



Before Installing Linux for the First Time

A Linux White Paper

Preface

Although I have used and supported operating systems for many years, Linux[®] was brand new to me. So it was with both eagerness and trepidation that I approached my assignment: Install Linux and write about my experiences as a Linux “newbie.” (A companion white paper, available from the same sources as this one, discusses installing Linux *applications*.) At first, it sounded simple. After all, I have installed every version of DOS, Microsoft[®] Windows[®], and IBM OS/2 that has ever existed up to this point (including quite a few beta versions along the way) and provided technical support for many of them. Yet they all have many things in common to simplify the learning process: command names, directory structures, file and directory naming conventions, and so on. The UNIX[®] world, however, was completely new to me.

Unlike Windows, which was all written by one company, Linux is a collaborative effort of thousands of programmers from around the world, each contributing ideas, programming skills, and much time and effort, with the goal of making Linux the most reliable, most powerful, and most flexible operating system available. Because there are so many people involved in the development of Linux, there are many different “flavors” of Linux available, each with its own unique combination of function, appearance, packaging, and installation tools.

Because of the “open source” nature of Linux, which allows each developer to produce customized versions of operating system utilities and related programs, no two Linux packages (called “distributions”) from different vendors are quite the same. To complicate matters further, a newer version of one distribution may be significantly different from an older version of the same distribution. In other words, Red Hat Linux version 6.2 is not only significantly different from Caldera OpenLinux 6.3, but also from earlier versions of Red Hat Linux.

This makes it difficult to write general Linux tutorials, because the instructions may differ dramatically from one distribution to another. I began this paper with the intention of writing a “How to Install Linux” paper but I quickly realized that it would not be practical. One paper covering several Linux distributions would not only be very large but also confusing, because some parts would apply to several—but not all—distributions, while others would be unique to only one distribution. I could write a number of different papers, each devoted to one distribution, but large portions of each would be redundant.

To top it all off, each distribution *already* includes an installation manual of some sort (printed, text file, PDF, HTML, etc.). Therefore, writing another installation manual would have been largely a waste of effort. However, from my experiences installing Linux, I realized that another type of paper would indeed be useful: something that explained what to do *before* you start

Preparing Today for Linux Tomorrow

the installation. In other words, what the installation manuals *should* tell you but do not. That is the subject of this paper.

Note: There probably are many terms used in this paper that are unfamiliar if you are new to Linux. You may find the white paper called “*A Brief Linux Glossary for Windows Users*”, available from the same sources as this paper, to be a useful reference.

Contents

Preface	1
Many Questions	4
The Answers.....	4
Hardware Support	4
Device Drivers	5
Partitioning.....	8
Installation Methods	10
Installation Types	11
Multiboot Setups.....	11
Beginning the Installation.....	12
After Installation – What then?	12
Appendix A - Creating a Bootable Disk to Install from a Non-Bootable CD-ROM Drive or from a Hard Disk Drive	14
Appendix B - Creating a Bootable Disk to Install from an External PCMCIA-Attached CD-ROM Drive	15
Appendix C - Creating a Bootable Disk to Install from a Network Drive (CD-ROM or Hard Disk)	16

Many Questions

Some of the issues that should be considered by the user *before* starting an installation include:

- Is my hardware supported by Linux? (Alternatively, has my system vendor tested Linux on my hardware and declared it compatible?) How do I find out before I start?
- Will I need separate device drivers for the audio and video chips, modem, USB port, CD-ROM drive, and other devices, or will Linux recognize them and install the drivers during installation. What about my printer? If installing on a notebook computer, will the built-in pointing device be recognized automatically? How about the PCMCIA (PC Card) slots and infrared port? Are there hot/warm-swap drivers available for the removable drive bay? If I need to find drivers, where can I look?
- How should I partition my hard drive to take the best advantage of Linux? (How many partitions do I need? How large should they be? Primary or Extended? Should I delete the Windows partitions and create new ones, or delete only the partition contents? Is one big partition for both Linux and applications/data files best, or should there be one for the OS and another for programs and data, or some other arrangement?)
- Will I be installing from a bootable internal CD-ROM drive, a bootable external PCMCIA-attached CD-ROM drive, a non-bootable CD-ROM drive using a boot diskette, a downloaded copy of Linux from my hard drive, or from a network drive? Will I need to create an installation floppy, and if so, how do I do it?
- What is the difference between a Server installation, a Workstation installation, and a Custom installation, and which installation method will best suit my needs?
- What should I do if I want to install Linux and Windows on the same system, or if I want more than one “flavor” of Linux installed? What sorts of “boot managers” are available to allow me to choose which OS to boot from each time?

