

Feature: Linux free, easy & secure

Linux is free to modify and distribute, safer than other systems and easier than ever to use. Find out how to get it, install it and use the open-source OS

Barry Shilliday, Personal Computer World 20 Jul 2006

Back in 1992, a Finnish student, Linus Torvalds, posted to some internet forums that he was creating a new operating system that would run on PCs.

The operating system would be free and open, and based around Unix – the commercial operating system used all over the world by companies, governments and universities.

The project received overwhelming support, and nobody could have known then that, in just a few years, it would become one of the most powerful operating systems in existence.

It would find itself running not just on PCs, but on hardware ranging from small PDAs to massive corporate mainframes, and that it would help to push an open software philosophy into the mainstream.

A long road

Linux has come a long way from its Unix roots. It has been many years since even the installation was a job best left to the experts, and the simplest of tasks had to be done using complex, obtuse text-based commands.

Today, Linux is just as likely to be running the software in your fridge as it is to be running on your desktop PC: it could be processing satellite data to predict weather patterns or acting as an internet games server.

The modern Linux desktop is comparable to Windows XP and Apple's OSX. Anybody familiar with a computer can sit down in front of Linux and have no trouble at all, and they may find that many tasks are simpler to perform.

Free

Unlike its counterparts from Apple and Microsoft, Linux is a free, open-source operating system. This means anybody can read and modify the underlying code, and anybody can distribute it without licensing restrictions.

The developer can contribute and learn, and the user at home can run a full operating system for next to nothing, in many cases completely free of charge.

Some of the best versions of Linux can be downloaded from the internet and installed onto your PC. There are no evaluation periods, no demands for payment, and no forced registration screens.

Drawbacks

So with all these benefits, surely there must be some major drawbacks?

Not many: the pace of development can be astounding, since people all over the world will contribute to Linux and its applications. Six months in the Linux world can see new features that would not be seen in several years with Windows.

However, things are not always perfect. Some exotic hardware may lack drivers for Linux although, except for a few cases, this is rare.

Another drawback is that Linux may lack certain specific applications you need. This is especially true for large-scale commercial applications such as Adobe Photoshop, where there are no real equivalents for Linux.

Fortunately, there are ways around this, and Linux can happily co-exist with Windows (or even OSX) on the same hard drive.

Here, we look at how to get hold of Linux and install it, and address some of the common pitfalls people experience. With some free space on your hard drive, or by using a spare PC, you can get Linux up and running without risking a thing.

It's easy to give Linux a try, and here you will see how.

Getting started

Notwithstanding the ever-growing list of great reasons to switch to Linux, installing a new and unfamiliar operating system can seem daunting for many people.

It can be confusing to know where to start and what to do. The first question to ask is whether you can run Linux. One of Linux's unique advantages is the wide range of hardware on which it runs, but like all evolving operating systems, hardware requirements can begin to increase over time.

Another question is, 'Why do you want to run Linux?' This can greatly affect the kind of machine and the amount of disk space you will need.

There is no doubt Linux is a hugely versatile operating system. It is not limited to a small niche, nor does it need to come in entirely different versions for different tasks.

You can run a full desktop environment, or maybe a busy dedicated web server, a small database back end or a fileserver for your work or home network.

Hardware requirements

Most modern Linux distributions are aimed at being able to deliver all these options. The desktop environment tends to need the biggest chunk of resources, and to run the latest versions there is a general minimal hardware requirement.

Any PC with a Pentium II or AMD K6 or above should be powerful enough to get by with, though a CPU of 1GHz or more is a better bet.

For lower end machines such as these, it is essential to have at least 128MB of Ram, and the more the better. There are some options for less powerful machines, as you will see later.

For machines that are not required to run a graphical desktop environment, the resource requirement is significantly lower: it is even possible to run Linux on an Intel 386 CPU with 32MB of memory for less intensive tasks, such as network routing or running a firewall.

As with processor and memory requirements, the disk space needed also varies. Linux can be installed alone on a single drive. It can run across several hard drives in a Raid configuration, or it can share a drive with other operating systems, including all versions of Windows.

As a rule, for a modern version of Linux, together with a large set of applications, two or three gigabytes is enough for the operating system. Space for personal files needs to go

on top. With large-capacity hard drives being the norm, disk space is rarely a problem any more.

