

Published in Queue Magazine, May 2004, under the title “Desktop Linux, Where Art Thou?”. Copyright 2004, Queue Magazine. All rights reserved. Contact Queue Magazine for republication or other licensing requests.

Linux on the desktop: Catching up, meeting new challenges, moving ahead

Linux on the desktop has come a long way – and it’s been a roller coaster ride. At the height of the dot-com boom, around the time of Red Hat’s Initial Public Offering, people expected Linux to take off on the desktop in short. A few years later, after the stock market crash and the failure of a couple of high profile Linux companies, pundits were quick to proclaim the stillborn death of Linux on the desktop. Then a funny thing happened: the GNOME and KDE desktop environments kept getting better, Mozilla and OpenOffice.org reached their 1.0 milestones and a new batch of Linux desktop companies sprang up (Lindows, Xandros, Lycoris). Rumors of the death of Linux on the desktop had been greatly exaggerated. Over the last year, there has been a steady stream of new developments and announcements that have added to the desktop Linux drumbeat. Perhaps most importantly, big IT companies, including Novell and Sun Microsystems, have placed bets on Linux on the desktop.

So will 2004 be “the year for Linux on the desktop”? This article takes stock of where the Linux desktop stands and looks at remaining obstacles to broader Linux adoption on the desktop, how Linux can meet the new challenges created by Microsoft’s plans for Longhorn, and opportunities to get ahead of the curve around collaboration-centric computing.

Linux' Long March to the Desktop

Sidebar

The Long Linux March to the Desktop

1991: Linux introduced

1991: Unix X11 GUI framework ported to Linux

1998-1999: KDE 1.0 & GNOME 1.0 – GUI environments for Linux

2001: Ximian Evolution 1.0 – Outlook look-alike that can be turned into a Microsoft Exchange client

2002: StarOffice 6.0 / OpenOffice 1.0 – real alternatives to Microsoft Office

2002: Mozilla 1.0 – a competitive browser

2002: KDE 3.0 & GNOME 2.0 – Linux GUIs become more polished

2003: Munich switches to Linux – The first big desktop success stories emerge

2003: Sun and Novell join the party – IT leaders throw their weight behind Linux on the desktop.

The epic communist Long March in China lasted about a year, from 1934-1935, during which time Mao and his band of communist fighters traveled almost 3000 miles through mountainous terrain and established their power structure. Here's the bad part: of the 200,000 fighters who started the March, only 30,000 or so made it to the end, and it would take them another 15 years to gain power.

We're now six years into the March of Linux on the desktop and there have been some setbacks along the way, the biggest one being the fading away of Netscape. While Netscape was not a Linux desktop company, Microsoft's control over 98% of the browser market continues to be one of the biggest threats, not just to the future of Linux on the desktop, but to the future of the Net. But Linux on the desktop has proven to be surprisingly resilient. The failure of companies like Eazel proved to be mere speed bumps. Out of the ashes of Netscape rose a lizard whose roar is becoming louder by the day, with Mozilla's new Firefox browser getting rave reviews before even reaching its 1.0 release.

Taking Stock

sidebar:

Desktop Linux: Making The Grade?

The product	
The desktop	B
Desktop developer platform	C-
Hardware support:	
- Computer	B-
- Peripheral devices	D
Applications	C+
Windows connectivity	A
The eco-system	
Linux distributions	A-
Desktop platforms	A-
Developers	B
ISVs	D
OEMs	C
Distribution channels	D
Consortia	A
News and support	C

Linux on the desktop is good enough for certain classes of users, but much work remains to be done. While there's a thriving community of open source developers, there is still a dearth of commercial third party applications and distribution channels remain inaccessible.

Linux on the desktop is now good enough for certain classes of users, but much work remains to be done. Two large, dynamic open source projects, KDE and GNOME, offer complete desktop environments that include a file manager, desktop administration tools, and a broad range of applications range from simple games to complete office productivity suites. Both desktops are reasonably mature, offering a fairly polished and stable desktop. They may not be as glitzy as Apple's Panther OS, but they compare reasonably well to Microsoft Windows. The main problem with the Linux desktop is that there are two desktop projects, making it hard for users and developers to pick a desktop to focus on and diluting the energies of the open source developers.

