layers of the cloud, combined with JBoss Enterprise Middleware, Red Hat Cloud delivers software as a service, thus helping customers build a complete cloud infrastructure on top of Red Hat software.

Installing Red Hat Enterprise Linux Server

There is a version of RHEL Server for almost any hardware platform. That means you can install it on a mainframe computer, a mid-range system, or PC-based server hardware using a 64- or 32-bit architecture. Currently, the 64-bit version of Red Hat Enterprise Linux is

the most used version, and that is why, in this chapter, you can read about how to install this software version on your computer. The exact version you need is Red Hat Enterprise Linux Server for 64-bit x86_64. If you don't have the software yet, you can download a free evaluation copy at www.redhat.com.

The ideal installation is on server-grade hardware. However, you don't have to buy actual server hardware if you just want to learn how to work with Red Hat Enterprise Linux. Basically, any PC will do as long as it meets the following minimum requirements:

- A CPU capable of handling 64-bit instructions
- 1GB of RAM
- 20GB of available hard disk space
- A DVD drive
- A network card



Make sure your computer meets these minimum requirements. To work your way through the exercises in this book, I'll assume you have a computer or virtual machine that meets them.

You can run Red Hat Enterprise Linux with less than this, but if you do, you'll miss certain functionality. For instance, you can install RHEL on a machine that has 512MB of RAM, but you'll lose the graphical user interface. You could also install RHEL on a 32-bit CPU or on a VMware or VirtualBox virtual machine, but within these environments you cannot configure KVM virtualization. Because this book includes some exercises that work directly on the hard disk of your computer and you don't want to risk destroying all of your data by accident, it is strongly recommended that you do not install a dual-boot RHEL and other OS configuration.



If you don't have a dedicated computer on which to install RHEL, a virtual machine is the second-best choice. RHEL can be installed in most virtual environments. If you want to run it on your own computer, VMware Workstation (fee-based software) or VMware Player (free software but with fewer options) works fine. You can download this software from www.vmware.com. Alternatively, you can use VirtualBox, a free virtualization solution provided by Oracle. You can download it from www.virtualbox.org.

You'll be working with Red Hat Enterprise Linux in a graphical environment in this book. RHEL offers some very good graphical tools, and for now, you'll need a graphical environment to run them. A typical Linux server that provides services in a datacenter does not offer a graphical environment. Rather, it runs in console mode. That is because servers in a datacenter normally are accessed only remotely. The administrator of such a server can still use graphical tools with it but will start them over an SSH session, accessing the server remotely. Later in this book, you will learn how to configure such an environment.

In Exercise 1.1, you will install Red Hat Linux on your computer.

EXERCISE 1.1**Installing Linux on Your Machine**

This procedure describes how to install Red Hat Enterprise Linux on your computer. This is an important exercise, because you will use it to set up the demo system that you'll use throughout this book. It is important that you perform the steps exactly as described here, to match the descriptions in later exercises in this book.

To perform this exercise successfully, you'll need to install on a physical computer that meets the following requirements:

- An entire computer that can be dedicated to using Red Hat Enterprise Linux
- A minimum of 1GB of RAM (2GB is recommended)
- A dedicated hard disk of 40GB or more
- A DVD drive
- A network card

Apart from these requirements, other requirements relate to KVM virtualization as well. The most important of these is that the CPU on your computer needs virtualization support. If you can enable virtualization from the computer BIOS, you are probably OK. Read Chapter 6, "Connecting to the Network," for more details about the requirements for virtualization.

1. Put the RHEL 6 installation disc in the optical drive of your computer, and boot from the installation disc. If the DVD drive is not in the default boot order on your computer, you'll have to go into the setup and instruct your computer to boot from the optical drive. After booting from the installation DVD successfully, you'll see the Welcome to Red Hat Enterprise Linux screen.



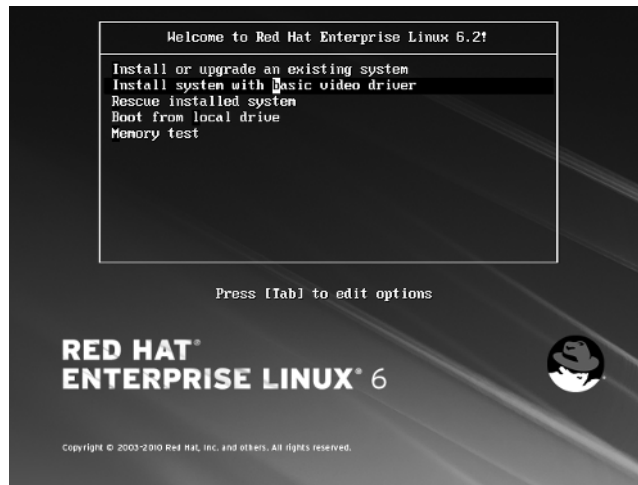
EXERCISE 1.1 (continued)

2. From the graphical installation screen, select Install Or Upgrade An Existing System. In case you're experiencing problems with the graphical display, you can choose to install using the basic video driver. However, in most cases that isn't necessary. The other options are for troubleshooting purposes only and will be discussed in later chapters in this book.
3. After beginning the installation procedure, a Linux kernel is started, and the hardware is detected. This normally takes about a minute.
4. Once the Linux kernel has been loaded, you will see a nongraphical screen that tells you that a disc was found. (Nongraphical menus like the one in the following image are referred to as *ncurses interfaces*. *Ncurses* refers to the programming library that was used to create the interface.)

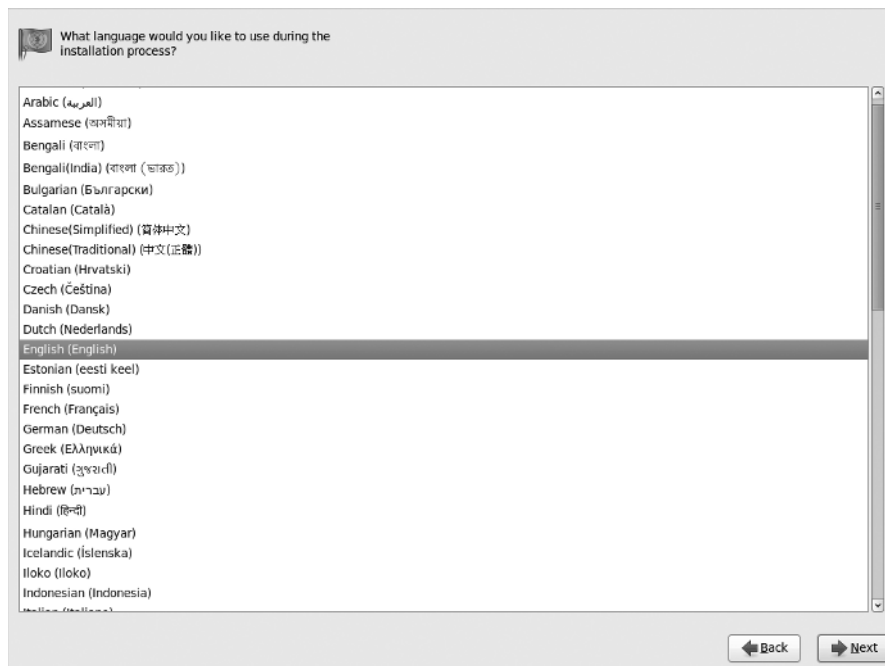


From this screen, you can start a check of the integrity of the installation media. Don't do this by default; the media check can easily take 10 minutes or more! Press the Tab key once to navigate to the Skip button, and press Enter to proceed to the next step.