Dual booting Linux and Windows

Most people currently running Windows who want to give Linux a try will install Linux alongside the same hard drive in what is known as a dual-boot configuration. As it's such a common requirement, almost all Linux distributions provide methods to help with this during the installation procedure.

Installing Linux onto a hard drive by itself is as simple as with Windows – the installation programs will usually offer to automate it for you. For dual booting, some disk partitioning steps are needed.

Partitions are separate logical areas on a hard drive that behave like individual drives themselves. A drive always has at least one partition, even if this fills the entire disk.

For a dual-boot (or multiboot) configuration, each operating system must be installed onto its own partition, so if Windows takes up an entire drive, its partition must be reduced in size before a partition for Linux can be created.

There are various commercial Windows tools, such as Partition Magic from Symantec, that can resize existing Windows partitions. Most modern Linux installations can do the same, although some may not allow the resizing of NTFS partitions, the filesystem commonly found in Windows XP installations.

Gparted is an alternative open-source partition management tool that supports NTFS partitions, and is available as a 'live' bootable CD ISO.

Linux requires at least one partition for its root filesystem: this is roughly the equivalent of a C: drive in Windows.

It is highly recommended to have an additional partition for swap space (also known as virtual memory). Traditionally, the recommendation was to have twice as much swap space as memory on the computer, though 1GB is almost certainly sufficient.

Before resizing any partitions, it is vital to make a backup of your existing data files. While the resizing techniques are completely safe, any physical problems such as hardware or power failure could result in all the partition's data being lost.

Distributions

Unlike operating systems such as Windows and Apple's OSX, Linux itself is not a single product. In fact, Linux refers only to the kernel, the most central part of the operating system.

The rest of the operating system, including all the applications, utilities and even the text and graphical environments, are all separate from the kernel and need to be provided alongside.

Because most of the fundamental parts of the system come from the separate open-source GNU project (GNU is a recursive acronym for 'GNU is Not Unix'), Linux as an operating system is regularly referred to as GNU/Linux, though this has never caught on in the mainstream.

The Linux kernel, together with the applications, is bundled up by an organisation and referred to as a distribution, or 'distro'. The organisation, be it commercial or otherwise, can decide what software to include, which versions, and how to put it together.

Typically a distribution will include an installation program, a mechanism for managing software packages, and utilities for maintaining the system – administration tasks, updates and so on.

Openness

Because of the open nature of Linux, anybody can create a distribution, and many have: there are hundreds of Linux distributions around, but fortunately just a handful of major ones.

With so many distributions, many are tailored for a certain kind of person, or a specific type of usage. Red Hat Enterprise Linux, for example, aims clearly at the corporate market and is not intended for home use.

Ubuntu Linux and its derivatives are aimed primarily at the desktop at home or work. Suse Linux from Novell has versions that are aimed at enterprise, and others for smaller home or office environments.

Each distribution has benefits and drawbacks, and some are more suited to advanced users (such as Debian) or beginners (such as Linspire). The bigger names, such as Suse, Xandros, Mandrake and Ubuntu, tend to be the most versatile and least specialised, and are generally more suitable for those new to Linux.

Some distributions are available through retail sales outlets, where you can buy a boxed set with manuals and CDs, and others are offered via the internet as a free or paid-for download.

Some distributions are only available commercially, and cannot be distributed freely; some are available only as a free download and can be distributed without restrictions.

The difference between the retail and free distributions is mainly documentation and support. Boxed retail versions often come with good-quality manuals, and a number of weeks or months of telephone or email support.

Downloadable distributions usually have good-quality online documentation and community support through internet forums, but no formal support.

Ubuntu is unique among distributors in that it will send you multiple pressed CDs of the current release free of charge – a great idea for giving out copies to friends.

A number of companies will sell pressed CD copies of the free distributions for a small charge. If you do not have a fast broadband connection this can be a good alternative.

32bit or 64bit?

Linux can run on many different processors, not just the ones in PCs. As such, support for 64bit processors (as supported by the Athlon 64 and most desktop Intel Pentiums) has been around for a long time and has matured well.

Almost all major distributions are supplied as 32 and 64bit versions. If you have a 64bit CPU, you have a free choice: you can run the 32bit version without any performance penalty over a native 32bit processor, or you can choose to run the 64bit version.

There are advantages and disadvantages to both: 64bit offers better memory handling for systems with lots of memory, which is important for servers.

