The Current State of Open Source

An Open Source Primer

Author: C.Beecher M.INSTIS
SSU/ITSA
Date: July 2003

Introduction

This paper is designed to appraise the reader of the current state of the Linux and Open Source software market and the business opportunities potential cost savings that are available by using this technology. It also aims to dispel some of the myths that have been spread about the Operating System (OS).

Linux has grown from its humble roots as first a University project, to a small hobbyist Operating System to the stable, scalable enterprise product it is today.

Before one delves too deeply in to the Open Source Software (OSS) world, it is worth covering a few of the basic terms that one is likely to encounter. The first and the one most people associate with Open Source and Free Software is Linux. Linux is pronounced with a short *i*, as LiH-nucks and takes its name from its founder, Linus Torvalds who started Linux as part of his University project looking at task scheduling at the University of Helsinki in Finland. In essence, Linux is a UNIX clone, free of AT&T proprietary source code.
The first official release of the Operating System was in October 1991, so Linux has been around for a lot longer than Microsoft Windows NT.

Other names that will crop up in the Open Source world are GNU, a project of the free software foundation in Cambridge Massachusetts. GNU reportedly stands for 'Gnu is Not Unix' and this project has been responsible for producing many of the UNIX like software (commands) which runs on the Open Source Operating Systems. GNU are also responsible for the GNU General Public License under which a large amount of the Open Source software is released.

The other big name one will encounter in Open Source is BSD, both variants of the Operating System (OpenBSD and FreeBSD) and the BSD license, which is considered one of the most 'free' licenses there is. BSD (Berkeley Software Distribution – from the University of California, Berkeley) can be found most commonly in the TCP network stack on Microsoft products and in the new Apple Mac system ten OS which is heavily based around the BSD OS.

Is Linux Ready for the Enterprise?

There have been many thousands of articles and white papers produced questioning if and when Linux would be ready for the Enterprise, when would it be ready for the desktop etc.

The answer is now, it is not only ready for the enterprise, it is in and running parts of the enterprise and has been for many years.
Linux is taking over the areas that were traditionally held by the high end UNIX Operating System and now in the S.M.E (Small Medium Enterprise) arena is challenging Microsoft on Total Cost of Ownership (TCO) availability and compliance to Open Standards.

Enterprise

In the early days Linux entered the Server room via the back-door. Many system administrators installed Linux Servers to replace the less mission critical servers providing Web and File and Print services. One of the most common Linux Distributions at this time was Slackware, still available today, it is considered one of the more secure and stable Linux distributions and has a very UNIX feel to it. As Linux gained ground and support, so it was 'ported' or made available under differing hardware platforms including Sun SPARC and ARM processors.

With Intel platforms taking a bigger share of the processing in many environments, Linux has really come to the fore. Originally written for the old 386 architecture Intel systems, Linux has grown with the advances in IA-32 technology.

Not only have Intel actively assisted the Linux Kernel team in making sure Linux takes advantage of the new processors, but so have their competitors like AMD.

In the last few years most if not all of the major I.T players have thrown their weight behind Linux and Open Source with companies like IBM, HP and Oracle.

Only recently Novell has stated that Netware 6.5 will be more Open Source friendly with a GroupWise client being available for Linux and the support of Apache Web services and the MY-SQL Relational Database Services (RDBMS)

Sun Microsystems has just announced that it will be partnering with RedHat (one of the major commercial Linux vendors) to supply Linux on Sun's Intel X86 systems.

Oracle the database vendor has supported Linux for several years now and that support is increasing with its “Unbreakable Linux” campaign with Oracle version 9i.
IBM has thrown its not inconsiderable weight behind Linux, not only its services and consulting arm but also products with DB2, Tivoli and Lotus products all being available on Linux.

Both SUN and IBM have and are still providing internal resources whose only role is to further Open Source software development, with Sun contributing to the Desktop (X-Windows) and IBM having over two hundred and fifty programmers world wide working on the Linux Kernel and 70 Linux based projects. If this seems a strange move for a company with IBM’s back-ground then this from IBM’s vice-president may explain things.
He claims that in the long term they (IBM) are getting a cheaper Operating System than if they built it themselves, plus they could not build a Linux class OS with just 250 programmers.