Through their tools and development libraries, Microsoft and especially Apple provide a more or less unified look and feel and a dominant development environment. Not so for Linux, where KDE, GNOME, Mozilla, OpenOffice and the WINE project all offer their own development environments, resulting in an often inconsistent end user environment.

But a consistent end user environment may be less critical than a strong selection of end user applications. After all, many of the leading Windows applications don't adhere to Microsoft's user interface guidelines. In most of the core applications areas, there is at least one strong Linux offering. Multiple strong web browser options are available, OpenOffice is a complete, mature office suite, and there are numerous strong Internet applications ranging from instant messaging to most of the popular multimedia plug-ins.

But significant gaps remain. While there are thousands of end user applications, most are unpolished and incomplete. There is no end-user Linux software industry to speak of. There's no strong personal finance application for Linux. If you need pixel-for-pixel compatibility with Powerpoint, or need to tap into an older Exchange server, Linux may not be right for you. You won't find educational games for Linux lining the shelves of your local computer superstore.

Similarly, the Linux eco-systems has its strengths and weaknesses. There's a wide range of Linux distributions to suit anyone's needs and a large, thriving community of open source developers each scratching their own itch. But while hundreds of white-box vendors are shipping Linux workstations, it's still all but impossible to purchase a Linux desktop PC in mainstream retail channels.

Three major obstacles mean that broad consumer adoption for Linux is years away, if it will happen at all:

- The lack of a critical mass of consumer applications: the types of applications that fill the aisles at your local Best Buys store – everything ranging from personal finance software to games to polished photo album software. While the open source community feels some of these gaps and the WINE project offers a partial solution, it's unlikely that independent software vendors will focus on the Linux market until there is established demand, creating a chicken-or-egg trap. The gradual shift of the software industry to web-based models may offer the best hope to overcoming this challenge, as long as these rich web applications aren't inextricably tied to the Windows operating system;
- The lack of major OEM support and consumer distribution channel access: while major OEMs such as Dell and software distributors stand among the most to gain if the Microsoft monopoly is broken down, they also have powerful financial incentives, in the form of joint marketing agreements with Microsoft, that make it almost impossible for these low-margin, commodity vendors to ship systems with alternative operating systems. Still, with the operating system and other Microsoft applications now making up the most expensive component in cheaper PCs, manufacturer have a major stake in advancing a lower-cost alternative;
- Resistance to change. Computer users invest tremendous energy learning to operate their computer and configuring it to meet their needs. They build up piles of data that make them reluctant to upgrade to a new computer, much less to an unfamiliar operating system. Of course, the successive waves of computer viruses, worms and Trojan horses have driven many users to despair and ready to explore safer alternatives.

While the consumer and knowledge worker markets, especially in the US, remain elusive, there are many markets where there are significant opportunities for Linux on the desktop, including “transactional workers”, price-sensitive markets outside of the US, and cheap, limited functionality PCs. These markets have the potential to provide desktop Linux with double-digit market share, making a serious dent in the current desktop mono-culture.

What's needed to seize the opportunities?

Strengthened desktop foundations

There are not just two Linux desktop platforms, there are at least five: KDE, GNOME, OpenOffice, Mozilla and WINE. Each of these big, successful open source projects includes a vibrant community of developers, is an umbrella to a significant number of projects, and is built on a different toolkit. And, of course, OpenOffice and Mozilla are cross-platform projects whose users mostly run the Windows operating system.

Much of the technical challenge ahead can be summed up, in the words of Red Hat's Havoc Pennington, as work that falls through the cracks between the existing open source projects or work that spans multiple projects. The latter problems are particularly intractable in the open source community, which is prone to endless heated technical arguments. Since code rules in the open source world and there is no one dictator or a single business imperative to help resolve issues, integration issues often remain unresolved. As a result, perhaps the single most critical challenge to the success of Linux of the desktop may be the need to create stronger communication links between the various projects, and, ultimately, a willingness for the various projects to make painful concessions in order to improve interoperability and jointly deliver a more competitive desktop environment.