For applications that involve heavy calculations, including graphics-rendering and games, 64bit code can run significantly faster. As 64bit processors have become cheap and almost standard, more work is going into optimising software, and these advantages will continue to increase.

The main disadvantage is that an application cannot mix 32 and 64bit code, which means 64bit applications cannot run 32bit plug-ins. This affects closed-source binary plug-ins such as Macromedia Flash, and Windows multimedia codecs, which can otherwise be used in Linux video players.

Installation

Having decided on a distribution, you need the installation discs. If you buy a retail box, that's all there is to it. If you get a distribution by downloading or copying from a cover disc, it will be in a format known as an ISO image, with the .iso letter file extension.

The name comes from the fact that the file is an image of a data CD, which is in the international standard ISO9660 format. The image is a raw copy of the CD's data structure, and needs to be written out as such: simply copying the file to a CD in the normal way will not work.

Working with ISO images

There are various ways to write a CD image to a blank disc. Linux and OSX systems can do the job without any extra software. Windows requires third-party CD-writing software (such as Nero) or a free downloadable tool such as ISO Recorder.

If you have a copy of Linux running already, K3b will write CD images (select Burn CD Image from the Tools menu, and in Gnome just right-click on an ISO icon. In OSX, run the Disk Utility and select 'Open Disk Image' from the File menu, and click the burn icon.

These images automatically create a bootable CD. In recent years, all PCs have supported booting directly from a CD. If necessary, enter the PC's Bios and set it to boot from CD before the hard drive. Insert the CD as soon as you power up the computer.

For older computers there are other options, such as booting from a floppy or network, but these are far more complex.

Installation

The first automated stage during installation checks the hardware of your PC and loads the appropriate drivers, such as network and graphics cards. Some distributions offer a text-based installation procedure while most of the major ones are graphical.

The next step is partitioning. Most distributions provide a lot of assistance here, with guided or automated partitioning.

This stage can vary between distributions, so an alternative is to partition your hard drive beforehand using a specialised distribution such as Gparted. This is just 30MB and is designed for partitioning hard drives using a graphical interface.

Remember to back up all important files to a separate location (such as another drive or a CD/DVD) before making changes.

Most distributions offer a software selection, often with preconfigured choices. You may be asked whether to install KDE, Gnome or both. These are the main desktop environments in Linux.

They each have their own native applications and a different look and feel, but it is possible to run Gnome applications under KDE and vice versa. Some distributions designed to be simple to install will not offer any choices and automatically install a predefined selection of applications.

Bootloaders

Once the system is installed, the bootloader must be placed onto the hard disk. This small program sits between the Bios and your hard drive, and allows you to select which operating system to boot.

The simplest place to put this is on the disk's MBR (Master Boot Record), but Windows will wipe this without warning if you later re-install the operating system.

For that reason it is advisable to install Windows first if you are setting up a dual-boot system on a new hard drive. If you are not dual booting you don't need to worry about this step.

After a reboot, everything should be ready to go. There will almost certainly be software updates to install, and this is usually handled by user-friendly tools that pop up to inform you and to perform the update automatically.

Network and internet

One of Linux's long-standing strengths is its networking capabilities. With its robust and high-performance network stack and its compliance to standards, Linux runs on routers and firewalls in the most critical of environments, and hosts many of the world's busiest websites.

On a modest PC, this same software gives you access to the internet through a local network or modem, and applications allow you to share files, browse the web, stream audio and video, and more.

Network cards

Almost all Ethernet network interface cards (Nics) are well supported under Linux. These include the onboard network interfaces found on most modern motherboards.

The same applies to external dial-up serial modems. With the world moving to broadband, these are becoming less common, but are still important for many.

Support for USB network hardware can vary: for example, the dial-up modem on most mobile phones will work without any issues, but some external USB modems will not.

With USB broadband modems, there are also problems; often, drivers can be downloaded and installed to make them work, but it is not the easiest of tasks.

The internal dial-up modems found in many PCs can also present problems, and support for these can vary between distributions. If you rely on USB hardware or internal modems to connect to the internet, it may be worth checking before installing, or trying alternative, up-to-date distributions.

Broadband

For broadband access, the easiest choice by a long way is to use an external broadband

modem/router. These units connect to your PC via an Ethernet connection and, as such, do not need to be supported by the operating system itself.