5. If the graphical hardware in your computer is supported, you'll next see a graphical screen with only a Next button on it. Click this button to continue. If you don't see the graphical screen at this point, restart the installation procedure by rebooting your computer from the installation disc. From the menu, select Install System With Basic Video Driver.

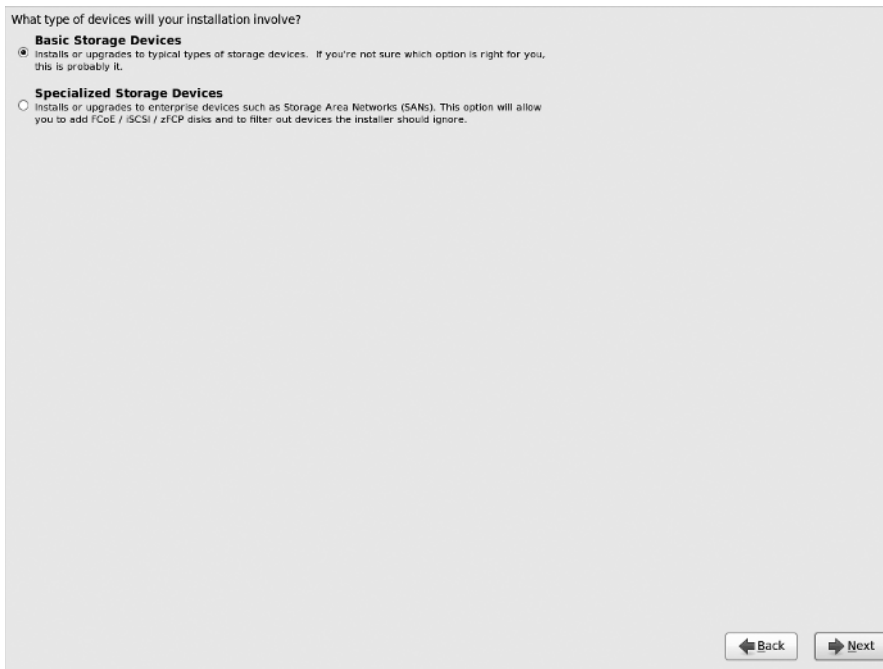
EXERCISE 1.1 (continued)

6. On the next screen, you can select the language you want to use during the installation process. This is just the installation language. At the end of the installation, you'll be offered another option to select the language you want to use on your Red Hat server. Many languages are supported; in this book I'm using English.



EXERCISE 1.1 (continued)

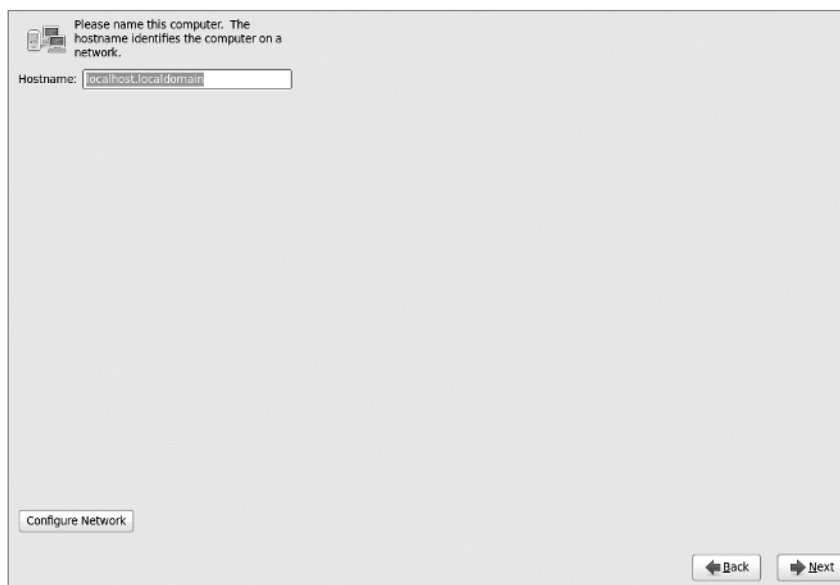
7. After selecting the installation language, on the next screen, select the appropriate keyboard layout, and then click Next to continue.
8. Once you've selected the keyboard layout you want to use, you need to select the storage devices with which you are working. To install on a local hard drive in your computer, select Basic Storage Devices. If you're installing RHEL in an enterprise environment and want to write all files to a SAN device, you should select the Specialized Storage Devices option. If you're unsure about what to do, select Basic Storage Devices and click Next to proceed.



9. After you have selected the storage device to be used, the installation program may issue a warning that the selected device may contain data. This warning is displayed to prevent you from deleting all the data on the selected disk by accident. If you're sure that the installer can use the entire selected hard disk, click Yes, and discard any data before clicking Next to continue.

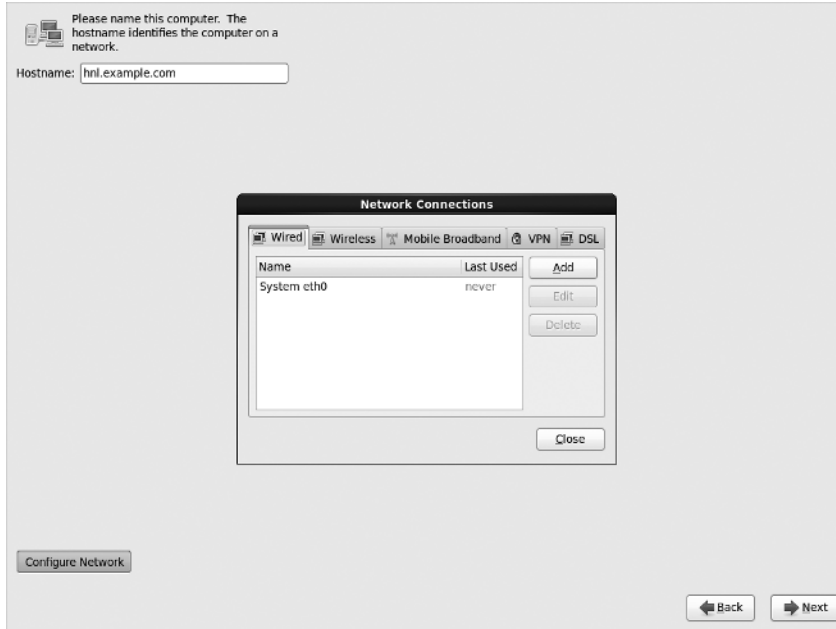
EXERCISE 1.1 (continued)

10. On the next screen, you can enter the hostname you want to use on the computer. Also on this screen is the Configure Network button, which you'll use to change the current network settings for the server. Start by entering the hostname you want to use. Typically, this is a fully qualified domain name that includes the DNS suffix. If you don't have a DNS domain in which to install the server, you can use `example.com`. This name is available for test environments, and it won't be visible to others on the Internet.
11. After setting the hostname, you have to click the Configure Network button on the same screen to change the network settings. If you don't do this, your server will be configured to get the network configuration from a DHCP server. There's nothing wrong with that if you're installing a personal desktop where it doesn't matter if the IP address it is using changes, but for servers in general, it's better to work with a fixed IP address. To set this fixed address, click Configure Network now.



EXERCISE 1.1 (continued)

12. You'll see the Network Connections window. This window comes from the NetworkManager tool, and it allows you to set and change all different kinds of network connections. In this window, select the Wired tab and, on that tab, click the System eth0 network card. Notice that depending on the hardware you are using, a different name may be used. Next click Edit to change its properties.



