

A subjective view of the Linux market

THE LITTLE DIFFERENCES...

Rudiger Berlich investigates the origins of the various Linux distributions, and asks the question: "Are they really all the same?"

Let's face it – most Linux distributions are built from very similar components. Once you've installed them, you'll be able to configure them in such a way that a mere user will hardly be able to notice the differences. That's not to say there are no differences at all, only that they are smaller than one might think. Differences on a political and economic level are bigger. The decision about which Linux Distribution to use should be based as much on criteria like the local market position or the availability of services and support, as on the technical merits of a particular brand.

The Free Software Foundation

The humble beginnings of Linux have their roots, not in 1991 with Linus Torvalds, but in the 1980s with Richard Stallman and the Free Software Foundation (FSF). At least, if you perceive Linux to be more than just the core Operating System kernel, that is. The fact is that Linux wouldn't be what it is today were it not for the plethora of programs provided by the FSF. These include compilers, editors (yes, the famous EMACS editor) and many of the standard utilities available in a Unix system.

Above everything else Richard Stallman and the FSF have contributed, however, is the GNU General Public License (GPL), which is the building block of Linux's success and is the core reason why there is an Open Source

movement today. If Linux wasn't possible without the FSF, then it's only fair to say that the FSF and its goals wouldn't be as widely known and accepted today were it not for Linux.

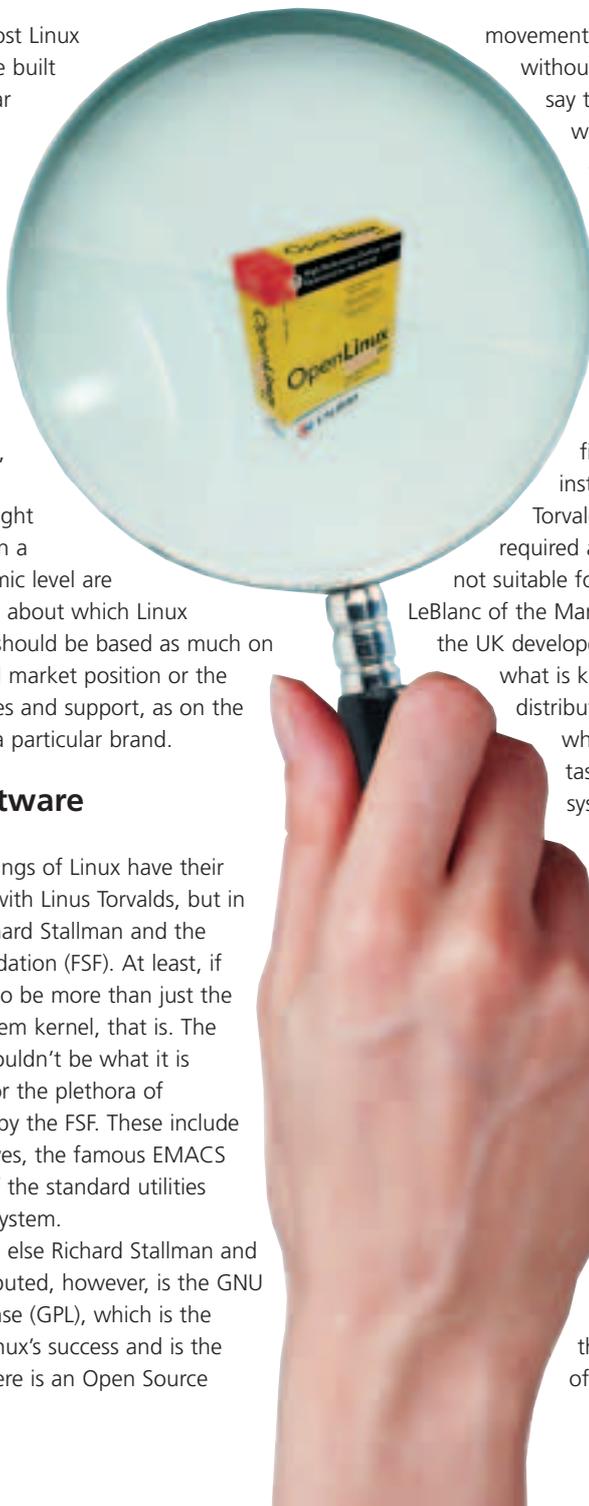
The early distributions

Linux quickly became widely accepted thanks to its free distribution under the GPL and the then emerging Internet. At first, the standard method of installing Linux was the Linus Torvald boot/root floppies – which required a lot of Unix expertise and was not suitable for a wider audience. Owen LeBlanc of the Manchester Computing Centre in the UK developed the first representative of what is known today as a Linux distribution: the MCC Interim Releases, which automated some of the tasks involved in installing a Linux system, such as copying software packages to your hard drive.

