



I

Prepare for Red Hat Hands-on Certifications

CERTIFICATION OBJECTIVES

- | | | | |
|------|------------------------------|------|---|
| I.01 | The RHCSA and RHCE Exams | I.06 | System Setup Options |
| I.02 | Basic Hardware Requirements | I.07 | Configure Default File Sharing Services |
| I.03 | Get Red Hat Enterprise Linux | ✓ | Two-Minute Drill |
| I.04 | Installation Requirements | Q&A | Self Test |
| I.05 | Installation Options | | |

The Red Hat exams are an advanced challenge. While this book covers the Red Hat Certified System Administrator (RHCSA) exam, it provides the foundation for those who want to earn the Red Hat Certified Engineer (RHCE) certification. Red Hat offers several courses to help prepare for these exams, as described in the front matter and in this chapter.

The focus of this chapter is installation, to create a common version of Red Hat Enterprise Linux (RHEL) as a test bed for future chapters. It assumes and describes hardware required to implement Red Hat's default virtual machine (VM) solution, the Kernel-based Virtual Machine (KVM). As rebuild distributions such as the Community Enterprise Operating System (CentOS) and Scientific Linux are essentially identical to RHEL, you should be able to use those solutions too. Just about the only difference between a rebuild and RHEL is the trademarks and the access to repositories, which will be described in Chapter 7.

Those of you familiar with earlier versions of the Red Hat requirements may note the recent changes to the Red Hat exams. Red Hat no longer gives a RHCT exam. It has now been replaced with the RHCSA. While the RHCSA is in many ways similar to the RHCT, there are significant differences. Most RHCSA objectives were covered on the former RHCT exam. However, the RHCSA is certainly not easier than the RHCT; it is just now a prerequisite to the RHCE. The RHCSA also includes a number of requirements that were formerly part of the RHCE objectives.

Nevertheless, Red Hat suggests that candidates for the RHCSA will have one to three years of experience with the bash shell, user administration, system monitoring, basic networking, software updates, and more. Details are described in the introduction to this book.

If you're new to Linux or Unix, this book may not be enough for you. It's not possible to provide sufficient detail, at least in a way that can be understood by newcomers to Linux and other Unix-based operating systems. If after reading this book, you find gaps in your knowledge, please refer to one of the following guides:

- *Linux Administration: A Beginner's Guide*, by Wale Soyinka (McGraw-Hill, 2008), provides a detailed step-by-step guide to this operating system.

- *Security Strategies in Linux Platforms and Applications*, by Michael Jang (Jones & Bartlett, 2010), gives you a detailed look at how you can secure your Linux system and networks in every possible way.
- *LPIC-1 in Depth*, by Michael Jang (Course Technology PTR, 2009), covers the certification many Linux professionals qualify for prior to working on the RHCSA and RHCE.

Before installing Red Hat Enterprise Linux (RHEL), you need the right hardware. Red Hat supports KVM only on physical systems with 64-bit CPUs capable of hardware virtualization. Details are discussed in the chapter. As such, while the RHCSA and RHCE exams are by and large not hardware exams, some basic hardware knowledge is a fundamental requirement for any Linux administrator. As for the operating system itself, you can purchase a subscription to RHEL, or you can use one of the “rebuild” distributions where the distribution is built by third parties from source code publicly released by Red Hat.

If you’re experienced with other Unix-type operating systems such as Solaris, AIX, or HP-UX, prepare to leave some defaults at the door. There are even significant differences between the Ubuntu and Red Hat distributions. When Red Hat developed its Linux distribution, the company made some choices that are not consistent with standard Unix practices. When I took Red Hat’s RH300 course, some students with these backgrounds had difficulties with the course and the RHCE exam.

For the purpose of this book, I’ll be running most commands as the Linux administrative user, root. Logging in as the root user is normally discouraged unless you’re administering a computer. However, since the RHCSA and RHCE exams test your administrative skills, it’s appropriate to run commands in this book as the root user. But you’ll also need to know how to set up regular users with partial or full administrative privileges.

INSIDE THE EXAM

A Virtual Host

The RHCSA assumes that you know how to “configure a physical machine to host virtual guests.” In other words, you need to be able to prepare a system to house VMs where other instances of RHEL (or even other operating systems such as Microsoft Windows) can be installed.

As this is RHEL, this is based on the Red Hat default VM system, KVM. As appropriate rebuild distributions such as CentOS and Scientific Linux use the same source code, they also use KVM. In this chapter, not only will you install RHEL, but also you will install those packages that support KVM.

Default File Sharing Configuration Services

RHCSAs are expected to know how to “Configure a system to run a default configuration HTTP server” and “Configure a system to run a default configuration FTP server.” The default Red Hat solutions for these services are the Apache Web server and the very secure FTP (vsFTP) server. While these services can be complex, the steps required to set up these

servers to share files are fairly simple. In fact, no changes are required to the default configuration files for these services. Some of the related steps described in this chapter depend on skills presented in future chapters.

The original release of the RHCSA objectives was worded slightly differently: “Deploy file sharing services with HTTP/FTP.” I believe this provides a significant clue to Red Hat’s intent with these objectives. To that end, you’ll examine how to set up these services as file servers, based on their default configurations.

Using Other Versions of Red Hat

For the purpose of this chapter, you can install RHEL 6 using a paid subscription or from a demonstration DVD. You can also use one of the rebuild distributions. However, while RHEL 6 is based in part on the work done on the Fedora Linux distribution, it’s based in part on both the Fedora 12 and 13 releases. So it’s possible if you use Fedora 12 or 13, the configuration files may in some cases be quite different from RHEL 6. Later versions of Fedora are likely to have features not found in RHEL 6.

CERTIFICATION OBJECTIVE 1.01

The RHCSA and RHCE Exams

Red Hat first started giving certification exams in 1999. Since that time, their exams have evolved. The former RHCT was a complete subset of the RHCE. Today, the RHCSA now covers topics separate from but closely related to the RHCE.

In addition, Red Hat has focused the exams more on hands-on configuration. Multiple choice questions were removed from the exam in 2003. More recently, in 2009, they simplified the exam by removing the requirement to install Linux on a “bare-metal” system. (However, the changes implemented in 2011 suggest that you need to know how to install Linux over a network on a VM.) In addition, there is no longer a separate troubleshooting portion of the exam. For more information, see www.redhat.com/certification/faq.

exam

Watch

Red Hat provides “pre-assessment” tests for Red Hat RHCSA and RHCE Exam Prep courses. They correspond to the RH134 and RH254 courses, respectively. These tests are

available through the Red Hat web pages for each course. Red Hat requires contact information before providing those preassessment tests.

The Exam Experience

Red Hat’s certification tests are hands-on exams. As such, they are respected throughout the industry as a sign of genuine practical knowledge. When you pass a Red Hat exam, you will stand head and shoulders above the candidate who has passed only a “standard” multiple-choice certification exam.

When time starts, you’ll be faced with a live system. You’ll be given actual configuration problems associated with the items listed in the exam objectives for each certification, shown at www.redhat.com/certification/rhcsa/objectives/ and www.redhat.com/certification/rhce/objectives/. Naturally, this book is dedicated to helping you gain the skills described on those web pages.

While you won't have Internet access during the exam, you will have access to online documentation such as man and info pages as well as documentation in the `/usr/share/doc/` directories, assuming appropriate packages are installed.

In addition, Red Hat provides the exam in electronic format. While the basic instructions may be in a local language such as English, the RHCSA and RHCE exams are available in 12 different languages: English, Simplified Chinese, Traditional Chinese, Dutch, French, Italian, Japanese, Korean, Portuguese, Russian, Spanish, and Turkish. If one of these alternatives is desired, you should contact Red Hat training to be sure, at training@redhat.com or 1-866-626-2994.

Red Hat also has prep courses for both exams. The outline for those courses are available from www.redhat.com. While this book is not intended as a substitute for such courses, they are consistent with the outline of those courses. This book covers the objectives associated with each of these exams.

exam

Watch

This book's coverage of the objectives can be found in the front matter items listed in the RHCSA and RHCE exam for this book, in Table 2, page xlvi.

The RHCSA Exam

The RHCSA exam allows you to demonstrate your ability to configure live physical and virtual systems for networking, security, custom filesystems, package updates, user management, and more. In essence, the RHCSA exam covers those skills required to configure and administer a Linux workstation in the enterprise.

The RHCSA exam lasts two and a half hours. When you sit down to take the exam, you'll have tasks to perform on a live RHEL system. Any changes that are made must survive a reboot. When you've completed the given tasks, the person grading the exam will see if the system is configured to meet the requirements. For example, if you're told to "create, delete, and modify local user accounts," it doesn't matter if the associated configuration file has been modified with the vi editor or the graphical User Manager tool. As long as you don't cheat, it's the results that matter.

The RHCE Exam

The RHCE exam tests your ability to configure live physical and virtual servers to configure network services such as Apache, vsFTP, the Network File System (NFS), Samba, remote logging, and more. It also tests your ability to handle complex configuration options associated with Security Enhanced Linux (SELinux), firewalls, networking, and more. In essence, if you pass the RHCE exam, hiring managers will know that you're qualified to help manage their enterprises of Linux systems.

The RHCE exam also lasts two hours. When you sit down to take the exam, you'll be given tasks to perform on a live RHEL system. As with the RHCSA, any changes that are made must survive a reboot. In any case, it doesn't matter if you've configured the associated configuration file with the nano editor or a GUI tool. As long as you don't cheat, it's the results that matter.

The topics in the Red Hat preparation courses in a few areas go beyond those listed in the Red Hat Exam Prep guide. While such topics are not currently part of the exam, they may be included in future versions of the Red Hat exams.

If You're Studying "Just" for the RHCSA Exam

Red Hat has been known to make minor changes to the requirements on occasion. Future changes may be based on topics covered in the Red Hat RHCSA Rapid Track course, RH199/RH200. So if you're not planning to take the RHCSA within the next few months, watch the outline for that course. It may in effect be a preview of where Red Hat wants to take the RHCSA exam in the future.

Evolving Requirements

Changes happen to the requirements for the Red Hat exams. You can see that in the differences between the RHCT and the RHCSA. You can see that in the changes to the exam format, where bare-metal installations are no longer required. In fact, that change happened over two years into the life of RHEL 5. Changes happened in the first month after RHEL 6 was released. So when you're preparing for the RHCSA or RHCE exams, watch the associated exam objectives carefully. In addition, Red Hat announced the removal of separate troubleshooting problems from the exams through the WordPress blog publishing site, at <http://redhatcertification.wordpress.com/>.

CERTIFICATION OBJECTIVE 1.02

Basic Hardware Requirements

Now it's time to explore in detail the hardware that Red Hat Enterprise Linux can handle. While some manufacturers now include their own Linux hardware drivers, most Linux hardware support comes from third parties, starting with the work of volunteers. Fortunately, there is a vast community of Linux users, many of whom produce drivers for Linux and distribute them freely on the Internet. If a certain piece of hardware is popular, you can be certain that Linux support for that piece of hardware will pop up somewhere on the Internet and will be incorporated into various Linux distributions, including Red Hat Enterprise Linux.

Hardware Compatibility

If you only have 32-bit systems available, be prepared to spend some money. KVM is the default VM solution for RHEL 6. Red Hat has configured it to operate only on 64-bit systems. Fortunately, most PCs and servers sold today are 64-bit systems. Even the lowly Intel i3 CPU can handle 64-bit operating systems. There are even 64-bit versions of the Intel Atom CPU common on netbook systems. Similar comparisons can be made for CPUs from Advanced Micro Devices.

Be careful when purchasing a new computer to use with Linux. Though Linux has come a long way the last few years, and you should have little problem installing it on most modern servers or PCs, you shouldn't assume Linux will install or run flawlessly on *any* computer, especially if the system in question is a state-of-the-art laptop computer. (And you do need 64-bit-capable hardware to prepare for the Red Hat exams.) Laptops are often designed with proprietary configurations that work with Linux only after some reverse engineering. For example, when I installed RHEL 6 on a brand-name business laptop built in 2010, I had to do a bit of extra work to make the wireless adapter work with RHEL 6.

The architecture of a server or PC defines the components that it uses as well as the way that they are connected. In other words, the architecture describes much more than just the CPU. It includes standards for other hardware such as the hard drive, the network card, the keyboard, the graphics adapter, and more. All software is written for a specific computer architecture.

Even when a manufacturer creates a device for a CPU platform, it may not work with Linux. Therefore, it's important to know the basic architecture of a computer. But strictly speaking, if you want hardware compatible with and supported by Red Hat, consult the hardware compatibility list at <http://hardware.redhat.com/hcl/>.

exam

Watch

While it is important to know how Linux interacts with your hardware, the Red Hat exams are not hardware exams. With the possible exception of the Internet Small Computer Systems Interface (iSCSI) targets in the RHCE objectives, no part of the exam

objectives directly addresses hardware. However, to practice for the exams, you do need to install RHEL 6. And to configure a system for KVM, which requires hardware-assisted virtualization, you need a system with 64-bit CPUs and related hardware.

Architectures

While RHEL 6 has been built for a variety of architectures, you can focus on the Intel/AMD 64-bit or x86_64 architecture for the RHCSA and RHCE exams. As of this writing, these exams are offered only on computers with such CPUs, so you need not worry about special architecture-specific issues such as specialty bootloaders or 64-bit module directories. Nevertheless, customized Red Hat distributions are available for a variety of platforms.

You can install RHEL 6 on systems with a wide variety of CPUs. But remember, to test a system for the required KVM virtualization solution, you'll need a 64-bit system. Red Hat supports four basic different CPU architectures:

- x86 (32-bit)
- Intel/AMD64 (x86_64)
- IBM Power Architecture (64-bit)
- IBM System Z Architecture

There's one significant omission from this list. Red Hat no longer builds RHEL 6 for the Intel Itanium CPU. To identify the architecture of a system, run the following command:

```
# uname -p
```

If you're planning to configure VMs on RHEL 6, be sure to choose an architecture that supports hardware-assisted virtualization, along with Basic Input/Output System (BIOS) or Universal Extensible Firmware Interface (UEFI) menu options that allow you to activate hardware-assisted virtualization. A configuration that supports hardware-assisted virtualization will have either the `vmx` (Intel) or `svm` (AMD) flags in the `/proc/cpuinfo` file.

Be aware, RHEL 6 installations on 32-bit systems with UEFI menus are not supported.



If you're not sure about a system, see if you can try it out at the "store" with one of the "Live CD" distributions. If allowed by store personnel, boot the system from that Live CD (or DVD). Once booted, you should be able to access the `/proc/cpuinfo` file from the command line. You may need to tinker a bit with the BIOS or UEFI menus. A system configured to support hardware-assisted virtualization will have the `vmx` or `svm` flags in that file, even when read from a Live CD distribution.