The Answers

Here are the answers I found to the preceding questions:

Hardware Support

The *minimum* requirements for Linux are pretty basic, so most likely your system, unless it is very old, will be up to the challenge. Yet, that is just the first step. Once you have determined that your system has at least the essentials covered (sufficient memory, disk space, etc.) it is time to move on to the more esoteric issues, such as which particular video chip or USB controller is in your computer.

First things first: To find out the various minimum hardware requirements for the different Linux distributions, look on the shrink-wrapped box, or in the installation manual that came with the distribution or, better yet, visit the vendor’s website. For example:

- Caldera Linux (www.calderasystems.com/support/hardware)
- Debian Linux (www.valinux.com/services/support/?session_hash=7d7ccc8f28c0f42cddc98da037d7875a)
- Red Hat Linux (www.redhat.com/support/hardware)
- SuSE Linux (www.suse.com/us/support/hardware/index.html)
- TurboLinux (www.turbolinux.com/products/syscompat.html)

Note: If you are reading this document online with Adobe Acrobat Reader, simply click on the Web addresses highlighted in blue to go to those Web sites via your Web browser.

Because it is possible that Linux will not automatically recognize all of your hardware, it is a good idea to find out what you have installed so that you can manually tell Linux what is in your box should the need arise. You should know what audio and video chips are in your system, as well as the type of CD-ROM or DVD drive and mouse (PS/2 or serial interface; two buttons or three), modem and USB controller chips, and the PCMCIA controller (if you are installing on a notebook computer). If your computer uses SCSI devices, you will need to know what brand and model of SCSI controller is installed as well.

If you do not know how to determine what components are being used, there are several ways to accomplish this:

- If the system is currently running Windows try (*prior* to installing Linux):
 - The System Information Tool (select *Start -> Programs -> Accessories -> System Tools -> System Information*, then click on *Components*)
 - The Windows Device Manager (select *Start -> Settings -> Control Panel*, then click on the *Hardware* tab and press the *Device Manager* button)
 - Any other similar utility on your computer
- The manuals that came with the computer and any options added later
- Vendor technical support for the system and/or options (the manuals should include phone numbers or websites for the vendors; see below for some Linux distribution vendors' websites)

Do not attempt to proceed with a Linux installation until you have collected all of this information and recorded it somewhere. There may be a place in the back of your computer's manual to write it down, or visit: members.tripod.com/~algolog/lxchk.htm for a checklist.

Note: Even if Linux automatically identifies all your hardware (or if you manually specify some components), it is possible that you still will have to locate and install some device drivers in order for some functions to work. These will be installed after Linux is up and running. You may also find that you have some hardware that will not work under Linux or which will require more work on your part than just installing a device driver. For example, if you discover that your computer uses a "winmodem", see the question "*How do I get my "winmodem" to work with Linux?*" in the *Linux Questions and Answers* white paper available from the same sources as this paper.

Device Drivers

Does Linux include device drivers for all your hardware, or will you need to locate and install some drivers yourself? If you will have to find them yourself, how do you find the drivers? Where do you look? Does it really matter if Linux does not support some of your hardware?

Starting with the last question first: Yes, it matters a great deal. The last thing you want to have happen in the middle of the installation (after you have already formatted the hard drive and deleted Windows) is to discover, when the screen goes blank, that your video chip is not supported or that it is only supported in VGA compatibility mode—at 640x480 resolution in 16 colors (ugly with a capital UGH!). Almost as bad is to find out after installation that your modem does not work. So how do you download drivers to correct the modem problem if your modem does not work, and you cannot download the drivers?

Preparing Today for Linux Tomorrow

There is no one simple answer to the other questions, but there at least are starting points and a strategy you can follow to simplify the process of locating device drivers.

First, contact the manufacturer of your PC, be it IBM, Compaq, Dell or whomever. No one knows your hardware better than the manufacturer. Their technical support department should be able to tell you whether your particular model has ever been tested with Linux and if it is supported for use with any particular distributions. (It may work with some distributions but not others, or it may be “certified” to work with some but not tested with others—which would leave you in limbo if you install an unsupported distribution and have problems. The system vendor probably would not help you in that situation—at least, not for free.)

If it is supported with one or more specific distributions (such as Red Hat 6.2 or Caldera OpenLinux 6.3), there may be new or updated device drivers available from the PC vendor that you will need to make particular components work correctly. If so, they should be available on the vendor’s website or from their technical support department, along with instructions for how to install the drivers. If one or more of the Linux distributions is supported, make sure you install one of those. That way if you run into trouble, at least the PC vendor’s tech support department will talk to you and may even be able to help you.

The *best-case* scenario is this: your PC is supported, the manufacturer offers device drivers and has a trained technical support staff. What do you do if this is *not* your situation?

The next approach is to contact the distribution vendor (Caldera, Corel, Red Hat, SuSE, TurboLinux, VA Linux, etc.) directly, via their tech support center or website. They may have a list of PC vendors’ systems that are known-compatible.