Linux has seen large take-up in the Film, Oil and Gas industries and now in the Financial sector where Linux can provide a cost effective alternative to the traditional UNIX.

Linux is no longer considered a ‘risk’, with the GNU software tools and the adaptation of the UNIX practice of modular development being common to both Linux and UNIX it is possible for UNIX administrators to support either Operating System.
It is worth noting that Linux, running on clusters of cheaper hardware has a higher performance than the proprietary UNIX (which are normally restricted to specific hardware platforms).

A recent white paper produced by Bloor Research for IBM on the “enterprise readiness of Linux” concludes that Linux is scalable (up to 6 way SMP with Kernel 2.4) on Intel, and extremely well in Grid and Cluster type environments as well as vertically on IBM Mainframe systems.
It also notes that Linux has a community of over 400,000 developer’s world wide all constantly improving the product.
They (Bloor) also state that from there findings over 90% of Linux is installed on an Intel platform and that the Linux platform has proved to be reliable.

S.M.E.

The majority of peoples only experience of Linux is from PC magazine coverage and CD-ROM based demonstrations.
This often provides an unfair bias, where Linux is put up against the competing Microsoft Operating Systems. How can a product which is essentially 'free' compete with a product costing several hundreds of pounds? Herein lies the problem for not only the 'public' but also Microsoft.

One of the issues with Linux is that it is 'free'. The cost of a desktop version of the OS from vendors such as RedHat or SuSE (the two bigger commercial Linux vendors) is about fifty pounds. Other distributions are 'free' often provided on Magazine cover disks or available for download over the Internet (as are RedHat and SuSE's offerings).
The common perception is that if something is free, then "it can't be any good, or else someone would charge for it". To understand this, first one must understand the Open Source philosophy.
Linux is not public domain software (and neither is the GNU tools and utilities) but released under the GPL (the GNU General Public License). See appendix A for an overview of GNU licenses.

The Linux Model did not follow the organised development, source-code controlled, statistical quality control that a commercial Operating System does. It was historically a 'hackers' (hacker in the real usage of the term, not a malicious person, bent on attacking computer systems) OS. This in part held back its adoption within the corporate market as it was perceived as un-stable and un-supported. This argument no longer holds water, as discussed in the Enterprise section of this paper, with the likes of Sun and IBM committing large amount of resources to Linux as well as many smaller companies and consulting firms.

Another area where most individuals find problems is with Linux hardware compatibility. This is still an issue today but is getting better with more and more hardware vendors now either supplying Linux drivers for their products or allowing Linux developers access to the product so they can write their own. In the past this was not always the case, and many of the drivers were the results of programmers 'reverse engineering' the Microsoft drivers that came with products. As Linux gains increasing acceptance in the market place and thus gains a larger share of the Intel base desktop, more and more manufactures will provide Linux drivers along-side the Microsoft ones.

Software and interoperability is the last hurdle for Linux on the desktop. In the enterprise this has not been an issue, it is almost transparent to the user for example, from the end users point of view it is immaterial if the database is a Microsoft SQL Server or a Linux MY-SQL one they are connected to. This is the

same with file and printer sharing, as long as the performance is similar (and often better on Linux) then they do not care.
Desktop office automation is one area that Linux has lagged behind in. This is not really the fault of Linux or the developers, but more to do with the dominance of Microsoft and the Office package. Microsoft Office (in its various incarnations) is the de-facto office automation package and has been for years. Due to its proprietary vendor controlled format, for another package to compete it must offer MS compatibility or it will be a non-starter, and historically this has been the problem. Sun has been Open Sources biggest ally in conquering this barrier with its commercial Star Office and Open Source version Open Office. These products provide very good Microsoft Office compatibility in many areas but fail in two. Open Office (and Star Office) can not provide the Visual Basic scripting macros MS does, nor can they offer the Access database plug-in flexibility many organisations require. They do offer their own version of Access in the commercial Star Office package though.
Thus, one of the bigger issues faced by organisations migrating to Linux/Open Source is “what to do with” the Office databases and macros they have built up over the years.
One option is to ‘bite the bullet’ and pay for a programmer(s) to covert their MS macros to Star Office, or to use a Windows Emulator (WINE) and stick with MS Office. The latter is at present the more popular option.