RAM Requirements

While it's possible to run RHEL 6 on less, the RAM memory requirements are driven by the needs of the Red Hat installer. For basic Intel/AMD-based 32- and 64-bit architectures, Red Hat officially requires 512MB of RAM and recommends at least 1GB of RAM per system. However, I've installed RHEL 6 on VMs with considerably less RAM. On a VM where the GUI wasn't installed, I didn't even need 200MB of RAM.

Of course, actual memory requirements depend on the load from every program that may be run simultaneously on a system. That can also include the memory requirements of any VMs that you might run on a physical RHEL 6 system. There is no practical maximum RAM, as theoretically, you could run 128TB (that's 128,000GB) of RAM on RHEL 6. But that's just theory. The maximum RAM supported by Red Hat on RHEL 6 is 16GB on 32-bit systems and 2TB on 64-bit systems.



If you're setting up Linux as a server, RAM requirements increase with the number of users who may need to log in simultaneously. The same may be true if you're running several different VMs on a single system. However, administrators typically "overbook" RAM on VMs configured with different functionality.

Hard Drive Options

Before a computer can load Linux, the BIOS or UEFI has to recognize the active primary partition on the hard drive. This partition should include the Linux boot files. The BIOS or UEFI can then set up and initialize that hard drive, and then load Linux boot files from that active primary partition. You should know the following about hard drives and Linux:

- The number of drives that can be installed on modern computers has increased. With port multipliers, it's relatively easy to configure 16 Serial Advanced Technology Attachment (SATA) drives on a system (assuming you can fit all of those drives).
Older PCs could handle only four Parallel Advanced Technology Attachment (PATA) drives.
- Depending on the SCSI (Small Computer Systems Interface) hardware available, you can attach up to 31 different SCSI hard drives.
- While you can install as many PATA, SATA, or SCSI drives as the hardware can handle, the Linux boot files from the /boot directory works only on one of the first two hard drives. If Linux is installed on a later drive, you'll need other boot media that are recognized by the BIOS / UEFI, such as a USB (Universal Serial Bus) or CD/DVD. (Red Hat no longer creates an image for a boot floppy.)
- RHEL 6 doesn't even have to be installed on a local drive. During the installation process, the specialized storage devices option supports the use of Storage Area Networks (SANs), Direct Access Storage Devices (DASDs), hardware RAID (Redundant Array of Independent Disk) devices, and more.

Networking

As Linux was originally designed as a clone of Unix, it retains the advantages of Unix as a network operating system. However, not every network component works with Linux. A number of manufacturers of wireless network devices have not built Linux drivers. In most such cases, Linux developers have been working furiously to develop appropriate drivers, and to get those drivers incorporated into the major distributions, including RHEL.

Virtual Machine Options

As virtualization makes it relatively easy to set up a large number of systems, it can help you configure a large number of systems, each dedicated to a specific service. To that end, virtualization can be divided into five different categories, as described in the aforementioned book: *Security Strategies in Linux Platforms and Applications*.

- **Application-level** Systems like Wine Is Not an Emulator (WINE) support the installation of a single application. In this case, WINE allows an application designed for Microsoft Windows to be installed on Linux. Loosely, that is one example of an application configured in a VM.
- **Platform-level VMs** Applications such as VMware Player and Virtualbox are open-source edition examples of platform-level VMs that emulate a complete computer system for the installation of a separate operating system.
- **Paravirtualization** While it's functionally similar to a platform-level VM, it works with fewer resources and usually requires a specialized kernel such as Xen.
- **Hardware-assisted virtualization** A hardware interface where VMs have access to the hardware features of a CPU, such as those described earlier with `vmx` or `svm` flags in the `/proc/cpuinfo` file.
- **Bare-metal virtualization** Some VM systems include a minimal operating system dedicated to VM operation. Two examples of bare-metal virtualization systems are VMware ESX and Citrix XenServer.

The KVM solution configured with RHEL 6 is known as a hypervisor, a VM monitor that supports the running of multiple operating systems concurrently on the same CPU. KVM replaces the previous default, Xen.



KVM has replaced Xen in many open-source distributions. XenSource is owned by Citrix, which started working with Microsoft after RHEL 5 was released.

CERTIFICATION OBJECTIVE 1.03

Get Red Hat Enterprise Linux

The RHCSA and RHCE exams are based on your knowledge of RHEL. When you take the RHCSA and/or the RHCE exams, it should be on a “standard” 64-bit computer. To get an official copy of RHEL, you’ll need a subscription. In some cases, trial subscriptions are available. However, if you don’t need the same “look and feel” of RHEL to prepare for an exam, third-party rebuilds are available. As such “rebuilds” use the same source code as RHEL, except for the trademarks and the connection to the Red Hat Network, they’re essentially functionally identical to RHEL.

Once you either purchase a subscription or get approved for an evaluation copy, you’ll be able to download RHEL 6 from the Red Hat Network (RHN) at <https://rhn.redhat.com/>. Downloads are available for the operating system in a format appropriate for a DVD. There’s also a download available for a network boot CD. You’ll even be able to download files with the source code for associated packages. These downloads are in ISO format, with a .iso extension. Such files can be burned to appropriate media, using standard tools such as K3b, Brasero, or even corresponding tools on Microsoft systems. Alternatively, you can set up a VM where the virtual CD/DVD drive hardware points directly to the ISO file, as discussed in Chapter 2. Unless you purchase an actual boxed subscription, the burning or other use of these ISO files is your responsibility.

Be aware, some of the installation options described in this part of the chapter have been subdivided into different sections. For example, the ways you can configure partitions using Red Hat’s Disk Druid tool are spread across multiple sections.



An ISO file is an image file that can be burned to CD/DVD media. The actual acronym is irrelevant, as it is based on a political compromise.

exam

Watch

While it is important to know how to get RHEL, that skill is not

a part of the objectives for the RHCSA or RHCE exam.

Purchase a Subscription

Different subscriptions are available for the desktop and the server. While the RHCSA is focused on workstations, it also does require the configuration of HTTP and FTP servers. Of course, the RHCE also requires the configuration of a variety of server services. So most readers will need a server subscription.

A variety of server subscriptions are available, depending on the number of CPU sockets and virtual guests. A system associated with a regular RHEL subscription is limited to two CPU sockets and one virtual guest. Each socket can have a multicore CPU. Significant discounts for academic users are available.

Get an Evaluation Copy

Red Hat currently offers a 30-day unsupported evaluation option for RHEL. Red Hat requires some personal information from such users. Once approved by Red Hat, you'll get instructions on how to download the distribution. For more information, see www.redhat.com/rhel/details/eval. The 30-day limit is simply a limit on access between your system and the Red Hat Network. The operating system continues to work, but without updates. No support is given.

Third-Party Rebuilds

You don't have to pay for operating system support to prepare for Red Hat exams. You don't have to live with a limited test subscription. To comply with the Linux General Public License (GPL), Red Hat releases the source code for just about every RHEL package at ftp.redhat.com. However, the GPL only requires that Red Hat release the source code. Red Hat's understanding is that it does not have to release the binary packages compiled from that source code.



The description of the GPL and trademark law in this book is not a legal opinion and is not intended as legal advice.

Under trademark law, Red Hat can prevent others from releasing software with its trademarks, such as its red fedora symbol. Nevertheless, the GPL gives anyone the right to compile that source code. If they make changes, all they need to do is release their changes under the same license. And several “third parties” have taken this opportunity to remove the trademarks from the released source code, and compiled that software into their own rebuilds, functionally equivalent to RHEL.

The source code is released in Source RPM package format, which means the RPM packages can be built using the `rpm` commands described in Chapter 7. However, the building of a distribution, even from source code, is a tricky process. But once complete, the rebuild has the same functionality as RHEL. It is true, rebuild distributions don’t have a connection to and can’t get updates from the Red Hat Network. However, the Red Hat Network is not part of the Red Hat Exam Prep guide. And the developers behind rebuild distributions also use the source code associated with new RHEL packages to keep their repositories up to date. Two options for rebuild distributions include:

- **Community Enterprise Operating System (CentOS)** The rebuild known as CentOS includes a number of experienced developers who have been working with RHEL source code since the release of RHEL 3 back in 2002. For more information, see www.centos.org.
- **Scientific Linux** This distribution is developed and supported by experts from the U.S. Government’s Fermilab and the European Organization for Nuclear Research, known by its French acronym, CERN. The people associated with these labs are among the smartest scientists around. For more information, see www.scientificlinux.org.



A number of the figures in this book are based on “snapshots” from a distribution that uses RHEL 6 source code, known as Scientific Linux. It was developed by two organizations with some of the smartest scientific minds in the world at Fermilab and CERN.

Check the Download

For downloads from the RHN, Red Hat provides checksums based on both the Message Digest 5 (MD5) and the 256-bit Secure Hash Algorithm (SHA256). You can check these ISO files to the given checksum numbers with the `md5sum` and

`sha256sum` commands. For example, the following commands calculate those checksums for the initial RHEL 6 DVD:

```
# md5sum rhel-server-6.0-x86_64-dvd.iso
# sha256sum rhel-server-6.0-x86_64-dvd.iso
```

While it's a good sign when a downloaded DVD passes these tests, such a result is not a guarantee that such DVDs are free of errors.

CERTIFICATION OBJECTIVE 1.04

Installation Requirements

According to the Red Hat certification blog, Red Hat now provides “pre-installed systems” for their exams. So you won't start from scratch, at least for the host physical system. But you'll still need to set up practice systems. The RHCSA objectives suggest that you need to do so with a network installation. On a pre-installed system, given the other requirements, that suggests that you need to know how to set up network installations on KVM-based VMs.

The installation requirements described in this section are suited to the creation of an environment for practice labs. That environment may also work as a baseline for other RHEL systems. On many real networks, new virtual systems are created or cloned from that baseline. Those new systems are then dedicated for a single service.

If you're creating a physical host for test VMs, make sure to have enough room available for the host physical system and the guest VMs. This section suggests that you create three VMs for test purposes. You should consider including a fourth VM as a backup, in case one of the VMs suffers a catastrophic failure due to a mistake or malfunction. For such purposes, 80GB of free space on a physical system would be sufficient. With some careful planning, you may be able to live with a smaller amount of free space and just two VMs. For more information how to configure RHEL 6 on a VM, see Chapter 2.



Since Linux filesystems are efficient, they can be filled to near capacity with minimal defragmentation issues. If you're configuring VMs on Microsoft-formatted partitions, considerable additional free space is required due to volume fragmentation.

You Won't Start from Scratch

Before installing RHEL 6, it may be helpful to review what is known about the latest RHCSA and RHCE exams. As described in the Red Hat blog announcement at <http://redhatcertification.wordpress.com/>, Red Hat now provides:

- Pre-installed systems
- Questions presented “electronically”

In other words, when seated for an exam, you'll see an installed copy of RHEL 6 on the test system, with questions in some electronic format. No public information is available on the format of the questions. This book will assume the most basic format for Red Hat exam questions, text files available in the root administrative user's home directory, `/root`.

The Advantages of Network Installation

Network installation means you don't have to use a full DVD on every system when installing RHEL 6. It means that every system is installed from the same set of installation packages. Network installations are faster than those from physical DVDs.

Network installations become especially powerful when combined with Kickstart files and the Pre-boot eXecution Environment (PXE). In that configuration, all you need to do to install RHEL 6 is boot a system, point the remote installer to the appropriate Kickstart file, and voila! After a few minutes, you'll have a complete RHEL 6 system.

Red Hat and Virtual Machines

The objectives associated with the RHCSA suggest that you need to know how to “configure a physical machine to host virtual guests.” It also suggests that you need to know how to perform a number of tasks with VMs, and “install Red Hat Enterprise Linux automatically using Kickstart.” That's consistent with the use of Kickstart files to set up RHEL 6 on a KVM-based VM.

One of the advantages of a VM is how it supports the use of an ISO file on a virtual CD/DVD drive. Files accessed from that virtual drive are not slowed by the mechanical speed of physical CD/DVD media. And as such access may not be slowed by network traffic, virtual CD/DVD drives may be as fast as network access from a host system.

Virtual and Physical Systems

Virtual systems can't stand alone. They require some connection to a physical system. Even “bare-metal” virtualization solutions such as VMware ESX and Citrix XenSource were built from or otherwise rely on specialized versions of the Linux kernel, which acts as the operating system on the physical host.

However, it's possible to install a substantial number of virtual systems on a single physical system. If those systems are dedicated to different services, they'll load the physical system at different times. Such loads makes it possible to “overbook” the RAM and other resources of the physical system.

For our purposes, there is no real difference whether the installation is performed on a physical or a virtual system. The software functions in the same way. As long as IP forwarding is enabled on the physical host system, networking on the virtual system works in the same way as well.

A Pre-installed Environment for Practice Labs

The baseline RHEL 6 system configured in this chapter is relatively simple. It starts with a 12GB virtual disk. Part of that disk will be organized as shown in Table 1-1. They will be configured as regular partitions. The remaining space on the hard drive will be left empty, for potential configuration during the exam as logical volumes.

Two additional virtual disks of 1GB each are included to facilitate the post-installation configuration of a logical volume. The 12GB hard disk and 8GB partition are arbitrary sizes that provide plenty of room for RHEL 6 software. If space is limited on your system, you might go as low as 8GB for a hard disk, as long as swap space is also appropriately limited. Swap space in Linux is used as an extension of local RAM, especially when that resource runs short.

The baseline minimum installation of RHEL 6 does not include a GUI. While it is fairly easy to install the package groups associated with the GUI after installation is complete, that process requires the installation of several hundred MB of packages. And that takes time. Since Red Hat provides a pre-installed system for the exam

TABLE 1-1

	Location	Size
Model Partitions	/boot	500MB
	/	8GB
	/home	1024MB
	Swap	1024MB

to reduce the time required for the exam, it is reasonable to suggest that the system provided by Red Hat includes the GUI. And the default GUI for Red Hat systems is the GNOME Desktop Environment.



GNOME is an acronym within an acronym. It stands for the GNU Network Object Model Environment. GNU is itself a recursive acronym, as it stands for GNU's Not Unix. Linux is filled with similar recursive acronyms, such as PHP: Hypertext Preprocessor (PHP).

The amount of RAM to allocate is more complex, especially on a VM. For the purpose of this book, I've configured VMs with 768MB of RAM to comfortably enable GUI-based illustrations of the RHEL installation process. If text-mode installations are acceptable, you can run a RHEL 6 GUI in 512MB, or possibly even less RAM. As different VMs rarely use the same RAM simultaneously, it's possible to "overbook" RAM; for example, it may be possible to set up three VMs, with 1GB of RAM each, on a physical host system with less than 3GB of physical RAM. Some RAM on the VMs will remain unused, available to the physical host system.

System Roles

Ideally, you can set up several systems, each dedicated to different roles. A network with a dedicated DNS (Domain Name Service) server, a dedicated DHCP (Dynamic Host Configuration Protocol) server, a dedicated Samba file sharing server, and so on, is more secure. In that situation, a security breach in one system does not affect any other services.