The preceding assumes that you are installing on an “off-the-shelf” or standard configuration as purchased from the manufacturer. Yet, what happens if you have added after-market components, such as a CD-ROM drive, or joystick, or Iomega Zip drive? Worse, what if you have built your own “Frankenclone” from a combination of scavenged parts of several older systems and newer store-bought components? Is there any hope of getting help in this situation?

Perhaps—although it is much trickier. No Linux distribution vendor is going to guarantee results on a combination of hardware that they have not tested themselves. But if you are willing to do a lot of footwork yourself tracking down what you need and serving as your own integration tester trying to get everything working together, you may very well be able to cobble something together that will work adequately. No mean feat for something assembled from scrap parts. Linux itself is fairly forgiving and will install on all sorts of hardware, but some of the non-essential components may need your personal touch to work properly.

Before you begin, it is important to realize that you may never be able to get some hardware working with Linux. It is possible that there simply are no drivers for some hardware and there never will be. In some cases, no one has bothered to write drivers for very old hardware, because of its age. Some devices that are not widely used may not have drivers because of the relative obscurity of the hardware. On the other hand, very recent products might not yet have drivers because the drivers are still in development. (At least in this case there is hope for the future.) In addition, drivers may be available for some Linux distributions but not others.

For best results, try to use mainstream products from well-known vendors that are fairly recent but which have been available for at least 6 months. These are the most likely to have drivers written for them. (In the future, whenever possible check for Linux device drivers *before* buying a product.)

First, try the distribution vendor. They will know which specific brands and models of video chips, audio chips, modems and so on have been tested and which ones they support. They

Preparing Today for Linux Tomorrow

also may have newly written or updated device drivers that were not ready in time to make the last distribution packaging.

To contact the distribution vendors, go to their technical support web pages, such as:

- Caldera Linux (www.caldera.com/support)
- Debian Linux
(www.valinux.com/services/support/?session_hash=7d7ccc8f28c0f42cddc98da037d7875a)
- Red Hat Linux (www.redhat.com/apps/support)
- SuSE Linux (www.suse.com/Support/index.html)
- TurboLinux (www.turbolinux.com/support)

In addition to vendor Web sites, they may also monitor Internet newsgroups devoted to their distributions, including:

- *alt.os.linux.caldera*
- *linux.debian.user*
- *redhat.rpm.general*

Even if the distribution vendors do not monitor a specific newsgroup, some of the users who read that newsgroup are likely to be knowledgeable about the issues you raise, and may have device drivers that they can post there for you.

Some, if not all, of the distribution vendors offer telephone support *for a fee*. Consult the documentation that came with the distribution (either printed or “readme” file) for the appropriate telephone number. Commercial application vendors may offer technical support for their products as well.

In addition to the distribution vendors, there are many Internet Web sites and newsgroups that offer assistance. The Web sites contain repositories of device drivers (and other useful Linux software) and may well have what you need. The newsgroups allow you to ask other users if they know of any sources for the particular drivers you need. In fact, because Linux is Open Source software, someone may even volunteer to write a customized driver for you by “tweaking” an existing driver to add support for unique features of your device. (Some people consider doing things like that to be “fun”!)

There are a number of sources for Linux device drivers, including the device vendors and web sites that specialize in Linux software. Some of these sites are:

- DLR Fresh Archive (www.go.dlr.de/fresh/linux/src)
- Freshmeat (freshmeat.net)
- Linux Apps (www.linuxapps.com)
- Linux Archives (home.linuxarchives.com/software.html)
- Linux MIDI + Sound Pages (www.xdt.com/ar/linux-snd/)
- Linux Software Encyclopedia
(stommel.tamu.edu/~baum/linuxlist/linuxlist/linuxlist.html)
- Linuxberg/Tucows Linux (linux.tucows.com)
- Linuxify! (www.linuxify.com)
- Slashdot (slashdot.org)

Preparing Today for Linux Tomorrow

The following Usenet newsgroups can provide a significant amount of user-to-user help:

- *alt.uu.comp.os.linux.questions* — General Linux questions and answers.
- *comp.os.linux.announce* — Announcements of new products, updates, bug fixes, etc.
- *comp.os.linux.development.system* — For programmers interested in developing Linux kernel and system software.
- *comp.os.linux.answers* — Linux FAQs, How-To's, READMEs and other documents that answer questions about Linux.
- *comp.os.linux.hardware* — Hardware questions, support, and how to get Linux working on certain hardware.
- *comp.os.linux.misc* — A catch-all for whatever does not fit in one of the other groups.
- *comp.os.linux.networking* — Networking and network hardware.
- *comp.os.linux.setup* — How to install and configure Linux and add-on software products.

There are also other Linux newsgroups provided for other topics that you might find helpful.