Soon after this, Peter McDonald brought the Softlanding Systems (SLS) distribution into existence.

This was followed by the Slackware Linux distribution by Patrick Volkerding, which was in large parts based on SLS. It's worth bearing in mind that this all happened in 1992, barely a year after Linux began.

Slackware was, and still is, semi-commercial – i.e. they fund their activities through the sales of Slackware on CD-ROM.



SuSE was founded in late 1992 in the Nuremberg area of Germany by four students of mathematics and computer science – Burchard Steinbild, Hubert Mantel, Thomas Fehr and Roland Dyroff. Other German distributions included LST (in Erlangen) and DLD (Stuttgart). SuSE was initially based on Slackware and incorporated various changes to make it more appropriate for the German market. Another Linux distribution – Jurix by Florian La Roche – was later incorporated into SuSE and the documentation was also translated into various languages, including English. SuSE is the oldest commercial distribution still available, and offers support for more hardware platforms than any of its commercial competitors.

Ian Murdock began the non-commercial Debian distribution in late 1993, in an attempt to provide free alternatives to the emerging commercial Linux distributions. It's arguably the most well known example of a free (as in both "free beer" and "free speech") Linux distribution, although in March 2000 Ian Murdock began work with Progeny on a commercial variant of Debian.

Red Hat and beyond

Bob Young and Mark Ewing founded Red Hat back in 1993. With the exception of Debian and its derived distributions, the Red Hat Package Manager is the standard amongst most Linux distributions. Red Hat was one of the first Linux companies to go public and it has subsequently bought various other companies. Among them is the German DLD distribution and Cygnus, the manufacturer of the embedded operating system, eCos. The Cygnus part of Red Hat's operations today contributes a significant proportion of its revenue stream. Red Hat initially made its Linux distribution available under the GPL, which enables other commercial vendors to build their own distributions based on Red Hat.

Utah based Caldera, Inc. was founded in 1994 by Ransom Love and Bryan Sparks. Caldera bought the

The market leader

If you want to make a statement about market leadership then you really need to define what you mean by the term market leader:

- 1 You could define market leadership in terms of the market's perception of this topic. For example, you could survey 1,000 people and see which Linux distribution gets mentioned most often. The problem with this method is that many people will tell you their perception of the market situation.
- 2 You could count the number of people that actually use a specific distribution. However this is also problematic as a single Linux distribution can be legally installed on any number of machines.
- 3 You could count the number of packages sold by a particular Linux distribution. This method suffers from the problem that few vendors may be willing to give you their exact sales figures (particularly if that vendor believes that they are not the market leader). Also, due to the ways that Linux can be distributed, not all of the installs of a distribution need to come about by someone purchasing the product.

Most vendors are the market leader in some way or another. Statements to this effect are printed all over press announcements and marketing material. What a specific vendor actually means by this varies greatly.

If you follow method number one for the definition of market leadership, then you'll probably come to the conclusion that Red Hat is the market leader, at least outside of Germany that is. If you do the same survey in Germany – SuSE's home turf – the situation is very different. The mood will change again in France, this time in favour of Mandrake.

There are independent surveys and online opinion polls that try to measure variables such as revenue, number of boxes sold or people using a specific Linux distribution. If you follow these (and thus use methods number two or three), the picture looks slightly different.

Within Europe, the UK used to be a Red Hat stronghold. However, a recent survey put Mandrake in front, with SuSE pipping Red Hat to second place. Within Germany, SuSE is consistently rated number one – a 1999 survey by Deutsche Bank even rated SuSE as the worldwide market leader. In the US, surveys frequently rate Red Hat as the market leader, though a recent survey instead placed Mandrake at the fore. Mandrake appears to be gaining ground in many areas worldwide. Turbo Linux has proven very strong in the Asian markets, where Red Hat has also had its successes. Caldera seems to be forging its own path by building a very business-oriented customer base, rather than positioning itself as a consumer-oriented company – as is evident in its acquisition of SCO.