However, that's not practical, especially during the Red Hat exams. Table 1-2 lists the roles appropriate for each of the three systems described in Lab 1.

TABLE 1-2

Roles for Test Systems

System	Roles
server1	Workstation and servers to be configured throughout the book, configured as server1.example.com on the 192.168.122.0/24 network. This book assumes a fixed IP address of 192.168.122.50.
tester1	Secure shell server that supports remote access, configured as tester1.example.com on the 192.168.122.0/24 network. May include servers for client testing, such as the Domain Name Service (DNS). This book assumes a fixed IP address of 192.168.122.150.
outsider1	Workstation on a third IP address, configured as outsider1.example.org. Some services should not be accessible from that workstation. This book assumes a fixed IP address of 192.168.100.100.

There's also another implicit fourth system in this network—the physical host for the virtual machines. When multiple networks are configured, that host will have virtual network adapters that connect to each network. For this book, I've set up a system named `maui.example.com`. The following excerpts from the `ifconfig` command display the `virbr0` and `virbr1` adapters, with connections to both networks:

```
virbr0 Link encap:Ethernet HWaddr 9E:56:D5:F3:75:51
       inet addr:192.168.122.1 Bcast:192.168.122.255 Mask:255.255.255.0
virbr1 Link encap:Ethernet HWaddr 86:23:B8:B8:04:70
       inet addr:192.168.100.1 Bcast:192.168.100.255 Mask:255.255.255.0
```

Of course, you can change the names and IP addresses associated with each of these systems. They are just the defaults to be used in this book. The `server1.example.com` system is the designated exam system, which will be used for exercises that address actual Red Hat exam requirements. For convenience, I've also set up some RHCE services on the physical host system.

The `tester1` system can be used to verify the configuration on the `server1` system. For example, if you've configured two virtual web sites with different names, you should be able to access both web sites from the `tester1` system. The Red Hat exams assume that you may connect a system as a client to servers such as Samba and LDAP. They also assume that a DNS server is configured with appropriate hostnames and IP addresses. While the configuration of some servers such as Kerberos is beyond the scope of the RHCSA/RHCE exams, they may be used during the exams by the other systems as clients.

Finally, the `outsider1` system is essentially a random system from an external network such as the Internet. Appropriate security settings mean that some services on the `server1` machine won't be accessible to `outsider1`.

As suggested earlier, it would be best to have a fourth virtual system available, in case of a failure in one of the three virtual systems described. Before following these recommendations, read Chapter 2. This chapter is focused on the configuration of a physical host system.

CERTIFICATION OBJECTIVE 1.05

Installation Options

Even most beginning Linux users can install RHEL 6 from a CD/DVD. While this section addresses some of the options associated with installation, it is focused on

the creation of that baseline system that can be used to set up other custom RHEL 6 systems.

In addition, the installation process is an opportunity to learn more about RHEL 6, not only boot media, but the logical volumes that can be configured after installation is complete. But as pre-installed physical systems are now the norm for Red Hat exams, detailed discussion of logical volumes have been consolidated in Chapter 6.

The steps described in this section assume a connection to the FTP server with RHEL 6 installation files created and configured in Lab 2. The steps are modestly different if you're installing RHEL 6 directly from a full DVD.

Boot Media

When installing RHEL 6, the simplest option is to boot it from the RHEL 6 DVD. While rebuild distributions may offer CDs, Red Hat only offers RHEL 6 on DVD. It also offers boot media in ISO format. From those media, there's no fuss about a separate boot disk, no worry about network connections. But in some organizations, you may not want to distribute the RHEL 6 DVD to everyone. In some cases, there is no CD/DVD drive. In essence, there are four methods available to start the RHEL 6 installation process:

- Boot from an RHEL 6 DVD.
- Boot from a network RHEL boot CD.
- Boot from a USB key.
- Boot from a Kickstart server using a PXE network boot card.

The last three options generally assume that you're going to install RHEL over a network. The installation and boot media are available from the Red Hat Network for users with a subscription. It should also be available from servers associated with rebuild distributions.

Some Red Hat documentation suggests that the network boot CD doesn't work on 64-bit systems with UEFI. For me, that was not true. It worked fine on both my UEFI-based laptop and server. For more information, see Red Hat Bugzilla item 661135. Nevertheless, if you need to create a boot USB key, find the `efidisk.img` file in the `images/` subdirectory of the RHEL 6 DVD. You can then write that image to a USB key. If that key is located on device `/dev/sdd`, you'd write that image with the following command:

```
# dd if=efidisk.img of=/dev/sdd
```

Be careful—if /dev/sdd is a drive with data, these commands will overwrite all data on that drive.



Know how to create the right boot disk for your system. If you have a problem, the installation boot CD or USB key can also serve as a rescue disk. At the boot prompt, the Rescue Installed System option will eventually bring you to a rescue mode that can mount appropriate volumes and recover specific files or directories.

CD/DVD or Boot USB Starts Installation

Now you can boot a target system from the network boot CD, the installation DVD, or the installation USB key. After a few files are opened and decompressed, an RHEL installation screen should appear with at least the following four options:

- Install or upgrade an existing system
- Install system with basic video driver
- Rescue installed system
- Boot from local drive

The first option should work for most users. If there is trouble with the graphics after the first option, try rebooting the system and work with the second option, which specifies a standard Video Electronics Standard Association (VESA) adapter associated with older Super Video Graphics Association (SVGA) monitors.

There are two modes associated with the Red Hat installation program, also known as Anaconda: text mode and graphical mode. While the minimum supported requirement is 512MB of RAM, you may be able to install graphically in a bit less. But if you have more than enough RAM and prefer text-mode installation, return to the RHEL installation screen. Highlight the Install Or Upgrade An Existing System option and press the TAB key. When you do, the following options are revealed on that screen, on one line:

```
> vmlinuz initrd=initrd.img
```

To force installation in text mode, add the word **text** to the end of this line.



While Red Hat published information suggests that a minimum of 512MB of RAM for a graphical installation, my experiments on KVM suggest that the minimum for a GUI installation on a VM is 652MB.

Basic Installation Steps

The basic RHEL installation is straightforward and should already be well understood by any Red Hat certification candidate. Most of the steps are described here for reference; it's useful to remember this process as you work on advanced configuration situations such as the Kickstart files described in Chapter 2.

The order of these steps vary depending on whether they're run directly from the CD/DVD or over a network. Variations occur depending on whether there's a previous version of Linux and Linux-formatted partitions on the local system. For this section, the following assumptions are made:

- Network installation based on the RHEL 6 network boot disk
- At least 652MB of RAM
- An available FTP server with the installation files, such as the one configured in Lab 2.
- RHEL 6 as the only operating system on the local computer

However, dual-boot situations are acceptable. In fact, I've written this book on an Intel I7 laptop system in a triple-boot configuration where RHEL 6 co-exists with Windows 7 and Ubuntu 10.04. If you're installing the system on a dedicated physical computer or a VM, the basic steps are the same. As a physical host is required for VMs, I assume you'll be first installing RHEL 6 on a physical system.

The most efficient, and thus (in my opinion) the most likely, way to install Red Hat Enterprise Linux is via a text or graphical installation from a remote server. For that purpose, Lab 2 configures an FTP server with the RHEL 6 installation files. Alternatively, it's possible to set up those installation files on an HTTP server such as the Apache Web server, as discussed later in this chapter.

The sequence of steps for the installation process varies, depending on whether you're installing from the DVD or the network installation CD, as well as whether you're installing in text or graphical mode. It also may vary if you're using

a rebuild distribution of RHEL 6. In fact, these instructions are something of a hybrid between two different types of installations, depending on what you do in Step 2. In fact, the developers behind rebuild distributions have modified the installation steps (slightly) as well. What you see will likely differ from the steps shown here. So be flexible when reading these instructions.

1. Boot your computer from the RHEL DVD, the RHEL network boot CD, or a boot USB key. This procedure assumes you're using the network boot CD. Five options are normally shown:

- Install Or Upgrade An Existing System
- Install System With Basic Video Driver
- Rescue Installed System
- Boot From Local Drive
- Memory Test

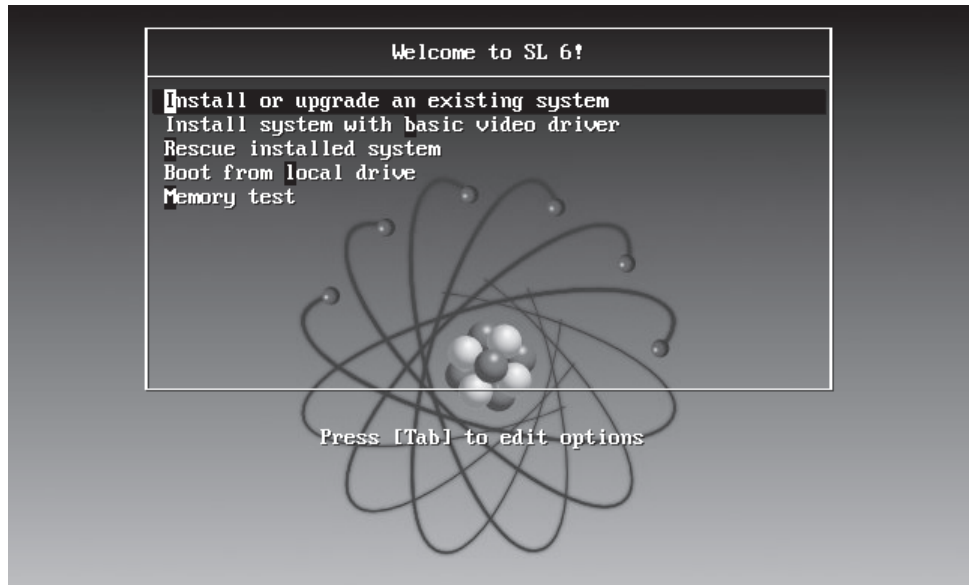
2. Figure 1-1 illustrates the options from the Scientific Linux rebuild. The options are the same as from a genuine RHEL 6 DVD. For installation, only the first two options matter. Try the first option. If there's trouble with the graphics part way through the installation, restart the system and try the second option listed.

If you're booting from the first RHEL 6 installation DVD, and want to be sure to review the options shown in Step 4, press tab, add a space and the word **askmethod** to the command line that appears, and then press enter. However, that first brings up the text-mode versions of the Choose A Language And Keyboard Type screens described in Steps 9 and 10. In addition, that would skip the Disc Found screen described in the next step.

3. The system responds with a Disc Found screen. The first steps of the network installation process starts in text mode (which is actually a low-resolution graphical mode), even if sufficient RAM is available. Choices are not "clickable." In this mode, use the TAB key to switch between options, and the SPACE or ENTER key to select or deselect an option.
4. If you like, you can use the Disc Found screen to test the integrity of the media. While it's good if the disc passes this test, it does not provide a guarantee that such media are free of errors. If you choose to test, be aware that the media is ejected after the test is complete. Accept the test or skip it and proceed to Step 4.

FIGURE 1-1

The Installation
Boot Screen



If you test the integrity of an ISO file used for CD/DVD media in a KVM-based VM, the system will “eject” that file without changing the status of the media in the VM. You’ll have to disconnect and then restore the media in the KVM Detail screen, as discussed in Chapter 2.

5. Choose your installation method. Four options are shown in Figure 1-2.


The options are straightforward:

- Local CD/DVD can be used if the CD/DVD (or equivalent ISO file) is loaded.
- Hard Drive assumes the RHEL 6 ISO file is available on a local hard drive partition. If you select this option, the program prompts you to specify the volume and directory with that file.
- NFS Directory assumes the installation files are available from a shared NFS directory.
- URL, short for Uniform Resource Locator, works with installation files stored on both an Apache Web server and an FTP server.

To point to the FTP server configured in Lab 2, select URL and click OK.

FIGURE 1-2


Select an
Installation
Method.



6. In the Configure TCP/IP screen, shown in Figure 1-3, choose how you want IP addressing configured. Your options are to enable support as a DHCP client for IPv4 and/or IPv6 addresses. (If the network DHCP server, such as a home router, does not support IPv6, enabling this option with DHCP slows the installation process.) As fixed IPv4 addresses will be needed for the test systems, select Manual Configuration, at least for IPv4 addresses. Fixed IPv4 addresses will be set up for all three systems from Table 1-2. In any case, KVM currently only supports IPv4.
7. Specify an IPv4 address for the system, the gateway, and the name server. If you're not sure what to do, this is an excellent time to plan a network as described in Lab 1. If you've configured a system on a VM, the gateway address is probably something like 192.168.122.1 or 192.168.100.1, on the same subnet as the IP address. For the purpose of this book, I've designated an IPv4 address for the server1.example.com system of 192.168.122.50. The name server, another name for the DNS server, is typically located on the same IP address as the gateway. Alternatively, for a physical system on a home network, the appropriate IP address for both the gateway and the name server is the IP address of the home router.

FIGURE I-3


Configure the type of network addressing.



- Direct your computer to the remote FTP server. As shown in Figure I-4, you can enter the hostname or IP address of the FTP server, as well as the shared directory.

FIGURE I-4

Enter the URL of the installation server.