Let's take a look at some of the specific foundational tasks that need to be tackled to ensure that the Linux desktop become more competitive. Thanks to Havoc Pennington for his input in this section.

Improved desktop "plumbing". The X Windows system is the right foundation for the Linux desktop, but innovation is needed to increase performance and usability, offer better multimedia support in the face of closed hardware specifications and patents, and generally design and implement new X extensions that allow the Linux desktop to offer the kind of graphical pizzazz that OS X and Longhorn provide.

Robustness and performance. You shouldn't have to know how to use the command line to maintain your Linux desktop system. Errors should be transparently handled or conveyed back to the UI with useful help for the user. Performance of the Linux desktop needs to be consistently tracked as it is for the server, monitoring metrics such as start-up time, opening a menu, or launching key applications.

Better support for peripheral devices. Linux needs a hardware abstraction library that makes peripheral devices available to applications and end-users. While the Linux kernel offers good support for USB and other peripheral device standards, that doesn't mean that you can just plug your digital camera into a Linux-powered desktop and edit pictures in your favorite image editing application. Linux doesn't have a universal hardware abstraction layer that makes these peripheral devices available to all applications. Last July, Havoc Pennington wrote a paper on this topic, entitled Making Hardware Just Work (<http://www.ometer.com/hardware.html>). Since then, a community volunteer has taken the lead on this and launched the HAL (Hardware Abstraction Layer) project (<http://pdx.freedesktop.org/Software/hal>).

A universal MIME system. There is no single repository of associations between file types and helper applications on Linux. Each of the platforms have their own MIME system, which means that things quickly get confusing if you use Evolution (a GNOME-based PIM) in KDE and are trying to figure out how to define, let's say, which web browser should handle links in email messages. The Shared

MIME Database project (<http://www.freedesktop.org/Standards/shared-mime-info-spec>) is making headway in tackling this challenge, with the GNOME project starting to adopt its specification approach.

A universal configuration/manageability framework. Each of the desktop platforms have their own configuration and management framework. Worse, lower pieces of the OS stack (such as the graphing layer aka the X windowing system, or the operating system proper) also have configuration and management frameworks. If you're lucky, you'll be able to change your monitor configuration using the GNOME control panel. More likely, you'll have to log out and log back in. It may also involve a trip to the command line. This problem is complicated by the fact that projects such as the GNOME desktop run not just on Linux but also on FreeBSD, various flavors of UNIX and even Windows.

A common end-user document filing system. The various desktop projects and Linux distributions don't always share the same filing system conventions for deciding where to place files that go on the desktop or in a user's home directory. The GNOME project has adopted KDE's convention for the user documents directory, but other projects and various Linux desktop distributions have their own conventions, so it's quite likely that your web browser will save documents to the desktop by default while your word processor may save them to your home directory and a third application may save files to a "My Documents" directory

A universal virtual file system so that the same URIs are known to all applications. Currently, the list of available file systems may differ from one application to another. On Linux, just because you are able to mount WebDAV or FTP volumes on the desktop doesn't mean that you'll be able to access those volumes from your word processor. Piecemeal progress is being made in this area, with, for example, Mozilla recently including support for the GNOME project's GNOME-VFS virtual file system.

A process spanning multiple desktop projects for interaction and UI-design. The GNOME project has invested the most effort into creating Human Interface Guidelines and encouraging the maintainers of GNOME applications to follow these guidelines. But what about user interface consistency between GNOME, KDE, OpenOffice and Mozilla, for instance?