One of the reasons organisations are moving away from Proprietary systems is Vendor Lock-in. They are ‘held’ in a continuous up-grade cycle of support, licensing and standards (as with Microsoft Office, where every new version differs in format from the previous). Microsoft (and others) has always used the argument that by paying for the support and product you are safe in the knowledge that issues will be resolved. This is true to a point and no large organisation would wish to rely on Internet news groups for software support as in the ‘old’ days of Linux. With the recent change in Microsoft Licensing many more organisations have been re-evaluating Open Source as a way of cutting TCO and vendor lock-in. Microsoft NT4 Operating System is a good example, it is now (just about) to become “end of life”, that means no more support, service patches or help. This leaves one with two options, upgrade, which will incur all the associated costs, not least a hardware upgrade (which is almost mandatory with MS) or stay on an un-supported platform. To be fair to Microsoft, RedHat has now also adopted this upgrade product life cycle, but only provides support for a year. There is however a big difference, with Linux the core of the OS, the kernel can be updated with new releases as can any or all of the distribution, (due to Linux’s modular approach) something one can not do with Microsoft products.

It is interesting to note that Microsoft has moved away from proprietary offering in the networking area (WINS for example) and ‘embraced’ more open standards like DNS and LDAP (in their Active Directory).

Government and Open Source

In the past, it would appear that Government (UK) in particular was wary of deploying Open Source, or more specifically Linux systems, why?
Possibly for the reasons given in the preceding sections perceived lack of support and compatibility issues.

Linux is gaining wide-spread acceptance in many Governments around the world, in India, China, Korea and especially in Germany, Linux is taking over not only on the 'back-end servers' but on the desktop too.
In Germany for example, the local government in Munich, (Germany's third largest city) has just announced it will migrate over 14,000 desktops to Linux. This move, they claim, will save not only money but provide greater commercial and technical flexibility for the council. This is interesting because the decision was reached in the face of tremendous pressure being brought to bare by Microsoft in the form of significant licensing reduction.

In the UK, the Department of Works and Pensions (DWP) has just implemented a Linux-based system to handle e-procurement, which will run over the GSI. To quote the chief executive of the Office of Government Commerce, Hugh Barrett, "the decision to use Open Source was based on its proven reliability, portability and lower licensing costs."
The O.G.C. is apparently also evaluating the Linux solution as part of its plan for an interoperable cross-government e-procurement system.

UK local city council in Nottingham has just implemented a Linux based email system. They have installed SuSE OpenExchange Server (which can migrate MS Exchange 5.5 data in to OpenExchange) to service over 7,500 users at a significantly lesser cost then by deploying a Microsoft solution.

Certain Companies have been accused of spreading FUD (Fear Uncertainty and Doubt) about Linux and OSS, being critical of the GPL license, portraying Open Source Software (OSS) as less secure and suggesting it might have an impact on national security.
One of the big advantages of Open Source is the peer review of code that constantly occurs. A proprietary Operating System or application will only represent the skill and direction of those employed by the vendor to undertake the project. The subsequent outcome is therefore dependant on the skills and ability of those employees. With O.S.S the code which stays is there on merit and no one individual or company can control it. Secondly, no one company but many (IBM, SUN etc) all contribute (and even their code is not always used) so the

best minds across the globe are helping to produce a feature rich, stable and secure OS.
With full access to the source code, organisations and Governments are free to modify and build-in their own security models (as the Banks have done). Although this would on the face of it, appear to be building a bespoke OS or application, the organisation would hold the source code and thus in years to come, if the product needed to be updated, one of any number of approved companies could update or modify the source code. This is something that can not be done with proprietary vendor products.
A big advantage of this is that items that are perceived as a security risk can be removed from the Kernel. One of the problems with the newer Microsoft OS is its ability to communicate with every bit of hardware on a PC. This includes floppy and CD drives USB devices and Wireless (Wi-Fi) and built-in Bluetooth devices as well as Infrared. With Linux, it is a simple matter to remove support for these devices from the Kernel, thus creating a more secure desktop system.