As a last resort, you can try the option vendors themselves (such as Iomega for Zip and Jaz drives, or Boca Research for one of its modems). “Standard” devices like keyboards and mice should work with Linux by default, but if they have any special features such as Internet launching buttons on the keyboard or a scroll wheel on the mouse, you probably will need special drivers to enable those features for Linux.

Partitioning

Ask ten different experienced UNIX/Linux users the best way to partition a large hard drive and you will probably get *at least* ten different suggestions. Here are some of the simplest:

1. For a new user or someone with simple needs, perhaps the best—and certainly easiest—solution is to create one large Primary partition (not counting the required swap file partition) for your operating system, programs and data files. However, some versions of some distributions may impose limits to partition sizes, or to the number of disk “cylinders” allocated to a partition (for example, 1024), so be sure to read the installation instructions that came with the distribution for this sort of information.
2. If a partition size or disk cylinder limitation restricts the size of the first Primary partition, or if you feel more comfortable putting the operating system in one partition and your applications and data on another, you can create two partitions. The second one, for the non-operating system stuff can be either Primary or Extended. (Disk cylinder limitations generally apply only to the first Primary partition, so you should be able to make the second one as large as you want.) The advantage of one large partition is simplicity. All your programs and data go in one place. You do not have to worry about where you put anything. The advantage of the two-partition method is that if you later wish to upgrade the operating system or want to format it and reinstall Linux from scratch for some reason, you do not have to worry about losing any data, because it is in a completely different partition.
3. If you think you will need to have both Linux and Windows (or another OS) on the same system, you can create two Primary partitions and use LILO (Linux LOader) as the means for selecting which operating system to run each time you boot up the system.
4. For a more sophisticated installation with multiple operating system partitions (such as for a programmer who needs to test software on Caldera, Debian, Red Hat, SuSE and TurboLinux, for example), several bootable Primary partitions can be set up with a boot

Preparing Today for Linux Tomorrow

manager utility, such as System Commander, acting as a traffic cop. (See *Multiboot Setups*, below.)

In each scenario above, remember to leave enough room for the swap partition. How much swap space will you need? See below.

Given a choice of creating swap space in a Primary or an Extended partition, which should you choose? Either will work. There is no technical advantage to doing it one way or the other, but there are a limited number of Primary partitions that can be created so if you plan on installing more than one operating system you might want to put the swap space in an Extended partition and save a Primary partition for other uses.

Note: If you have Microsoft Windows NT installed, it may not boot if you have more than one Primary partition installed on the bootable disk drive (what Windows sees as C:). In this case, Linux would *have* to be installed in an Extended partition, or in a Primary partition on a second hard disk drive. For more information, refer to the lilo (Linux LOader) documentation on a Linux system using the command: `man lilo` (“man” is short for manual). If you do not yet have a Linux system set up where you can read the lilo documentation there are a couple of other options available: 1) If you are installing from a CD set, look for a directory or separate CD (as in Red Hat 6.2) containing documentation, possibly in HTML format. 2) Visit Linux websites, such as linux.ctyme.com and www.linuxdoc.org, for online documentation.

After you have decided what type of partition to use, the next question is how large a swap space partition to create. Although it can be smaller, for best results the partition size should be *at least* equal to the amount of memory installed in the system—preferably *twice* the amount of physical memory. In other words, if you have 64MB of RAM installed use double that amount for your swap space partition (64MB * 2 = 128MB).

If you need more than 128MB of swap space, but you are using an older distribution that does *not* support a swap file larger than 128MB, there are ways around the 128MB swap file limitation: 1) Linux allows more than one swap partition on a hard drive, and 2) Linux allows swap files to be put in *data* partitions (i.e., in the same partitions as your programs and data files). For instructions on how to set up either or both of these, refer to the white paper entitled, *Linux Questions and Answers*, available from the same sources as this paper.

Now that you have decided how large your swap partition should be, the next question is how big to make your operating system and/or data partitions. (By “data partitions” I am referring to anything you as the user will be installing or creating other than the operating system itself—in other words, application software and data files.) The simplest method, of course, is to subtract the swap partition size from the total available disk space and make the remainder one big partition. But if you have decided to have separate operating system and data partitions or multiple operating system partitions and one or more data partitions, things get a lot muddier.

It is difficult to give a concrete answer to this because it depends on so many variables: How many partitions you will be creating, the minimum requirements for each operating system, and the total size of your hard drive(s). The best advice I can offer is to look at the minimum requirement for the operating systems first. If the installation instructions say to allow 500MB or 700MB for the OS, I would allow 1GB to allow for growth (in case later versions are larger). Better too much than too little, unless it leaves you with insufficient space for your data partition. A few hundred megabytes one way or the other are not much if you are dealing with a 10GB hard drive, for example. Once you have your operating system partition requirements out of the way, the rest is available for your program/data partition(s). Unless you have multiple operating system partitions and wish to have corresponding data partitions for testing purposes, I would recommend one single partition for all user files (applications and data) for the sake of simplicity.