It's difficult to draw any conclusion from all this other than that Red Hat, SuSE and Mandrake are arguably at the forefront of the Linux revolution. Where they exist in relation to one another is left to the judgement of the reader.

German LST distribution, which today forms the company's German arm. In 2001 Caldera also purchased the assets of the Santa Cruz Operation, and so now owns the rights for SCO Unix and Unixware. (The only part of the former SCO organisation that remains independent is the Tarantella division). This development symbolises one of the effects of the Linux movement – a consolidation phase in the whole Unix industry. It is particularly noteworthy that Caldera has decided to branch out from being solely a Linux company – it now provides a customised version of UnixWare (now



**There are
literally
hundreds of
Linux
distributions**

called OpenUNIX 8) to the former SCO community. The UnixWare kernel now includes a "Linux personality", which in simple terms means that it's capable of running Linux programs. In fact, one large component of OpenUNIX 8 is a complete Linux distribution. OpenUNIX could therefore be described as a Linux distribution with a SCO Unix kernel.

Turbo Linux was founded in 1992 under the name of Pacific Hitech. However, it was only later, as a Red Hat-based Linux distribution, that it became known throughout America and Europe. Turbo Linux recently attempted a merger with Linux Care, a company providing professional services mainly to the US market.

Mandrake was founded in 1998 and is another of the "big" European Linux distributions. It quickly gained a large base of followers due to technical reasons and its unique distribution model – Mandrake is sold through the MacMillan publishing house in a franchise style agreement. Mandrake has recently gone through a mini-IPO, which has furnished them with 4.3 million Euros in additional funds, to help them through the hard months ahead. Although the company



offers all the standard services, it's still regarded more as targeting the desktop user rather than businesses.

Although Caldera, Turbo Linux and Mandrake were all initially based on Red Hat, they can today be thought of as wholly independent (though compatibility is being maintained). Caldera in particular is heavily based on the former LST distribution, to the extent that much of the development is being undertaken in Erlangen.

The list goes on

After bringing its WordPerfect office suite to the Linux market, Corel later decided to test the waters with its own Debian-based Linux distribution. Having failed to generate much market awareness they now seem to restrict themselves to merely selling Linux software, such as WordPerfect and Corel Draw.

Connectiva is a Brazilian-based distribution for the Latin American market and is mainly available in Portuguese and Spanish. Although relatively new, the distribution seems to be gathering significant momentum.

There are literally hundreds of other Linux distributions, both commercial and non-commercial, but very few of these have gained a wider international acceptance.

Many other companies are active outside the realm of Linux distribution creation. Examples include Linux Care, who started to provide professional services to commercial entities in 1998. VA Linux started in 1993 and for a long time positioned itself as the worldwide market leader in Linux hardware solutions. Following the decline of the whole IT industry in early 2000, VA has undergone extensive restructuring to become a services and software engineering company – based upon its development platform, SourceForge. VA Linux went public at the end of 1999 and alongside Red Hat has undergone the most successful IPO of the Linux industry.

IBM has a special role to play in the Linux market, as this huge corporation has integrated Linux into its strategic planning and invests heavily in new technology. IBM's engagement has also marked a turning point in the adoption of Linux by large corporations. Other blue chip companies such as Compaq, Dell, Oracle and Fujitsu-Siemens also invest in and develop the Linux market.

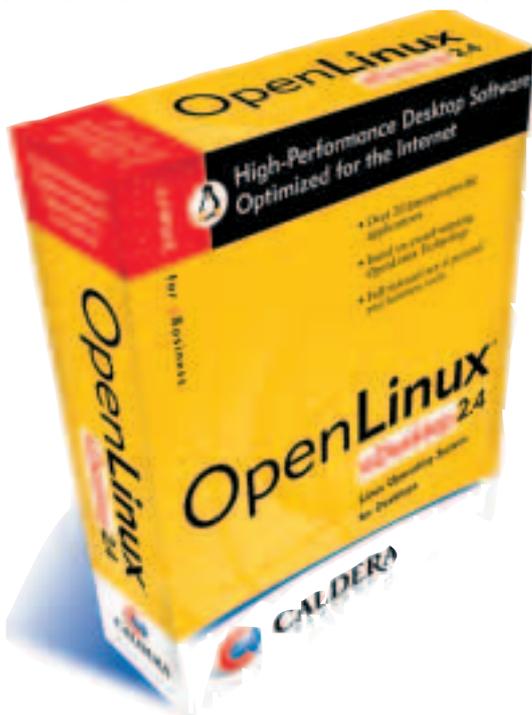
Similarities

Linux distributions all use very similar components. To start with, in order to be called a Linux distribution the Linux operating system kernel must be included. While this particular piece of software (now totalling over a million lines of source code) is being developed and maintained with the assistance of Linux distributors (they employ some of the core developers, for starters) the kernel itself is not bound to any specific company. The final control rests in the