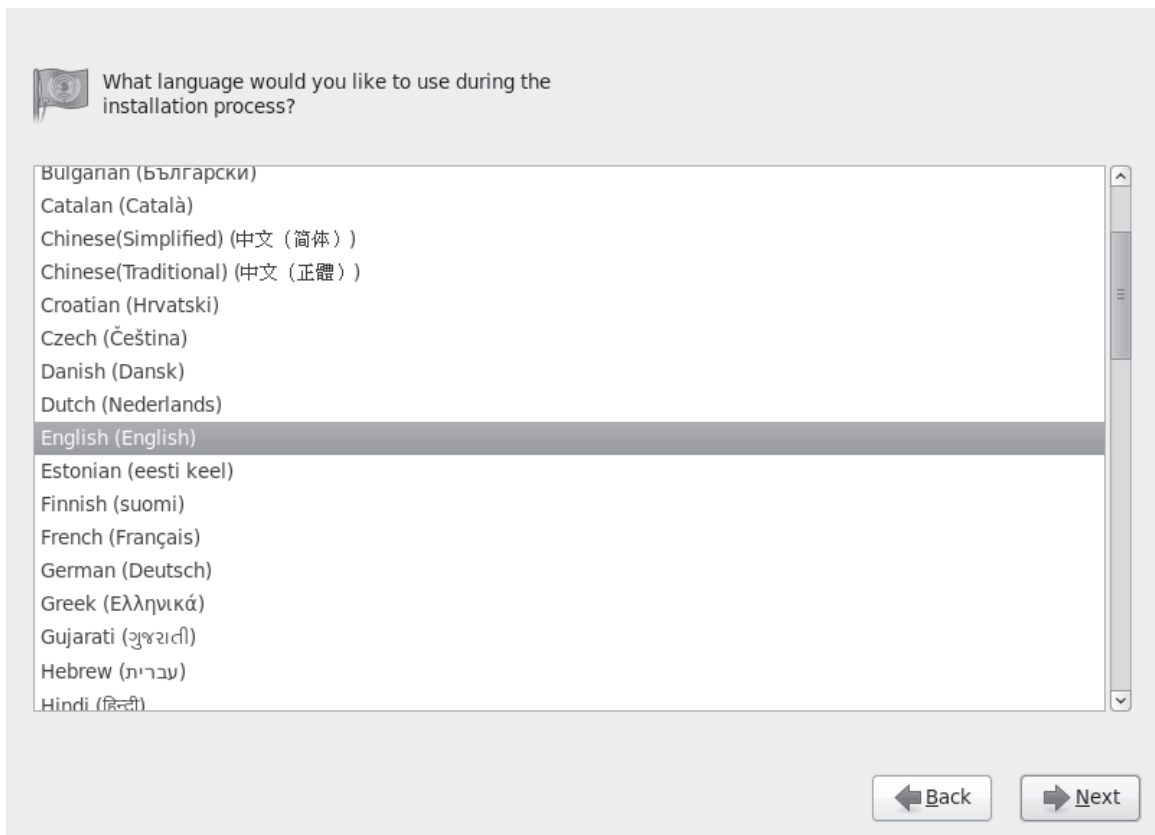


9. If a proper connection is made, and the files are available, you'll see the first installation screen. Assuming sufficient RAM, it's a graphical screen. Click Next to continue.
10. Select a language to use during the installation process, as shown in Figure 1-5. English is the default; over 50 options are available.
11. Select a keyboard type; the default depends on the language selected.



If you encounter problems, examine the messages in the third, fourth, and fifth consoles; to do so, press ALT-F3, ALT-F4, or ALT-F5. A command line is available by pressing ALT-F2. To return to the GUI screen, press CTRL-ALT-F6. If in text-mode installation, you can return to that screen by pressing ALT-F1.

FIGURE 1-5 Select a language for installation.



12. The next screen relates to local or specialized storage devices. If you have one or more “regular” local hard drives (SATA, PATA, or a virtual hard drive on a KVM system), select Basic Storage Devices, and click Next.
13. If there are new hard disks being applied to the system, you’ll be asked to confirm changes to those disks.
14. If there’s a previous version of RHEL (or a similar recognized Linux operating system) installed, you may see a Fresh Installation option. If it appears, select it. Otherwise, Anaconda skips to the next step.
15. You’re now asked to give a hostname for the local system, such as server1.example.com. Click Configure Network to open the Network Connections tool described in Chapter 5. In some cases, the network settings configured in Step 6 may not be reflected in this tool. Make appropriate changes and click Next to continue.
16. Now you’ll see a world map, where you can select the time zone of the local system. The System Clock Uses UTC option is a reference to the local hardware clock and the atomic realization of Greenwich Mean Time. (UTC is a non-English language acronym also based on a political compromise.) While incompatible with Microsoft Windows, the UTC option supports changes for daylight saving time. Make appropriate changes and click Next to continue.
17. The next step is to enter the password for the root administrative user, twice. Do so and click Next to continue.
18. The next step, shown in Figure 1-6, determines how space on configured hard disks, local and remote, is used. The options are fairly well explained in the figure. To summarize:
 - **Use All Space** Removes all partitions on all configured devices, including those created by non-Linux operating systems.
 - **Replace Existing Linux System(s)** Removes all Linux-formatted partitions, including those created for other Linux distributions such as Fedora and Ubuntu Linux.
 - **Shrink Current System** Takes account of unused space on available existing partitions, shrinks those partitions, allowing that free space to be used for the new installation.
 - **Use Free Space** Uses existing free space for the new installation.
 - **Create Custom Layout** Supports custom configuration using Disk Druid, the Red Hat disk partitioning tool, during the installation process.

- **Encrypt System** Supports encryption of the partitions created during the process. You'll learn to encrypt and manage an existing partition in Chapter 6.
- **Review And Modify Partitioning Layout** Starts the Disk Druid utility described in the next section.

For the purpose of this installation, select Create Custom Layout. Note how that blanks out the deselection of the Encrypt System and the selection of the Review And Modify Partitioning Layout options. Click Next to continue.

19. If there's more than one hard drive installed, you'll have to select those drives on which RHEL 6 will be installed. One example is shown in Figure 1-7. Per the baseline discussion earlier in this chapter, choose the 12GB drive. The actual size will be shown in MB and will vary by a few percent. That variance is irrelevant on the Red Hat exams.
20. In this case, select the 12GB drive by clicking the arrow pointing right. You should now see that 12GB drive in the right-hand column, listed in the

FIGURE 1-6

Partition layout options

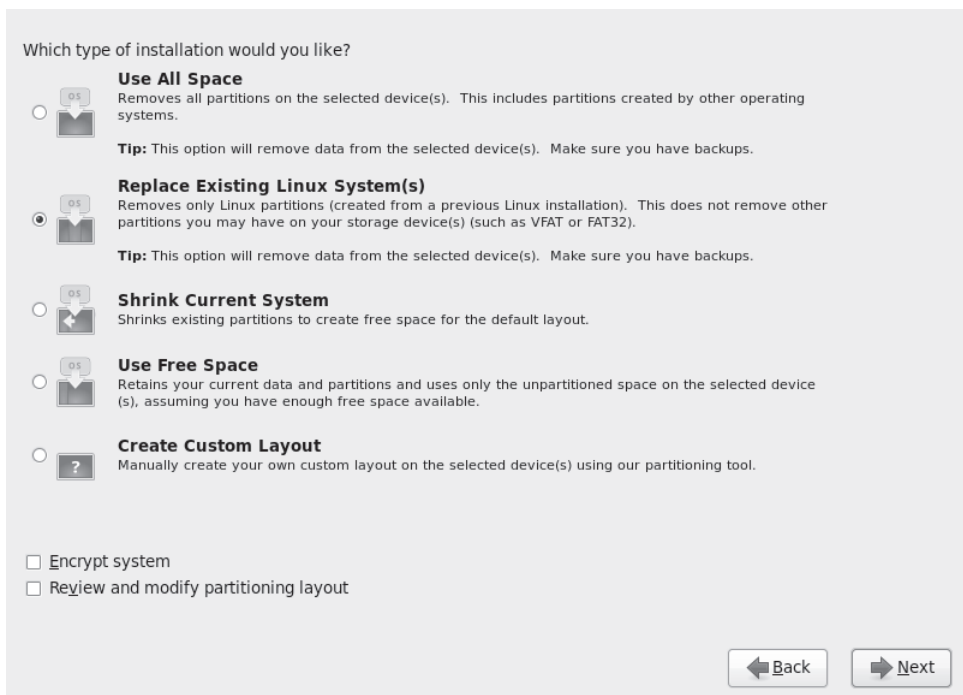
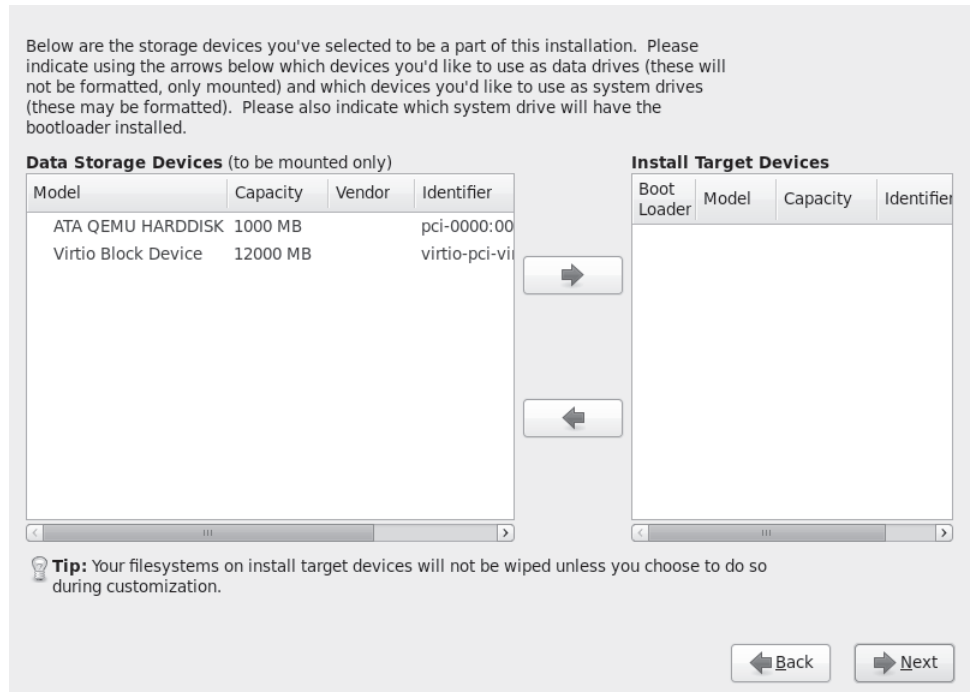


FIGURE 1-7

Select from available local hard drives.



Install Target Devices box. Click Next to continue. You'll see the Please Select A Device screen shown in Figure 1-8, the main Disk Druid screen. Detailed discussion of this utility continues in the next section.

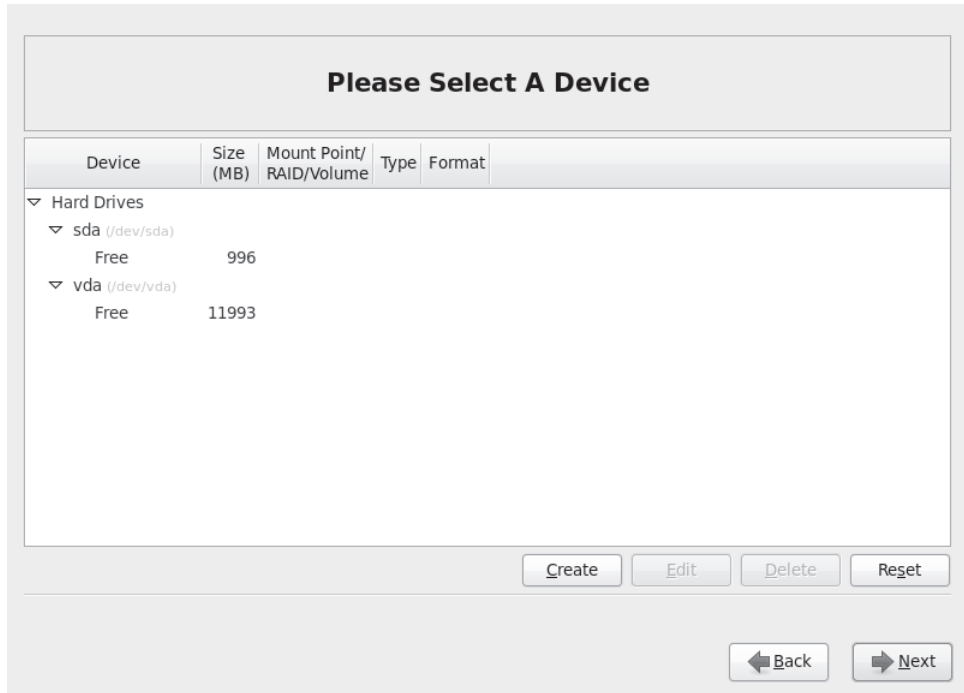
The Installation Perspective on Partitions

Although it's possible to create more, RHEL will recognize only up to 16 partitions on any individual SATA, SCSI, PATA, or virtual hard drive. Once a partition is created, you can configure Linux to mount a directory directly on that partition. Alternatively, that partition can be designated as a RAID device or as part of a logical volume.

To define a partition, you may need some background on naming conventions, the configuration of different filesystems, uses of swap space, logical volumes, and RAID arrays. This is just an overview. Detailed information is available in Chapter 6, including tasks that may be required during the Red Hat exams, and on real systems.

FIGURE I-8

Partition configuration at the Please Select A Device screen.



Naming Conventions

Linux has a simple naming standard for disk partitions: three letters followed by a number. The first letter identifies the type of drive (*s* is for PATA, SATA, or SCSI; and *v* is for virtual disks on KVM-based VMs). The second letter is *d* for disk, and the third letter represents the relative position of that disk, starting with *a*. For example, the first SATA drive is *sda*, followed by *sdb*, *sdc*, and *sdd*.

The number that follows is based on the relative position of the primary, extended, or logical partition. Primary partitions can contain the boot files for an operating system. Hard drives can also be configured with one extended partition, which can then contain a number of logical partitions.

Hard disks are limited to four primary partitions. When four partitions are not enough, an extended partition can be substituted for the last primary partition. That extended partition can then be subdivided into logical partitions. So when planning a partition layout, make sure that extended partition is big enough.

Each partition is associated with a Linux device file. At least that is straightforward; for example, the device filename associated with the first logical partition on the first SATA drive is `/dev/sda5`.

A volume is a generic name for a formatted segment of space that can be used to contain data. Volumes can be partitions, RAID arrays, or those logical volumes associated with Logical Volume Management (LVM).

A filesystem is the way a volume is formatted to allow it to store files. For example, Red Hat uses the fourth extended filesystem (`ext4`) as the default format for its volumes. The standard way to access data in Linux is to first mount that filesystem onto a directory. For example, when the `/dev/sda1` partition is formatted to the `ext4` filesystem, it can then be mounted on a directory such as `/boot`. It is common to say something like, “The `/dev/sda1` filesystem is mounted on the `/boot` directory.” For more information, see Chapter 6.

Separate Filesystem Volumes

Normally, you should create several volumes for RHEL 6. Even in the default configuration, RHEL is configured with at least three volumes, for a top-level root directory (`/`), a `/boot` directory, and Linux swap space. Additional volumes may be suitable for directories such as `/home`, `/opt`, `/tmp`, and `/var`. They’re also suitable for any custom directories such as for web sites, dedicated groups of users, and more. While it’s important to configure the `/boot` directory on a regular partition, other directories can readily be configured on logical volumes or RAID arrays.

Dividing the space from available hard drives in this manner keeps system, application, and user files isolated from each other. This helps protect the disk space used by the Linux kernel and various applications. Files cannot grow across volumes. For example, an application such as a web server that uses huge amounts of disk space can’t crowd out space needed by the Linux kernel. Another advantage is that if a bad spot develops on the hard drive, the risk to your data is reduced, as is recovery time. Stability is improved.

While there are many advantages to creating more volumes, it isn’t always the best solution. When hard drive space is limited, the number of partitions should be kept to a minimum. For example, if you have a 4GB hard drive and want to install 3000MB of packages, a dedicated `/var` or even a `/home` volume could lead to situations where disk space runs out far too quickly.