Note: Just to make your life more difficult, even though hard drives are labeled something like “10GB” they really mean “10 billion bytes”, which is *not* the same as “10 gigabytes”. Ten gigabytes technically is $10 \times 1024 \times 1024 \times 1024$, or 10,737,418,240 bytes. However, a “10GB” drive may have only 10,550,587,164 bytes (9.826GB), or something else between 10,000,000,000 and 10,737,418,240 bytes. This was *not* done in an effort to cheat customers but to simplify the labeling of the drive by rounding the numbers. (Would you really prefer that the drive be called a 9.826GB drive?) [10GB of *memory*, on the other hand, is actually the full 10,737,418,240 bytes, in case you were wondering.] The reason the distinction between “billion bytes” and “gigabytes” is important is that when you are partitioning your hard drive, if you think you have 10 gigabytes available when you actually have 9.826GB, your last partition will be smaller than you expect by almost 200MB. This is something to keep in mind when doing your partition sizing.

Of course, if you are starting with a Windows system that you will be installing Linux onto, you already have one or more partitions created. If you decide to remove Windows (rather than installing Linux in addition), you have the choice of reformatting the existing partition(s) and using them as is, versus deleting them and creating new partitions from scratch—or do you? Remember that Linux requires a Primary *Linux* partition (type 83, to be precise) from which to boot. A “normal” DOS/Windows type Primary partition (types B, C, F or others) will not do the job. Therefore, you really are better off just replacing the Primary Windows partition with a Primary Linux partition. (For more on partition types, refer to the related white paper “*Linux Questions and Answers*” available from the same sources as this paper.)

Once you begin the Linux installation, at some point (depending on the installation program used by your distribution) you will probably have a choice of how to do the disk partitioning. The distribution may offer a graphical tool, such as Disk Druid, it may drop you into the text mode *fdisk* program, or it may let you choose between the two. Unless you have used the Linux *fdisk* program before and are comfortable with it, you would probably be best served by choosing the graphical tool, which does most of the work for you. (This rule applies even if you are familiar with *fdisk* from Windows or OS/2. Amazingly, the Linux version is even *less* user-friendly than the Windows and OS/2 versions.)

Installation Methods

The simplest approach involves installing Linux from a distribution CD. Most, if not all, Linux distributions will boot from the CD, however this requires that the CD-ROM drive itself be bootable. If it is, simply shut down the system, insert the CD-ROM and reboot. (This requires that you have previously set up the system BIOS to boot from the CD-ROM drive *before* the hard drive. If you do not know how to do this, refer to the user’s guide that came with the system.)

If your CD-ROM drive is *not* bootable, or if you are installing from an external PCMCIA (PC Card)-attached CD-ROM drive, or from files copied to an internal hard disk drive, or from a network drive (CD-ROM or hard disk drive), or an intranet site (via FTP or HTTP), you will need an installation diskette from which to boot. The files contained on the installation floppy will differ depending on the type of drive from which you will be installing. A boxed copy of Linux probably already includes one or more installation disks; if not, or if you downloaded the operating system or are installing from a network drive or intranet site, you will have to create the installation disk. (See *Appendix A*, *Appendix B*, or *Appendix C*, below, for instructions on creating the necessary bootable installation disk, depending on your installation method.)

Note: If you will be installing Linux on more than one system and need the ability to install from different types of drives, rather than making all the different types of boot disks, just create the one for PCMCIA-attached CD-ROM drive installation (*Appendix B*). *This one boot*

disk contains the files necessary not only for installing from a PCMCIA CD-ROM drive, but also from a local hard drive, network drive, FTP and HTTP.

Installation Types

The use to which you intend to put your system will determine which type, or class, of installation you should choose. If you are setting up a system to function as a server of some sort, then your choice is simple: select *Server*. If you will be using it as a desktop or laptop workstation, then you will probably want to select a *Workstation* option (you may be given choices of a *KDE Workstation*, a *Gnome Workstation* and possibly others—these determine which graphical user interface (GUI), also called a desktop environment, Linux will install). If you are an experienced, or “power”, user you might prefer the *Custom* installation option if it is offered.

Warning: A *Server* or *Workstation* installation will erase all information from all existing *Linux* partitions on all hard disk drives. If you do not wish this to happen, you must either install via the *Custom* installation option or use the *Upgrade* option. Notice that I said all existing *Linux* partitions. This does *not* mean your Windows partitions will be erased! Only those partitions that already contain the Linux operating system, applications or data files will be affected.