13. You'll now see the properties of the eth0 network card. First make sure that the option Connect Automatically is selected. If it isn't, your network card won't be activated when you boot the server.

EXERCISE 1.1 (continued)

Editing System eth0

Connection name: System eth0

Connect automatically

Wired | 802.1x Security | **IPv4 Settings** | IPv6 Settings

Method: Manual

Addresses

Address	Netmask	Gateway	
192.168.1.99	24	192.168.1.1	<input type="button" value="Add"/>
			<input type="button" value="Delete"/>

DNS servers: 192.168.1.1

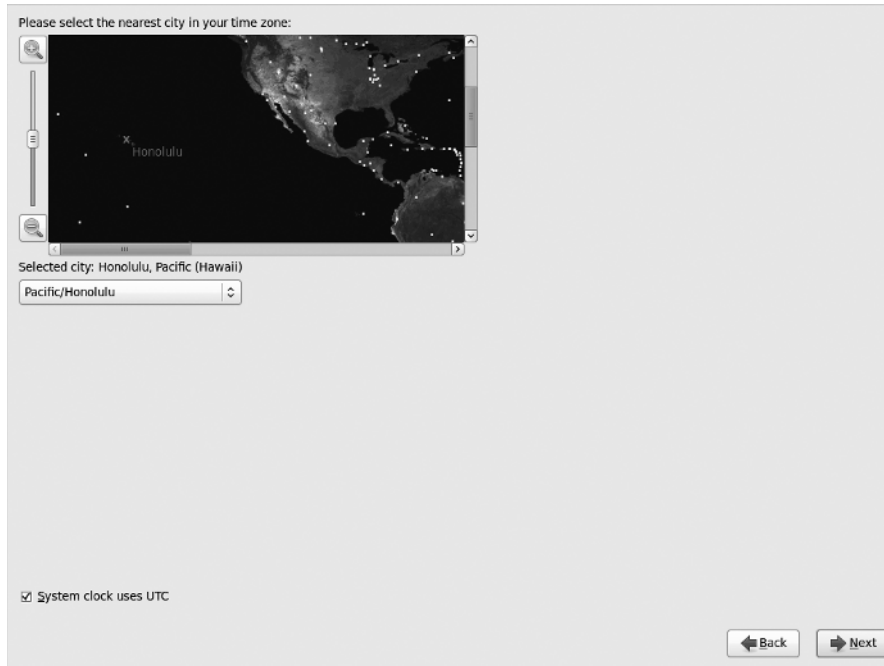
Search domains:

DHCP client ID:

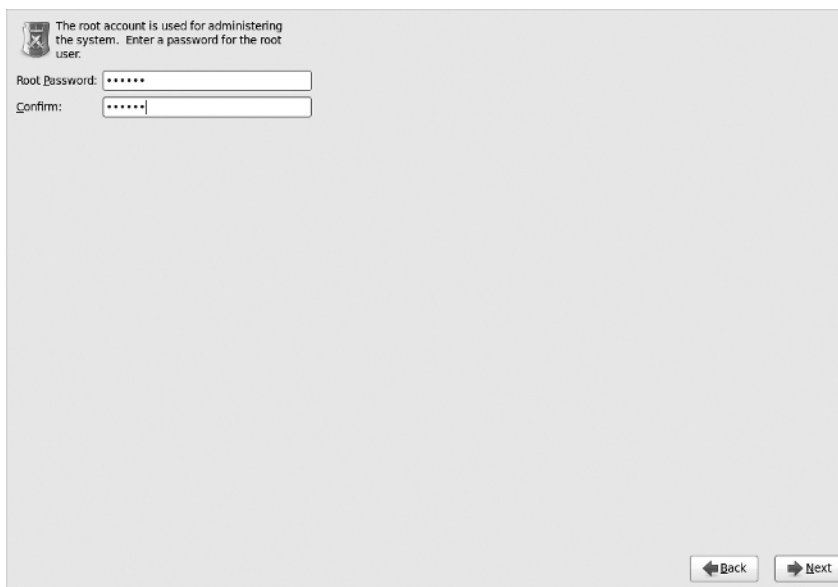
Require IPv4 addressing for this connection to complete

Available to all users

14. Select the IPv4 Settings tab, and in the Method drop-down list, select Manual.
15. Click Add to enter the IP address you want to use. You need at least an IP address and a netmask. Make sure that the address and netmask you're using here do not conflict with anything else that is in use on the network to which you are connecting. In this book I'll assume your server uses the IP address 192.168.0.70. If you want to communicate with other computers and the Internet, you'll have to enter the address of the gateway and the address of at least one DNS server. You need to consult the documentation of the network to which you're connecting to find out which addresses to use here. For the moment, you don't have to enter anything here.
16. After entering the required parameters, click Apply to save and apply these settings.
17. Click Close to close the NetworkManager window. Back on the main screen where you set the hostname, click Next to continue.
18. At this point, you'll configure the time settings for your server. The easiest way to do this is just to click the city nearest to your location on the world map that is displayed. Alternatively, you can choose the city that is nearest to you from the drop-down list.

EXERCISE 1.1 (continued)

19. You'll also need to specify whether your computer is using UTC for its internal clock. UTC is Coordinated Universal Time, a time standard by which the world regulates clocks and time. It is one of several successors to Greenwich Mean Time, without Daylight Saving Time settings. Most servers have their hardware clocks set to UTC, but most PCs don't. If the hardware clock is set to UTC, the server uses the time zone settings to calculate the local software time. If your computer has its hardware clock set to UTC, select the option System Clock Uses UTC, and click Next to continue. If not, deselect this option and proceed with the installation.
20. Next you'll specify the password that is to be used by the user root. The root account is used for system administration tasks, and its possibilities are nearly unlimited. Therefore, you should set the root password to something that's not easy for possible intruders to guess.

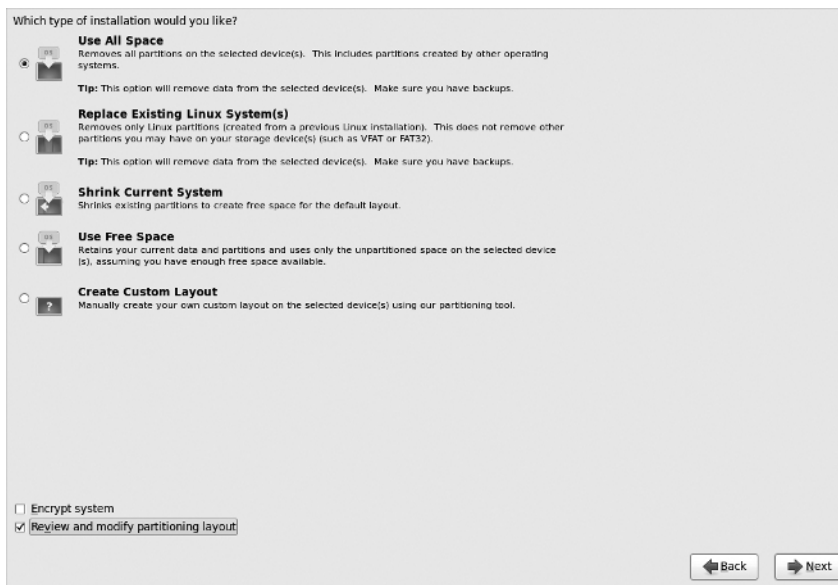
EXERCISE 1.1 (continued)

The root account is used for administering the system. Enter a password for the root user.