Linux Swap Space

Linux swap space is normally configured either on a dedicated partition or a logical volume. Such space is used to extend the amount of effective RAM on a system, as virtual memory for currently running programs. But you can't just buy extra RAM and eliminate swap space. Linux moves infrequently used programs and data to swap space even if you have gigabytes of RAM. As such, RAID arrays of swap space make little sense, as why would anyone back up fragments of data from RAM?

The way Red Hat assigns default swap space is based on the amount of RAM on a system and the space available in local hard drives. For systems of up to 2GB, the default swap space size is twice the amount of installed RAM. Above 2GB, it's the amount of RAM + 2GB. But those are not "hard and fast" rules. Workstations with several GB of RAM frequently use very little swap space. On my home server, I have 8GB of RAM and 4GB of swap space. That swap space is rarely used, but it may be used more frequently on systems that aren't rebooted for months at a time or have heavy demand from certain services. In any case, the default installation configures swap space not in a dedicated partition, but as a logical volume.

Basic Information on Logical Volumes

The creation of a logical volume from a partition requires the following steps. Details on these concepts as well as the actual commands required to execute these steps are described in Chapter 6. Some of these steps are run automatically if you create a logical volume during the installation process.

- The partition needs to be labeled as a logical volume.
- The labeled partition can then be initialized as a physical volume.
- One or more physical volumes can be combined as a volume group.
- A volume group can be subdivided into logical volumes.
- A logical volume can then be formatted to a Linux filesystem or as swap space.
- The formatted logical volume can be mounted on a directory or as swap space.

Basic Information on RAID Arrays

RAID was an explicit requirement on the RHCT/RHCE exams up to the release of RHEL 6. As it is no longer found in either the RHCSA/RHCE objectives or the outlines of the prep courses for these certifications, you can relax a bit on that topic.

In any case, the RAID configured on RHEL 6 is software RAID. The acronym, Redundant Array of Independent Disks, is somewhat misleading, as software RAID is usually based on independent partitions. Redundancy comes from the use of partitions from different physical hard drives.

Partition Creation Exercises

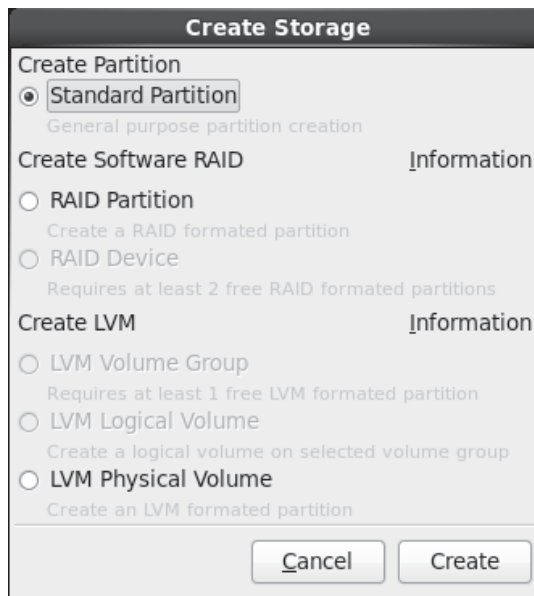
Now return to the installation process. If you followed the steps described so far in this chapter and the system has sufficient RAM, you should see the Disk Druid Please Select A Device screen shown in Figure 1-8.

At this screen, you have the opportunity to configure partitions, logical volumes, and RAID arrays.

1. Configure standard partitions as described earlier in Table 1-1. Larger partitions are acceptable if you have the space. They would be necessary if you're creating the physical host system that will contain the VMs. The Create button supports the creation of standard partitions, logical volumes, and RAID arrays, as shown in Figure 1-9. Select Standard Partition and click Create to continue.

FIGURE 1-9

You can create a variety of storage devices.



2. You should now see the Add Partition window shown in Figure 1-10. It supports a number of choices:
 - Allowable drives. For the purpose of this installation, limit the partitions to be configured to the virtual hard drive, labeled as the Virtio Block Device.
 - The Mount Point, which is the directory (such as /boot) whose files will be stored on the partition. You can type in the mount point; alternatively, it's a drop-down text box that provides options for typical mount point directories.
 - The File System Type; the default ext4 filesystem is sufficient. Click the box; options for other formats, along with configuration as a physical volume, RAID array component, or swap space, are also available.
 - The Size of the partition in MB; in this case, the partitions to be configured for this baseline system are defined in Table 1-1.

Now it's time for some exercises. First, examine how you can create and configure partitions during the installation process. You'll also look at how to allocate a filesystem to a partition or a logical volume.

FIGURE 1-10

Add a partition.

The screenshot shows the 'Add Partition' dialog box with the following configuration:

- Mount Point:** A text input field with a dropdown arrow.
- File System Type:** A dropdown menu showing 'ext4'.
- Allowable Drives:** A list box containing two entries:

<input type="checkbox"/>	sda	1000 MB	ATA QEMU HARDDISK
<input checked="" type="checkbox"/>	vda	12000 MB	Virtio Block Device
- Size (MB):** A text input field containing '200'.
- Additional Size Options:**
 - Fixed size
 - Fill all space up to (MB):
 - Fill to maximum allowable size
- Force to be a primary partition
- Encrypt

Buttons for 'Cancel' and 'OK' are located at the bottom right of the dialog.

EXERCISE 1-1**Partitioning During Installation**

This exercise is based on an in-process installation of RHEL 6. Mistakes are easy to recover from on VMs, as it's fairly easy to restart the installation process. This exercise starts with the Please Select A Device screen shown in Figure 1-8 and continues with the windows shown in Figures 1-9 and 1-10. In addition, it assumes sufficient RAM to work with the graphical installation.

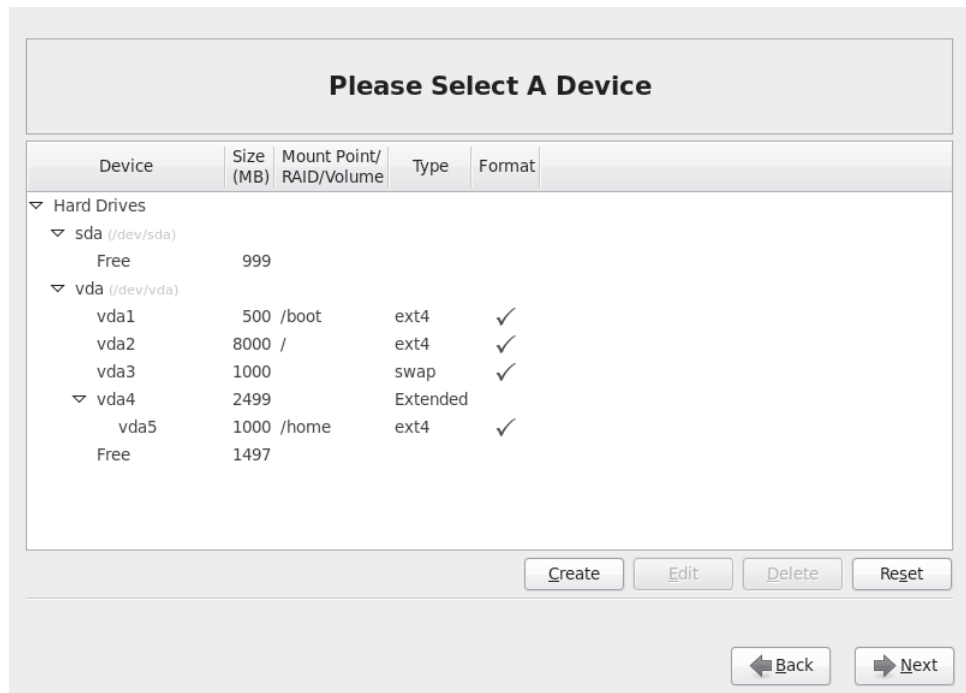
1. Start creating a custom layout. If you're starting with blank hard disks, no partitions will be configured. Delete configured partitions if no space is available.
2. Try to create a regular partition. Click Create to open the Create Storage window. Select Standard Partition and click Create to open the Add Partition window shown in Figure 1-10. Set up an appropriate mount point, such as /home/user. Click the File System Type drop-down text box and review the available formats. If more than one allowable drive is available, make sure an appropriate drive is selected.
3. Retain the default ext4 File System Type, and click OK to continue.
4. Create one additional partition, using the steps just described. For the purpose of this exercise, the default 200MB is good enough, assuming sufficient free space is available.
5. Now click Create again, select LVM Physical Volume (LVM), and click Create. Note how it opens the Add Partition window with the Physical Volume (LVM) File System Type. If more than one allowable drive is available, make sure an appropriate drive is selected. Click OK.
6. Repeat the preceding step to create a second LVM partition. If more than one allowable drive is available, make sure to select a drive different from that chosen in Step 5.
7. Click Create. You should now be able to select LVM Volume Group. Do so and click Create.
8. In the Make LVM Volume Group window, click the Physical Extent drop-down text box. Review the available Physical Extents, which are units associated with volume groups. Typically, no changes are needed.
9. Make sure all available Physical Volumes To Use are active.

10. Click Add; this opens the Make Logical Volume window.
11. Enter an appropriate mount point such as /home/volume. The Logical Volume Name shown is just the default; you can use any legal filename for your logical volume. Set a size that does *not* use all available space. Click OK.
12. Review the result in the Make LVM Volume Group window. Click OK and review the result in the original partition window.
13. Now it's time for some clean-up. Click Cancel to return to the main partitioning window. Click Reset to restore the original Disk Druid configuration before the start of the exercise.

Now that the exercise is complete, the partition configuration should reflect at least the minimums shown in Table 1-1. One version is shown in Figure 1-11. If a mistake is made, highlight a partition and click Edit. The Edit Partition window that appears includes the same options shown. Different partitions and modest variations in size are not relevant on the Red Hat exams.

FIGURE 1-11

Sample partition configuration



To complete this part of the process, click Next. If you see a Format Warnings window, there may be existing data on the partitions that have been created. This is your last chance to cancel before proceeding. Assuming you're satisfied, click Format to continue, and in the next window, click Write Changes To Disk.

Configure the Bootloader

The standard Linux bootloader is GRUB, the GRand Unified Bootloader. While RHEL 6 uses a slightly older version of GRUB, version 0.97, it's a "tried and true" bootloader. The settings shown in Figure 1-12 are reasonable defaults for a dual-boot configuration. On a VM, the screen will be simpler. In most cases, no changes are required.



The terms "boot loader" and "bootloader" are interchangeable. Both are frequently found in Red Hat documentation.

FIGURE 1-12

Configure the
bootloader.

Install boot loader on /dev/sda.

Use a boot loader password

Boot loader operating system list

Default	Label	Device
<input type="radio"/>	Other	/dev/sda1
<input checked="" type="radio"/>	Exam Prep System	/dev/sda8

exam**Watch**

The Linux Loader (LILO) hasn't been supported by Red Hat for nearly a decade. In addition, while

GRUB 2.0 is available, it is also not supported for RHEL 6, at least as of the initial release.

- The Install Boot Loader On /dev/sda option would install the bootloader on the master boot record of the noted hard drive. Unless another bootloader is installed, that is appropriate.
- The Use A Boot Loader Password can help secure the system. As you'll see in Chapter 5, without a bootloader password, anyone with access to the boot menu can access Linux with full administrative privileges just by booting into runlevel 1.
- The Boot Loader Operating System List specifies a list of detected operating systems. While you might see a second entry for an operating system such as Microsoft Windows in real life, RHEL 6 should be the only operating system installed here for the Red Hat exams.

While it's possible to add an entry from this menu for other operating systems, it won't work for other Linux installations on the same system. For that purpose, it's more effective to modify the GRUB bootloader configuration file directly after installation is complete.

While you may choose to set up a bootloader password, the defaults should be acceptable on any system where RHEL 6 is the only operating system installed on the local machine. Make any appropriate changes and click Next to continue.

Wow, Look at All That Software!

There are over 2500 packages available just from the RHEL 6 installation DVD. That number does not include a number of packages available only through the Red Hat Network. With so many packages, it's important to organize them into groups. After configuring the GRUB bootloader, you'll see the options shown in Figure 1-13, which allows you to configure the local system to a desired functionality. The selection depends on your objective. If you're installing on a physical system to set up KVM-based virtualization, select Virtual Host. If you're setting up virtual guests

FIGURE 1-13

Functional
installation
options

The default installation of Red Hat Enterprise Linux is a basic server install. You can optionally select a different set of software now.

Basic Server
 Database Server
 Web Server
 Virtual Host
 Desktop
 Software Development Workstation
 Minimal

Please select any additional repositories that you want to use for software installation.

High Availability
 Load Balancer
 Red Hat Enterprise Linux
 Red Hat Enterprise Linux

You can further customize the software selection now, or after install via the software management application.

Customize later Customize now

(or other dedicated servers), select Basic Server. During a Red Hat exam, you'll be installing most additional software after basic operating system installation is complete. Other options are listed in Table 1-3. Depending on the rebuild distribution, the options may vary significantly.

TABLE 1-3

Installation Software Categories

Category	Description
Basic Server	Installs basic packages for Red Hat as a server
Database Server	Includes MySQL and PostgreSQL database packages
Web Server	Sets up a system with the Apache Web server
Virtual Host	Configures a system with the KVM VM system
Desktop	Includes desktop productivity software
Software Development Workstation	Adds tools to modify and compile software
Minimal	Includes a minimal list of packages for the operating system

For a truly secure baseline in a production environment, consider the minimal installation. Fewer packages means fewer vulnerabilities. You can then add just the packages needed for the desired functionality. Any software that isn't installed can't be exploited by a cracker.



In the world of Linux, the term “hacker” refers to good people who want to create better software. The term “cracker” refers to people who want to break into other systems with evil intent.

The repositories listed can be useful for additional server functionality. As suggested by their names, the unselected repositories shown in Figure 1-13 support the installation of software in a number of categories, including storage clusters, systems that require high availability, and systems that balance the load for high-traffic services. As these are Red Hat repositories, maintained by Red Hat, some third-party repositories may not duplicate the availability or functionality of some of the repositories shown in the figure.

If you want to redirect systems to local or internal repositories, to control the packages others install on their systems, click Modify Repository. That allows you to specify where the local system looks for new packages and updates.



Some groups of developers enable the installation of software not supported by Red Hat. Two examples can be found at <http://atrpms.net> and <http://rpmrepo.org>.

For the purpose of this chapter, retain the default Red Hat Enterprise Linux repository (or the defaults for a rebuild distribution such as CentOS/Scientific Linux). Make sure to select Customize Now; otherwise, you won't be able to follow along with the next section. Click Next to continue.

Baseline Packages

In this section, you'll get a basic overview of what's available during the RHEL 6 installation process. During the exams, you may refer to one of these package groups with the Red Hat Add/Remove Software tool. You can also find a list of available package groups with the **yum grouplist** command. More information is available in Chapter 7.