Of course, if you are a power user or a network administrator it is unlikely that you will be reading this paper, which is intended for Linux novices. Therefore, realistically your choices are limited to one of the workstation options. A workstation installation is almost entirely automated, with only a few questions for you to answer. *Which* desktop environment you choose is really a matter of taste. KDE and Gnome are both very popular (as are other freely available user interfaces). You might ask your Linux-using friends which GUI they prefer, to help you decide. Fortunately, should you later decide you would rather have something else, all is not lost. You can download and install a variety of desktop environments to replace whichever one you installed initially.

Note: One scenario in which you might be using the *Custom* install method would be where you are creating a dual-boot environment with Linux and Windows NT. You would have to configure LILO (the Linux boot manager program) to install itself in the *boot partition*, not in the default location: the Master Boot Record (MBR), which is where the Windows NT boot loader is installed. If you install LILO in the MBR, it will write over the Windows NT boot loader, rendering Windows NT non-bootable. After installing LILO, you would then configure the Windows NT system loader to boot LILO from the boot partition (otherwise Linux will not boot). Again, only an experienced Linux user should attempt to use the *Custom* installation option.

Multiboot Setups

If you decide that you will need to have several different operating systems or Linux distributions installed in different partitions on your system, you will require a way to select which operating system to load each time you turn on the system or reboot. The type of tool you will need for this is called a boot manager. This is essentially a program that loads before any other and provides you with a menu of choices of which operating system/partition to boot from. There are any number of such tools available, both free and commercial. The more sophisticated of these programs do more than just provide a list of partitions from which to choose. They even help you create the partitions and install the operating systems to those partitions.

If you are pursuing a multiboot configuration, it is essential that you decide on the boot manager program before installing Linux on your system, because the tool you choose will determine how you go about installing the operating systems. Be sure to read and follow the instructions that accompany the boot manager program.

Some of the free boot manager programs include: *KbootSelector*, *Klilo*, and *Smart Boot Manager*. (You can download these programs from various websites, such as those listed in the *Device Drivers* section, above.) If you will be installing Linux to share space with Windows NT specifically, a useful tool is BootPart (members.aol.com/gvollant/bootpart.htm).

Commercial boot manager programs include: *BootMagic*, *Power Boot*, and *System Commander*, available either directly from the vendor or from stores or web sites that sell Linux software. Some of these commercial web sites include:

- BlueSky Innovations (www.boot-manager.com) — PowerBoot vendor.
- Indelible Blue (www.indelibleblue.com) — Online store.
- Linux Mall (www.linuxmall.com) — Online store.
- PowerQuest (www.powerquest.com/bootmagic/index.html) — BootMagic vendor.
- The Linux Store (www.thelinuxstore.com) — Online store.
- V Communications (www.v-com.com/product/index.html) — System Commander vendor.

In Australia, Linux users can go to:

- Everything Linux (www.everythinglinux.com.au) — Online store.

Beginning the Installation

If you have gone through the preceding sections, obtained the information you need about your hardware and decided how you plan to partition your disk drive, you should be ready to begin the installation. As mentioned previously, there are a number of ways to install Linux: from CD; from an Internet, intranet, or network download; or by automated remote install. Not only that, no two Linux distributions will have exactly the same installation process. One distribution may differ from the next in terms of which installer program is used. In addition, one distribution may install a different default user interface and utilities than another would. Installation procedures may even differ between versions of the same distribution.

As I said in the Preface, above, I will not attempt to describe in this paper the procedures for installing every version of every Linux distribution. Instead, I attempted to highlight in this paper some important issues to consider prior to installing, as well as to point out “gotchas” to be wary of, and explain some choices that might be confusing. These aspects will be the same, regardless of whether you install from CD or from a downloaded package.

If you go the download route, be sure to read the instructions that came with the package for information on how to create the installation disk and other information unique to download installs. If you install a “shrink-wrapped” copy, be sure to read the printed installation manual that is included.

If you will be using an automated remote installation procedure, the network administrator or whoever else is handling the centralized installation will provide the necessary installation disk and/or other instructions. A remote installation will be preprogrammed using a script, to select all the necessary choices during installation, so you are essentially done at this point except to start the installation and then remove the installation disk at the end and reboot the system.

After Installation – What then?

Getting the operating system installed and running is a major step, but not the final one. Linux is not yet as user-friendly as Windows, at least as far as recognizing and supporting some

Preparing Today for Linux Tomorrow

hardware options. In the case of my IBM ThinkPad® 600E notebook computer, audio and modem drivers were not installed, leaving me with no sound/music or dial-up access to the Internet.

If this happens to you, it means that you will have to track down drivers that will be compatible with these devices. (See the *Device Drivers* section, above, for driver sources.) Once you get the modem support working, you will still have to configure the system for dial-up access to the Internet. Your ISP should be able to supply you with the correct settings. (With any luck, they will have prior experience with Linux.) For information on the steps required to set up your Linux system for dial-up Internet access, refer to the “*Linux Questions and Answers*” white paper, available from the same sources as this paper.