Standardisation efforts

Whilst technical differences exist between Linux distributions, these are becoming increasingly less important due to the ongoing standardisation efforts pursued by the Linux community. Some of this standardisation happens silently, such as the adoption of the RPM package format by most commercial Linux vendors. Others are a consequence of the fact that key components tend to be identical across different Linux distributions. Other efforts are more proactive and are steered by a committee of companies and private participants. Amongst these are the Linux Standard Base (LSB), which aims to promote a set of standards to increase compatibility, and the Filesystem Hierarchy Standard, which is consolidating the filesystem layout.

On the training side, the specifications of the Linux Professional Institute could be regarded as a standard – particularly as most Linux vendors observe them. However, Red Hat has also brought out its own course material and test specifications, so the impact of the LPI is not as big as could be desired. Another standardisation that goes beyond the Linux market can be seen on the side of vendors of proprietary Unix systems.



hand of Linus Torvalds, the original author of Linux.

All Linux distributions will aim at providing the latest stable version of the kernel in order to remain competitive, particularly as driver support is usually provided by the kernel (with the exception of the graphics subsystem). All this means that, with certain restrictions, all recent Linux distributions compiled for the same architecture (Intel, for example) will be able to run the same programs. At the very least, installing and compiling a source package of another distribution on your preferred version will be possible, and in most cases you'll even be able to use the binary packages.

On Linux a distinction is being made between the graphical interface and the underlying technical infrastructure. The part that communicates directly with the graphics card – and also has certain networking responsibilities – is called the X Window system. The arguments raised above for the Linux kernel apply just as well to the X Window system, which is called Xfree86 under Linux. What this means is that commercial Linux companies will all tend to provide the newest version of the X Window system, and there also have the same hardware support for graphics cards.

Today there are two major graphical interfaces on Linux that sit on top of the X Window system: GNOME and KDE. KDE is the older development (GNOME started a year later) but the two systems are arguably now on the same technical level. Both GNOME and KDE provide the functionality of a Window Manager (they draw a frame around a window and let the user move and resize it with the mouse, for example) but they also offer all the usual features of an integrated desktop environment (drag

and drop, unified look and feel of all applications through the use of the same widget set, etc.). While some vendors, such as Red Hat, put more emphasis on GNOME, and others (SuSE and Mandrake) on KDE, every vendor will provide both environments.

So, all distributions are the same?

Although the similarities between the different distributions are strong, there are also some differences.

In the following example, the package zsh, compiled for Red Hat 7.1, was chosen randomly from the Red Hat ftp server (well, it was the last one on the list, so choosing it was easy although maybe not all that random...) and an installation was attempted on a default SuSE Linux 7.2 system. SuSE is based on the RPM format and also uses the newest versions of the core packages (kernel, libraries), as does Red Hat. The first attempt to install the package results in the following message:

```
# rpm -i zsh-3.0.8-8.i386.rpmerror: failed 2
dependencies: libtermcap.so.2 is needed by 2
zsh-3.0.8-8
```

libtermcap was subsequently installed from the SuSE CDs and a second attempt was made to install the software. The error message now encountered was:

```
# rpm -i zsh-3.0.8-8.i386.rpmfile /etc/zshrc 2
from install of zsh-3.0.8-8 conflicts with 2
file from package aaa_base-2001.5.15-2
```

It should be noted that zsh was not installed in the system before, and yet there was still a file `/etc/zshrc`. It is still possible to install the Red Hat zsh package, however you have to force RPM to overwrite the existing file or resolve this conflict in some other way, which can be dangerous. After the problem with `/etc/zshrc` was resolved, zsh ran without problems.