Root Password:

Confirm:

21. The next screen you'll see is used to specify how you'd like to use the storage devices on which you'll install Red Hat Enterprise Linux. If you want to go for the easiest solution, select Use All Space. This will remove everything currently installed on the selected hard disk (which typically isn't a bad idea anyway). Table 1.1 gives an overview of all the available options.



Which type of installation would you like?

- Use All Space**
Removes all partitions on the selected device(s). This includes partitions created by other operating systems.
Tip: This option will remove data from the selected device(s). Make sure you have backups.
- Replace Existing Linux System(s)**
Removes only Linux partitions (created from a previous Linux installation). This does not remove other partitions you may have on your storage device(s) (such as VFAT or FAT32).
Tip: This option will remove data from the selected device(s). Make sure you have backups.
- Shrink Current System**
Shrinks existing partitions to create free space for the default layout.
- Use Free Space**
Retains your current data and partitions and uses only the unpartitioned space on the selected device(s), assuming you have enough free space available.
- Create Custom Layout**
Manually create your own custom layout on the selected device(s) using our partitioning tool.

Encrypt system

Review and modify partitioning layout

EXERCISE 1.1 (continued)**TABLE 1.1:** Available storage options

Option	Description
Use All Space	Wipes everything that is currently on your computer's hard disk to use all available disk space. This is typically the best option for a server.
Replace Existing Linux System(s)	Removes existing Linux systems only if found. This option doesn't touch Windows or other partitions if they exist on your computer.
Shrink Current System	Tries to shrink existing partitions so that free space is made available to install Linux. Using this option typically results in a dual-boot computer. Using a dual-boot computer is a bad idea in general, and more specifically, this option often has problems shrinking NTFS partitions. Don't use it.
Use Free Space	Use this option to install Linux in the free, unpartitioned disk space on your computer. This option assumes that you've used external tools to make disk space available.
Create Custom Layout	The most difficult but also the most flexible option available. Using this option assumes you'll manually create all the partitions and logical volumes that you want to use on your computer.

22. To make sure you're using a setup that allows you to do all exercises that come later in this book, you'll need to select the Create Custom Layout option.
23. After selecting the Create Custom Layout option, click Next to continue. You'll now see a window in which your hard drive is shown with a name like sda or hda on old IDE-based computers below it. Under that appears one more item with the name Free that indicates all available disk space.

EXERCISE 1.1 (continued)

24. To configure your hard disk, you first have to create two partitions. Click Create to start the Create Storage interface. For the first partition, you'll select the Standard Partition option. Select this option, and click Create.



EXERCISE 1.1 (continued)

25. You'll now see the Add Partition interface in which you have to specify the properties of the partitions you want to create. The first partition is a rather small one that is used for booting only. Make sure to use the following properties:

Mount Point: /boot

File System Type: ext4

Size: 200 MB

Additional Size Options: Fixed size

Force to be a primary partition



26. After creating the boot partition, you'll need to create a partition that's going to be used as an LVM physical volume. From the main partitioning screen, click Create, and in the Create Storage options box, select LVM Physical Volume. Next click Create.

EXERCISE 1.1 (continued)

At this point, the purpose is to get you up and running as fast as possible. Therefore, you'll read how to configure your disk, without overwhelming you with too many details on exactly what it is you're doing. In Chapter 5, "Configuring and Managing Storage," you'll read more about partitions and logical volumes and what exactly they are.

27. In the Add Partition window, you now have to enter the properties of the physical volume you've just created. Use the following values:

File System Type: Physical Volume (LVM)
 Size: 40000
 Additional Size Options: Fixed size
 Force to be a primary partition

EXERCISE 1.1 (continued)

Please Select A Device

Device	Size (MB)	Mount Point/ RAID/Volume	Type	Format
▼ Hard Drives				
▼ sda (sda/sda)				
sda1	200	/boot		
Free	20279			

Add Partition

Mount Point: <Not Applicable>

File System Type: physical volume (LVM)

Drive	Size	Model
<input checked="" type="checkbox"/> sda	20480 MB	VMware, VMware-Virtual S

Allowable Drives:

Size (MB): 10000

Additional Size Options

Fixed size

Fill all space up to (MB):

Fill to maximum allowable size

Force to be a primary partition

Encrypt

Cancel OK

Create Edit Delete Reset

← Back Next →

28. At this point, you have created an LVM physical volume, but you can't do anything useful with it yet. You now need to create a volume group on top of it. To do this, click Create, and under the Create LVM option, select LVM Volume Group. Next click Create.

Please Select A Device

Device	Size (MB)	Mount Point/ RAID/Volume	Type	Format
▼ Hard Drives				
▼ sda (sda/sda)				
sda1	200	/boot	ext4	✓
sda2	10000		physical volume (PV)	✓
Free	10279			

Create Storage

Create Partition
General purpose partition creation

Create Software RAID Information
Create a RAID formatted partition

RAID Partition
Create a RAID formatted partition
Requires at least 2 free RAID formatted partitions

Create LVM Information

LVM Volume Group
Requires at least 3 free LVM formatted partition

LVM Logical Volume
Create a logical volume on selected volume group

LVM Physical Volume
Create an LVM formatted partition

Cancel Create

Create Edit Delete Reset

← Back Next →

EXERCISE 1.1 (continued)

29. You'll now see the properties of the LVM volume group. The only relevant parameter is the name, which is set to `vg_yourhostname`, which is perfectly fine. Change nothing, and click Add to add logical volumes in the volume group. The logical volumes are what you're going to put your files on, and you'll need three of them:

- One 20GB volume that contains the root directory
- One 512MB volume to use for a swap
- One 2GB volume that contains the `/var` directory

To start creating the logical volumes, click Add.



30. You need to add three logical volumes using the following parameters:

The root volume:

Mount Point: `/`

File System Type: `Ext4`

Logical Volume Name: `root`

Size: `20000`

The swap volume:

File System Type: `swap`

Logical Volume Name: `swap`

Size: `512`

The var volume:

Mount Point: `/var`

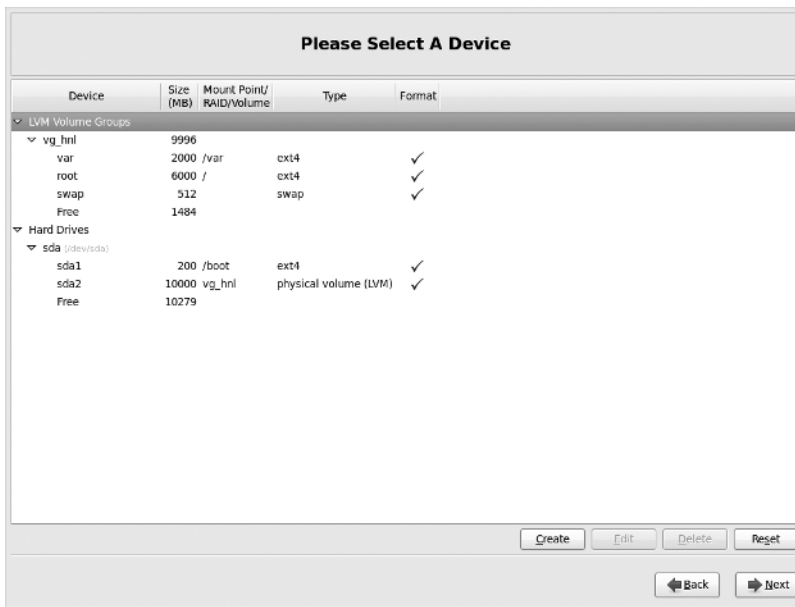
File System Type: `Ext4`

Logical Volume Name: `var`

Size: `2000`

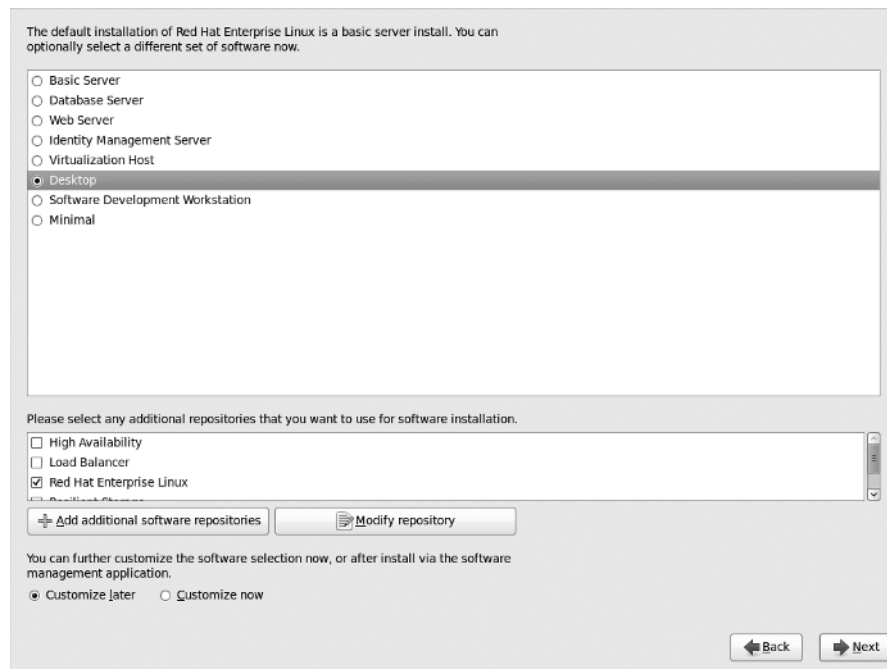
EXERCISE 1.1 (continued)

Once you've finished configuring storage devices on your computer, the disk layout should look like this:



EXERCISE 1.1 (continued)

31. Now click Next to continue. In the Format Warning window that you now see, click Format to start the formatting process. Next, confirm that you really want to do this by selecting the Write Changes To Disk option.
32. At this point, the partitions and logical volumes have been created, and you're ready to continue with the installation procedure. On the following screen, the installer asks what you want to do with the boot loader. Select the default option, which installs it on the master boot record of your primary hard drive, and click Next.
33. You now have to specify what type of installation you want to perform. The only thing that counts at this moment is that you'll need to select the Desktop option. If you don't, you'll end up with a server that, by default, doesn't have a graphical environment, and that is hard to fix if you're just taking your first steps into the world of Red Hat Enterprise Linux. After selecting the Desktop option, click Next to continue.



34. The installation process is now started, and the files will be copied to your computer. This will take about 10 minutes on an average system, so it's now time to have a cup of coffee.
35. Once the installation has completed, you'll see the Congratulations message telling you that your server is ready. On this screen, click Reboot to stop the installation program and start your server.

EXERCISE 1.1 (continued)

36. Once the server has successfully started for the first time, you'll see the Welcome screen that guides you through the remainder of the installation procedure. From this screen, click Forward. Next you'll see the License Information screen in which you have to agree to the license agreement. After doing so, click Forward to proceed.
39. Now you'll see the Set Up Software Updates screen where you can connect to the Red Hat Network.
 - a. If you have credentials for Red Hat Network, you can connect now.
 - b. If you don't and just want to install a system that cannot download patches and updates from Red Hat Network, select the No, I Prefer To Register At A Later Time option, and click Forward.

In this book, RHN access is not required, so select No, I Prefer To Register At A Later Time. You'll see a window informing you about all the good things you'll miss without RHN. In this window, click No Thanks, I'll Connect Later to confirm your selection. Now click Forward once more to proceed to the next step.

EXERCISE 1.1 (continued)

Welcome
License
Information
› Set Up Software Updates
Create User
Date and Time
Kdump

Set Up Software Updates

This assistant will guide you through connecting your system to Red Hat Network (RHN) for software updates, such as:

- Your Red Hat Network or Red Hat Network Satellite login
- A name for your system's Red Hat Network profile
- The address to your Red Hat Network Satellite (optional)

Why Should I Connect to RHN? ...

Would you like to register your system at this time? **(Strongly recommended.)**

Yes, I'd like to register now.

No, I prefer to register at a later time.

Back Forward



If you don't connect your server to RHN, you cannot update it. This means it's not a good idea to use this server as a production system and provide services to external users; you'll be vulnerable if you do. If you need to configure a Red Hat system that does provide public services, you have to purchase a subscription to Red Hat Enterprise Linux. If you don't want to do that, use Scientific Linux or CentOS instead.

40. At this point, you'll need to create a user account. In this book, we'll create the user "student," with the full name "student" and the password "redhat" (all lowercase). You can safely ignore the message that informs you that you've selected a weak password.

EXERCISE 1.1 (continued)

Welcome
License Information
Set Up Software Updates
▶ **Create User**
Date and Time
Kdump

Create User

You must create a 'username' for regular (non-administrative) use of your system. To create a system 'username', please provide the information requested below.

Username:

Full Name:

Password:

Confirm Password:

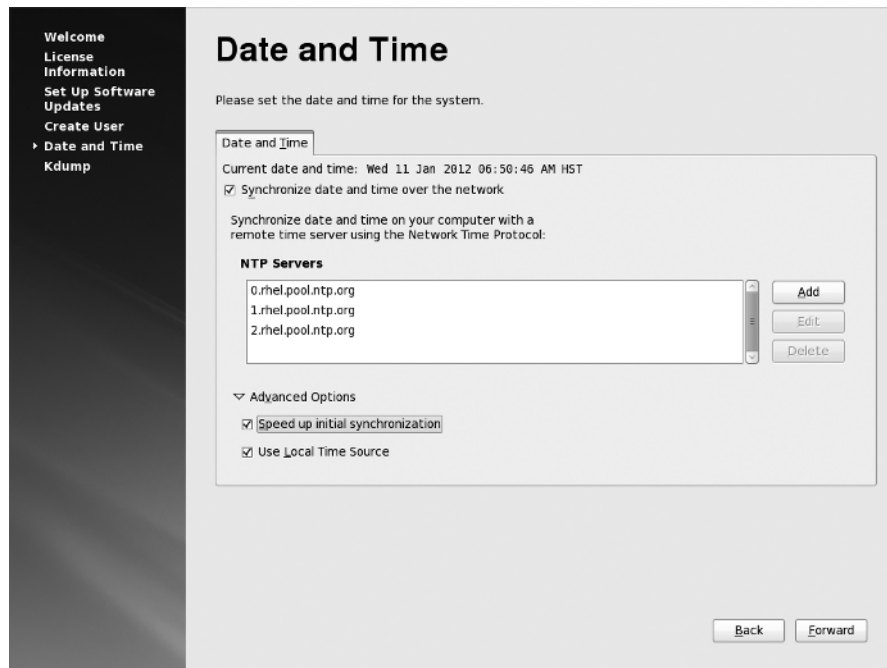
If you need to use network authentication, such as Kerberos or NIS, please click the Use Network Login button.