Red Hat package groups are organized logically; for example, packages associated with a specific language can be found in the Languages package group. It's important

to choose only the package groups you need. Fewer installed packages means more room for personal files, as well as the log files needed to monitor systems.

Package Groups

This section includes the briefest possible overview of each of the package groups available during the RHEL installation process. As you can see from Figure 1-14, there are high-level groups in the left-hand pane, such as Desktops, and regular package groups in the right-hand pane, such as the X Window System. The details of the RPMs associated with each package group are stored in an XML file. To review that file, go to the RHEL installation DVD and read the compressed *-comps-rhel6-Server.xml.gz file in the /repodata directory.

For an example of the details within a package group, select Desktop and click Optional Packages. This opens the Packages In Desktop window shown in Figure 1-15. Compare this list to the aforementioned XML file. Mandatory packages aren't shown in the associated window, as their installation is required for the package group. As

FIGURE 1-14

Red Hat
Enterprise Linux
package groups

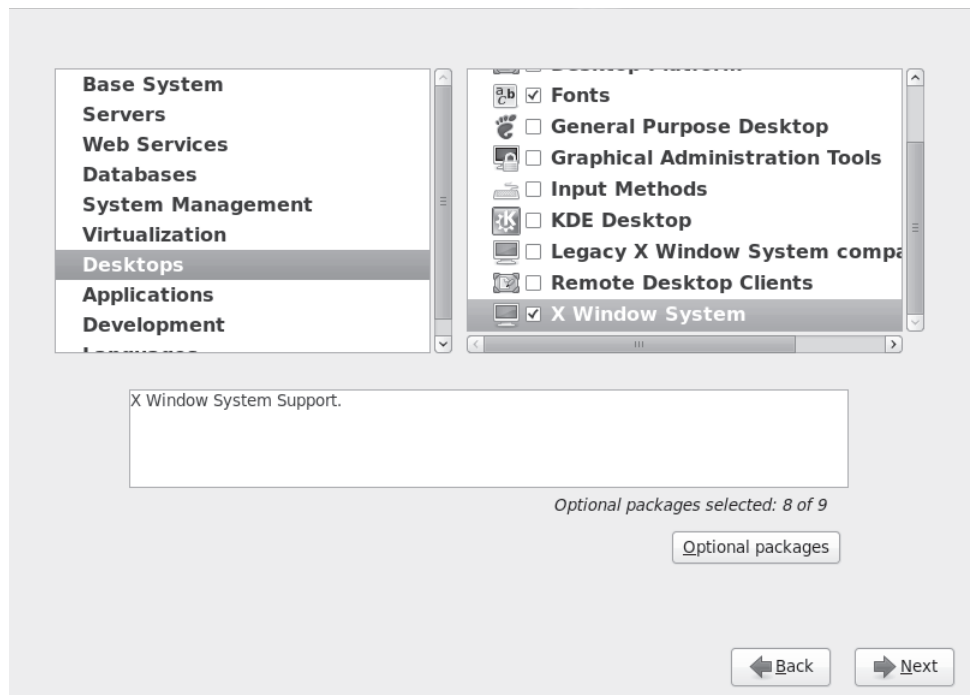
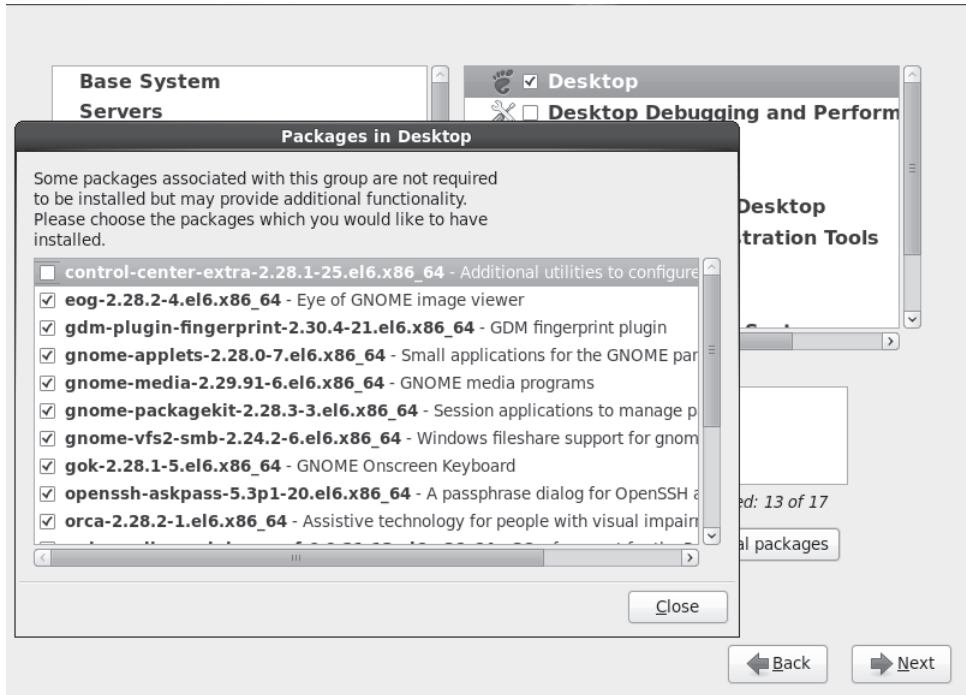


FIGURE I-15

Red Hat
Enterprise Linux
Desktop package
group details



the XML labels suggest, default packages are selected by default; optional packages are not selected.

Take some time studying this screen. Examine the packages within each package group. You'll learn about the kinds of packages that are installed by default. If you don't add them during the installation process, it isn't the end of the world. You can still add them with the **rpm** and **yum** commands or the Add/Remove Packages tool described in Chapter 7.

If the XML file is too confusing, just make a note of the name of a package group. From that name, you can find a list of associated packages after installation is complete. For example, the following command identifies mandatory, default, and optional packages for the **base** package group:

```
$ yum groupinfo base
```

For the purpose of this book, I've created two different baseline installations. One baseline is suited as a host system for VMs; the other is suited to be installed in a VM. Both include a relatively minimal GUI installation, as most Linux

administrators appreciate the convenience of multiple open command line consoles, side by side. To that end, I selected the following package groups during the installation process.

- **Desktops – Desktop** Includes the basic GNOME Desktop Environment.
- **Desktops – Fonts** Installs fonts that improve system readability in the GUI.
- **Desktops – X Window System** Adds the server packages associated with the GUI.
- **Applications – Internet Browser** Includes the Firefox Web browser.

The package groups fall into different categories; for example, to select the X Window System package group, I first highlighted the Desktops category in the left-hand pane. I could then select the X Window System package group in the right-hand pane.

In addition, for the physical host system configured for KVM-based VMs, I made sure the Virtualization package groups were included. They should be automatically included if you selected the Virtual Host system category earlier in the installation process.

Once desired packages are selected, click Next. Anaconda then proceeds to the installation process. When the process is complete, you'll see a final screen that confirms the installation, along with an option to reboot the system.



On the system used to write this book, I also installed the OpenOffice.org suite and The GIMP (The GNU Image Manipulation Program).

On Reboot

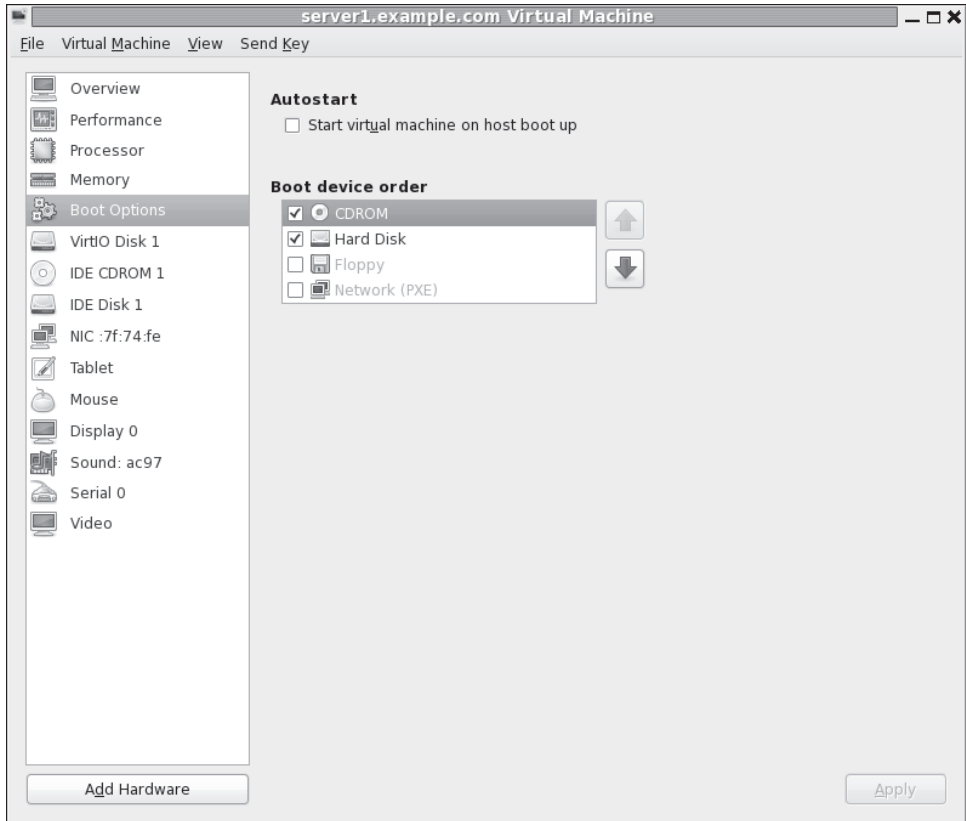
When installation is complete, you'll see a final message to that effect, with an option to reboot the system. If you've installed RHEL 6 on a physical system, don't forget to eject or remove the boot disk and/or the installation DVD.

For RHEL 6 installed in a KVM-based VM, you'll need to change the boot device. If it isn't already open, click Applications | System Tools | Virtual Machine Manager. Enter the root administrative password if prompted, and double-click on the desired VM. Then change the boot device with the following steps:

1. In the window associated with the VM, click View | Details.
2. In the window that appears, shown in Figure 1-16, click Boot Options.

FIGURE I-16

Boot devices
in the Virtual
Machine Manager



3. If the default boot device is set to CDROM, you should be able to change it to Hard Disk (or vice versa) using the arrow buttons. Once your selections are complete, click Apply. Be aware, the changes aren't implemented until the next time the KVM VM is powered on.
4. Click View | Console. If you haven't already done so, click Reboot to complete the RHEL 6 installation process.
5. Click Virtual Machine | Shut Down | Force Off. Confirm if prompted.
6. Click Virtual Machine | Run. The system should then boot normally into the newly installed system.



As of this writing, the Virtual Machine | Shut Down | Reboot and Virtual Machine | Shut Down | Shut Down menu options do not work.

In most installations, RHEL 6 will start the first boot process discussed later in this chapter.

CERTIFICATION OBJECTIVE 1.06

System Setup Options

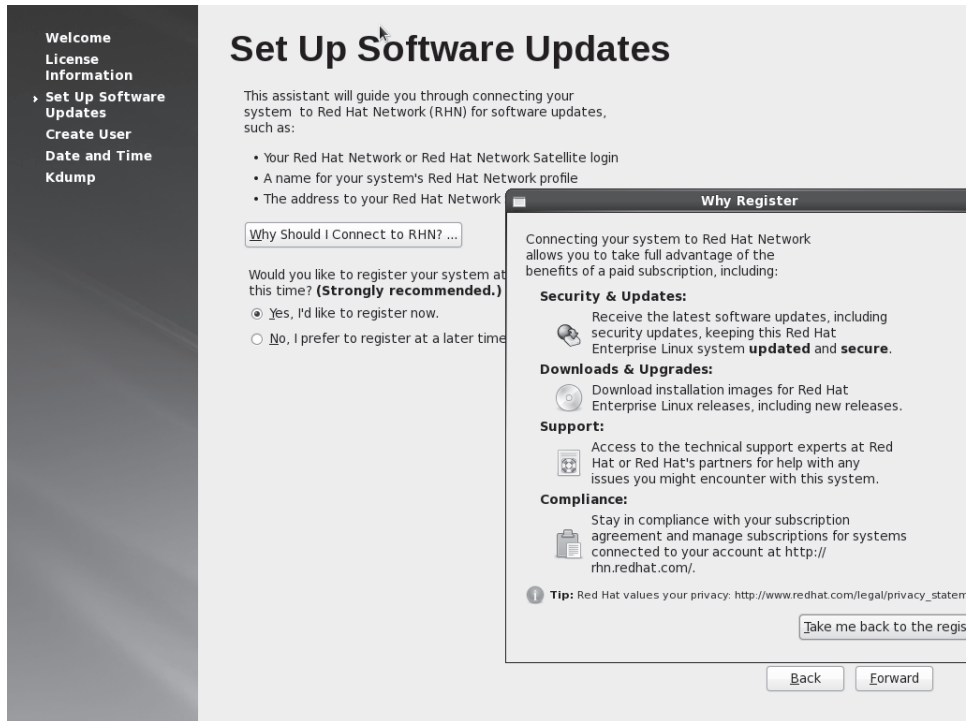
Baseline configurations are important. Once configured, you can clone that baseline to set up as many systems as needed. On a real network, a good baseline can be used to create systems dedicated to specific services. To enable remote access, it will have a Secure Shell (SSH) server, configured with a regular user.

For the boot process, RHEL 6 includes an implementation of the Upstart system, which replaces the Unix-based SysVInit system. It determines the consoles, services, and displays, as well as the runlevel that starts when a system is booted. Some systems use remote authentication, configured to connect to remote servers for username and password verification. While these systems are covered in other chapters, enough information is provided in this section to set up a baseline system.

The First Boot Process

But first, there's the process that starts the first time most RHEL 6 systems are booted after installation. And that's the first boot process. While there's a functionally similar text-mode version of the process, the default installation described earlier in this chapter leads to the GUI version of the process. The steps described in this section are based on an installation of the actual RHEL 6; the steps associated with a rebuild distribution will vary.

1. It starts with a welcome screen. Click Forward to continue.
2. The first boot process continues with a license agreement, which varies depending on whether this is RHEL 6 or a rebuild distribution. (Scientific Linux 6 doesn't even include this step.) If you refuse the license agreement, you're prompted to shut down and remove RHEL 6 from the local system. If you can accept this agreement, select Yes and click Forward to continue.
3. You're prompted to connect the system to the RHN. The illustration includes the window that appears if you click Why Should I Connect To RHN.