Tackling these and other remaining challenges to offering a competitive, integrated Linux desktop may be best taken on by a new effort that brings together people from the different desktop-related projects. The good news is that the open source community is moving in exactly this direction and two projects have recently broadened their charter to be able to more effectively tackle these challenges: the X Consortium has morphed into the X.org Foundation and the Freedesktop.org project has broadened its charter and now hosts a growing number of desktop foundation-related projects, including several mentioned above.

Fighting Microsoft's lock-in strategies

Microsoft continues to systematically pursue strategies that lock customers in to their solutions, using the office document format, the Windows Media format, its Exchange servers and Active Directory as ways to lock users into the Windows desktop. Let's take a look at one of the key areas of lock-in.

The OpenOffice document filters are already quite good and continue to get better (witness the improvements made in OpenOffice 1.1 and StarOffice 7). Developers at OpenOffice, KWord and other open source productivity applications are already collaborating informally to solve remaining problems. Improved, collaboratively maintained documentation by the open source community of the Microsoft Office DTDs, similar to the OpenOffice project's documentation of the Excel file format (<http://sc.openoffice.org/excelfileformat.pdf>), could further improve this situation.

The availability of a rich, open and XML-based file format standard for productivity applications will also be a significant boost for Linux desktop adoption. An effort is under way at the OASIS standards consortium to develop a formal standard, largely based on the OpenOffice file format. Development, adoption, and implementation of the standard by OpenOffice and other Linux productivity applications will offer librarians, data archivists, policy makers and other interested parties an open alternative that they can adopt as a standard for document exchange in government, educational institutions and elsewhere.

More commercial end-user applications for Linux

As indicated above, the lack of a critical mass of commercial end-user applications is a major barrier to broader consumer adoption of Linux. This chicken and egg problem is unlikely to be resolved within the next few years, but the WINE project, an open source implementation of the Windows APIs on top of X and Unix, may provide a partial answer. Thanks to the efforts of the WINE community and CodeWeavers, many Windows applications, including Microsoft Office, Internet Explorer, Windows Media Player, Quicktime and Adobe Photoshop already run unmodified under Linux. TransGaming, the other company that builds on the WINE libraries, boasts that it allows 250 of the world's hottest games, including EverQuest, Battlefield 1942 and SimCity, to run unmodified under Linux. Additionally, the WINE libraries offer an easy way for ISVs to port applications to Linux. For instance, a few years ago MusicMatch was ported to Linux using WINE. In the not so distant future, further improvements to WINE may allow end-user Linux distributions to credibly claim that Linux runs Windows applications.

User-friendly documentation and support options

Most Linux documentation, and the most popular mailing lists and support forums are hard to parse for non-technical users, making it a daunting task to solve the inevitable problems that arise.

Absence of a standard Linux

The multitude of Linux distributions and hardware platforms supported by Linux makes it hard for ISVs to package applications for Linux. Efforts such as the Linux Standards Base and easy-to-use software installers such as Linux's Click'N'Run Warehouse help to address this challenge.

Brian Proffitt, Managing Editor of Linux Today, gives a great example of the end user cost of this lack of a standard definition of Linux, at <http://linuxtoday.com/mailprint.php3?action=pv<sn=2004-01-30-023-26-OP-DT-SW>. The GNOME desktop boasts beautifully anti-aliased fonts, so Proffitt was horrified when he installed the latest version of the Mozilla browser on his Red Hat system and ended up with jagged fonts. It turns out that anti-aliased fonts in GNOME require two libraries that don't

come pre-installed on all the major Linux distribution, which is why XFT is not pre-compiled in every release of Mozilla.

Better support for developers

Defining a standard Linux specification would not just benefit users, it would also be a boon for developers who wish to build Linux desktop applications. But more is needed. The Linux desktop projects are still best suited for programmers who code in C or C++ and who think of Emacs as an IDE, making the Linux desktop hard to approach for the millions of Java and Visual Basic programmers who program in Visual Studio. Offering primarily C and C++ strongly limits the pool of developers. Open source implementations of Java and C# and python and the other open source “p” languages offer great promise but it continues to be a challenge for the Linux community to offer developers the easy to use tools that Microsoft provides.