What's more, the internet connection is made by the unit rather than the PC, so your PC does not have to remain switched on to maintain the connection: useful when more than one computer needs to access the internet.

As long as the hardware is recognised and supported, configuration is usually automatically performed by the Linux distribution.

If your router provides network addresses by DHCP, normally you will not have to do anything to get up and running. Most provide tools to configure network access, including setting a static IP address or configuring multiple interfaces.

Linux supports many network protocols and applications. You can access Windows network shares using the Samba package, which comes installed on most distributions.

In the KDE desktop, support for Samba is built into the file/web browser, allowing you to browse shares on Windows, OSX, or other Unix and Linux systems.

Support for traditional Unix file services is also provided, including NFS (network file system), most commonly used by commercial Unix systems.

On the application side, there are multiple web browsers, FTP clients, instant messaging and VoIP (Voice over IP) programs, Bit Torrent downloaders, and so forth.

Wireless networking

Wireless connectivity has until lately been rather a struggle because of the lack of drivers from manufacturers. Some have made specifications or drivers available, including Intel and Ralink, and these drivers have filtered through to distributions over the past couple of years.

A good way to see if your wireless network card or USB device is supported is to plug it into a live CD-based distribution. If your wireless network uses encryption, bear in mind that some distributions still only offer graphical configuration for Wep encryption and not WPA.

These can be configured to work with WPA, but may require manual changes. Support for Bluetooth file transfers is good, and both KDE and Gnome have utilities to manage the protocol.

Linux has enterprise-standard routing and firewall capabilities and, as such, it can be tricky to configure and set up without assistance. There are a number of tools for configuring these; some are supplied by the distribution vendor and others by a third party.

Some of the larger vendors, including Suse and Mandriva, provide central administration tools with easy-to-use controls for setting up networking. For other distributions, a utility such as the free Firestarter may come in handy.

Applications

A typical Linux distribution differs from Windows in that, while Windows is supplied as a bare operating system with a few token basic utilities, Linux comes with a wealth of applications, games and server software.

What you get depends on the distribution, as does how to install extra supported software. Some software is almost a given. Open Office is the omnipresent office suite. Compatible with Microsoft Office, it is being adopted by businesses and governments.

Recently its file format, an open equivalent of Word's doc format, became an international standard. For web browsing, Firefox is available and, as with Open Office, differs from the Windows version only in its appearance.

Other applications will vary depending on whether the desktop is Gnome or KDE. Kopete and Gaim, two instant-messaging applications, are compatible with just about all the IM protocols.

Amarok can manage all your music files and leaves iTunes far behind – the latest version can manage iPods and other digital players. Digikam and f-spot can organise images.

For server software, the list includes all the standard software on which the internet runs: the Apache web server, Perl, PHP and Python, the MySQL database, and major commercial applications such as Oracle and DB2.

Software installation used to be laborious and frustrating, but package management has developed considerably in recent years. The 'apt' system used by Debian, Ubuntu and some other distributions lets you install any supplied package (and there are hundreds fully supported) with just a couple of mouse clicks.

With such large collections of software provided by the distributions, it is usually unnecessary to go elsewhere for downloads.

Due to US software patents, distributions rarely provide applications or libraries affected by them, even in areas where software patents do not exist (such as Europe). This is especially true of multimedia codecs, including the omnipresent Mp3 and video codecs such as Mpeg2.

Although the distributions will not supply the software preinstalled, most give clear instructions on how to install the software yourself.

DVD playback also presents an issue, because video DVDs are encrypted so only licensed players can play them back, which is impossible with open-source video players.

Circumvention is available through the libdvdcss library; this is illegal in the US, but not in Europe. The American-led legal case against the code was rejected in Norway, and again on appeal, and the library is available to download free.

Note that libdvdcss does not circumvent any DVD copy protection; it merely allows you to play back region-coded DVDs. You should also be aware that you will be legally responsible for paying any necessary royalty and/or licensing fees required by the patent holders for the playback of mpeg-encoded material.

Hardware support for Linux

Linux supports almost all types of hardware and peripherals you will find on a typical PC. These include all the standard things, such as DVD writers, graphics, network and sound cards, floppy and Zip drives and USB storage devices, digital cameras, and lots more.

The majority of drivers are already contained in the Linux kernel (the core operating system) and will be automatically recognised and configured without any downloads or manual intervention.