If you need more control when creating the user (specifying home directory, and/or UID), please click the Advanced button.

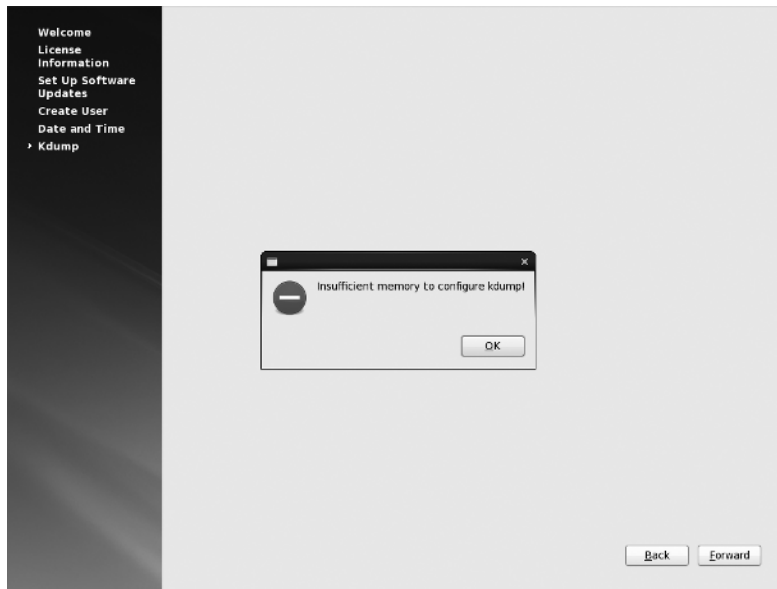
- 41.** During the installation, you already indicated your time zone and whether your server is using UTC on the hardware clock. At this point, you need to finalize the Date And Time settings.
- Specify the current time.
 - Indicate whether you want to synchronize the date and time over the network.
 - Because time is an essential factor for the functioning of many services on your server, it is a very good idea to synchronize time with an NTP time server on the Internet. Therefore, on the Date And Time screen, select Synchronize Date And Time Over The Network. This will show a list containing three NTP servers on the Internet. In many cases, it doesn't really matter which NTP servers you're using, as long as you're using some NTP servers, so you can leave the servers in this list.

EXERCISE 1.1 (continued)

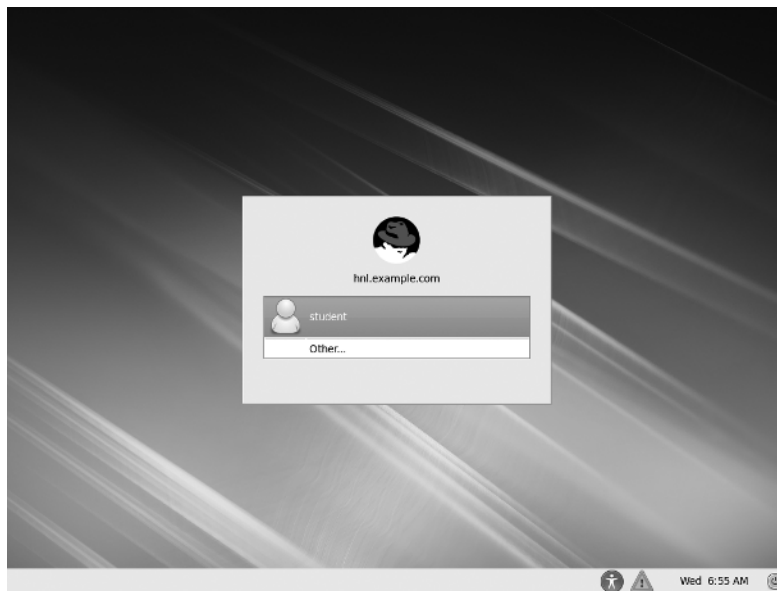
- d. Open Advanced Options, and select the Speed Up Initial Synchronization and Use Local Time Source options. The first option makes sure that, if a difference is detected between your server and the NTP time server it is synchronizing with, your server will synchronize its time as fast as it can. If you are installing your server in a VMware virtual environment, it is important to use this option to prevent problems in time synchronization. The second option tells your server to use the local hardware clock in your server as a backup option. It is a good idea to enable this option on all servers in your network, because it creates a backup in case the connection to the Internet is lost for a long period of time.
- e. After enabling the advanced options, click Forward to continue.



42. In the final part of the configuration, you can enable the Kdump settings. *Kdump* refers to crash dump. It allows a dedicated kernel to activate on the rare occasion that your server crashes. To use this feature, you need at least 2GB of available RAM. If you're using less, you'll see an error message indicating that you have insufficient memory to configure Kdump. You can safely ignore this message.

EXERCISE 1.1 (continued)

- 43.** On the next and final screen of the installation program, click Finish. This completes the installation procedure and starts your system. You'll now see a login window where you can select the user account you'll use to log in.



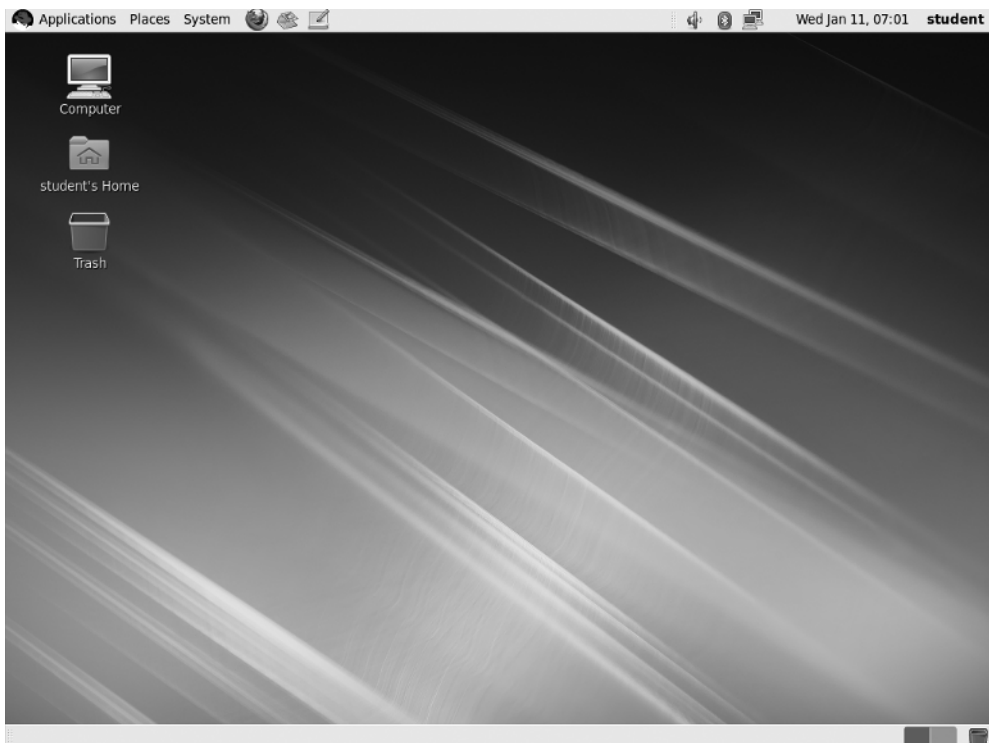
Exploring the GNOME User Interface

Now that your server is installed, it's time to get a bit familiar with the GNOME user interface. As indicated, on most servers, the graphical user interface (GUI) is not enabled. However, to get familiar with RHEL, it is a good idea to use the GNOME interface anyway.