On the subject of cost, if staff are familiar with one OS and not Linux, setting up a server with Linux will probably cost more as savings on the OS price are taken up by time on the learning curve. But once installed and setup, Linux systems require low maintenance, and subsequent systems will require less time and therefore save money. But beyond saving a few pounds, the licensing issue goes further.
Many system administrators prefer the thin server approach, where services are spread across multiple low cost machines rather than centered on big central boxes, with the load being split by service rather than users. This approach is more easily scalable and limits downtime. It also makes it easier to upgrade and maintain individual system elements. Linux not only makes the thin server model financially viable, the high efficiency of the OS means that desktop machines no longer considered viable to run the latest desktop OS may be recycled as non critical servers and routers.

Skills in the Linux arena are more costly than for a Microsoft certified person but, the Linux person will generally be better multi-skilled and be able to operate in both a Linux and UNIX environment and often have Microsoft knowledge as well.

Conclusions

With the advent of Windows Dot Net Server 2003 product, there is probably little difference between Linux and Microsoft from a Security point of view.
Linux has a reasonable security record but one must take care if a big swing to Linux occurs as with Microsoft at present being the dominant OS attackers will turn their full attention to Linux as the system to attack. This is already starting to happen in the Web Server area.
It should be noted that no system can be totally secure – as always, it is the users and developers who hold the key.

On the desktop Linux in a Web enabled or Net centric world is a viable option today. With more and more applications requiring the user to interface with nothing more than a web browser, Linux can provide a cost effective desktop solution, if all that is required is email and web browsing plus an industry standard compatible word processor, then Linux with Open Office is one solution.

In the back-office, Linux can also fair very well. It does appear to run faster than other Operating Systems, often on lower powered hardware and with such applications as SAP, DB2, WebSphere, Lotus, Oracle 9, Apache, MY-SQL, Sybase Tivoli, SuSE OpenExchange and Netware (V7) to name a few all available natively on Linux.

For Linux to continue to grow users need to actively deploy it distributors must continue to enhance their installation, training and support and Independent Software Vendors (ISV) must aggressively extend applications to support it.
By migrating parts of the IT infrastructure to Linux, organisations will enhance their flexibility, gain the ability to port applications across platforms and increase leverage of hardware and software.
It would appear that more and more organisations are migrating there low end systems (UNIX and NT) to Linux and also consolidate their applications on Linux virtual machines, based on Linux compatible Mainframes.

Linux is not the answer to all I.T problems and for example, in the back office, Microsoft NT/2000 is a far better Application Server than Linux, but when it comes to file and print Serving, Linux wins. Microsoft SQL Server (2000) is also a lot faster than Oracle and has advantages over DB2.

Microsoft Server 2003 has moved the game on; it will be interesting to see if Linux can catch up in this Mid-range space.

As a foot note, it is interesting to observe that IBM appear happy to allow AIX skilled employees to cross-pollinate Linux with their skills and experience, maybe Microsoft should follow suit, how much longer will it be before we see Microsoft Office for Linux I wonder?

Appendix A (Licensing)

The GNU General Public License (GPL) is intended to guarantee ones freedom to share and change free software and ensure that it is free for all users. Items which are covered under the GPL are the Linux Kernel, the KDE and GNOME (X-Windows) desktops and most of the GNU utilities.

What the GPL does do, is to stop a commercial company from taking GPL code and then making parts of it proprietary by adding to it. It also ensures that patches and upgrades are fed back in to the Open Source community.

BSD also has a license, which is considered to be the freest of the lot. This is because there are no source redistribution clauses in it. This is why parts of the Microsoft Operating system and the Apple Mac System X Operating system use BSD code. They are under no obligation to feed the changes back in to the user community, thus the code is freely modifiable but the community will befit less than from the GPL system.

Microsoft also has a shared source license, as an individual you can download and modify any Microsoft (MSSL) code and distribute it! The catch, you can only do this on a non-commercial basis and there is at present very little MSSL covered code.

SUN also have a couple of free type source code license agreements. The first is called Sun Community Source License which allows code modification but no distribution, and the second is the Sun Public License which is free.