4. To register, you'll need a RHN account, with an available subscription. As that is not required in the RHCSA/RHCE objectives, select No and click Forward to continue.
5. You're prompted with the reasons shown in the illustration for registering. To avoid registering at this time, click No Thanks, I'll Connect Later to continue. You can register on the RHN later with the `rhn_register` command.
6. You're prompted with a window that the local system is not set up for software updates. As you'll perform this task in Chapter 7, click Forward to continue.
7. While not required for installation, the next step allows you to create a regular user for the system. While you can also connect to a remote database of users, a regular local user is required in this step. If you click Advanced, it opens the Red Hat User Manager, which can help customize user details,

as discussed in Chapter 8. Set up at least a local user, and click Forward to continue.


8. Now you can set up a date and time for the local system. The date and time that appears reflects that configured in the local hardware. If you select the Synchronize Date And Time Over The Network option, the window changes to show available NTP servers, based on the options associated with the Date/Time Properties tool briefly described in Chapter 5 and covered in detail in Chapter 17. Make desired changes and click Forward to continue.
9. Unless you're working with a system with a lot of memory, a message about "Insufficient memory to configure kdump" appears. That system collects data associated with kernel crashes. Click OK or make appropriate changes, and click Finish.

If you haven't installed GUI options such as the GNOME Desktop Environment and the X Window System, you'll see the text-mode alternative to the First Boot tool, as shown in Figure 1-17. Even if the noted GUI options are configured, you can start the equivalent of that tool with the `setup` command.

It includes six tools. The functionality associated with each of these tools is discussed in several different chapters.

FIGURE 1-17

The Text-Mode Setup Tool



Default Security Settings

When RHEL 6 is installed, there are default settings associated with SELinux and `iptables`-based firewalls. As this section just summarizes default settings, it does not include much detail. For more information on these security options, see Chapters 4, 10, and others.

First, SELinux is enabled in enforcing mode by default. You can confirm the setting with the `sestatus` command, which should lead to the following output:

```
SELinux status:                enabled
SELinuxfs mount:              /selinux
Current mode:                  enforcing
Mode from config file:        enforcing
Policy version:                24
Policy from config file:      targeted
```

You'll learn more about SELinux and the RHCSA exam in Chapter 4. If you go for the RHCE, you'll also learn to configure SELinux in Chapters 10 and on to support a wide variety of services.

If you want detailed information about the current `iptables` command, examine the `/etc/sysconfig/iptables` file. That file is used by the `iptables` service, courtesy of the `/etc/init.d/iptables` script. The following line from that file allows traffic sent through port 22 access from the outside.

```
-A INPUT -m state --state NEW -m tcp -p tcp --dport 22 -j ACCEPT
```

Port 22 is the default port for the Secure Shell (SSH) service, which supports remote administration of the local system. If there's a good network connection, you'll be able to connect remotely to this system. If the local IP address is 192.168.122.50, you can connect remotely to user `michael`'s account with the following command:

```
# ssh michael@192.168.122.50
```

The SSH server can be configured to enhance security even further. For more information, see Chapter 11.

Special Setup Options for Virtual Machines

On a KVM-based virtual host, you may notice additional firewall rules. For example, in the `/etc/sysconfig/iptables` file of the virtual host on my system, I see the following additional rule, which accepts traffic over a physical bridged network device.

```
-I FORWARD -m physdev --physdev-is-bridged -j ACCEPT
```

It works for IPv4 networking with the help of active IP forwarding in the dynamic file, `/proc/sys/net/ipv4/ip_forward`. It's a binary file; if it's set to 1, IPv4 forwarding is active. To set it on a permanent basis, open the `/etc/sysctl.conf` file, and make sure the following directive is set to 1:

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

To implement the changes immediately on the local system, run the following command:

```
# sysctl -p
```

This information is also covered briefly in Chapter 5 in the context of network configuration. A detailed discussion of the related `/proc` filesystem is a RHCE topic covered in Chapter 12.

CERTIFICATION OBJECTIVE 1.07

Configure Default File Sharing Services

One of the requirements of the RHCSA guide is to “Configure a system to run a default configuration HTTP server,” as well as “Configure a system to run a default configuration FTP server.” Let's interpret those statements. The default HTTP server on RHEL 6 is the Apache Web server. The corresponding default FTP server is vsFTP.

While detailed configuration of these services is the province of the RHCE, these systems include basic functionality in their default installations. You'll confirm the operation of the default installation. Next, you'll take this process one step further, to set up these services to share files, specifically the files copied from the installation DVD.

The configuration of these services for file sharing is fairly simple. No changes are required to the main configuration files. Assuming SELinux is enabled (as it should be during the exams), the basic steps are:

- Copy the contents of the RHEL 6 installation DVD to the appropriate directory.
- Make sure the contents of the noted directory are configured with the right SELinux contexts.
- Configure the noted service to point to the specified directory, and to start when the system is booted.

Naturally, the steps vary by service. The details described in this chapter are rudimentary and may not be good enough if the related commands and services are new to you. For more information on the **mount** command, see Chapter 6. For more information on SELinux, see Chapter 4. For more information on the Apache Web and vsFTP services, see Chapters 14 and 16, respectively.

Mount and Copy the Installation DVD

You may already know that the **mount** command is used to connect a device such as a partition or a DVD drive to a specified directory. For example, the following command mounts the standard DVD drive onto the `/media` directory:

```
# mount /dev/dvd /media
```

RHEL 6 includes a number of similar device files in the `/dev` directory. With the **ls -l** command, you can confirm that these files are all linked to the `/dev/sr0` device file:

```
# ls -l /dev/dvd
# ls -l /dev/dvdrw
# ls -l /dev/cdrom
```

If the DVD is properly configured, it should automatically find the appropriate filesystem format from the `/etc/filesystems` file. In this case, it's shown in the file as `iso9660`, which is the ISO standard for CD and DVD formats. It is reasonable to expect that the DVD is properly configured. If there's a problem, you'd see the following error message from the noted **mount** command:

```
mount: you must specify the filesystem type
```

The most likely cause is an error either in the ISO file or in the way that file was burned to the physical media. Speaking of the ISO file, it's easy to mount it directly on a directory, without wasting a physical DVD. The following command would work on the RHEL 6 DVD:

```
# mount -o loop rhel-server-6.0-x86_64-dvd.iso /media
```

The next step is to copy the contents of the DVD to the directory that's shared with the file server of your choice, FTP or HTTP. Generically, the following command makes sure to copy files in archive (-a) mode, recursively (-r). The dot makes sure to include hidden files. The files are copied to the noted directory:

```
# cp -ar /media/. /path/to/dir
```

The directory that should be used depends on the server. Of course, servers can be configured to use directories in other than default locations.

Set Up a Default Configuration Apache Server

The Red Hat implementation of the Apache Web server uses the `/var/www/html` directory for HTML files. Subdirectories that you create are suitable for file sharing. It works with the default Apache configuration, as installed from Red Hat packages. Nevertheless, you'll have to make sure port 80 is open in any existing firewall.

The steps required to configure Apache as a RHEL installation server are similar to those required to configure vsFTP. In Exercise 1-2, you'll take the steps required to configure Apache as an installation server. But first, you need to make sure Apache is installed with the following command:

```
# yum install httpd
```

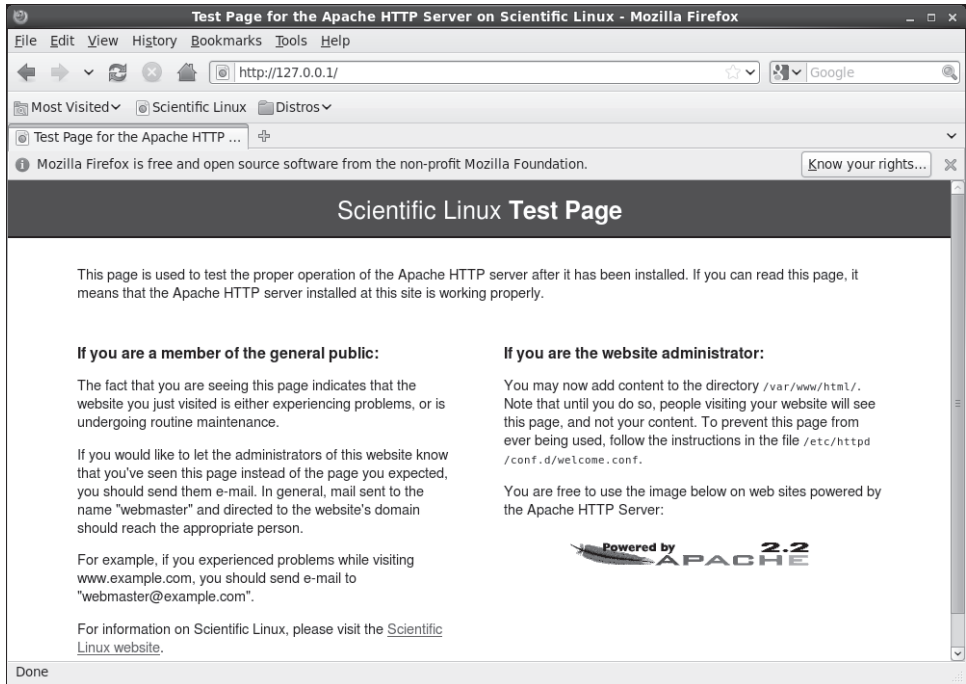
If the command is successful, you can find the main Apache configuration file, `httpd.conf`, in the `/etc/httpd/conf/` directory, as well as the main data directory in `/var/www/html`. But remember, don't touch that configuration file! To make sure the default installation works, first start the Apache service with the following command:

```
# /etc/init.d/httpd start
```

Next, use a browser on the system where Apache is installed and navigate to the localhost IP address with the following URL: `http://127.0.0.1/`. An example from the Scientific Linux rebuild of RHEL 6 is shown in Figure 1-18.

FIGURE I-18

Proof of a
working default
Apache server



If port 80 is open in an existing local firewall, this page should be accessible from remote systems as well. In addition, you'll need to make sure Apache starts automatically the next time RHEL 6 is booted. One way to do so is with the following command:

```
# chkconfig httpd on
```

For more information on how services like Apache are controlled during the boot process, see Chapter 11. While that is an RHCE chapter, the `chkconfig` command shown here is simple.

EXERCISE 1-2**Configure Apache as an Installation Server**

In this exercise, you'll install and configure the Apache Web server as a file server, suited for RHEL 6 installations. You'll need a copy of the RHEL 6 DVD, or the associated downloaded file in ISO format. In this exercise, you'll take the steps needed to create an appropriate directory, copy the installation files, set an appropriate SELinux context, open port 80 in any existing firewall, and restart the Apache service. These are basic steps; detailed Apache configuration is discussed in Chapter 14.

1. Mount the RHEL 6 DVD on an empty directory. You might use one of the following two commands. While the first mounts an actual physical DVD, the second mounts the ISO file:

```
# mount /dev/dvd /media
# mount -o loop rhel-server-6.0-x86_64-dvd.iso /media
```

2. Create an appropriate directory for the installation files. Since the standard directory for Apache Web server files is `/var/www/html`, it's simplest to create a subdirectory there with the following command:

```
# mkdir /var/www/html/inst
```

3. Copy the files from the mounted DVD to the new directory:

```
# cp -ar /media/. /var/www/html/inst/
```

4. Make sure the files have the right SELinux context with the `chcon` command. The `-R` applies the changes recursively through the copied installation files. The `--reference=/var/www/html` switch applies the default SELinux context from that directory.

```
# chcon -R --reference=/var/www/html/ /var/www/html/inst
```

5. Open port 80 associated with the Apache Web server; it's fairly easy to do with the utility that you can start with the `system-config-firewall` command. If that's too complex, just run the following command to disable the firewall on the current system. You'll learn to configure firewalls in more detail in Chapters 4 and 10.

```
# iptables -F
```

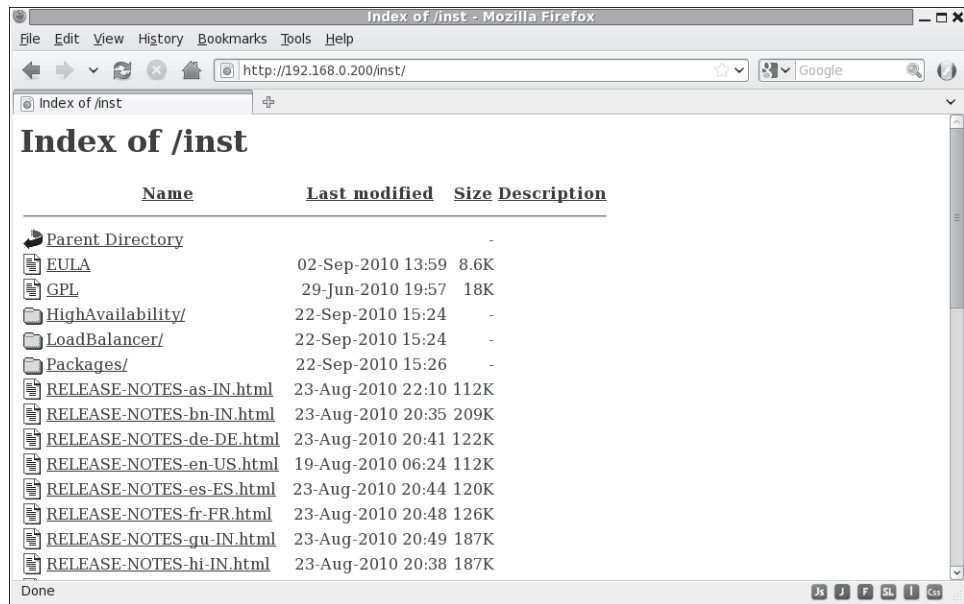
6. Make sure the Apache Web server is running with the following command:

```
# /etc/init.d/httpd restart
```

The Apache Web server should now be ready for use as a file server, sharing the installation files from the RHEL 6 DVD. To confirm, point your browser to the IP address for the server and to the `inst/` subdirectory. If that IP address is 192.168.0.200, you'd navigate to:

```
http://192.168.0.200/inst
```

If successful, you'll see a page of clickable and downloadable files as shown here:



Share Copied Files via FTP Server

The Red Hat implementation of the vsFTP server includes the `/var/ftp/pub` directory for published files. For the purpose of the installation files, you can create the `/var/ftp/pub/inst` directory. To make the system compatible with SELinux, you'll then change the security contexts of each of those files with one command. When you then start or restart the vsFTP server, it will be ready for use as an installation server.

The process is documented in Lab 2. Assuming an appropriate connection to remote repositories using the RHN or from a rebuild distribution, you can make sure the latest version of the vsFTP server is installed with the following command:

```
# yum install vsftpd
```

If successful, you can find the main vsFTP configuration file, `vsftpd.conf`, in the `/etc/vsftpd` directory, as well as the main data directory in `/var/ftp/pub`. Just don't touch that configuration file, as the requirements point to a default configuration server. Just make sure to start the vsFTP service with the following command:


```
# /etc/init.d/vsftpd start
```

As web browsers can access FTP servers, you can confirm the default FTP server configuration on the local system by navigating to `ftp://127.0.0.1/`. The default result in the Firefox Web browser is shown in Figure 1-19. The `pub/` directory shown is actually the `/var/ftp/pub` directory.