To make yourself known to your Red Hat server, you can choose between two options. The best option is to click the name of the user account that you've created while installing the server and enter the password of that user. It's easy to pick the username—a list of all user accounts that exist on your server is displayed on the graphical login screen. Selecting a username from the graphical login screen connects you with normal user credentials to the server. That means you'll enter the server as a nonprivileged user, who faces several restrictions on the server.

Alternatively, from the graphical login screen, you can click Other to enter the name of another user you want to use to log in. You can follow this approach if you want to log in as user root. Because there are no limitations to what the user root can do, it is a very bad idea to log in as root by default. So, at this point, click the name of the user that you've created, and enter the password. After successful authentication, this shows the default GNOME desktop with its common screen elements, as shown in Figure 1.1.

FIGURE 1.1 The default GNOME graphical desktop



In the GNOME desktop, there are a few default elements with which you should be familiar. First, in the upper-left part of the desktop, there is the GNOME menu bar. There are three menu options: Applications, Places, and System.

Exploring the Applications Menu

In the Applications menu, you'll find a limited number of common desktop applications. The most useful applications are in the System Tools submenu. The Terminal Application is the single most important application in the graphical desktop because it gives you access to a shell window in which you can enter all the commands you'll need to configure your server (see Figure 1.2). Because it is so important, it's a good idea to add the icon to start this application to the panel. The panel is the bar which, by default, is at the top of the graphical screen. The following procedure describes how to do this:

1. Open the Applications menu, and select System Tools. You see the contents of the System Tools submenu.
2. Right-click the Terminal icon, and select Add This Launcher To Panel.
3. You'll now see a launcher icon that enables you to start the Terminal application in a quick and easy way from the panel.

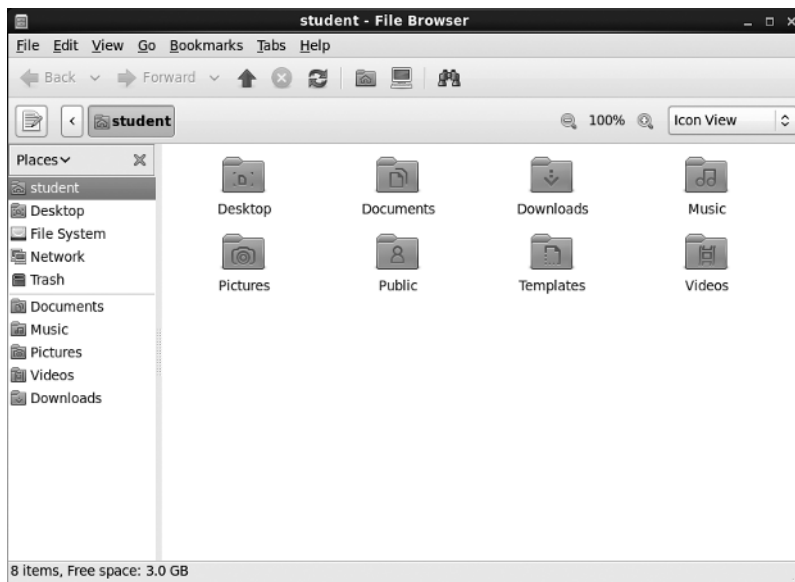
FIGURE 1.2 The Terminal application gives access to a shell interface.



Another rather useful application in the System Tools submenu of the Applications menu is the file browser. Selecting this application starts Nautilus, the default file browser on a Red Hat system. Nautilus organizes your computer in Places, which allow you to browse the content of your computer in a convenient way.

After opening Nautilus, you'll see the contents of your home directory, as shown in Figure 1.3. This is your personal folder where you can store your files so that other users have no access. By using the Places sidebar, you can navigate to other folders on your computer, or by using the Network option, you can even navigate to folders that are shared by other computers on the network.

FIGURE 1.3 After opening Nautilus, you'll get access to your home folder.



The file system is among the most useful places that you'll see in Nautilus. This gives you access to the root of the Linux file system, which allows you to see all the folders that exist on your computer. Be aware that, as an ordinary user without root permissions, you won't have access to all folders or files. To get access to everything, you should run Nautilus as root.

From Nautilus, you can access properties of files and folders by right-clicking them. This gives you access to the most important properties, including permissions that are assigned to a file or folder. However, this is not the way that you would normally change permissions or other file attributes. In subsequent chapters of this book, you'll learn how to perform these tasks from the command line.

Exploring the Places Menu

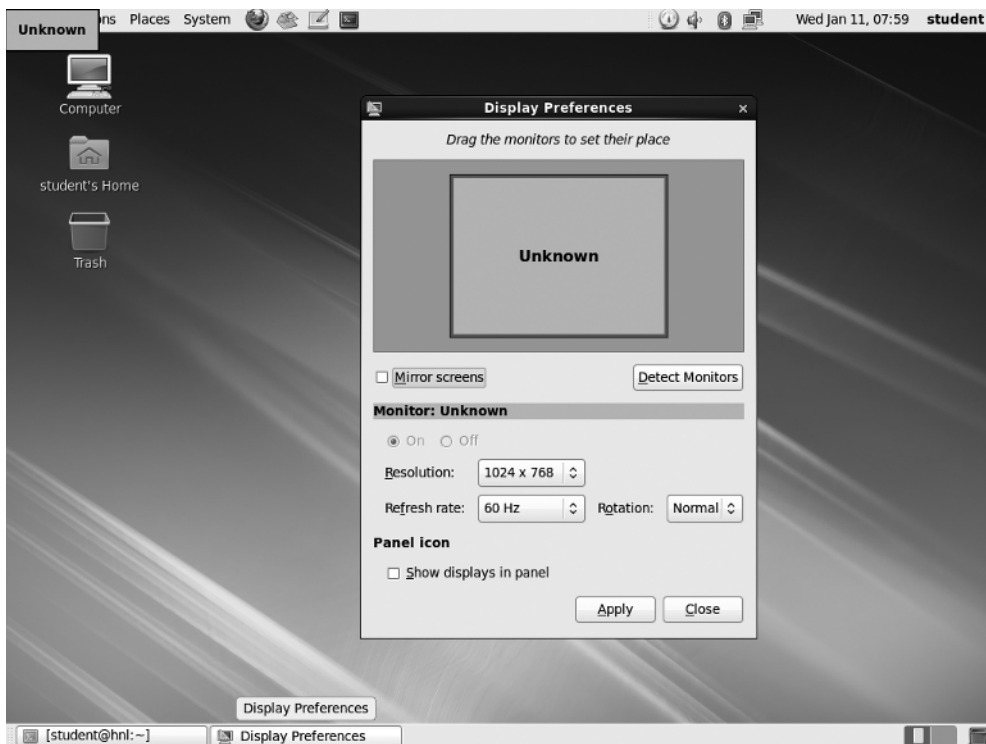
Now let's get back to the main menus in the GNOME interface. There you'll notice that the name of the second menu is Places. This menu, in fact, shows more or less the same

options as Places in Nautilus; that is, it includes all the options you need to connect to certain folders or computers easily on the network. It also includes a Search For Files option, which may be useful for locating files on your computer. However, you will probably not be interested in the Search For Files option once you've become familiar with the powers of the Find command.

Exploring the System Menu

The third of the default GNOME menus, the System menu, gives you access to the most interesting items. First you'll find the Preferences submenu, which has tools such as the Screensaver and Display tools. You'll use the Display Preferences window (see Figure 1.4) to change the settings of the graphical display. This is useful in configuring external monitors or projectors or just to correct the screen resolution if the default resolution doesn't work for you.