Much of the driver support comes from third-party developers rather than manufacturers themselves, which regularly release drivers only for Windows.

This can mean that some older hardware is supported under Linux but not under Windows XP; unfortunately, it can also mean that some hardware is not supported under Linux at all.

This generally only occurs when a manufacturer refuses to release drivers for Linux or even the basic information needed by others to write them. The main problem area is with internal 'winmodems' (or soft modems) and USB broadband modems.

A broadband Ethernet-connected router is recommended for use with Linux, as these will work perfectly. Or, if you need to use a dial-up modem, ideally use an external serial one. Wireless network cards have also been a problem in the past, but the situation here is rapidly improving.

Specialised distributions

The major distributions are all geared around running Linux on a typical home or office PC or on a PC server, and come with a large selection of available applications and server software.

Of course, not everybody wants this: there are different reasons for running Linux. As a result there are specialised distributions that pop up all the time, and some become popular and endure as niche products.

The most common of these are the 'live CDs', which allow you to boot Linux without installing anything onto a hard drive; indeed, a hard drive is not even necessary to get them running.

Some of the larger ones use heavy compression on the CD to give quite a substantial system, others are small enough to fit onto a business card sized CD or a bootable USB pen drive.

Similar to the live CDs are the rescue distributions. These are designed to recover PCs from problems such as booting or file-system corruption. Since Linux can access Windows partitions, they can be useful in recovering data from an unbootable Windows PC, too.

Common installation problems

Problem: Windows takes up the whole disk and the installation program cannot resize the partition.

Solution: Repartition the disk using Gparted or another 'live' CD. Alternatively, repartition with Windows software. If Windows is not needed, delete the Windows partition altogether (but this will destroy all files on it).

Problem: The installation programs allow me to resize the partition, but I can reduce it only by a small amount.

Solution: You have too many files on your Windows partition. Delete some large files and try again.

Problem: Should I put the bootloader on the MBR or the root partition?

Solution: If the root partition is a primary partition (numbered 1 to 4) you can safely put the bootloader onto it. You must then mark that partition as active in order for the Bios to locate and start the bootloader.

Problem: I don't know which software to choose during installation.

Solution: Pick the standard desktop installation option. You can always add software later. Or use a distribution such as Ubuntu, which comes with carefully selected applications chosen for you.

Is Linux really safer than Windows ?

A common incentive to use Linux over Windows is the extra security. But does this really exist?

Linux and Windows are architecturally very different. The way programs are run and designed tends to make the environment generally quite safe in Linux.

The open nature of Linux also means that security problems are regularly recognised and fixed, often within hours, whereas in Windows, these problems can go unfixed for months.

Many of the Windows trojans and viruses that exist are designed to use features of the Windows operating system that do not exist in Linux and never will.

It is a common belief that Linux's security comes from it being far less widespread than Windows on the desktop, making it an unappealing target.

Arguably, it could be said that virus developers do not think this way: a security hole in Linux or its applications is just as much a target.

Most Linux systems run only the internet services you choose. With some, such as Ubuntu, that means as standard there are no services running at all. This makes the need for a firewall somewhat redundant; an analogy is someone with a house with no windows worrying about installing window locks.

Common services that do run are regularly scrutinised and updated for any security exploits, and when vulnerabilities are found they are often just theoretical. Ultimately, while it is never wise to be complacent, running Linux means leaving the worry of viruses and spyware behind.

Top 10 distributions

With so many distributions available, which one should you choose? Everybody has their own preferences as to what makes a good distribution, and a glance at the Linux-related websites will show these disagreements can get heated.

One way to gauge the popularity of a distribution is to look at the distrowatch.com website, which orders the top 100 distributions based on the average number of page hits per day.

The list below shows the standings at the time of writing. Ubuntu has become the most talked about distribution, and is a good choice whether you're new to Linux or an expert.

1. Ubuntu
2. Suse
3. Mandriva

4. Fedora
5. Mepis
6. Debian
7. Damn Small Linux
8. Knoppix
9. Gentoo
10. Slackware

Linux resources

Info on distributions

Pressed CD retailers/distributors

<http://linuxiso.co.uk>

<http://linuxexporium.co.uk>

<http://www.tuxdiscs.co.uk>

ISO Recorder (Windows)

Firestarter Firewall

Linux News

<http://linux.slashdot.org>

<http://lwn.net>