The filesystem layout used to be a big problem in the interoperation of Linux distributions, for example Red Hat places its start-up scripts under `/etc/rc.d`, while SuSE had for some time put its scripts under `/sbin/init.d`. Documentation will also find different homes. This is changing thanks to the efforts of the FHS (Filesystem Hierarchy Standard), but even when these differences still existed, vendors tried to make their distributions as compatible as possible. SuSE, for example, provided links in the appropriate places in order to make sure that Red Hat RPMs could be installed without problems.

The example above illustrates the following points:

- Different vendors have different ideas about what software should be installed by default. This means that you can't always rely on the fact that all software needed by a particular package by vendor

The filesystem layout used to be a big problem in the interoperation of linux distributions

Crystal ball? Eno Crystal Ball!

- A is already installed in the system of vendor B. In the case of SuSE, libtermcap could be post-installed without problems from the SuSE CDs, which will be the usual situation.
- Different vendors might place certain files in different packages, which means you may run into some conflicts with already existing files, if you install a “foreign” package.
 - The Red Hat zsh binary was dynamically linked with three libraries: libc, libtermcap and libnsl. All of these were available in the correct version on the SuSE system or could be installed without problems (in the case of libtermcap).

The last point is probably the most important. Even though there were some minor difficulties in installing the Red Hat package on the SuSE system, the necessary infrastructure was available or could be easily installed. Incidentally, similar problems to those described above will probably happen when you install a SuSE package on a Red Hat system.

When you have a look at the reviews of Linux distributions in computer magazines, you’ll see a strong focus on the installation rather than the everyday usage of the tested products. This is unfortunate, as it is not the installation alone that determines the quality and long-term success of a Linux distribution. On the other hand, this is the area where you’ll see the biggest differences between different distributions, although these differences are more on a cosmetic rather than a functional level. As the installation covers only one hour in the life cycle of a system of ideally a year or longer, differences on this level are largely unimportant.

Another difference in the various distributions is the configuration utilities provided. Of course, you can always configure your system manually, and none of the distributions restrict you to using their specific configuration utilities, though they certainly play a role as no mere desktop user can be expected to fathom the ins and outs of a Unix system.

Finally, let’s have a look at the licensing side of the picture. The majority of components in a Linux distribution are covered by the GNU General Public License, which basically allows everybody to do (almost) everything with it, as long as he/she doesn’t try to take this right away from anybody else.

Occasionally, in-house developments of Linux vendors are covered by licenses other than the GPL, which then prevents the copying of the CDs.

Generally, differences between different versions of a specific Linux distribution can be at least as big as the differences between the products of different vendors (on the same technical level), at least as far as programming libraries, compilers and kernel versions are concerned.

Looking into the crystal ball

Linus Torvalds once started an interview by looking into an imaginary crystal ball, which subsequently broke. He commented this with “Crystal Ball? Eno Crystal Ball”. ‘Eno’ refers to the way kernel error messages are labelled. He then had to replace it with an imaginary “I can’t believe it’s not crystal” ball made out of plastic. It gave less accurate results, but didn’t break.

Looking into the future isn’t easy, especially when taking into account the current market conditions. Some vendors will emerge from the current recession stronger than they were before,

as they absorb market shares from other vendors that didn’t make the grade.

It’s evident that the core business of Linux vendors is shifting away from pure operating system design, to providing services and software for the Linux operating system. This development has already been evident over the past couple of years. Standardised systems increase the available market segment for everybody. Only a marginal amount of money will be earned in the future through the sales of Linux distributions alone. Already the technical differences between the distributions are small, and people debating the question “which particular brand of Linux should be used” should put at least as much emphasis on local market conditions and the availability of services and support for that brand as on the increasingly unimportant technical differences.

The key message here is that differences between Linux distributions on the same technical level are smaller than differences between different versions from the same vendor. It should also be kept in mind that no one can be interested in a market that is dominated by just one Linux vendor. We’re not yet in a situation where Linux has gained the majority of the market share in the computer market. Linux has a stronghold in the server market, but it’s still weak on the desktop. With the advent of more handheld devices and the immense improvements taking place with regards to Linux’ graphical interface, a large increase in its share of the desktop market is not unlikely – we should keep in mind that Linux still only “owns” six per cent of this segment. It is in the interest of all of us to join forces, not to encourage fragmentation. ■



Info

(1)Rebel Code – Glyn Moody (Penguin Books)

(2)Just For Fun – Linus Torvalds and David Diamond (Texere Publishing)