FIGURE 1.4 The Display Preferences menu helps you optimize properties of the graphical display hardware.



In the Screensaver tool, you can set the properties of the screensaver, which by default activates after five minutes of inactivity. It will lock the screen so that you get access to it again only after entering the correct password. This is very useful in terms of security, but

it can also be annoying. To disable the automatic locking of the screensaver, select System > Preferences > Screensaver and make sure the option Lock Screen When Screensaver Is Active option is unchecked.

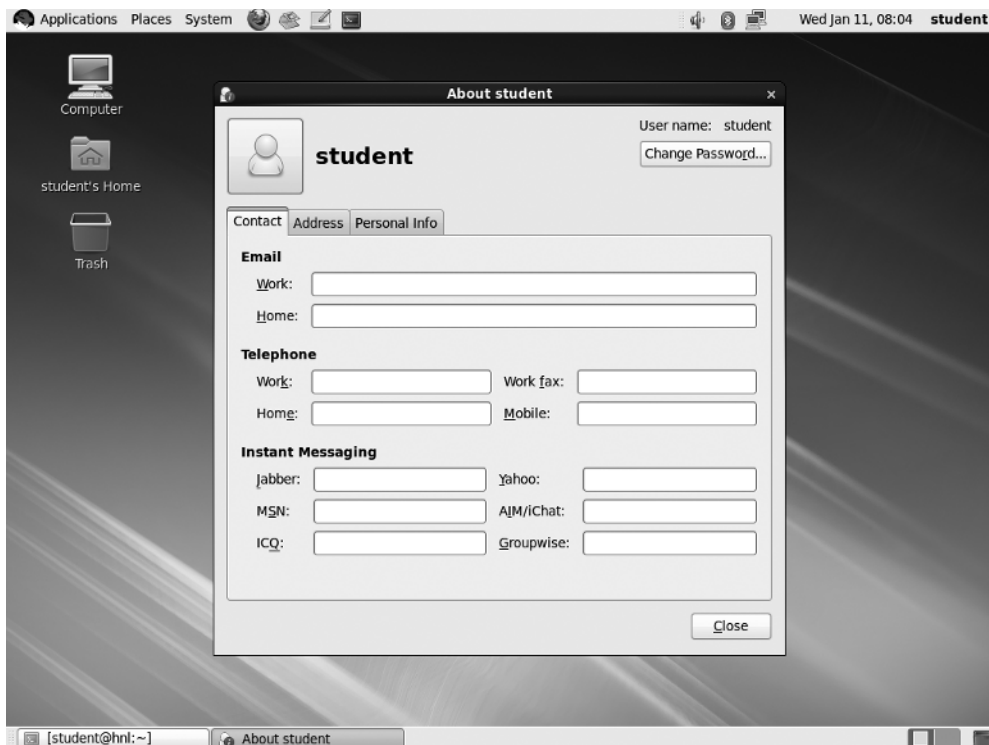
In the Administration submenu under System, you'll get access to some common administration utilities. These are the *system-config* utilities that allow you to perform common administration tasks in a convenient way. These tools relate more to system administration tasks than the tools in any of the other GNOME submenus.



You'll learn how to use the system-config utilities in later chapters.

The upper-right part of the GNOME panel displays some apps that give access to common tools, including the Network Manager utility, which gives you easy access to the screens that help you configure the network cards in your computer. You'll also find the name of the current user in the upper-right corner of the screen. You can click on it and then Account Information to get access to personal information about this user, as well as the option to change the user's password (see Figure 1.5).

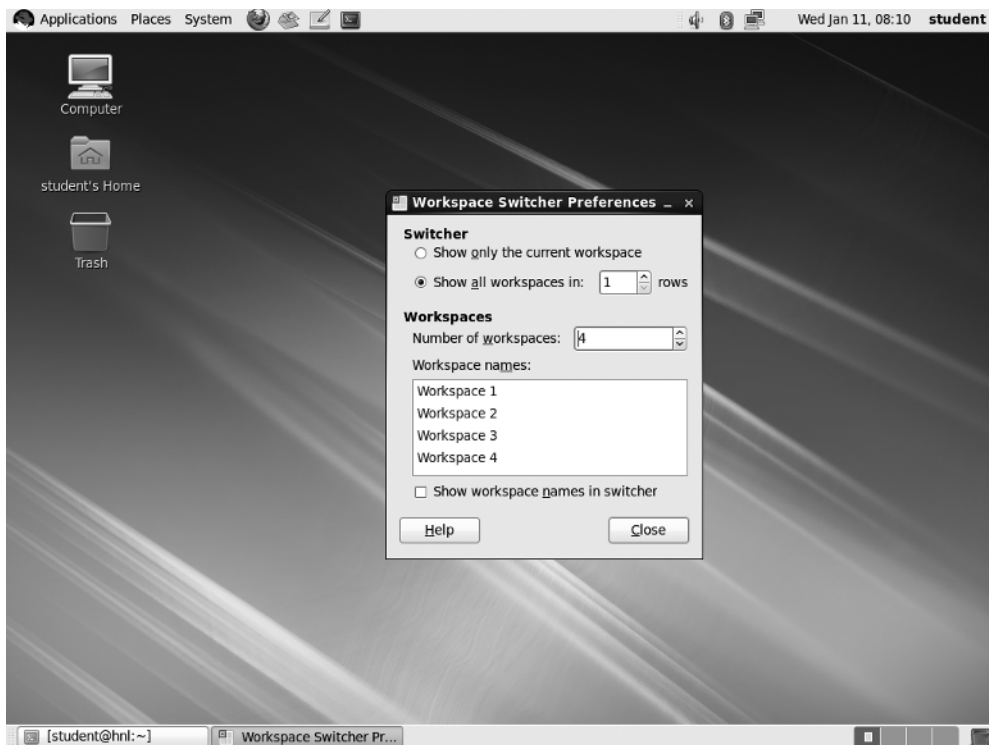
FIGURE 1.5 Click the name of the current user to get access to account information about that user.



The menu associated with the current user also gives you access to the Lock Screen tool. Use it whenever you walk away from the server to lock the desktop in order to make sure that no one can access the files on the server without your supervision. Another useful tool is Switch User, which allows you to switch between two different user accounts that are both logged in.

The last part of the screen gives access to all open applications. Just click the application that you want to use to access it again. A very useful element in this taskbar is the Workspace Switcher (see Figure 1.6). This screen is one of the two workspaces that are activated by default. If you want to open many applications, you can use multiple workspaces to work in a more organized way. You can put specific application windows on those workspaces where you really need them. By default, Red Hat Enterprise Linux shows two workspaces, but you can increase the number of workspaces to an amount that is convenient for you. To activate another workspace, just click the miniature of the workspace as it is shown in the taskbar.

FIGURE 1.6 Increasing the number of workspaces



Summary

In this chapter, you became familiar with Red Hat Enterprise Linux (RHEL). You learned about what Linux is and where it comes from. You read that Linux comes from a tradition of open source software, and it is currently in use in most of the Fortune 500 companies. Next you will read about the Red Hat company and its product offerings.

You then learned how to install Red Hat Enterprise Linux on your computer. If all went well, you now have a usable version of RHEL that is available to you while working your way through this book.

Finally, the chapter introduced you to the GNOME graphical desktop. You learned that using it makes the process of learning Linux easier. You also saw where some of the most interesting applications are located in the different menus of the GNOME interface.