If you are using a corporate gateway server to get to the Internet (and/or if you need access to an internal LAN), you will have to contact your internal help desk or IT department for assistance. They can provide the information regarding proxy servers, ports, and so on that you will need.

Mark T. Chapman
IBM Server Group
September 29, 2000

Appendix A - Creating a Bootable Disk to Install from a Non-Bootable CD-ROM Drive or from a Hard Disk Drive

If your system BIOS does not support booting from a CD-ROM drive, you will need to create a bootable diskette from which to start the installation process.

Note: Details will differ from one Linux distribution to another and perhaps even from one version to the next, so these instructions may not apply to any installation but *Red Hat Linux 6.2*. Also, if you will boot from an external PCMCIA-attached CD-ROM drive (*Appendix B*) or a network drive (*Appendix C*), you will have a slightly different procedure.

To create the installation disk for installing from a non-bootable CD-ROM drive:

1. If you are creating the disk on a DOS, Windows or OS/2 system:
 - a. Insert a pre-formatted disk in the floppy drive. (This process will delete any existing files on the diskette, so be sure it does not contain anything you need.)
 - b. The boot image can be found in the **IMAGES** directory on the Red Hat 6.2 CD-ROM. Windows and OS/2 users can use the respective file manager graphical program to open the **DOSUTILS** directory and then double-click on the RAWRITE icon to execute it. (DOS users can simply type: `\DOSUTILS\RAWRITE` and press the **Enter** key.)
 - c. When prompted for the disk image file name, type: `\IMAGES\BOOT.IMG` and press **Enter**. When asked for the target, type: `A:` (or whichever drive letter is correct) and press **Enter**. This will create a bootable disk for a CD-ROM installation.
 - d. When the drive stops clicking and whirring, remove the diskette, close the DOS session window and label the diskette something like "Red Hat Linux 6.2 Installation Disk".
2. If you are creating the boot disk from a Linux or other UNIX system:
 - a. Insert a pre-formatted disk in the floppy drive. (This process will delete any existing files on the diskette, so be sure it does not contain anything you need.)
 - b. The boot image can be found in the **/images** directory on the Red Hat 6.2 CD-ROM. Make sure the CD-ROM drive is mounted (with the command: `mount /mnt/cdrom`), go to the **/images** directory (`cd /mnt/cdrom/images`) then use the command:

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

to copy the boot image file to the floppy drive. (`if=` specifies the input file name, `of=` represents the output file name, and `bs=` identifies the boot disk size; `/dev/fd0`, of course, is the floppy drive.)
 - c. When the drive stops clicking and whirring, remove the diskette, close the command session window and label the diskette something like "Red Hat Linux 6.2 Installation Disk".

Note: This boot disk can also be used to install from a local hard disk drive. (In other words, if you copy the entire CD-ROM drive contents to your hard drive or download the operating system from a web site, this same installation diskette will allow you to install from the hard drive or the CD-ROM drive.) The only difference in the procedure is in Step 2b. Simply go to the directory on the hard drive instead of on the CD-ROM drive using the command: `cd /images` from whichever hard drive the directories were copied to.

Appendix B - Creating a Bootable Disk to Install from an External PCMCIA-Attached CD-ROM Drive

If you will be installing from an external CD-ROM drive that is connected to your system by a PCMCIA (PC Card) adapter, use the following instructions.

Note: Details will differ from one Linux distribution to another and perhaps even from one version to the next, so these instructions may not apply to any installation but *Red Hat Linux 6.2*. Also, if you will install from an *internal/local* hard disk drive or non-bootable CD-ROM drive (*Appendix A*) or a network drive (*Appendix C*), you will have a slightly different procedure.

To create the installation disk for a PCMCIA-attached CD-ROM drive:

1. If you are creating the disk on a DOS, Windows or OS/2 system:
 - a. Insert a pre-formatted disk in the floppy drive. (This process will delete any existing files on the diskette, so be sure it does not contain anything you might need.)
 - b. The boot image can be found in the **IMAGES** directory on the Red Hat 6.2 CD-ROM. Windows and OS/2 users can use the respective file manager graphical program to open the **DOSUTILS** directory and then double-click on the RAWRITE icon to execute it. (DOS users can simply type: `\DOSUTILS\RAWRITE` and press the **Enter** key.)
 - c. When prompted for the disk image file name, type: `\IMAGES\PCMCIA.IMG` and press **Enter**. When asked for the target, type: `A:` (or whichever drive letter is correct) and press **Enter**. This will create a bootable disk for a CD-ROM installation.
 - d. When the drive stops clicking and whirring, remove the diskette, close the DOS session window and label the diskette something like "Red Hat Linux 6.2 Installation Disk".
2. If you are creating the boot disk from a Linux or other UNIX system:
 - a. Insert a pre-formatted disk in the floppy drive. (This process will delete any existing files on the diskette, so be sure it does not contain anything you need.)
 - b. The boot image can be found in the **/images** directory on the Red Hat 6.2 CD-ROM. Make sure the CD-ROM drive is mounted (with the command: `mount /mnt/cdrom`), go to the **/images** directory (`cd /mnt/cdrom/images`) then use the command:

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

to copy the boot image file to the floppy drive. (`if=` specifies the input file name, `of=` represents the output file name, and `bs=` identifies the boot disk size; `/dev/fd0`, of course, is the floppy drive.)
 - c. When the drive stops clicking and whirring, remove the diskette, close the command session window and label the diskette something like "Red Hat Linux 6.2 Installation Disk".

Note: This boot disk contains the files necessary not only for a PCMCIA-attached CD-ROM drive installation, but also for local hard drive, NFS, FTP and HTTP installations.

Appendix C - Creating a Bootable Disk to Install from a Network Drive (CD-ROM or Hard Disk)

If you will be installing from a remote CD-ROM drive or hard disk drive on a network, use the following instructions.

Note: Details will differ from one Linux distribution to another and perhaps even from one version to the next, so these instructions may not apply to any installation but *Red Hat Linux 6.2*. Also, if you will install from an *internal/local* hard disk drive or non-bootable CD-ROM drive (*Appendix A*) or a PCMCIA-attached external CD-ROM drive (*Appendix B*), you will have a slightly different procedure.

To create the installation disk for a network CD-ROM drive:

1. If you are creating the disk on a DOS, Windows or OS/2 system:
 - a. Insert a pre-formatted disk in the floppy drive. (This process will delete any existing files on the diskette, so be sure it does not contain anything you need.)
 - b. The boot image can be found in the **IMAGES** directory on the Red Hat 6.2 CD-ROM. Windows and OS/2 users can use the respective file manager graphical program to open the **DOSUTILS** directory and then double-click on the RAWRITE icon to execute it. (DOS users can simply type: `\DOSUTILS\RAWRITE` and press the **Enter** key.)
 - c. When prompted for the disk image file name, type: `\IMAGES\BOOTNET.IMG` and press **Enter**. When asked for the target, type: `A:` (or whichever drive letter is correct) and press **Enter**. This will create a bootable disk for a CD-ROM installation.
 - d. When the drive stops clicking and whirring, remove the diskette, close the DOS session window and label the diskette something like "Red Hat Linux 6.2 Installation Disk".
2. If you are creating the boot disk from a Linux or other UNIX system:
 - a. Insert a pre-formatted disk in the floppy drive. (This process will delete any existing files on the diskette, so be sure it does not contain anything you need.)
 - b. The boot image can be found in the **/images** directory on the Red Hat 6.2 CD-ROM. Make sure the CD-ROM drive is mounted (with the command: `mount /mnt/cdrom`), go to the **/images** directory (`cd /mnt/cdrom/images`) then use the command:

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

to copy the boot image file to the floppy drive. (`if=` specifies the input file name, `of=` represents the output file name, and `bs=` identifies the boot disk size; `/dev/fd0`, of course, is the floppy drive.)
 - c. When the drive stops clicking and whirring, remove the diskette, close the command session window and label the diskette something like "Red Hat Linux 6.2 Installation Disk".

Note: *This boot disk can be used to install from a network hard disk drive. (In other words, if someone copies the entire CD-ROM drive contents to a network hard drive, this same installation diskette will allow you to install from either a network hard drive or a network CD-ROM drive.) The only difference in the procedure is in Step 2b. Simply go to the directory on the hard drive instead of on the CD-ROM drive using the command: `cd /images` from whichever hard drive the directories were copied to. This boot disk contains the files necessary not only for a network drive installation, but also for FTP and HTTP installations.*



© Copyright IBM Corporation 2000

IBM Server Group
3039 Cornwallis Road
Dept. LO6A
Research Triangle Park, NC 27709

Produced in the United States of America
9-00
All rights reserved

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM reserves the right to change specifications or other product information without notice.

IBM, the IBM logo, and ThinkPad are trademarks of IBM Corporation in the United States and/or other countries.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

All other trademarks and registered trademarks are the property of their respective owners.

THIS PUBLICATION MAY INCLUDE TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES AND MAY BE CHANGED OR WITHDRAWN AT ANY TIME. THE CONTENT IS PROVIDED AS IS, WITHOUT WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. SOME JURISDICTIONS DO NOT ALLOW DISCLAIMER OF EXPRESS OR IMPLIED WARRANTIES IN CERTAIN TRANSACTIONS, THEREFORE THIS DISCLAIMER MAY NOT APPLY TO YOU.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.