Note the security associated with the vsFTP server. Click the [Up To Higher Level Directory](#) hyperlink. The current directory does not change. Users who connect to this FTP server are unable to see, much less download, files from anything above the

FIGURE 1-19

Access to the default FTP server



`/var/ftp` directory. This is a security concept known as a chroot jail. It's often used on other systems such as the Domain Name Service (DNS) server to increase security on the host system.

If port 21 is open in the local firewall, this FTP server should also be accessible from remote systems. In addition, you'll need to make sure the vsFTP server starts automatically the next time RHEL 6 is booted. One way to do so is with the following command:

```
# chkconfig vsftpd on
```

The `d` at the end of vsFTP refers to its daemon. For more information on how services like vsFTP are controlled during the boot process, see Chapter 11. While that is an RHCE chapter, it should not be hard to remember this one **chkconfig** command.

CERTIFICATION SUMMARY

The RHCSA and RHCE exams are not for beginners. This chapter helps you install a basic RHEL system, with the packages and settings suitable for the remainder of this book. Both exams are practical, hands-on exams. When you sit for either exam, you'll be faced with a live RHEL system with a series of problems to solve and systems to configure. The RHCSA covers core system administration skills.

The Red Hat exams now implicitly assume the use of a 64-bit system. Red Hat configures the default VM system, KVM, only in the 64-bit build of RHEL 6. And you're required to configure RHEL 6 as a virtual host for the RHCSA.

With a subscription to the RHN, you can download RHEL installation ISO files from the associated account. Since RHEL software is released under open-source licenses, third parties such as CentOS and Scientific Linux have used that source code without Red Hat trademarks. You can also use such rebuild distributions to study for the RHCSA and RHCE exams.

It will be helpful to create multiple installations of RHEL 6 to practice the skills you'll learn in later chapters. To that end, I recommend the configuration of three systems. While many users don't have three spare physical computers to dedicate to their studies, VMs make it possible to set up these systems on a single physical computer.

As the installation of RHEL 6 is relatively easy even for newer Linux users, not every detail is covered in this chapter. However, it is based on the network installation source created in Lab 2. After installation comes the First Boot process, which varies depending on whether you've installed a GUI.



TWO-MINUTE DRILL

Here are some of the key points from the certification objectives in Chapter 1.

The RHCSA and RHCE Exams

- ❑ The RHCSA is a separate exam from the RHCE.
- ❑ Red Hat exams are all “hands-on”; there are no multiple-choice questions.
- ❑ If you’re studying for the RHCSA, focus on Chapters 1–9. If you’re studying for the RHCE, while you’re responsible for the information in the entire book, focus on Chapters 1 and 10–17.

Basic Hardware Requirements

- ❑ While RHEL 6 can be installed on a variety of architectures, you’ll need hardware that can handle the 64-bit version with hardware-assisted virtualization for the Red Hat exams.
- ❑ Red Hat supports RHEL 6 installations on systems with at least 512MB of RAM. Less is possible, especially on systems without a GUI. More is required for a GUI installation within KVM.
- ❑ RHEL 6 can be installed on local or a variety of network hard drives.

Get Red Hat Enterprise Linux

- ❑ The Red Hat exams use Red Hat Enterprise Linux.
- ❑ Subscriptions and test copies of RHEL 6 are available.
- ❑ Since Red Hat releases the source code for RHEL 6, third parties are free to “rebuild” the distribution from the Red Hat source code (except for the trademarks).
- ❑ Third-party rebuilds of RHEL 6 are functionally identical, except for access to the Red Hat Network.
- ❑ Reputable third-party rebuilds are available from CentOS and Scientific Linux.

Installation Requirements

- Red Hat has stated that exams are presented on “pre-installed systems” with questions presented “electronically.”
- The RHCSA requires the configuration of a physical machine as a virtual host.
- The native RHEL 6 VM solution is KVM.
- It’s useful to set up multiple VMs to simulate network communications.

Installation Options

- You can start the installation process from a variety of boot media.
- RHEL 6 can be installed from DVD, from a local drive, from an NFS directory, from an Apache web server, or from an FTP server.
- RHEL 6 should be configured on separate volumes for at least the top-level root directory (/), the /boot directory, and Linux swap space.
- RHEL 6 includes installation package groups in a number of categories.

System Setup Options

- The first post-installation step is the First Boot process.
- SELinux and `iptables`-based firewalls are enabled by default.

Configure Default File Sharing Services

- The RHCSA exam includes requirements to deploy HTTP and FTP servers in their default configurations.
- The default services associated with the HTTP/FTP protocols are the Apache Web server and the vsFTP server.
- One way to deploy a default HTTP or FTP server is to configure it with the installation files from the RHEL DVD.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. As there are no multiple-choice questions on the Red Hat exams, there are no multiple-choice questions in this book. These questions exclusively test your understanding of the chapter. It is okay if you have another way of performing a task. Getting results, not memorizing trivia, is what counts on the Red Hat exams.

The RHCSA and RHCE Exams

1. How many multiple-choice questions are there on the RHCE exam? And on the RHCSA exam?

Basic Hardware Requirements

2. Assuming Intel-based PC hardware, what's the default virtualization technology for RHEL 6?
3. How many bits are required of a CPU that can be used to configure KVM on RHEL 6?

Get Red Hat Enterprise Linux

4. Name one third-party distribution based on RHEL 6 source code.

Installation Requirements

5. How much time is allocated for installation during the RHCSA and RHCE exams?

Installation Options

6. Name two different options for installation media that will boot the RHEL 6 installation program.

62 Chapter 1: Prepare for Red Hat Hands-on Certifications

7. Name three types of volumes that can be configured and formatted during the RHEL 6 installation process to store data.

8. Say you've mounted the RHEL 6 DVD on the /media directory. There's an XML file on that DVD with a database of packages and package groups. In what directory can you find that XML file?

System Setup Options

9. What command starts the same menu as the text-mode First Boot process?

10. What service is allowed through the default firewall?

Configure Default File Sharing Services

11. What is the standard directory for file sharing for the RHEL 6 implementation of the vsFTP server?

12. What is the standard directory for HTML files on the Apache Web server?

LAB QUESTIONS

The first lab is fairly elementary, designed to get you thinking in terms of networks and networking. The second lab should help you configure an installation server. The third lab suggests that you look at the requirements associated with the Linux Professional Institute, for a different perspective on system administration.

Lab 1

In this lab, you'll plan the systems required for a Red Hat network. You have three computers configured with RHEL 6. Two of these computers are to be configured on one domain, example.com. These computers will have hostnames server1 and tester1. The third computer is to be configured on a second domain, example.org, with a hostname of outsider1.

One of the computers in the example.com domain, server1, may be configured with two network cards. Alternatively, if these systems are configured as guests on a KVM virtual host, IP Forwarding will make it possible for these systems to communicate, even though they're set up on different networks. The focus will be on IPv4 addressing.

- Systems on the example.com domain will be configured on the 192.168.122.0/24 network.
- Systems on the example.org domain will be configured on the 192.168.100.0/24 network.

Ideally, you should set up the server1.example.com system as a workstation and a server. The basic instructions described in this chapter should suffice, as it will be up to you to install and configure required services after installation is complete. It will be the primary system you use for practice. You'll install RHEL 6 on this system in Chapter 2, and you'll clone it for different chapters as well as the sample exams at the end of this book.

The tester1.example.com system will be a system that allows remote access only from the SSH service. In some cases, server services not necessarily required for certification may be configured on the outsider1.example.org network. That will allow you to test clients that are required for certification.

Lab 2

This lab assumes you've downloaded the DVD-based ISO for RHEL 6 or a rebuild such as CentOS or Scientific Linux. The DVD-based ISO is important, as it will serve two purposes. It will be the installation repository used earlier in this chapter as well as the package repository to be configured in Chapter 7. This lab simply includes those commands required to configure the noted files on the vsFTP server.

While the Red Hat exams are given on a pre-installed system, the associated requirements do suggest that you need to know how to install systems over a network, and configure Kickstart installations. And as you don't have Internet access during the exam, you won't have access to the Red Hat Network or any other Internet repository.

1. Create a directory for your installation files. With the following command, create the /var/ftp/pub/inst directory. (If you get an error message, vsFTP may not be properly installed.)

```
# mkdir /var/ftp/pub/inst
```

64 Chapter 1: Prepare for Red Hat Hands-on Certifications

2. Insert the RHEL 6 installation DVD into its drive. If not automatically mounted, do so with a command such as `mount /dev/cdrom /media`. (If all you have are the ISO files, say in the Downloads/ subdirectory, substitute `mount -ro loop Downloads/rhel*.iso /media`.)
3. Copy the required files from the RHEL 6 installation DVD. Use the `cp -ar /source/. /var/ftp/pub/inst` command, where *source* is the mount directory (such as /media/). Don't forget the dot (.); it makes sure to copy hidden files, including the .discinfo file.
4. Make sure there's nothing blocking access to your vsFTP server. Use a tool like the Firewall configuration tool to open up port 21 on the local system. For more information on firewalls and SELinux, see Chapter 4.
5. If SELinux is enabled on the local system, run the following command to apply appropriate SELinux contexts to the files on the new directory:

```
# chcon -R -t public_content_t /var/ftp/
```

6. Now activate the FTP server with the following command:

```
# service vsftpd restart
```

7. Test the result. On a remote system, you should be able to use the `lftp` command to connect to the local FTP server, using its IP address. Once connected, you'll be able to find the installation files in the `pub/inst/` subdirectory.

Lab 3

The Red Hat exams are an advanced challenge. In this lab, you'll examine the Red Hat exam prerequisites from a slightly different perspective. If you're uncertain about your readiness for this exam, the Linux Professional Institute has Level 1 exams that test basic skills in more detail. In addition, they cover a number of related commands which I believe are implied prerequisites for the Red Hat certifications.

To that end, examine the detailed objectives associated with the noted exams 101 and 102. Links to those objectives are available from www.lpi.org. If you're comfortable with most of the files, terms, and utilities listed in the objectives for those exams, you're ready to start your studies for the Red Hat exams.

SELF TEST ANSWERS

The RHCSA and RHCE Exams

1. There are no multiple choice questions on any Red Hat exams. It's been nearly a decade since the Red Hat exams had a multiple-choice component. The Red Hat exams are entirely “hands-on” experiences.

Basic Hardware Requirements

2. The default virtualization technology for RHEL 6 is KVM. While there are many excellent virtualization technologies available, KVM is the default option supported by Red Hat on RHEL 6.
3. 64. To configure KVM on RHEL 6, you need a system with one or more 64-bit CPUs.

Get Red Hat Enterprise Linux

4. There are several different distributions available built on RHEL 6 source code. Two options are Oracle Linux and Scientific Linux. There may be additional correct answers.

Installation Requirements

5. There is no correct answer to this question. While the Red Hat exams are now presented on pre-installed systems, it's possible that you'll have to install RHEL 6 on a VM within an existing RHEL 6 installation.

Installation Options

6. Options for installation boot media for RHEL 6 include a CD, a DVD, and a USB stick. Red Hat no longer creates boot media in a floppy format.
7. You can configure and format regular partitions, RAID arrays, and logical volumes during the installation process to store data.
8. You can find the specified XML file under the noted conditions in the /media/repodata directory.

System Setup Options

9. The command is `setup`, which requires root administrative privileges.
10. The default RHEL 6 firewall allows access to the Secure Shell service, SSH for short.

Configure Default File Sharing Services

11. The standard directory for file sharing for the RHEL 6 implementation of the vsFTP server is `/var/ftp/pub`.
12. The standard directory for HTML files for the RHEL 6 implementation of the Apache Web server is `/var/www/html`.

LAB ANSWERS

Lab 1

When configuring a network connected to the Internet, you'll want to allow access to some systems and deny them to others. To that end, this lab provides a framework for the systems you'll want to set up to study for the RHCSA/RHCE exams.

As the RHCSA is in many ways an exercise in configuring a workstation, it may seem less important to set up a network to study for that exam. However, there are server elements to that exam, such as the configuration of FTP and HTTP file servers, so networks can't be neglected for the RHCSA.

With the development of VMs, the cost of hardware should be less of a handicap even for home users who are studying for the Red Hat exams. However, the RHCSA requires the configuration of VMs. And the default Red Hat solution, KVM, can only be configured for RHEL 6 on 64-bit systems.

While dynamic IPv4 addresses are used for most workstations, static IPv4 addresses are more appropriate in many cases, such as for DNS and e-mail servers. Clients are easier to configure when the IP address of such servers are known. So it's appropriate to set up static IPv4 addresses for such systems.

Three systems is a suggested minimum, as the rules associated with firewalls are typically not applied to a local system. The second system is a remote client that should have access to local server services. The third system is a remote client that should not have such access. All of these systems can be configured on VMs.

Of course, "real-life" networks are much more complex. And you're welcome to set up a network with more systems.

In Chapter 2, when you install RHEL 6 systems on KVM-based virtual machines, you'll want to clone that system to support configuration from a baseline. And, in fact, that's what happens in many enterprises. VMs makes it practical to dedicate one or more RHEL 6 systems to a specific service, such as the Apache Web server.

Lab 2

During the Red Hat exams, you won't have access to the Internet. However, many installations and updates require Internet access to get to those files associated with installation and updates.

When you configure the files from the RHEL 6 installation DVD on a remote system, you're configuring an effective substitute for the purpose of installation of additional packages. In addition, those files support network installation, which is still a RHCSA requirement.

The steps described are associated with the configuration of the vsFTP server, protected by SELinux. Do not fear SELinux. As suggested by the steps in this lab, the configuration of the vsFTP server is fairly simple. While the use of SELinux may seem intimidating to the RHCSA candidate, it is a requirement. The commands described in this lab show how you can live with SELinux on a vsFTP server. Chapter 4 will explain how you can make live with SELinux work for you in a number of other situations.

Lab 3

This lab may seem odd, as it references the requirements for a different Linux certification. However, many Linux administrators take the exams of the Linux Professional Institute (LPI) seriously. LPI creates excellent certifications. Many Linux administrators study for and pass the LPIC Level 1 exams. Passing the LPIC 101 and 102 exams provides an excellent foundation for the RHCSA and RHCE exams.

If you feel the need to get more of a grounding in Linux, refer to some of the books described at the beginning of this chapter.

The Red Hat exams are an advanced challenge. Some of the requirements for the RHCSA and RHCE exams may seem intimidating. It's okay if some of them seem beyond your capabilities at the moment, as that is the purpose of this book. However, if you're uncomfortable with basic command line tools such as **ls**, **cd**, and **cp**, you might need more of a grounding in Linux first. But many candidates are successfully able to fill in the gaps in their knowledge with some self-study and practice.