Longhorn: the next challenge

The technical agenda described above, ambitious as it may be, is really a catch-up agenda that will allow the Linux desktop to make headway in areas where the current state of the desktop is *good enough* and decision makers focus mostly on cost issues, such as transactional workers, cost-sensitive markets outside of the US and consumers looking for limited functionality Internet PCs. But Microsoft is planning to raise the bar considerably with its upcoming Longhorn OS.

Perhaps most significantly for the prospects of Linux on the desktop, Longhorn will come with an XML mark-up language called XAML that allows developers to build rich Internet applications that are seamlessly integrated into the desktop, potentially relegating the web browser to a legacy system and inextricably integrating the open APIs of the web with the Longhorn APIs. If XAML takes off, users may be able to interact with Net content in rich new ways, but could become more locked in to their Windows desktops than ever before.

But XAML didn't pop up out of nowhere. The language bears some striking similarities to Mozilla's XUL (XML User interface Language). XUL is a UI mark-up language that can be used to build Gecko-based applications (Gecko is the name of the Mozilla rendering engine). The Mozilla 1.x suite, Firefox and Thunderbird are all XUL applications. More intriguingly, there are a number of standalone XUL applications, such as the Mozilla Amazon Browser (<http://mab.mozdev.org>) which runs as a client application that lets you locally manipulate Amazon data without the latencies of web page updating – just like the XAML Amazon demo shown to great fanfare at a recent Microsoft developer conference. The difference, of course, is that XUL is available now and will run on any platform.

The Mozilla people believe that the combination of XUL, the Mozilla Gecko rendering engine, SVG support and web services capabilities are the keys to delivering rich Internet applications to the desktop, especially because XUL is relatively easy to learn for the millions of HTML coders. Mozilla developers are building prototype applications that combine the benefits of web-based application deployment and management with the responsiveness and usability of desktop applications.

In sum, XUL has the potential to be a cross-platform, open answer to XAML. If the Mozilla developers work closely with Linux desktop efforts such as GTK (the toolkit used by the GNOME project) to combine the cross platform strengths of XUL with native widget support on Linux, Linux could offer rich Internet applications that are every bit as compelling as XAML applications and offer the benefit of running on any platform.

Moving ahead

The open source community has a plan for catching up on the desktop and may be able to best Microsoft in its plans for bringing rich Internet content to the desktop. But how can the open source community deliver on the original promise of the PC as “bikes for the mind” (Steve Jobs) or the promise of the Internet as a creative enabler? Can the open source community articulate its own vision for tapping into the creative power enabled by combining a rich client environment with the network effects of the Net? Are there areas of opportunity that the dominant provider isn’t pursuing because they don’t fit into their strategy for maximizing shareholder return on investment?

Let’s take a look at three projects that aim to tackle this larger challenge.

Chandler

Today’s PIMs force users to adapt to technology. Information is locked into silos such as e-mail, calendar entries, address book items and to-dos.

Frustrated by this inability of existing PIMs to adapt to his way of organizing his work, Mitch Kapor, the creator of Lotus 123, turned to the open source development process, invested millions of his own money and created the Open Source Applications Foundation, where a team of two dozen developers are creating Chandler, a next generation PIM. Chandler, slated to be released sometime next year, breaks down the walls between the various PIM applications, and the wall between what’s mine and what’s yours. All data is organized into one central repository (an idea Microsoft is picking up on in their plans for WinFX in Longhorn) and users can organize their information however they want. The calendar becomes just another way to view appointment data – you could just as easily view your data according to which project it relates to. The other major emphasis in Chandler is collaboration – the ability to share information using flexible policies. User and presence management and notifications are handled not through central servers but using the peer-to-peer Jabber instant messaging platform.

“Both the PC and the Internet revolution started off as decentralizing forces, putting power with the end user,” notes Mitch Kapor, “but were then embraced by enterprises, which gained control over them. Software investment in general, including on Linux, is largely driven by the needs of large enterprises, who are mainly interested in managing centralized computing resources, enterprise-wide application development, limiting the ability of the user to modify their desktop environment. The end user computing experience matters less to them.” But the command and control approach taken by many large enterprises, and around which the business models of the large software vendors are built, is at odds with the increasing interconnection enabled by the Internet. People are working together on projects in ad-hoc teams that transcend organizational boundaries. Microsoft Exchange is designed to work within the confines of a large enterprise, but what if you’re a small business

owner who works with a fluid group of people dispersed around the world on a variety of different projects? Even within large organizations, bottom-up and ad hoc collaboration within the enterprise and across the corporate firewall is becoming increasingly important. That's where Chandler comes in: "A big part of Chandler's mission is to help small groups of people collaborate across organizational boundaries," says Mitch Kapur.

The dashboard

Nat Friedman, founder of Ximian (recently acquired by Novell) also believes that the future of the Linux desktop lies in collaboration, and he's putting his money *and* his fingers where his mouth is. Recently, he launched a Desktop Integration Bounty Hunt (<http://www.gnome.org/bounties/>) for the GNOME project, offering bounties for making more than 50 improvements to GNOME aimed at improving the experience of collaboration in the desktop environment.

Nat also coordinates the GNOME dashboard, a project that aims to bring relevant information to users, rather than requiring people to dig around for information. With the dashboard, if you're chatting with a friend using instant messaging, and your friend asks you about your recent trip to Rio, the dashboard will helpfully figure out that you were assembling your pictures from that trip last night, dig up your shots, and allow you to drag and drop them into your chat software.

"While you read email, browse the web, write a document, or talk to your friends on IM, the dashboard does its best to proactively find objects that are relevant to your current activity, and to display them in a friendly way, saving you from digging around through your stuff like a disorganized filing clerk," as Nat puts it.

The dashboard is not just a single piece of software you can install. It requires some modifications to front-end applications such as the web browser, PIM software, chat software etc, to allow them to send "clue packets" to the dashboard. So volunteers in the GNOME project are now retrofitting GNOME applications to support the dashboard concept. For Longhorn, Microsoft is working on a similar concept, which it calls "implicit query."

Mozilla

Written off by many just a few years ago, the newly established Mozilla Foundation is not just thinking about the future of XUL, the organization is also getting ready to roll out its much-anticipated next generation web browser and email client (Firefox and Thunderbird), slated to reach 1.0 this Summer.

Looking further ahead, leaders in the Mozilla project are exploring ways to give users easier access to all the information they care about, whether it's a web page they have visited before, a bookmark, an email message or an attachment they have filed away, using a personal data server that offers Google-like search of both local and remote files. The Mozilla Foundation is also interested in tapping into friend-of-a-friend small world networks to enable sharing of bookmarks and web page annotations, for example. An overview of some of the ideas the Mozilla Foundation is exploring can be found at <http://www.mozilla.org/events/dev-day-feb-2004/mozilla-futures/>.

So what do these three projects have in common? The projects share a focus on integrating the Internet into a rich client environment in innovative ways. They are breaking information out of its traditional silos. They are working to build applications that adapt to the user's needs, rather than the other way around. They bring information closer to the user and peer to peer communication and collaboration. In sum, these projects tap into the power of the desktop while also understanding that a network application framework, which opens up new possibilities for users and offers important cost advantages for developing, deploying and maintaining software, is a key to success. Most importantly, they move beyond trying to copy what Microsoft did a few years ago. Each of these projects is breaking new ground and is seeing some of its ideas copied by Microsoft, putting to shame the claim by Microsoft's OS chief Jim Allchin that open source is a threat to innovation.

What's next? These ideas feeding off each other, and linking up with web innovations such as wikis, RSS and weblogs, and social networks. Stay tuned. The open source community is bringing innovation back to the desktop.

About the author:

Bart Decrem works for the Open Source Applications Foundation and currently spends most of his time helping the Mozilla Foundation with marketing and business affairs. His bio is at http://www.decrem.com/bart/bio_03.html.